Dimensions of Learning in Agent-based Design

Dan L. Grecu & David C. Brown

Department of Computer Science,
Worcester Polytechnic Institute, Worcester, MA 01609, U.S.A.

Email: dgrecu@cs.wpi.edu, dcb@cs.wpi.edu
URL: http://cs.wpi.edu/~dgrecu, http://www.wpi.edu/~dcb

Keywords:

Design, multi-agent systems, SiFAs, dimensions of learning.

1.0 Introduction

This research originates in the work started several years ago at Worcester Polytechnic Institute dedicated to the investigation, modelling and evaluation of multi-agent based design. The main thrust behind our approach was the idea of finding the elementary patterns of agent problem-solving and interaction in design tasks. To achieve this goal we introduced and defined the concept of Single Function Agents (SiFAs) [Dunskus et al, 1995]. SiFAs are agents specialized to perform one single *function* during the design process. Some typical functions would be to select, evaluate, and provide critique. The SiFA concept is generic and agents can be instantiated for different, particular design domains. This introduces the other two important characteristics of SiFAs: the unique *target* on which they operate and the unique *point of view* from which they perform their task. In parametric design the targets are represented by design parameters. For example, a Selector (*function*) can be specialized to choose materials (*target*) with good corrosion resistance *(point of view*).

SiFAs helped us analyze what functions are needed in specific design situations, what types of knowledge are associated with each function, and what kind of information is exchanged between various pairs of agents. They also allow new approaches to conflict resolution and to negotiation over design decisions. Conflicts might occur when agents have different preferences over the same issues. Single Function Agents have proven to be particularly useful in detecting the basic factors which influence design decisions and design solutions in a concurrent engineering approach.

Learning adds additional power for agents to deal with real design problems. In this context learning is needed for several important reasons. First, it can adapt a system to the domain and to the design problems it has to solve. And, no less important, it can adapt the agents to each other, to make the problem-solving process smoother. Therefore, learning targets two levels in each agent: the level of its domain knowledge and the level of its interaction knowledge (fig.1).

In this paper we intend to analyze some of the ingredients of learning specific to a multi-agent design system. The considerations presented so far suggest that it is necessary to establish a set of dimensions under which to consider possible learning patterns. Modelling multi-agent design in terms of primitive agent functionality (SiFAs) provided us with the means to accurately explore cooperation in design. In a similar way, these learning dimensions are intended to allow for clear research orientation or classification in the very large and heterogeneous space of adaptability.

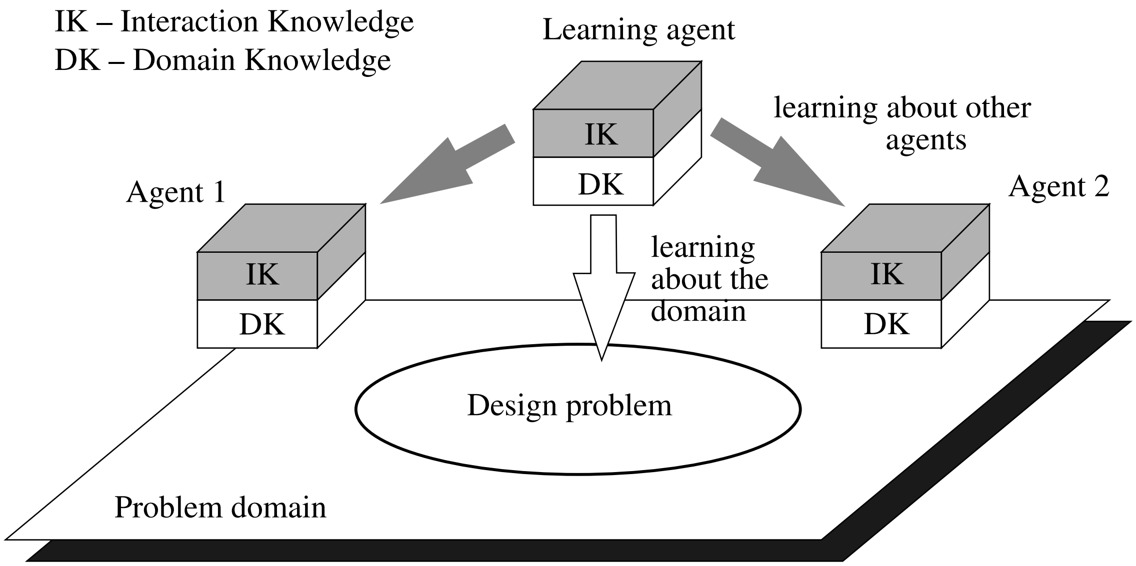

Figure 1: Learning in a multi-agent design context

2.0 Dimensions of learning

The dimensions are concerned with: what can trigger learning, what information supports learning, what might be learned, how information is made available, the learning methods, the scope of learning, and the consequences.

