Federated Meta-Learning: Democratizing Algorithm Selection Across Disciplines and Software Libraries

-- Proposal, v1.1 –

Joeran Beel
Trinity College Dublin
School of Computer Science and Statistics
Artificial Intelligence Discipline
ADAPT Centre
beelj@tcd.ie

KEYWORDS

Meta-Learning As-A-Service, Interoperable Meta-Learning, Decentralized Meta-Learning, Decentralized Algorithm Selection, Federated Algorithm Selection, Algorithm Selection Problem, Automated Algorithm Configuration, Hyperparameter Optimization, Federated Learning

Introduction

There is an ever-growing number of tools for automating the machine learning pipeline, both commercial and open source. Auto-sklearn [11, 15], Auto Weka [14], ML-Plan [18], and H2O.ai are only some examples. In addition, the *auto* ∗ movement reached other areas of data science in the wider sense, for instance, the recommender-system community with LibRec-Auto [17]. A key component of *auto*∗ and algorithm selection tools is the algorithm selection and configuration process, which is subject to intensive research [1–3, 6, 7, 9, 10, 12, 13, 16, 20, 22, 24]. In some disciplines, like the SAT community, automated algorithm selection tools like SATzilla achieved remarkable improvements compared to ´standard´ algorithms [25].

Meta-learning is one of the most promising techniques to warm-starting the algorithm selection and configuration process [12]. With meta-learning, a machine learning model is trained to predict how algorithms will perform on a given task. That predicted algorithm can then either be applied directly to solve the task, or further be optimized. The meta-learning model is built based on the past performance of algorithms on a large number of tasks (datasets). The tasks are typically described through meta-features, and for new unseen tasks, the most performant algorithms can be predicted through the meta-learner. Meta-learning is not only used for predicting the performance of machine learning algorithms but for almost any kind of algorithm selection in various disciplines (SAT [25]; recommender systems [5, 8]; reference parsing [21]…).

Research Problem

The prediction accuracy of meta learning for algorithm selection varies strongly among disciplines. A challenge is the (non) availability of data in some disciplines to build the meta-learning model. This is due to the typical workflow of machine learning, data science etc. Typically, software libraries – be it machine learning libraries like (Auto) sklearn, or recommender system libraries like LibRec (-Auto) are used in isolation, either locally or in the cloud. Either way, the information how different algorithms and their configurations performed on the particular dataset, is neither published nor shared with others.

Research Goal

Our goal is to facilitate the algorithm selection process by leveraging historic performance data that was on various devices by various software libraries.

Federated Meta Learning

We propose "Federated Meta-Learning" (FML), a concept that allows everyone to benefit from the data that is generated through software libraries including machine learning and data science libraries as well as the *auto*∗ tools and tools from other domains such as the SAT community. We envision a peer-to-peer or client-server architecture that allows the exchange of meta-data and models for the purpose of meta-learned algorithm selection and configuration across disciplines.

The input to Federated Meta Learning is a description of the task, and the output is a recommendation for the potentially best performing algorithm(s) to solve that task. This recommendation could consist simply of a list of the best algorithms, or their predicted performance values. The list could also consist of multiple sub-lists created with different meta learners.

In its simplest form, federated meta-learning is a knowledge base or directory of algorithms-data performance measures. Ultimately, Federated Meta-Learning would be able to predict algorithm performance for unseen tasks.

To the best of our knowledge this concept is novel. The term "Federated Meta Learning" has only been used once before by Chen et al. but in a different context [4]. Federated Machine Learning has no central instance that stores and controls all raw data. Instead, the learning is performed on the local devices who remain the data owners and in full control. This makes Federated Machine Learning different from *distributed* machine learning [19] or central repositories like OpenML [23]. Federated Meta Learning has similarities with "federated machine learning", which was recently introduced by Google: "*Federated [Machine] Learning enables [devices] to collaboratively learn a shared prediction model while keeping all the training data on device, decoupling the ability to do machine learning from the need to store the data in the cloud."*[1]

However, Federated *Machine* Learning focuses on learning one machine learning task across multiple devices. Federated *Meta* Learning focuses on learning algorithm performance for arbitrary tasks across devices (or one generic model applicable to all tasks). We envision federated meta learning as an ecosystem where the raw data is kept on the original devices. However, either a) the meta data would be shared among the devices, b) the meta data would be stored on a central device, or c) the learning would be distributed, and the created model would be shared among the devices.

