

A Taxonomy of Computational and Social Learning

Koji Morikawa (morikawa@crl.mei.co.jp)

Advanced Technology Research Laboratories
Matsushita Electric Industrial Co., Ltd.
3-4 Hikaridai, Seika, Soraku, Kyoto 619-0237 Japan

Sameer Agarwal (sagarwal@cs.ucsd.edu)

Charles Elkan (elkan@cs.ucsd.edu)

Garrison W. Cottrell (gary@cs.ucsd.edu)

Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Drive, La Jolla, California, 92093-0114 USA

Abstract

This paper presents a first attempt at explaining the relationship between the psychological and artificial intelligence points of view of learning with a special focus on social learning. A two dimensional classification methodology is proposed that classifies learning behaviors in intelligent agents on the basis of agent structure and of information acquisition modes. We make a fundamental distinction between active and passive modes of learning, depending on the agent's ability to affect the information source. A preliminary classification of learning behaviors in this scheme is presented and areas of future work are identified.

Introduction

There exists a mismatch in how psychologists and computer scientists view learning. Traditionally computer scientists classify learning into roughly three distinct areas of study: supervised learning, unsupervised learning and reinforcement learning. This classification ignores issues like where the learner is situated and how the training data is acquired. Psychologists on the other hand consider learning as a social activity involving other agents. Recently there has been interest in developing agents and learning algorithms which take into consideration these issues and can learn in multi-agent environments. This is referred to as imitation learning or more generally social learning [?, ?, ?, ?, ?]. Unfortunately even within the psychological community there is little agreement as to what exactly constitutes imitation. In this paper we aim to survey the various kinds of social learning behaviors and present a computational view of social learning. The benefit we seek is to generalize the computational view while also considering the algorithmic requirements of social learning.

The task of designing self developing autonomous agents such as robots or intelligent agent software requires us to have a clear understanding how such agents can learn by interacting with and observing other agents and/or humans in their environments. More specifically, it requires us to be able to classify the various kinds of learning behavior that an agent like this can be expected to exhibit, characterizing them in terms of their computational requirements. By computational we mean a characterization in terms of the goal of the learning behavior, and the type of the information source from which learning data is acquired.

This paper presents a first attempt at constructing a taxonomy of learning in agents and relating it to existing literature on social learning in humans and animals. We identify learning behaviors as tasks defined in terms of the particular component of the agent that is modified and the specific mode of information acquisition it uses to do so. We identify behaviors for which implementation mechanisms are known and point out gaps which need to be filled in future research.

The rest of the paper is organized as follows. The next section presents an overview of definitions and distinctions made in the psychological literature dealing with learning in multi-agent environments. We begin our analysis in the third section by defining an agent in terms of its components. Section four presents the various modes of information acquisition and how they relate to various kinds of learning behaviors. Finally we end with a discussion and ideas for future work.

Social Learning

Humans and animals learn continuously by interacting and observing others around them. Learning in human and animal societies is an extensively studied topic in psychology, ethology and anthropology. Consequently there are a number of definitions of social learning in the literature. A modern definition which can be applied to both psychological social learning and computational one is provided by [?]:

Social learning is the phenomenon by means of which a given agent (learner) updates its own knowledge base (adding to, or removing from it a given information, or modifying an existing representation) by perceiving the positive or negative effects of any given event undergone or actively produced *by another agent* on a state of the world which the learning agent has as a goal. [emphasis added]

[?] makes a distinction between two major kinds of social learning behaviors, namely social facilitation and imitation. In social facilitation a learner's acquisition of a piece of information maybe caused by another agent but that does not necessarily imply that information transmission occurred from the mind of the latter to the former. Hence the learner need not attribute mental states

and goals to the agent he is observing. Contrast this with the case of imitation, where the agent being imitated has to be attributed intent as well as mental states. Thus Conte distinguishes between the two learning behaviors on the basis of nature of the information source.

