



Proposal for the 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)

Joeran Beel¹(✉) and Lars Kotthoff²

¹ School of Computer Science and Statistics, Artificial Intelligence Discipline,
ADAPT Centre, Trinity College Dublin, Dublin, Ireland

beelj@tcd.ie

² Department of Computer Science, Meta Algorithmics,
Learning and Large-Scale Empirical Testing Lab, University of Wyoming,
Laramie, USA

larsko@uwyo.edu

Abstract. The algorithm selection problem describes the challenge of identifying the best algorithm for a given problem space. 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 little attention to the algorithm selection problem, although the problem is highly relevant in information retrieval. This workshop will bring together researchers from the fields of algorithm selection and meta-learning as well as information retrieval. We aim 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. In particular, we will explore to what extent existing solutions to the algorithm selection problem from other domains can be applied in information retrieval, and also how techniques from IR can be used for automated algorithm selection and meta-learning.

1 Motivation

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

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number 13/RC/2106.

© Springer Nature Switzerland AG 2019

L. Azzopardi et al. (Eds.): ECIR 2019, LNCS 11438, pp. 383–388, 2019.

https://doi.org/10.1007/978-3-030-15719-7_53

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. Particularly in artificial intelligence, impressive performance achievements have been enabled by algorithm selection systems. A prominent example is the 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 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 [8]. 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 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]. 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 Objectives, Outcomes and Vision for the Workshop

We propose the 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)¹. We aim at achieving the following goals:

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

The expected outcome is a workshop proceedings book, which we aim to publish at <http://www.ceur-ws.org/>. Our vision is to establish a regular workshop at ECIR or

¹ <http://amir-workshop.org/>.

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 Topical Outline

We will explore (a) how existing solutions for algorithm selection and meta-learning can be applied to identify the best algorithm for a given information retrieval problem and (b) how information retrieval techniques may be applied to solve the algorithm selection problem in IR and other domains.

More precisely, topics relevant for the workshop are

- Automated Machine Learning
- Algorithm Selection
- Algorithm Configuration
- Meta-Learning
- Hyper-Parameter Optimization
- Evaluation Methods and Metrics
- Benchmarking
- Meta Heuristics
- Learning to Learn
- Recommender Systems for Algorithms
- Algorithm Selection as User Modeling Task
- Search Engines for Algorithms
- Neural Network Search

4 Planned Format and Structure

We envision a half-day workshop with the following submission types.

- Research Papers, Position Papers, Case Studies (6 or 12 pages, LNCS format) with 5–15 min presentations at the Workshop
- Posters, Demonstrations, Nectar², Datasets (4 pages, LNCS format) with a 1–2 min teaser presentation, and a poster session.

The tentative schedule is as follows.

8:30	Poster setup	10:00	Coffee Break and Posters
9:00	Welcome and	10:30	Paper Presentations
9:10	Keynote Talk	12:00	Lunch, Outlook, Discussion
9:45	Poster Pitches		

² A Nectar track allows well known researchers to present a summary of their recent work relating to the workshop topic. Compare <http://www.ecmlpkdd2018.org/nectar-track/>.

Depending on the number of submissions, part of the paper session may be replaced by a panel discussion with experts from IR and machine learning to discuss the algorithm selection problem in the context of information retrieval in depth.

5 Expected Audience and Attendees

Workshops on the algorithm selection problem in the field of machine learning attracted significant attention. For instance, the *Meta-Learning and Algorithm Selection Workshop* at ECMLPKDD in 2015 resulted in 15 publications [9]. The *Workshop on Meta-Learning* (MetaLearn 2017) at NIPPS resulted in 29 publications, including posters [10]. Given that the algorithm selection problem is less known in the IR community, we expect around 5 publications (short and full paper) plus a few posters.

We are confident to receive a significant number of manuscripts and expect a high number of attendees as the algorithm selection problem is relevant for everyone in information retrieval, particularly for everyone who wants to deploy a real-world system. It is an easy-to-understand problem that every researcher has faced himself/herself; and has attracted a lot of attention in other communities already.

6 Organizers

6.1 Joeran Beel

Joeran Beel³ 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 center that cooperates with industry partners including Google, Deutsche Bank, and Huawei. Joeran is further a Visiting Professor at the National Institute of Informatics in Tokyo. His research focuses on information retrieval, recommender systems, user modeling 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 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.

6.2 Lars Kotthoff

Lars Kotthoff⁴ 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 \$400 K in external funding to date. Lars is also the PI for the Artificially Intelligent Manufacturing center (AIM) at the University of Wyoming. He co-organized multiple workshops on meta-learning and automatic machine learning

³ <https://www.scss.tcd.ie/joeran.beel/>.

