A FRAMEWORK FOR AUTONOMOUS AGENTS BASED ON THE CONCEPT OF ANTICIPATORY SYSTEMS

PAUL DAVIDSSON, ERIC ASTOR, and BERTIL EKDAHL*

Department of Computer Science, Lund University, Box 118
Lund, S–221 00, Sweden

ABSTRACT

This paper presents a new framework for autonomous agents that is based on the concept of anticipatory systems. It is a hybrid approach that synthesizes low-level reactive behavior and high-level symbolic reasoning. According to this framework, an agent, i.e. an anticipatory agent, consists of three main entities: a reactive system, a world model, and a meta-level component. The world model should, in addition to the description of the agent's environment, also include a description of the reactive part of the agent. The basic idea is that the meta-level component makes use of the world model to make predictions of future states. These predictions are then used by the meta-level to guide the agent's behavior on a high-level, whereas the low-level behavior is controlled by the reactive component.

1 INTRODUCTION

Anticipation is a mental activity that humans, but also some other living organisms, are capable of and one which is frequently practiced. A tennis player, for instance, has to anticipate the trajectory of the ball in order to make a good hit. A stockbroker makes forecasts of stock prices in order to make a good deal. In short, they use knowledge of future states to guide their current behavior. To do this they make use of an internal model of the particular phenomenon. The experienced tennis player's use of his model is probably on an unconscious level and has been learned through tedious sensorimotor training. A novice, on the other hand, has to use his model of the ball's trajectory in a more conscious manner. Similarly, the stockbroker's model is probably on a conscious level, learned through theoretical studies and experience of previous stock prices.

It is the authors' opinion is that this kind of anticipatory reasoning has not been sufficiently studied and, as a consequence, is not well-understood within the field of intelligent autonomous systems. This is probably due to the strong influence from the more traditional sciences (e.g., physics) that have essentially been limited to the study of causal systems which, in contrast to anticipatory systems, do not take knowledge of future states into account. We believe that autonomous systems with the ability to anticipate as described above would exhibit novel, interesting and possibly unexpected properties that might enhance the capacity of autonomous systems.

*Also affiliated with the University of Karlskrona/Ronneby, Department of Economics and Computer Science, S–372 25 Ronneby, Sweden

1.1 Background

In the last decade, there has been a growing interest in the field of autonomous agents* among computer scientists. The reason for this renewed interest is probably due to the emergence of a new paradigm radically different from the traditional one. Agents constructed following this new approach are often referred to as *reactive*, or *behavior-based*, agents.

According to the earlier approach, the two main components of an agent are the world model (i.e., an internal description of the agent's external environment) and some kind of planner. The working of traditional agents can be described as following a sense-model-plan-act cycle. The sensors sense the environment and produce sensor-data that are used to update the world model. The world model is then used by the planner to decide which actions to take. These decisions serve as input to the effectors that actually carry out the actions. However, when actually embodying these kinds of agents (e.g., Shakey[18]) it was noticed that even if they were able to do some "advanced" cognitive tasks such as planning and problem solving, they had problems with more basic tasks like routine reaction that require fast action but no extensive deliberation.[6]

As a reaction to the relative failure of this approach, the first reactive agents emerged in the mid-eighties. They were inspired by the idea that most of our every-day activities consist of routine action rather than of abstract reasoning. So, instead of doing world-modeling and planning, it was argued that the agents should only have a collection of simple behaviors that react to changes in the environment in a stimulus-response fashion. Some of the most influential agents of this kind are Brooks' robots based on the subsumption architecture,[5] Pengi,[1] and situated automata.[20] Probably the most controversial element of this new approach concerns the representation of knowledge. Brooks, in particular, argues that explicit representations of the world are not only unnecessary but also get in the way when implementing actual agents. Instead the agent should use "...the world as its own model — continuously referring to its sensors rather than to an internal world model."[5]

Behavior-based agents have been shown to be superior (in some situations) to traditional in doing a limited number of simple tasks in real world domains. However, in addition to not being particular versatile, they have problems with handling tasks that require knowledge about the world that must be obtained by reasoning or from memory, rather than perception. According to Kirsh[11] some possible candidates for such tasks are activities which require: (1) response to events beyond the agent's current sensory limits, (2) some amount of problem solving, (3) understanding a situation from an objective perspective, or (4) prediction of other agents' behavior.