[?], using the more specialized term “cultural learning”, distinguishes between three kinds of social learning behaviors by focusing on the relationship between the learner and those around it. The authors differentiate between **imitative**, **instructive** and **collaborative** learning. Where the primary distinction is the direction of information flow. Imitation has a unidirectional information flow with no direct interaction between the demonstrator and the learner. Instructive learning corresponds to a student teacher relationship where the information flow is bi-directional, but asymmetric with the instructor being the primary source of information. Finally collaborative learning is the situation where none of the agents involved have enough knowledge to complete the task involved. Learning happens through mutual cooperation and discovery and information flow is symmetrical and bi-directional.

Tomasello also points out that these learning behavior are correlated with the developmental stages of a child. Each stage embraces the previous one, that is, instructive learning employs imitative learning, and collaborative learning employs both instructive and imitative.

While Tomasello and Conte choose to differentiate on the basis of information sources and flow, [?] classifies social learning on the basis of *what* is learned. She makes the distinction between non-imitative social learning, which corresponds to information acquisition about objects in the environment, and imitative social learning, which is behavior acquisition by observing conspecifics.

Thus, there are two ways in which existing approaches categorize learning behaviors, on the basis of the source of information (inputs) and on the basis of what part of the agent is improved (output). We believe that these two approaches represent an orthogonal characterization of learning behaviors, and this is the approach we take in the rest of the paper. The dual specification of a behavior in terms of input and output allows for a more functional, implementation-friendly view.

Structure of an Agent

Agent structure is the first dimension along which we base our analysis, it is useful to parameterize our notion of a self-developing agent. An agent is an entity characterized by:

1. **Input Selector** Every agent has a set of sensors which inform it about the external state of the environment and its own internal states. But what really matters are the inputs streams and features that the agent is devoting attention to at a given moment. Even though it might be sensing a multitude of inputs from different sources, it is using only a few for reasoning and learning. Learning input selection then corresponds to learning to choose a particular input stream to attend

to and/or learning to select particular features from an input stream. Computationally this corresponds to the task of signal detection and feature selection.

2. **Knowledge Base** This serves as a store of experiences and factual knowledge about the environment and self. It includes, for example, simple statements that indicate qualitative properties of objects, and laws of physics that hold true. Knowledge about the environment may be explicitly presented as preprogrammed or learned rule bases or it may be implicitly represented inside other components, e.g. inside the decision function as constraints on output actions.
3. **Goals** The intent to perform a certain task or to achieve a certain target is defined to be a goal of an agent. An agent without intent will not do anything at all. It is assumed that every agent starts out with a set of high-level goals innate to it, which govern its overall behavior. As part of its learning the agent learns to decompose goals into subgoals or change the salience of existing ones. The process of subgoaling is important because it restricts the problem specification and constrains the search space, thereby converting the complex problem corresponding to a high-level goal satisfaction into pieces that can be solved more easily. The exact structure of the goal is also of importance as it decides what is the output of the learning process [?].

It is instructive here to compare our definition of a Goal with an existing framework. The BDI framework proposed by [?] views intelligent agents as rational agents with certain mental attitudes: Belief, Desire and Intentions (BDI). The three attitudes refer to the informational, motivational and deliberative components of the agent. The informational component is what we refer to as knowledge base, the motivational component is responsible for assigning values, payoffs or priorities to tasks and the deliberative component is responsible for defining the criterion for choosing the best course of action, i.e. whether the agent should maximize expected returns or minimize the likelihood of errors.

Mathematically, goals translate into error functions that need to be minimized, reward functions that need to be maximized or first order predicate logic statements that need to be satisfied. Here the function that assigns values to the actions is the agent’s desire and the exact form of the error function that should be minimized, its intent.

