

jPL: A Java-based Software Framework for Preference Learning

Pritha Gupta, Alexander Hetzer, Tanja Tornede,
Sebastian Gottschalk, Andreas Kornelsen, Sebastian Osterbrink,
Karlson Pfannschmidt, and Eyke Hüllermeier

Intelligent Systems Group
Paderborn University

Preference learning (PL) is an emerging subfield of machine learning, which deals with the induction of preference models from observed preference information [3]. Such models are typically used for prediction purposes, for example to predict context-dependent preferences of individuals on various choice alternatives. Depending on the representation of preferences, individuals, alternatives, and contexts, a large variety of preference models and problems are conceivable.

We developed a software framework offering tools and algorithms for solving preference learning problems.¹ While software frameworks for core machine learning problems such as classification abound, we are not aware of any comprehensive library of tools for preference learning. In fact, existing libraries are essentially restricted to one or two types of PL problems (e.g. [2], [6], [5], [4], [1]).

Our framework, called jPL, is implemented in Java. It is based on a unified data format, the Generic Preference Representation Format (GPRF), which is suitable for modeling data related to different kinds of preference learning problems. This also includes a dataset transformer, which converts data from several existing formats to GPRF. As problem classes, the framework currently supports collaborative filtering, instance ranking, label ranking, multilabel classification, object ranking, ordinal classification, and rank aggregation out of the box, with at least two algorithms being implemented for each problem. It provides a convenient command line interface as well as an API, both allowing one to configure the system using json files. The whole framework was developed in a quite generic way, so as to allow other problems and algorithms to be added easily.

Our framework also supports the evaluation and comparison of different methods in terms of standard validation techniques, and includes a set of commonly used loss functions. Just like the framework as a whole, the evaluation component is easily extensible by new evaluation techniques and loss functions.

References

1. V Dang. The lemur project-wiki-ranklib. lemur project.
2. Vincent E Farrugia, Héctor P Martínez, and Georgios N Yannakakis. The preference learning toolbox. *arXiv preprint arXiv:1506.01709*, 2015.

¹ <https://github.com/intelligent-systems-group/jpl-framework>

3. Johannes Fürnkranz and Eyke Hüllermeier. Preference learning: An introduction. In *Preference learning*, pages 1–17. Springer, 2010.
4. Nicholas Mattei and Toby Walsh. Preflib: A library for preferences <http://www.preflib.org>. In *International Conference on Algorithmic Decision Theory*, pages 259–270. Springer, 2013.
5. Jesse Read, Peter Reutemann, Bernhard Pfahringer, and Geoff Holmes. Meka: a multi-label/multi-target extension to weka. *The Journal of Machine Learning Research*, 17(1):667–671, 2016.
6. Grigorios Tsoumakas, Eleftherios Spyromitros-Xioufis, Jozef Vilcek, and Ioannis Vlahavas. Mulan: A java library for multi-label learning. *Journal of Machine Learning Research*, 12(Jul):2411–2414, 2011.