2.1 What can trigger learning?

Four major situations generate learning:

1. *Failure:* Failures can occur either during the design process or by testing/simulating the designed artifact. Both situations represent learning opportunities.
2. *Success*: A successful design or design decision is a learning opportunity, as far as the result represents an improvement compared to previous cases.
3. *Differences*: Differences, especially between points of view, are an important source of conflicts. Learning triggered by noticing differences helps anticipate and alleviate conflicts.
4. *Need to improve abilities*: The need to improve designs or the design process can be translated into a requirement to improve agent abilities. This means that an agent uses the learning about design situations or other design agents as a possibility to take more informed decisions in the future. Observing patterns, dependencies or other relations can be useful even though it is not motivated by events falling in any of the previous three categories.

2.2 What are the elements supporting learning?

Design agents learn by using information that becomes available in the following ways:

- knowledge provided through critique, praise, estimation, evaluation and advice;
- sequences of design decisions, whether stored or in progress;
- traces of knowledge exchanges and negotiation sessions between agents;
- analysis of failures and conflicts;
- feedback provided after completing the design task.

2.3 What might be learned?

While the range of the learnable is fairly wide, there are areas where learning might have a serious impact for design and for agent interaction. Some of the following types of learning target the level of design reasoning, while others target the meta-level of agent interaction.

- *Constraints related to parameters or, in general, to elements of the design*: By recording constraint information from other agents, it becomes possible to avoid conflicts through anticipatory evaluations.
- *Dependencies between design parameters*: Dependencies act as soft constraints. They are critical for deciding on which agents critique and praise should make an impact.
- *Support in favor of or against a decision*: Rationale can primarily help in weighing proposals of different agents and in deciding which agent should revise its decision.
- *Design rules*: As agents have different functionality and points of view, sharing of design rules is desirable among agents having the same target or at least the same domain.
- *Preferences in selection tasks*: If agent preferences are known, they help set up better initial proposals made by one agent to another.
- *Preconditions and postconditions for rules, actions and tasks*: This learning is important for a better adaptation of decisions to specific contexts. Agents decide by themselves when to act, and therefore should be able to recognize favorable situations to get involved in the design.
- *Prediction of decisions of other agents*: For SiFAs this amounts to building a behavioral model of other agents, by relating factors which influence an agent with the agent's decisions [Grecu & Brown, 1996].
- *Types of conflicts*: Conflicts can be seen, indexed and classified differently by the different types of agents. Enabling agents to classify and recognize conflicts enhances the power of the conflict anticipation and resolution mechanism.
- *Heuristics to solve conflicts and to negotiate*: These heuristics, being less dependent on the specific functionality of an agent, can be learned or transferred from one agent to another and refined according to local necessities [Cross 1978; Grecu & Brown 1994].

2.4 Availability of knowledge for learning

The way that agents make information available to the other agents determines where the learning opportunities arise. Varying the representation and the context in which the knowledge is communicated favors some of the learning types in comparison to the others.

- *Direct communication* implies transfer of knowledge between agents through direct channels. The transferred knowledge is not altered in this process.
- *Indirect communication* assumes the presence of an intermediary agent, conveying information from one agent to another. Knowledge may be pruned or abstracted in this process and may also be submitted to ontological transformations.
- *The design board* provides all the agents with the design decisions taken so far. Various attributes attached to these decisions can make the learning more powerful.
- *Repositories of design and interaction histories* store traces usable for later processing and compilation by the agents.

2.5 Methods of learning

The set of learning methods enumerated below is not claimed to be exclusive with respect to their application in the SiFA design framework. Only some of the potential applications are given in this brief traversal of the learning methods. Most scenarios lend themselves to on-line learning, as agents would make changes the latest after the current design session. The use of design repositories is the most typical off-line learning situation.

1. *Explanation-based learning.* Targets primarily the possibility to generalize strategies and plans, and to make them available across agents [Mitchell et al 1986].
2. *Induction (*[Fisher 1987]*).* The application area is extremely large, from classifying design situations to building behavioral models of other agents [Grecu & Brown, 1996].
3. *Knowledge compilation (*[Brown 1991]*).* Agents use plans and strategies, which represent opportunities to compile sequences of steps into macro-entities or to refine general operations in more domain-specific rules [Gordon & Subramanian 1993]. Histories offer the most compelling resource for this approach.
4. *Case-based learning.* Many atypical design contexts, which are not suitable for generalization, can be “remembered” as design cases [Sycara & Lewis, 1992]. Case-based learning can also prove itself useful for retrieving past interaction experiences between agents.
5. *Reinforcement learning.* The existence of critics offering immediate feedback on decisions, provides the most straightforward example of acquiring new decision knowledge through reinforcement techniques [Tan 1993; Whitehead 1991].