4. **Decision Function** This is the mapping from stimuli to actions. It is assumed that the agent has a fixed set of primitive actions that it can perform. One of the tasks decision functions perform is constructing high-level behavior by choosing a sequence of primitive actions for a given situation. Another task that they perform is reasoning and planning. [?] refer to this as the *program* part of an intelligent agent. Computationally

this corresponds to a classifier or a regression function, represented for example by a neural network or a decision tree or a procedure which takes the map of a maze and produces a possible path out of it.

5. **Learning Method** A learning agent possesses mechanisms by which it can construct new decision functions, modify existing ones to satisfy new or modified goals or adapt to changing conditions in the environment. Computationally this corresponds to various kinds of Learning/Training algorithms.

Treating each of the above as a variable, an agent is a particular instantiation of a combination of these variables. In general, traditional machine learning research assumes all but one of the above variables is given and attention is focused on learning the one which is left out. Learning has a very narrow meaning and is restricted to function induction in one form or another. However, to construct agents which autonomously improve their skills over time, one must allow for modifications or improvements to be made to more than one of these variables simultaneously.

The notion of learning for psychologists, however, is a much more general one. They include acquisition of variables 1-3 within the scope of learning. Hence psychologists use the term *learn* when an agent learns to focus its attention on particular objects in the environment or acquires a new goal.

[?] distinguishes learning how to learn i.e. learning new learning methods from other kinds of learning behaviors by referring to it as second order learning. In the current paper we focus our attention on what Bateson refers to as first order learning, and do not deal with acquisition of learning methods. This will be the subject of future work.

Information Acquisition

The second dimension in learning is the source of information and how it is acquired. We make a primary distinction between active and passive modes of information acquisition based on the role of the learner in collecting the training data.

Passive

Passive or observational information collecting modes are characterized by a one-way transfer of information. The learner collects information by observation and has no control over the examples that it is exposed to. This is a commonly occurring model of training set availability in machine learning, where examples come from a fixed, and in most cases an unknown distribution. A problem with learning in this mode is that the data available may be biased or skewed, resulting in models which are not robust to noise in the inputs.

Environment Passive information collection from the environment implies that an agent observes the environment without disturbing it. The agents learn about the behavior of the various objects in the environment and

how to relate their behavior to events in the environment. Both traditional supervised and unsupervised learning techniques can be used to learn from data from the environment.

Agent-agent & agent-environment interaction Observing an agent interact with the environment and other agents present therein, and using this information to acquire new behavior, is commonly referred to as imitative learning. Imitative learning is attractive due to a number of reasons, it allows an agent to learn novel behaviors without incurring the cost of making mistakes which it would if it were to experiment by itself, an agent whose explicit role is that of a teacher is not required and finally since the process is passive no interaction between the learner and the teacher is needed [?]. Here it is important to note that there is a large class of behaviors that are referred to as imitative, including acquisition of high-level actions, attention, and new goals, which may not correspond to some researchers' definition of *true* imitation learning. [?] present an analysis of three learning behaviors which have traditionally been attributed to imitation but can be explained by a much simpler mechanism. While the exact definition of true imitation learning remains a subject of lively debate [?], for our purposes we choose to follow Conte's definition:

Imitation is a behavior ruled by the goal that a given agent **O** (which stands for observer) be-like or act-like another agent **M** (which stands for model) as long as **M** is (perceived as) a suitable model under a given circumstance.

Notice that the above definition requires the learner not only to learn from another agent, but to also evaluate the suitability of that agent to serve as a model.

The principal characteristic that differentiates true imitation learning from other imitative behaviors is that the agent learns to copy and perform a novel behavior which is not part of its existing repertoire.

The existing research in computational mechanism of imitation learning [?, ?, ?, ?, ?], is limited to acquiring decision functions. We are not aware of work on computational models of attention acquisition (stimulus enhancement) or goal acquisition without acquiring the related method of goal satisfaction (emulation).

