

A Value Theory of Meta-Learning Algorithms

Abraham Bagherjeiran
University of Houston
Department of Computer Science
Houston, TX 77204-3010
abagherj@cs.uh.edu

Abstract

We use game theory to analyze meta-learning algorithms. The objective of meta-learning is to determine which algorithm to apply on a given task. This is an instance of a more general problem that consists of allocating knowledge consumers to learning producers. Solving this general problem in the field of meta-learning yields solutions for related fields such as information retrieval and recommender systems.

Introduction

The problem that is the subject of this work arises in three distinct and seemingly unrelated areas of research within Artificial Intelligence: meta-learning, information retrieval, and recommender systems. The main objective of a meta-learning algorithm is to determine which learning algorithm is best to apply to a given learning task. In information retrieval, one seeks to retrieve the documents most relevant to a user's query. Finally, the goal of a recommender system is to provide the user with a set of items for potential purchase.

These three fields solve the same problem: the producer-consumer problem. A consumer has a utility function defined over items it can consume, which we call products. A producer, which creates the products, needs to know what to produce to give the consumer maximal utility. Producers and consumers interact only through an intermediary agent, which we call the mediator.

Although we concentrate on meta-learning, we ask whether a solution common to all three related fields exists. Answering this question is the subject of our research. The current focus is on the use of game theory to model how producers and consumers would interact as collaborative users that do not necessarily trust each other.

Producer-Consumer Model

The goal of a meta-learning algorithm is to map a sequence of tasks onto a set of hypothesis spaces (Baxter 2000). A learning algorithm learns a hypothesis $h \in \mathcal{H} \in \mathbb{H}$ to maximize expected reward on a task:

$$E[r(h)]_P = \int_{X \times Y} r(h(x), y) dP(x, y)$$

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

where P is the example distribution, h is the hypothesis output by the learning algorithm, and r is a reward function. A meta-learning algorithm operates one level above a learning algorithm to learn hypothesis spaces for a sequence of tasks. We think of a hypothesis space as a set of parameters for an algorithm (the parameters bias the algorithm toward some hypothesis space).

The producer-consumer model of meta-learning decomposes the meta-learning problem into autonomous agents. A producer produces a hypothesis space. A consumer applies a hypothesis space on a sequence of tasks. The solution to the meta-learning problem is a mapping of tasks to hypothesis spaces—consumers to producers. In this model, producers and consumers must agree on a *single*, shared mapping. Having a single mapping increases the speed with which a consumer can find the best product. Since consumers and producers will have to compete for their contributions to the model, their respective utility functions and strategy choices are of particular interest. Considering the utility of their respective choices allows us to examine the *value* of a task or a hypothesis space.

Meta-Learning as a Game

For autonomous agents to form this shared mapping, they must collaborate. When one producer's utility depends on a consumer's choice and *vice versa*, the collaboration becomes competitive. We analyze this competitive collaboration as a game in which each producer and consumer is a player.

A producer is a learning algorithm that selects a parameter combination for an algorithm. Since parameters bias a learning algorithm to a hypothesis space, we can interpret selecting parameters as a form of specialization: selecting a smaller portion of the hypothesis space. The producer's utility function depends on the number of consumers that prefer this parameters combination and its own preference for the parameters. For example, a producer of neural networks prefers smaller networks—*a priori*—because they are faster to evaluate and less likely to over-train the data.

A consumer evaluates a learning algorithm over a sequence of tasks from the same distribution. A consumer's strategy is to select a learning algorithm—a producer. A consumer's utility depends on the average performance of the algorithm on its tasks. For example, a consumer can select

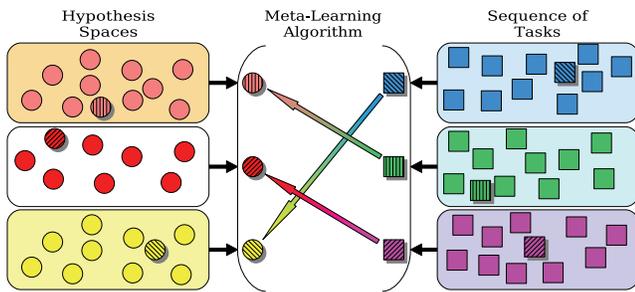


Figure 1: In the producer-consumer model of meta-learning, producers select a parameter combination for an algorithm (hypothesis space). Consumers—modeled as prototype tasks—select a specialized algorithm for a sequence of tasks.

the producer with the best average performance using cross-validation.