Challenges

Federated Meta Learning is a highly challenging research project. If the learning is distributed, as with Federated Machine Learning, the key challenge lies in the question of how to effectively train models across the devices? Other challenges include the creation of a generic data description language that is powerful enough to satisfy the needs of different disciplines. Architectural challenges include the question whether to use a peer-to-peer or client-server architecture. In the long run, social questions need consideration such as preventing manipulation (developers of algorithms may have an interest that their algorithms are ´recommended´) and free-rider problems (users benefiting from the system without sharing their data). Another challenge would arise if the system should not only focus on the globally best algorithm for a task (an entire dataset) but if per-instance algorithm selection should be learned. This would make the whole system even more complex.

REFERENCES

[1] Bischl, B., Kerschke, P., Kotthoff, L., Lindauer, M., Malitsky, Y., Fréchette, A., Hoos, H., Hutter, F., Leyton-Brown, K., Tierney, K. and others 2016. Aslib: A benchmark library for algorithm selection. *Artificial Intelligence*. 237, (2016), 41–58.

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

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

[4] Chen, F., Dong, Z., Li, Z. and He, X. 2018. Federated Meta-Learning for Recommendation. *arXiv preprint arXiv:1802.07876*. (2018).

[5] Collins, A. and Beel, J. 2019. A First Analysis of Meta-Learned Per-Instance Algorithm Selection in Scholarly Recommender Systems. *Workshop on Recommendation in Complex Scenarios, 13th ACM Conference on Recommender Systems (RecSys)* (2019).

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

[7] Cunha, T., Soares, C. and Carvalho, A.C. 2017. Metalearning for Context-aware Filtering: Selection of Tensor Factorization Algorithms. *Proceedings of the Eleventh ACM Conference on Recommender Systems* (2017), 14–22.

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

[9] Edenhofer, G., Collins, A., Aizawa, A. and Beel, J. 2019. Augmenting the DonorsChoose.org Corpus for Meta-Learning. *Proceedings of The 1st Interdisciplinary Workshop on Algorithm Selection and Meta-Learning in Information Retrieval (AMIR)* (2019), 32–38.

[10] Ferrari, D.G. and De Castro, L.N. 2015. Clustering algorithm selection by meta-learning systems: A new distance-based problem characterization and ranking combination methods. *Information Sciences*. 301, (2015), 181–194.

[11] Feurer, M., Klein, A., Eggensperger, K., Springenberg, J., Blum, M. and Hutter, F. 2015. Efficient and Robust Automated Machine Learning. *Advances in Neural Information Processing Systems 28*. C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, eds. Curran Associates, Inc. 2962–2970.

[12] Hutter, F., Kotthoff, L. and Vanschoren, J. 2019. Automatic machine learning: methods, systems, challenges. *Challenges in Machine Learning*. (2019).

[13] Kotthoff, L. 2016. Algorithm selection for combinatorial search problems: A survey. *Data Mining and Constraint Programming*. Springer. 149–190.

[14] Kotthoff, L., Thornton, C., Hoos, H.H., Hutter, F. and Leyton-Brown, K. 2017. Auto-WEKA 2.0: Automatic model selection and hyperparameter optimization in WEKA. *The Journal of Machine Learning Research*. 18, 1 (2017), 826–830.

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

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

[1] https://ai.googleblog.com/2017/04/federated-learning-collaborative.html

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

[18] Mohr, F., Wever, M. and Hüllermeier, E. 2018. ML-Plan: Automated machine learning via hierarchical planning. *Machine Learning*. 107, 8-10 (2018), 1495–1515.

[19] Peteiro-Barral, D. and Guijarro-Berdiñas, B. 2013. A survey of methods for distributed machine learning. *Progress in Artificial Intelligence*. 2, 1 (2013), 1–11.

[20] Romero, C., Olmo, J.L. and Ventura, S. 2013. A meta-learning approach for recommending a subset of white-box classification algorithms for Moodle datasets. *Educational Data Mining 2013* (2013).

[21] Tkaczyk, D., Gupta, R., Cinti, R. and Beel, J. 2018. ParsRec: A Novel Meta-Learning Approach to Recommending Bibliographic Reference Parsers. *Proceedings of the 26th Irish Conference on Artificial Intelligence and Cognitive Science (AICS)* (2018), 162–173.

[22] Tu, W.-W. 2018. The 3rd AutoML Challenge: AutoML for Lifelong Machine Learning. *NIPS 2018 Challenge* (2018).

[23] Vanschoren, J., Van Rijn, J.N., Bischl, B. and Torgo, L. 2014. OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*. 15, 2 (2014), 49–60.

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

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