⁴ <http://www.cs.uwyo.edu/~larsko/>.

(e.g. [9]) 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 ~800 citations and his research has been supported by funding agencies and industry in various countries.

References

1. Beel, J., Gipp, B., Langer, S., Breiterger, C.: Research paper recommender systems: a literature survey. *Int. J. Digit. Libr.* 305–338 (2016)
2. Rice, J.R.: The algorithm selection problem (1975)
3. Xu, L., Hutter, F., Hoos, H.H., Leyton-Brown, K.: SATzilla: portfolio-based algorithm selection for SAT. *J. Artif. Intell. Res.* **32**, 565–606 (2008)
4. Kotthoff, L., Thornton, C., Hoos, H.H., Hutter, F., Leyton-Brown, K.: Auto-WEKA 2.0: automatic model selection and hyperparameter optimization in WEKA. *J. Mach. Learn. Res.* **18**, 826–830 (2017)
5. Lindauer, M., van Rijn, J.N., Kotthoff, L.: The algorithm selection competition series 2015–17. arXiv preprint [arXiv:1805.01214](https://arxiv.org/abs/1805.01214) (2018)
6. Tu, W.-W.: The 3rd AutoML challenge: AutoML for lifelong machine learning. In: *NIPS 2018 Challenge* (2018)
7. Brazdil, P.: Metalearning and algorithm selection. In: *21st European Conference on Artificial Intelligence (ECAI)* (2014)
8. Hoos, H.H., Neumann, F., Trautmann, H.: Automated algorithm selection and configuration. Report from Dagstuhl Seminar 16412, vol. 6 (2016)
9. Vanschoren, J., Brazdil, P., Giraud-Carrier, C., Kotthoff, L.: Meta-learning and algorithm selection workshop at ECMLPKDD. In: *CEUR Workshop Proceedings* (2015)
10. Calandra, R., Hutter, F., Larochelle, H., Levine, S.: Workshop on meta-learning (MetaLearn 2017) @NIPS (2017). <http://metalearning.ml>
11. Miiikkulainen, R., Le, Q., Stanley, K., Fernando, C.: Metalearning symposium @NIPS (2017). <http://metalearning-symposium.ml>
12. Vanschoren, J., Van Rijn, J.N., Bischl, B., Torgo, L.: OpenML: networked science in machine learning. *ACM SIGKDD Explor. Newsl.* **15**, 49–60 (2014)
13. Ahsan, M., Ngo-Ye, L.: A Conceptual model of recommender system for algorithm selection. In: *AMCIS 2005 Proceedings*, p. 122 (2005)
14. Collins, A., Tkaczyk, D., Beel, J.: A novel approach to recommendation algorithm selection using meta-learning. In: *Proceedings of the 26th Irish Conference on Artificial Intelligence and Cognitive Science (AICS), CEUR-WS*, pp. 210–219 (2018)
15. Cunha, T., Soares, C., de Carvalho, A.C.: Metalearning and recommender systems: a literature review and empirical study on the algorithm selection problem for collaborative filtering. *Inf. Sci.* **423**, 128–144 (2018)
16. Cunha, T., Soares, C., de Carvalho, A.C.: CF4CF: recommending collaborative filtering algorithms using collaborative filtering. arXiv preprint [arXiv:1803.02250](https://arxiv.org/abs/1803.02250) (2018)
17. Cunha, T., Soares, C., de Carvalho, A.C.P.L.F.: Selecting collaborative filtering algorithms using metalearning. In: *Frasconi, P., Landwehr, N., Manco, G., Vreeken, J. (eds.) ECML PKDD 2016. LNCS (LNAI)*, vol. 9852, pp. 393–409. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46227-1_25

18. Matuszyk, P., Spiliopoulou, M.: Predicting the performance of collaborative filtering algorithms. In: Proceedings of the 4th International Conference on Web Intelligence, Mining and Semantics (WIMS 2014), p. 38. ACM (2014)
19. Misir, M., Sebag, M.: ALORS: an algorithm recommender system. *Artif. Intell.* **244**, 291–314 (2017)
20. Romero, C., Olmo, J.L., Ventura, S.: A meta-learning approach for recommending a subset of white-box classification algorithms for Moodle datasets. In: Educational Data Mining (2013)
21. Vartak, M., Thiagarajan, A., Miranda, C., Bratman, J., Larochelle, H.: A meta-learning perspective on cold-start recommendations for items. In: Advances in Neural Information Processing Systems, pp. 6907–6917 (2017)