► Latent Learning
► Machine Learning
► Online Learning
► Reinforcement Learning
► Schema-Based Architectures of Machine Learning
► Supervised Learning

## References


Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *The Behavioral and Brain Sciences, 27*, 377–396.

Hoffmann, J. (2003). Anticipatory behavioral control. In M. V. Butz, O. Sigaud, & P. Gérard (Eds.), *Anticipatory behavior in adaptive learning systems: Foundations, theories, and systems* (pp. 44–65). Berlin: Springer.

Jeannerod, M. (2006). *Motor cognition*. Oxford: Oxford University Press.

Pezzulo, G., Butz, M. V., Castelfranchi, C., & Falcone, R. (Eds.). (2008). *The challenge of anticipation: A unifying framework for the analysis and design of artificial cognitive systems, LNAI 5225*. Berlin: Springer.

Stock, A., & Stock, C. (2004). A short history of ideo-motor action. *Psychological Research, 68*, 176–188.

Tolman, E. (1932). *Purposive behavior in animals and men*. London: Century/Random House UK.

Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks, 11*, 1317–1329.


# Anticipatory Learning Mechanisms


Filipo Studzinski Perotto
Constructivist Artificial Intelligence Research Group, Toulouse, France


## Synonyms

Adaptive systems; Constructivist agents; Predictive model learning algorithms

## Definition

In the artificial intelligence domain, *anticipatory learning mechanisms* refer to methods, algorithms, processes, machines, or any particular system that enables an autonomous agent to create an anticipatory model of the world in which it is situated. An *anticipatory model of the world* (also called *predictive environmental model*, or *forward model*) is an organized set of knowledge allowing inferring the events that are likely to happen. For cognitive sciences in general, the term *anticipatory learning mechanism* can be applied to humans or animals to describe the way these natural agents learn to anticipate the phenomena experienced in the real world, and to adapt their behavior to it.

## Theoretical Background

When immersed in a complex universe, an agent (natural or artificial) needs to be able to compose its actions with the other forces and movements of the environment. In most cases, the only way to do so is by understanding what is happening, and thus by anticipating what will (most likely) happen next. Therefore, a predictive model of the world can be very useful to an agent as a tool to guide its behavior; the agent has a perception of the current state of the world, and it decides what actions to perform according to the expectations it has about the way the situation will probably change.

Butz et al. (2003) have made a comprehensive review of the different forms of anticipatory behavior present in artificial intelligence, including simple implicit anticipatory systems (in which the anticipatory behavior is programmed into the agent as useful predefined action rules), payoff anticipatory systems (in which the agent estimates the future rewards), and model-based anticipatory systems (in which the agent effectively anticipates the world transformations using a predictive environmental model). Very often, these artificial anticipatory systems are inspired by real adaptive mechanisms found in nature.

Nevertheless, even if most of researchers in artificial intelligence agree that anticipation is a fundamental characteristic of intelligent behavior, there is no consensus on what kind of model a "strong" intelligent agent should possess, much less on how it can learn it in a feasible way. Moreover, for a wide range of complex or realistic problems, it is very hard to provide the agent with a complete model of the world in advance; the agent has no alternative but to incrementally learn how the universe evolves, from its own experiences, in order to adapt its behavior to it. This is the advantage of being endowed with an anticipatory learning mechanism.

The necessity of such a mechanism is more evident when the agent is fully situated and completely autonomous; that means, when the agent is by itself, interacting with an unknown, dynamic, and complex world, through limited sensors and effectors, which give it only a local point of view of the state of the universe and only partial control over it. In other words, the agent is not omniscient (it is not aware of the complete state of the universe), and is not omnipotent (it is just one among other possible sources of perturbation affecting the environment). In this case, it is very hard to predefine static solutions (like automatic behaviors) designed to deal with all possible situations the agent can face throughout its existence.

An autonomous and situated agent is necessarily self-motivated; it is a creature that has goals. Sometimes these goals are implemented as explicitly defined states to be reached or specific tasks to be accomplished; but in general, the agent is just motivated by sporadic reward signals, or intrinsically evaluative sensations that it wants to experience or avoid (like pleasure and pain). In any case, predefined reactive behaviors can properly work only in a restricted set of problems where the important variables are fairly known and controllable. The remaining problems can only be successfully faced by cognitive agents, who will be compelled to discover the regularities that govern the universe, understand the causes and the consequences of the phenomena, identify the forces that influence the observed changes, and especially master the impact of its own actions over the ongoing events. So, in the machine learning community, it is common to consider two subproblems: on the one hand, the construction of a predictive model of the world (i.e., structured knowledge that allows the agent to anticipate the environment dynamics); on the other hand, the definition of a policy of actions (i.e., a behavioral strategy that guides the agent in its plans and decisions according to its objectives). Generally, for a situated agent, there is no separate training phase; the learning mechanism needs to create both the model of the world and the policy of actions online (while the agent is already performing its activities).