2.6 Local vs. Global Learning

There are two distinct ways in which learning can happen. The traditional way confines learning at the level of one agent. That means that the adaptive process happens at the level of one agent, in a multi-agent setting. Whenever the agents applies the learned knowledge or makes use of the changed reasoning parameters (such as weights and probabilities), the effects will be observable. A different kind of learning is of organizational type [Hutchins 1991; Shoham & Tennenholtz 1993; Weiβ 1993]. Several agents perform adaptive changes at their local level. However, the result of learning can be observed only through the concurrent use of these changes by the agents. The effect will be only at the global level. The learned items generating the effect will be stored by several agents.

2.7 Consequences of learning

Learning processes can be classified by the areas in which they generate change:

- *Design improvement*: One of the most desirable effects of learning is to improve the quality of designs. However, it is often hard to relate design quality change with the design decisions or cooperation decisions that caused them. SiFAs, being very specialized agents, also allow for an extension of the search space for design solutions, which is not affordable in agents with large complexity. In this case, the resulting improvement is reflected in a potentially larger number of solutions
- *Improvement of the design process*: A new multi-agent design system will cause a reasonable amount of decision revision during its first problem-solving sessions. This is due to the fact that the system is not "tuned" to the domain and the agents have little knowledge about each other. Adaptation can provide a fair amount of support in changing this situation, by avoiding conflicts and reducing the interaction needed to find a design solution.

3.0 Conclusions

The methodology described in this paper can be regarded as both a theoretical and a practical tool. On a theoretical level it offers a basis for evaluating and comparing multi-agent learning methods in design. From a practical point of view it creates a search space of learning opportunities in multi-agent design. By combining coordinates along different dimensions we are able to imagine new learning scenarios. The ranges given for each dimension could still be extended and refined. Despite that, they already provide us with valuable information about making design agents adaptive in various ways.

4.0 References

D. C. Brown, Compilation: The Hidden Dimension of Design Systems. In: *Intelligent CAD, III*, (eds.) H.Yoshikawa & F.Arbab, North-Holland, 1991.

J. Cross, Negotiation as a Learning Process. In: W. Zartman (ed.), *The Negotiation Process – Theories and Applications*, Sage Publications, 1978, pp. 29-54

B. Dunskus, D. Grecu, D. Brown & I. Berker, Using Single Function Agents to Investigate Negotiation. To appear in: *AI in Engineering Design, Analysis and Manufacturing – Special issue on Conflict Management*, I. Smith (ed.), 1995

D. Fisher, Knowledge Acquisition Via Incremental Conceptual Clustering, *Machine Learning*, Vol. 2, 1987, pp. 139-172

D. Gordon & D. Subramanian, A Multistrategy Learning Scheme for Agent Knowledge Acquisition, *Informatica*, Vol. 17, 1993, pp. 331-344

D. Grecu & D. Brown, Learning by Design Agents During Negotiation, *AI in Design Conference – Workshop on Machine Learning,* Lausanne, Switzerland, 1994

D. Grecu & D. Brown, Learning by Single Function Agents During Spring Design, *4th International AI in Design Conference,* Stanford, CA, 1996

E. Hutchins, Organizing Work by Adaptation, *Organization Science*, Vol. 2, No. 1, February 1991, pp. 14-39

T. Mitchell, R. Keller & S. Kedar-Cabelli, Explanation-Based Generalization: A Unifying View, *Machine Learning*, Vol. 1, No. 1, 1986

Y. Shoham & M. Tennenholtz, Co-Learning and the Evolution of Social Activity, *Tech. Rep. Stanford Univ.*, 1993, 36 pp.

K. Sycara & M. Lewis, Cooperative Negotiation in Concurrent Engineering Design, *Cooperative Engineering Design*, Springer Verlag, 1990, pp. 269-297

M. Tan, Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents, *Workshop on Reinforcement Learning, Machine Learning Conference*, 1993, 8 pp.

G. Weiβ, Learning to Coordinate Actions in Multi-Agent Systems, *IJCAI-93*, Vol. 1, pp. 311-316

S. Whitehead, A Complexity Analysis of Cooperative Mechanisms in Reinforcement Learning, *AAAI-91*, pp. 607-613