There is also an evolution-based critique of the reactive paradigm that has seldom been expressed. It seems reasonable to compare reactive agents to, for example, reptiles in the sense that their cognition is based on stimulus-response behavior. For instance, Sjölander[22] writes: "There is thus no true intermodality in the snake, just a number of separate systems, involving specific behavioral patterns connected to a specific sensory input." Humans and other more developed species, on the other hand, are equipped with central representations and Sjölander suggests that this could be an explanation of why reptiles were superseded

*By autonomous agents we mean systems that are able to interact independently with its environment through its own sensors and effectors.

by mammals and birds. He concludes: "To go from monosensorially governed constructions of several internal representations to a centralized intermodal one must surely be one of the most important breakthroughs in the evolution of mind." This suggests that agents based on the reactive paradigm will never reach human-level performance.

Recently, several researchers have pointed out that an intelligent agent must have both high-level reasoning and low-level reactive capabilities.[8, 10, 14, 15, 17] The rationale behind this hybrid approach is to utilize the reaction ability of reactive agents necessary for routine tasks, but still have the power of planning necessary for more advanced tasks. In line with this we will suggest a novel way of combining reactive and deliberate agents into hybrid agents. This approach, called *anticipatory agents*, is based on the idea of anticipatory planning[3] which, in turn, is inspired by the concept of *anticipatory systems*.[19]

2 ANTICIPATORY AGENTS

According to Rosen[19], an anticipatory system is " ... a system containing a predictive model of itself and/or of its environment, which allows it to change state at an instant in accord with the model's predictions pertaining to a latter instant." (p.339) Thus, such a system uses the knowledge concerning future states to decide which actions to take in the present.

In more formal terms, the next state of an anticipatory system would be a function of past and future states:

$$s_{n+1} = f(s_1, s_2, ..., s_n, s_{n+1}, ..., s_k), \quad k > n$$

whereas a causal system only depends on past states:

$$s_{n+1} = f(s_1, s_2, ..., s_n)$$

However, since an agent cannot normally have true knowledge of future states,[†] it is, of course, not possible to implement an anticipatory system in this strict sense. The best we can do is to approximate such a system by using *predictions* of future states. Thus, we have:

$$s_{n+1} = f(s_1, s_2, ..., s_n, \hat{s}_{n,1}, ..., \hat{s}_{n,k-n}), \quad k > n$$

where $\hat{s}_{n,i}$ is the predicted value of s_{n+i}.

2.1 Computational Framework

In the suggested framework, an anticipatory agent consists mainly of three entities: an object system S, a meta-level component M, and a world model W. S is an ordinary (non-anticipatory) dynamic system. W is a description of the environment *including* S, but excluding M.[‡] M should be able to make predictions using W and to use these predictions to change the dynamic properties of S. Although the different parts of an anticipatory agent certainly are causal systems, the agent taken as a whole will nevertheless behave in an anticipatory fashion.

[†]An agent acting in a closed world and having a perfect world model, on the other hand, would be able to have true knowledge of future states.

[‡]The importance of having an internal model that includes both the agent as a part of the environment and (a large portion of) its abilities has been stressed by, for instance, Zeigler[25] and Kohout.[12]

To get this approach working, it is essential that the sequence of states of W are parameterized by a time variable that goes faster than real time. That is, if W adequately describes the environment (and S) at some time t_0, then after an arbitrary time interval Δt, W's sequence of states will have proceeded $t_0 + \Delta t$.

2.2 Implementational Issues

When implementing an anticipatory system, what should the different components (S, M, and W) correspond to, and what demands should be made upon these components? To begin with, it seems natural that S should correspond to some kind of reactive system similar to the ones described above. It must be a fast system in the sense that it should be able to handle routine tasks instinctively and, moreover, it should have an architecture that is both easy to model and to change.

