Towards Understanding Learning Behavior (Abstract)

Joaquin Vanschoren and Hendrik Blockeel

Declarative Languages and Artificial Intelligence research group,
Computer Science department, K.U.Leuven

For time-efficient learning, each machine learning (ML) algorithm must make some assumptions about the underlying structure of the data, and the utility of an ML algorithm largely depends on how well these assumptions match the concept hidden in the data at hand. Still, in the current state of the art, not much is known about the conditions under which one inductive method is better than another, making it hard to make a rational choice. The objective of meta-learning is to assist ML practitioners by learning from previous learning tasks. Learning how characteristics of data and methods relate to the behavior of inductive methods can provide useful knowledge for tackling new learning problems.

However, the effectiveness of current meta-learning approaches remains limited, because of a lack of publicly available datasets (on the meta-level each training example requires a completely new dataset), because learning methods are usually not characterized (inhibiting generalization over different settings or modified versions of the algorithm), because it offers no causal explanations for an algorithm's behavior and because preprocessing is often not considered.

To learn more effectively from previous ML experiments, we propose constructing an experiment database, stating all involved parameters of past experiments, making the results generalisable, reproducible and reusable for further research. We could then investigate (by querying) or discover (by data mining) a wide range of hypotheses. To cover a broad spectrum of data characteristics, we propose using synthetic datasets (validated against natural ones) to augment our meta-learning experiments. This requires a fine-grained dataset generator able to hide a wide range of different concepts in the data, to introduce possibly complex relations, correlations and different value distributions in the attributes, and to add artificial noise, missing values, irrelevant attributes,...

To generalize over algorithms (and to explain their behavior in terms of their properties) we need an adequate characterization consisting of parameter settings, used techniques, representational model,... Also, a bias-variance decomposition of the predictive error can provide a deeper understanding by linking poor performance to e.g. inappropriate modeling or bad parameter settings, and advise corrective measures. Finally, since preprocessing has a large impact on the utility of a learning algorithm, we should build a separate experiment database measuring the effect of preprocessing techniques on data characteristics, so we may learn which techniques may benefit specific ML algorithms on a specific problem. As such, we could provide truly practical advice on the utility of data mining techniques, together with the required preprocessing techniques.