Self An agent capable of inspecting its own internal states can use this information to learn about itself and its relationship with things around it. Introspection allows the agent to reason about and make explicit facts and deductions that might not be obvious to begin with. The result of these deductions can then be stored for later reference and use saving on time and effort. Mathematical theorem proving is an example of this kind of activity. Proving a theorem does not provide us with any new knowledge which is not already contained in the system of axioms that imply it, however keeping track of existing deductions allows us to reuse them in subsequent theorem proving tasks. Explanation based learning (EBL) is an example of such learning behavior [?].

Active

Active/Interactive information collection modes are characterized by two-way transfer of information, the learner having some control over the process generating the training data. The learner can ask questions from an all-knowing oracle which could be a parent or a teacher, or perform experiments to verify the hypothesis or models that he has constructed. Learning in such query models is much more robust and less susceptible to errors due to noise in the inputs [?, ?]. Of course this model has its own set of complications, since now the agent must also possess the ability to generate queries based on what he has learned and what he wants to learn.

Instead of lumping all agent-agent interactions in one category we choose to make a finer distinction between interaction with other agents such as parent-like agents who are aware of the correct model and will answer the learner's questions truthfully, and interaction with sibling like agents who may or may not be aware of the correct answers or may choose not to reply truthfully. This is similar to a distinction made by Tomasello, who differentiates between instructive learning and collaborative learning. An interesting consequence is that now the agent needs to possess language skills to communicate with its parents and peers.

Interaction with the environment Being able to manipulate objects allows an agent to learn all that it could learn by simply observing them and more. Scientific discovery is a prime example of learning in this mode. Reinforcement learning is one way mechanisms for learning from the environment can be implemented. Here the agent tries to predict the result of his experiments, and self-reinforces if the results of experiments agree with its hypothesis.

Interaction with parent Interaction with parents is an example of interaction with a trusted agent who has perfect or near perfect knowledge of how to perform a task and is willing to share its knowledge. Information from the parent can come in a number of ways. The parent provide labels which tell the learner what is the correct decision to make or just indicate whether the learner's action was right or wrong. The parent can also help the learner by providing scaffolding in the form of breaking up the task into easier subtasks, allowing the learner to decompose his original goal into subgoals. Most forms of supervised and reinforcement learning would fall within this category.

Interaction with sibling This is the case where two of more agents are working together on a task but none of them has enough knowledge or skill to perform it on its own. Hence they work by sharing information and working on each other's suggestions. This introduces the additional complication of the agent responses not being reliable or complete anymore. Tomasello refers to this as collaborative learning.

[?, ?] are examples of this kind of behavior in the field of cooperative robotics. Parker uses the term *intentional*

cooperation for this kind of behavior, distinguishing it from swarm behavior.

Self Experimentation with self allows humans to construct models of themselves and their physical abilities. An example is babbling, where newborns start by uttering gibberish in an attempt to construct forward models of their speech system. Once acquired this model can be used to reason and plan the control sequence that the child needs to send to its muscles in order to utter a particular sound. e.g. a robot with a manipulator arm can start out with no knowledge of how to plan and move its arm to reach out to specific points in 3D space, but can learn to do so by experimenting with its motors and using self-reinforcement to indicate levels of success [?]

Discussion

Computer scientists while designing agents have the luxury of picking a subset of human cognitive abilities and implementing them without any physiological restrictions, e.g. natural language understanding. So while this approach produces functioning applications, the agents produced lack the ability to interact with the user and adapt to his needs. Psychologists have no such luxury and are constrained to explain human/animal behavior "as it is". Figure 1 presents an example of classification of learning behaviors on the basis of our scheme. Its important to note that none of the cells in the table exists in isolation, and there are close interactions between behaviors along the same column and rows. The dotted boxes in the figure correspond to what we believe are gaps in our understanding of computational mechanisms underlying these behaviors. We hope that future work in machine learning will focus on these subject areas.

Conclusions

In this paper we have presented a characterization of learning in intelligent agents, basing it on the two dimensions of agent structure and information acquisition. We pointed out classes of behavior for which computational implementations exist and where there are gaps in our understanding. Our classification also provides a way of assessing the capabilities of an agent by looking at the area covered by it on the table presented in Figure 1.