M would then correspond to a more deliberative meta-level component that is able to "run" the world model faster than real time. When doing this it must be able to reason about the current situation compared to the predicted situations and its goals in order to, among other things, detect undesirable situations that cannot be handled by the reactive component. If such a situation is detected it should decide whether (and how) to change the reactive component or to issue commands directly to the effectors. There is a large body of work concerning different aspects of meta-levels, but none of this work seems readily applicable to the outlined approach. The closest is perhaps the studies on reflective architectures[16] and some works on meta-reasoning architectures in the context of autonomous agents.[13]

Since a world model is based on a representation, it can only approximately describe any given subset of the real world. Thus, since world models are abstractions of reality, we must decide on which level of abstraction is the most appropriate for the world model. Gat[9] argues that the debate concerning traditional agents versus reactive agents is really an argument about the proper use of a world model (i.e., internal state information). He writes that "... internal state should be maintained at a high level of abstraction and that it should be used to guide a robot's action but not to control these actions directly." Thus, local sensor information is necessary for the immediate control. He provides an example to show that this is probably also the way humans work: We are able to find our house and belongings because we have a world model at a high level of abstraction. We do not know the exact location of our house or our belongings, but we use sensor data to fill in the details that the world model does not provide. We thus suggest that W should be on a rather high level of abstraction. Moreover, there is certainly a trade-off between how detailed the world model is and the accuracy of the predictions made using it; the higher the level of abstraction, the more accurate (i.e., probable) are the predictions.

It is important to make a distinction between two parts of the world model.[21] We have the *environment model* that describes the current state of the agent's surroundings (i.e., some kind of map), and the rest of the world model that includes more general knowledge about other possible states and ways of achieving these states. An environment model is often some kind of 3-D model, typically used for navigation tasks. It contains dynamic and situation-dependent knowledge, whereas a world model contains more stable and general knowledge, such as knowledge about objects, properties of objects, relationships between objects, events, processes, and so on.

The resulting architecture is illustrated in Figure 1. (It is also possible to regard W as a part of M.) To summarize, the sensors receive input from the environment. This data

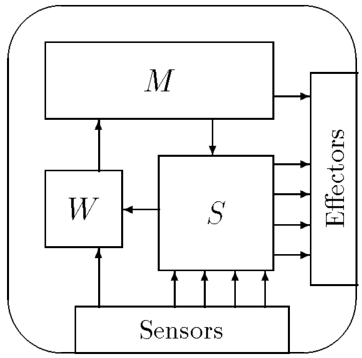

Figure 1: The basic architecture of an anticipatory agent.

is then used in two different ways: (1) to update W and (2) to serve as stimuli for S. S reacts to these stimuli and provides a response that is forwarded to the effectors, which then carry out the desired action(s). Moreover, M uses W to make predictions and on the basis of these predictions M decides if, and what, changes are necessary. Either it commands the effectors directly to carry out some action(s), or it make some changes on the dynamic properties of S. Every time S is modified, W should, of course, be updated accordingly. Thus, the working of an anticipatory agent can be viewed as two concurrent processes, one reactive at the object-level and one more deliberative at the meta-level.

2.3 Learning

Since an autonomous agent is supposed to act in environments that are continually changing, it is important that it can adapt its behavior according to these changes. In the approach suggested there are two ways of making such adaptions: (1) by changing S, and (2) by changing W.

As mentioned earlier, one way of changing S is to let M modify it directly based on M's predictions. It is, however, also desirable that S itself has a learning capacity for acquiring low-level stimuli-response schemas. This kind of learning seems closely related to what has been studied under the label of *reinforcement learning*.[23] In addition, by letting M control the training, the suggested framework seems to have a potential for performing this kind of learning in a more systematic manner than in existing systems.

As M will rely heavily on W, it is important that both the environment model and the world model are continually updated as the environment changes. While a considerable amount of work has been done on building and maintaining environment models,[2, 21] less research has been carried out on updating the more general parts of a world model.