## Important Scientific Research and Open Questions

Over the last 20 years, several anticipatory learning mechanisms have been proposed in the artificial intelligence scientific literature. Even if some of them are impressive in theoretical terms, having achieved recognition from the academic community, for real-world problems (like robotics) no general learning mechanism has prevailed. Until now, the intelligent artifacts developed in universities and research laboratories are far less wondrous than those imagined by science fiction. On the other side, neuroscientists, psychologists, and philosophers have been working hard to try to explain how intelligence works, in particular how animals and humans learn, and how things are modeled and represented in their brains and minds. Even if some important findings have been done, we are still far from being able to explain in detail the main part of intelligent processes, and, in the current state of the art, we are not able to present a complete and definitive model neither of the intelligence in general, nor of the faculty of learning in particular. Within the artificial intelligence community, it is possible to highlight at least four lines of research more or less explicitly related to the conception of anticipatory learning mechanisms: constructivist AI, automata learning, model-based reinforcement learning, and anticipatory classifiers systems.

Drescher's book (1991) can be considered the first impacting work published on the subject of constructivist models. He presented the *schema mechanism*, an algorithm conceived to reproduce in machine some aspects of the human cognitive development as described by ► Piaget's learning theory, representing anticipatory knowledge as (computational) ► schemas, in the form [context] + [action] → [result], similarly to the classical Fikes and Nilsson's *STRIPS* system. The schema mechanism inaugurated an interesting line of research called *constructivist artificial intelligence*. After Drescher, some other authors tried to follow the same way proposing a variety of constructivist learning mechanisms, often focused on *abstract concept creation* (i.e., how the agent can develop its own representational vocabulary beyond its basic sensorimotor signals). Guerin (2011) present a good review about these algorithms, including Chaput's *CLA*, Holmes and Isbell's *PST*, and Perotto, Buisson, and Alvares's *CALM*.

The *automata learning* research community also played an important role in the development of model-based anticipatory learning algorithms. The problem of finding the structure of an automaton (a finite-state machine) from examples is similar to that in which an agent has to learn a model of the environment from the observation. Another essential reference is the *reinforcement learning* research community (in AI), concerned with the *decision-making* problem. Reinforcement learning algorithms are generally designed to estimate the utility of state-actions pairs, and to establish a policy of actions to maximize the rewards received by the agent over time. This problem is popularly modeled as a *Markovian decision process*. The classical MDP model is represented as a state machine; at each time step, the machine is in some state *s*, and the agent may choose some action *a* to carry out; at the next time step, according to some (nondeterministic) transition function, the process changes into a new state *s*', giving the agent a corresponding reward *r*. This formalism has been extended to deal with partial observability; in this case, the agent does not know *s*, only perceiving an observation *o*, which works as an indirect and incomplete indication to the underlying state of the process. Several algorithms have been proposed to solve MDPs and POMDPs (i.e., to find the optimal or near-optimal policy, to maximize the average or cumulative discounted reward over time), and a good overview about them can be found in the Feinberg and Shwartz's book (2002).

Another important line of research related to anticipatory learning mechanisms was generated within the *evolutionary computation* (or genetic algorithms) community, from where the *Anticipatory Behavior in Adaptive Learning Systems* conference series emerged. Sigaud et al. (2009) present the *anticipatory learning classifier systems* framework, including representative algorithms like Stolzmann's *ACS*, Butz's *ACS2* and *XACS*, and Gerard's *YACS* and *MACS*, comparing it with other related models. In recent years, a convergent movement of all these research branches toward the use of factored MDPs have been noticed; in a factored MDP the state space is decomposed into a set of variables or properties, which permits to avoid an exhaustive enumeration of states.

Thus, an MDP can be extended to become at the same time factored and partially observable, and it is so called FPOMDP. In order to be factored, the original set of states $S$ is decomposed and replaced by the set $X = \{X_1, X_2, \ldots X_n\}$ of properties or variables; each property $X_i$ is associated to a specified domain, which defines the values the property can assume. Furthermore, in order to be partially observable, the set $X$ is divided into two subsets, $X = P \cap H$, where the subset $P$ represents the observable properties (those that can be accessed through the agent sensory perception), and the subset $H$ represents the hidden or non-observable properties. The set $C = \{C_1, C_2, \ldots C_m\}$ represents the controllable variables, which compose the agent actions; $R = \{R_1, R_2, \ldots R_k\}$ is the set of (factored) reward functions, in the form $R_i : P_i \rightarrow \mathbb{R}$; and $T = \{T_1, T_2, \ldots T_n\}$ is the set of transformation functions, in the form $T_i : X \times C \rightarrow X_i$, defining the system dynamics (which can be nondeterministic).

When the agent is immersed in a system represented as a FPOMDP, the complete task for its anticipatory learning mechanism is both to model the transformation function and to define a sufficiently good policy of actions. The transformation function can be described in the form of a *dynamic Bayesian network*, i.e., an acyclic, oriented, two-layers graph, where the first layer nodes represent the environment situation in time *t*, and the second layer nodes represent the next situation, in time $t + 1$. A stationary policy $\pi : X \rightarrow C$ defines the action to be taken in each given situation in order to optimize the rewards received by the agent over a potentially infinite time horizon. Certain algorithms create stochastic policies, and in this case the action to take is defined by a probability.