References

- [Alissandrakis et al., 2000] Alissandrakis, A., Nehaniv, C. L., and Dautenhahn, K. (2000). Learning how to do things with imitation. In *AAAI Fall Symposium on Learning How to Do Things*. AAAI.
- [Bateson, 1972] Bateson, G. (1972). *Steps to an Ecology of Mind*. Ballantine Books.
- [Billard, 2000] Billard, A. (2000). Learning motor skills by imitation: a biologically inspired robotic model. *Cybernetics & Systems*.
- [Bratman et al., 1988] Bratman, M. E., Israel, D., and Pollack, M. E. (1988). Plans and resource bounded

Information Source What is Learned	Passive/Observational			Active/Experimental			
	Environment	Agent-Environment	Self	Environment	Parent-Environment	Sibling-Environment	Self
Input Selector	Relation Discovery (Clustering)	Attention Imitation (Stimulus enhancement)	Introspection	Reinforcement Learning (for Input Selection)	Instructive Learning	Collaborative Learning	Experiments by itself
Knowledge Base	Discovery of facts about environment	Discovery of Agent-Environment Relations		Reinforcement Learning			
Goals	Subgoal Discovery	Goal imitation (Emulation)		Subgoal Discovery			
Decision Function	Predictive Model Production	Action imitation (Imitation)		Reinforcement Learning			

Figure 1: Classification of learning behaviors

practical reasoning. *Computational Intelligence*, 4:349–355.

[Byrne and Russon, 1998] Byrne, R. W. and Russon, A. E. (1998). Learning by imitation: A hierarchical approach. *Behavioural and Brain Sciences*, 16(3).

[Cohn et al., 1995] Cohn, D. A., Ghahramani, Z., and Jordan, M. I. (1995). Active learning with statistical models. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems*, volume 7, pages 705–712. The MIT Press.

[Conte, 2000] Conte, R. (2000). Intelligent social learning. In *Proceedings of the AISB'00 Symposium on Starting from Society - the Application of Social Analogies to Computational Systems*. AISB.

[Demiris and Hayes, 1996] Demiris, J. and Hayes, G. (1996). Imitative learning mechanisms in robots and humans. In Klingspor, V., editor, *Proceedings of the 5th European Workshop on Learning Robots*.

[Heyes, 1993] Heyes, C. M. (1993). Imitation, culture and cognition. *Animal Behaviour*, 46.

[Heyes and Galef, 1996] Heyes, C. M. and Galef, B. G., editors (1996). *Social Learning and the roots of culture*. Academic Press.

[Matarić, 1997] Matarić, M. J. (1997). Learning social behavior. *Robotics and Autonomous systems*, 20.

[Matarić, 2000] Matarić, M. J. (2000). Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics. *Cybernetics & Systems*, 32.

[Morimoto and Doya, 1998] Morimoto, J. and Doya, K. (1998). Reinforcement learning of dynamic motor sequence: Learning to stand up. In *Proceedings*

of IEEE/RSJ International Conference on Intelligent Robots and Systems, volume 3.

[Parker, 1998] Parker, L. E. (1998). Alliance: An architecture for fault tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*.

[Plutowski et al., 1994] Plutowski, M., Cottrell, G., and White, H. (1994). Learning mackey-glass from 25 examples, plus or minus 2. In *Advances in Neural Information Processing Systems*, volume 6. Morgan Kaufmann.

[Russell and Norvig, 1995] Russell, S. and Norvig, P. (1995). *Artificial Intelligence: A modern Approach*. Prentice Hall.

[Segre and Elkan, 1994] Segre, A. M. and Elkan, C. P. (1994). A high-performance explanation-based learning algorithm. *Artificial Intelligence*.

[Tomasello et al., 1993] Tomasello, M., Kruger, A., and Ratner, H. (1993). Cultural learning. *Behavioral and Brain Sciences*.