2.4 Perception

Even though the same sensors may be used both to update W and to provide stimulus for S, they might be processed in rather different ways. The updating of W would probably rely on the methods from the traditional computer vision (i.e., perception) paradigm where the goal is to construct a detailed symbolic representation of the world independent of the tasks under consideration.

For S, as we have hinted earlier, it is probably a better choice to follow some reactive approaches[1, 5] and rely on task-driven perception, often referred to as *purposive* vision. The goal for this kind of vision is to solve particular visual tasks. To do this it is necessary to develop "fast visual abilities which are tied to specific behaviors and which access the scene directly without intervening representations."[4]

3 RELATED RESEARCH

The only experiment on computer-based anticipatory systems known to the authors has been carried out by Tsoukalas.[24] He applied the approach to the diagnosis and control of a nuclear reactor. His work is, however, not directly comparable to our framework, mainly because it is not addressing autonomous agents. In addition, he makes several simplifications. For instance, the system has no potential for updating its world model or for learning in general. Moreover, his system has no explicit meta-level component which we regard as a highly important feature for achieving a more powerful and flexible system.

However, some systems in traditional AI-planning and control theory can, although in a limited sense, be said to show anticipatory behavior. More closely related to our notion of anticipation is some work done within the field of biology.

3.1 Anticipation in Traditional Agents

It can be argued that the planning component in a traditional system provides some amount of anticipation by internally testing different actions (or sequences of actions) to see which future states can be achieved. This kind of anticipating is, however, limited in the sense that it covers only state changes that are caused by the pursuing of a particular goal. Changes that have other causes, such as other agents or independent physical processes, are typically not taken into account. The present framework, on the other hand, has the potential to take into consideration all kinds of changes in the environment, and in particular, it will be able to detect unwanted situations in advance. It is a kind of passive and general anticipation rather than the active and goal-directed anticipation of planning systems.

Goal-directed anticipation can (will) in some cases also be performed by an anticipatory agent. For example, when an unwanted situation has been detected implying that M must change the properties of S, or when the agent actively pursues a goal, M will be forced to perform reasoning related to traditional planning.

3.2 Anticipation in Control Systems

Some control systems use a method called *feedforward* to eliminate disturbances that can be measured. According to Åström and Wittenmark[26], "the basic idea is to use the measured disturbance to anticipate the influence of the disturbance on the process variables and to introduce suitable compensating control actions." (p.166) To do this, the feedforward method requires an adequate model of the process to be controlled.

Thus, the feedforward method resembles the suggested approach in that a model is used to predict future states and that these predictions are used to guide the behavior of the system. However, the anticipation is limited to the prediction of just one future state (i.e., the error in output that the disturbance would have caused) whereas anticipatory systems are able to make predictions of arbitrary "time-steps" into the future. Moreover, the process to be controlled is characterized by just a single numerical value, whereas in autonomous agent contexts there is a need for qualitatively more powerful models (i.e., symbolic descriptions).

3.3 Anticipation in Biological Systems

It has been suggested that anticipation, in addition to causal relationships, should be taken into account when analyzing biological phenomena.[7] Sjölander[22] provides an example of such an analysis: a dog hunting a rabbit uses its internal world model to predict the rabbit's future positions. These predictions can then be used to focus the attention on the relevant aspects of the situation. This anticipatory behavior can also help in recognition tasks. For example, the dog need not see the whole rabbit all of the time. Since it can predict the rabbits current position, it needs only glimpses of parts of the rabbit to confirm its predictions. To make this possible within the suggested architecture requires, of course, that M is able to control the sensors.

4 CONCLUSIONS AND FURTHER RESEARCH

As should be clear from the above, a shift of paradigm seems desirable for the further development of autonomous agents. We believe that the framework of anticipatory agents is a promising candidate for such a paradigm.

A prototype of an anticipatory agent is currently being developed. While initial experiments, where very simple reactive components have been used, have shown some promising results, we need to test more sophisticated reactive components in order to evaluate the approach. Moreover, we have in this paper ascribed many different functions to M, such as using the world model to make predictions and interpret the results; modifying the dynamical properties of S; issuing commands directly to the effectors; controlling learning and experimentation, and controlling the sensors. Since we have only studied most of these functions superficially, further studies of these topics are necessary.