In an meta-learning algorithm as a game—as illustrated in Figure 1—a consumer selects a producer which selects a parameter of a learning algorithm. The consumer applies the specialized algorithm on the set of tasks and receives its average performance as a utility. The producer receives a utility proportional to the number of consumers that the specialized algorithm satisfies. A producer benefits when it finds an algorithm combination that satisfies many consumers just as a consumer benefits when it selects a producer that provides good performance.

When performed simultaneously, the strategy combination results in a shared mapping that provides a utility value to each player. An interesting feature of game theory is that we can find an *optimal* solution to the game: the Nash equilibrium. The set of strategies such that no player has any interest to change its strategy is a Nash equilibrium strategy (Rasmusen 2001). We can investigate the meta-learning maps that result from playing equilibrium strategies.

Given this game-theoretic analysis framework, the *value* of a task or hypothesis space is the value its strategy provides a player. The value of the mapping is the value to each player. The purpose of meta-learning, thus, is to find a mapping that maximizes the value for all players.

Applications

Although we focus on meta-learning, the idea has applications to other fields. Interestingly, this producer-consumer model is an extension of a model of insurance markets in which insurance companies must specialize their policies for different classes of buyers (Wilson 1977). As a model of meta-learning, it allows us to compare meta-learning algorithms based on their approximation of equilibrium strategy combinations.

In meta-learning, the use of game theory allows us to extend the idea of landmarking learners. Landmarking represents a task as the performance of simple learning algorithms on it (Bensusan & Giraud-Carrier 2000). In the game theory approach, players with identical preferences of the strategy set are interchangeable. We use an existing similar consumer

as an approximation to a new consumer. This allows us to reuse the recommendation of the old consumer on the new one.

In information retrieval, the set of products is the finite set of documents in a database at any point in time. Consumers are modeled as a query and a relevance measure which is derived from the query. Producers are groups of related documents. The mediator is a retrieval algorithm that receives the query and retrieves the relevant documents. Most work in information retrieval focuses on the rapid retrieval of documents, but we focus on learning a consumer’s relevance measure. We successfully used the ideas for content-based image retrieval (Bagherjeiran, Vilalta, & Eick 2005).

In recommender systems, the set of products is the set of actual goods and services. Consumers are potential buyers modeled as a history of recent purchases. Producers actually produce the goods, which we model as groups of related products. As a recommender system, the mediator recommends items for purchase through indirect communication such as advertisements or user comments.

Conclusion

We discussed the use of game theory to analyze a problem common to the seemingly unrelated fields of meta-learning, information retrieval, and recommender systems. This producer-consumer problem consists of learning producers and knowledge consumers as players in a game that compete to find a single, shared solution to their respective problems. The solution to the game maps consumers to the producers that can best satisfy them.

Using game theory to analyze this problem, we intend to compare existing algorithms with respect to the optimal solutions to the game. In future work, we will determine which equilibria best describe the solutions of existing algorithm. We will also develop approximation methods to quickly find good solutions to the game.

References

- Bagherjeiran, A.; Vilalta, R.; and Eick, C. F. 2005. Content-based image retrieval through a multi-agent meta-learning framework. In *Proc. 17th Int’l Conf. on Tools with Artif. Intell.*, 24–28. Hong Kong, China: IEEE Computer Society.
- Baxter, J. 2000. A model of inductive bias learning. *J. Artif. Intell. Res.* 12:149–198.
- Bensusan, H., and Giraud-Carrier, C. G. 2000. Discovering task neighbourhoods through landmark learning performances. In *Proc. 4th European Conf. on Principles of Data Mining and Knowledge Discovery*, 325–330. London, UK: Springer-Verlag.
- Rasmusen, E. 2001. *Games and Information: An Introduction to Game Theory*. Oxford, UK: Basil Blackwell, 3rd edition.
- Wilson, C. 1977. A model of insurance markets with incomplete information. *J. Economic Theory* 16:167–207.