Degris and Sigaud (2010) present a good overview of the use of this representation in artificial intelligence, referring several related algorithms designed to learn and solve FMDPs and FPOMDPs, including both the algorithms designed to calculate the policy given the model (like Boutilier's *SVI* and *SPI*, Hoey, St-Aubin, Hu, and Boutilier's *SPUDD* and APRICODD, Guestrin, Koller, and Parr's *FALP* and *FAPI*, Poupart's *VDCBPI*, Sim and Kim's *SHSVI*, and Shani, Brafman, and Shimony's *FSVI*) and the algorithms designed to discover the structure of the system (like Degris and Sigaud's *SDYNA*, *SPITI*, and *UNATLP*, Strehl, Diuk, and Littmann's *SFL*, and Jonsson and Barto's *VISA*).

Despite the growing interest in anticipatory learning mechanism within the artificial intelligence community, some questions have not yet been convincingly answered. How can an agent enrich its perception with high-level, conceptual, or abstract understanding? How can it consistently solve the exploration–exploitation dilemma (find the good compromise between exploring new possibilities in order to learn new things, and taking profit of the knowledge already learned)? How can the agent correctly identify the relevant properties of the situations and discover the causal relations of the world? How can it efficiently deal with continuous, nondeterministic, nonstationary, noisy, large, and complex universes (like most of real-world problems)? These important questions still remain open.

## Cross-References

- ▶ Adaptation and Anticipation: Learning from Experience
- ▶ Adaptive Learning Systems
- ▶ Anticipatory Learning/Anticipation and Learning
- ▶ Anticipatory Schema(s)
- ▶ Belief-Based Learning Models
- ▶ Computational Models of Human Learning
- ▶ Developmental Robotics
- ▶ Incremental Learning (Definition)
- ▶ Learning Algorithms/Machine Learning
- ▶ Mental Models of Dynamic Systems
- ▶ Piaget´s Learning Theory

## References


Butz, M. V., Sigaud, O., & Gérard, P. (2003). Internal models and anticipations in adaptive learning systems. In M. V. Butz, O. Sigaud, & P. Gerard (Eds.), *Anticipatory behavior in adaptive learning systems: Foundations, theories, and systems* (Vol. LNAI 2684, pp. 86–109). Berlin-Heidelberg: Springer.

Degris, T., & Sigaud, O. (2010). Factored markov decision processes. In O. Buffet & O. Sigaud (Eds.), *Markov decision processes in artificial intelligence*. Vandoeuvre-lès-Nancy: Loria.

Drescher, G. (1991). *Made up minds: A constructivist approach to artificial intelligence*. Cambridge, MA: MIT Press.

Feinberg, E. A., & Shwartz, A. (2002). *Handbook of markov decision processes: Methods and applications*. Norwell: Kluwer.

Guerin, F. (2011). *Learning like baby: A survey of AI approaches. Knowledge engineering review*, 26, Cambridge. (to appear).

Sigaud, O., Butz, M. V., Kozlova, O., & Meyer, C. (2009). Anticipatory learning classifier systems and factored reinforcement learning. In G. Pezzulo, M. V. Butz, O. Sigaud, & G. Baldassarre (Eds.), *Anticipatory behavior in adaptive learning systems: From psychological theories to artificial cognitive systems* (Vol. LNAI 5499, pp. 321–333). Berlin-Heidelberg: Springer.


# Anticipatory Schemas


Norbert M. Seel
Department of Education, University of Freiburg, Freiburg, Germany


## Synonyms

Anticipatory behaviour; Schema-based expectations

## Definition

Anticipatory schemas direct how adaptive learners (such as humans and animals) explore their environments. What happens to them is rarely completely unexpected. Actually, one important distinction among schema-based architectures is their reactive or anticipatory nature. Reactive schema-based architectures of cognition assume quick and automatic responses to dynamic environments whereas anticipatory schema-based architectures include anticipatory mechanisms, which generate and exploit expectations about the next stimuli to be processed. These anticipatory aspects are inspired by psychological theories of action control, indicating that anticipated effects of (possible) actions play a fundamental role in regulating the agent's behavior.

## Theoretical Background

Research on anticipatory schemas (and behavior, respectively) in adaptive learning systems gains increasingly more recognition and appreciation in various research disciplines, such as cognitive psychology, neuroscience, linguistics, artificial intelligence, machine learning, robotics, and others. However, the idea of anticipatory schemas can be traced back to the seminal work on productive thinking from Otto Selz in 1913. It is strongly associated with *complex completion* due to an activated schema. This can be illustrated with the example of a word schema.

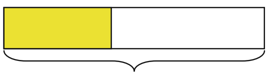


"The consciousness of the particular word to be found transcends from the consciousness of not specified word to the consciousness of a word which begins with 'Mel —'. We should conceive the development of this consciousness in a way that an empty schema of a concrete word will be filled through the insertion of a spoken series of phonemes at its beginning, i.e. a combinatory process" (Selz 1913, p. 113 f.).

The development of anticipatory schemas facilitates the completion of a complex, i.e., and is based on three regularities:

- A given totality, which functions as a part of a complex, tends to initiate the reproduction of the entire complex.