References

1. P.E. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In *AAAI-87*, pages 268–272, 1987.

2. M. Asada. Map building for a mobile robot from sensory data. *IEEE Trans. on Systems, Man, and Cybernetics*, **20**(6):1326–1336, 1990.

3. E. Astor, P. Davidsson, B. Ekdahl, and R. Gustavsson. Anticipatory planning. Technical Report LU–CS–TR: 90–69, Dept. of Computer Science, Lund University, Lund, Sweden, 1990. Also in Advance Proceedings of the European Workshop on Planning, 1991.
4. M.J. Black et al. Action, representation, and purpose: Re-evaluating the foundations of computational vision. In *IJCAI-93*, pages 1661–1666, 1993.
5. R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, **47**(1):139–159, 1991.
6. R.A. Brooks. New approaches to robotics. *Science*, **253**:1227–1232, 1991.
7. J.M. Burgers. Causality and anticipation. *Science*, **189**:194–198, 1975.
8. I.A. Ferguson. Touring machines: Autonomous agents with attitudes. *IEEE Computer*, **25**(5):51–55, 1992.
9. E. Gat. On the role of stored internal state in the control of autonomous mobile robots. *AI Magazine*, **14**(1):64–73, 1993.
10. J. Hendler. Abstraction and reaction. In J. Hendler, editor, *AAAI Spring Workshop on Planning in Uncertain, Unpredictable or Changing Environments*, Stanford, 1990.
11. D. Kirsh. Today the earwig, tomorrow man? *Artificial Intelligence*, **47**(1):161–184, 1991.
12. L. Kohout. *A Perspective on Intelligent Systems.* Chapman and Hall Computing, 1990.
13. D.R. Kuokka. MAX: A meta-reasoning architecture for "X". *SIGART*, **2**(4):93–97, 1991.
14. L.M. Lewis. A time-ordered architecture for integrating reflexive and deliberative behaviour. *SIGART*, **2**(4):110–114, 1991.
15. D.M. Lyons, A.J. Hendriks, and S. Mehta. Achieving robustness by casting planning as adaption of a reactive system. In *IEEE Int. Conf. on Robotics and Automation*, 1991.
16. P. Maes and D. Nardi. *Meta-Level Architectures and Reflection.* Elsevier Science, 1988.
17. T.M. Mitchell. Becoming increasingly reactive. In *AAAI-90*, pages 1051–1058, 1990.
18. N.J. Nilsson. A mobile automaton: An application of artificial intelligence techniques. In *IJCAI-69*, pages 509–520, 1969.
19. R. Rosen. *Anticipatory Systems – Philosophical, Mathematical and Methodological Foundations.* Pergamon Press, 1985.
20. S. Rosenschein and L. Kaelbling. The synthesis of digital machines with provable epistemic properties. In J.F. Halpern, editor, *The 1986 Conference on Theoretical Aspects of Reasoning and Knowledge*, 1987.
21. Y. Roth-Tabak and R. Jain. Building an environment model using depth information. *IEEE Computer*, **22**(6):85–90, 1989.
22. S. Sjölander. Some cognitive break-throughs in the evolution of cognition and consciousness, and their impact on the biology of language. *Evolution and Cognition*, 1994. In print.
23. R.S. Sutton. (ed.), Machine Learning, **8**(3/4), Special Issue on Reinforcement Learning, 1992.
24. L.H. Tsoukalas. *Anticipatory Systems Using a Probabilistic-Possibilistic Formalism.* PhD thesis, Dept. of Nuclear Engineering, University of Illinois at Urbana-Champaign, 1989.
25. B.P. Zeigler. *Object-Oriented Simulation with Hierarchical, Modular Models – Intelligent Agents and Endomorphic Systems.* Academic Press, 1990.
26. K.J. Åström and B. Wittenmark. *Computer Controlled Systems: Theory and Design.* Prentice-Hall, 1984.