Not Everything We Know We Learned

Mihai Nadin

Program in Computational Design,
University of Wuppertal, Germany;
President, MINDesign, USA/Germany
nadin@acm.org

Abstract. This is foremost a methodological contribution. It focuses on the foundation of anticipation and the pertinent implications that anticipation has on learning (theory and experiments). By definition, anticipation does not exhaust all the forms through which the future affects human activity. Accordingly, guessing, expectation, prediction, forecast, and planning will be defined in counter-distinction to anticipation. The background against which these distinctions are made is explicit in the operational thesis advanced: Anticipation and reaction can be considered only in their unity. The interrelation of anticipation and reaction corresponds to the integrated nature of the physical and the living. Finally, an agent architecture for a hybrid control mechanism is suggested as a possible implementation.

Context and Reference

Einstein [1] observed that, "No problem can be solved from the same consciousness that created it. We must learn to see the world anew." As my own work in anticipatory computing evolved, I have constantly faced attitudes varying between skepticism and sheer enmity from outside the small community of researchers dedicated to the study of anticipation. Every example of anticipation my colleagues or I advanced was whittled down to the reactive explanations of the deterministic cause-and-effect sequence, that is, to a particular form of causality. Even the reviewers of this paper could not agree among themselves on a line of argument that in the final analysis suggests that there is more to causality than what the Cartesian rationality that we learned in school and have practiced since then preaches. (A good source of information on this topic is http://www.culture.com.au/brain_proj/Descartes.htm.) Given this difficult, but not surprising situation, I shall proceed in a more didactic manner than I would otherwise be inclined. The intention is to clarify terminology before working with the concepts—a practice that Charles S. Peirce [2] defined as the "ethics of terminology," a prerequisite of any scientific endeavor.

From among the many definitions of anticipation advanced since the pioneering work of Robert Rosen [3, 4, 5] and my own early work [6], I would like to focus on the following operational definitions:
1.   An anticipatory system is a system whose current state is defined by a future state. Eventually, a change was introduced in this definition [7] so to imply that in

addition to the future state, a current or even past state could affect the state of the system (cf. Definition 9 in [8]).

2. 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 later instant, in faster than real time (cf. Definition 3 in [8]).

3. Anticipation is the result of the competition among a number of mind models. Reward mechanisms explain the dynamics of this competition [cf. Definition 2 in [8]; see also [9]).

In the end, I adopted the viewpoint according to which anticipation is a characteristic of the living. At the foundation of this perspective lie the work of Rosen, especially [4] and [5] and Elsasser [10]. Moreover, in discussing the neural basis of deciding, choosing and acting, Jeffrey D. Schall [11] distinguishes between external forces (explaining the "movements of physical bodies, such as rocks") and "reasons" explaining "many human movements" (actions directed towards a goal). Anticipatory behavior—such as the recognition of the prey's trajectory—is not the only expression of anticipation. We can mention design, creative activities, and conceptual elaborations as pertaining to the rich forms expressing anticipation.

In respect to learning, it is rather difficult to limit oneself to a reduced number of definitions. The richness of the forms of learning and of the expressions of learning makes the attempt tenuous at best. However, given the focus of this volume— adaptive learning systems—a first pruning of the rich tree of learning definitions becomes possible.

Computational learning theory (COLT) focuses on "the design and analysis of algorithms for making predictions about the future based on past experiences" (cf. Freund and Schapire, www.learningtheory.org). This definition falls within the reactive paradigm. An applied definition, originating in studies in time series prediction [12] sees learning as an emulation of the structure of time series, i.e., of the dynamics of processes we intend to understand, control, and automate. In this case, the knowledge about the unfolding in time of a process is eventually used in order to produce some desired behavior or to avoid undesired behavior.

Lastly, for the purpose of this paper, learning as the pursuit of regularities or patterns in processes involving the physical/inanimate as well as the living is of special interest to me. This extraction (in forms such as data-mining, perception, knowledge mapping, etc.), by humans, by machines, or by hybrid entities, of regularities or patterns can be subjected to further analysis, i.e., further learning. Learning is a necessary condition of adaptive processes in the living, in machines, and in hybrid systems. In what follows, this methodological section remains the only reference.

Not everything we know was learned. This holds true whether "knowledge is seen as socially situated or whether it is considered to be an individual extraction (cf.

Ernest [13]). Taken at face value, it maintains that the process of knowledge acquisition is complemented by knowledge production. Moreover, in addition to what we learn, there is innate knowledge—that space that "has to exist before data," as it was defined by some researchers (cf. Novak *et al* [14]) focusing on language and its active role in learning. There are genetically defined processes (such as language acquisition, (cf. Nadin [15] pp. 77ff) or seeing [15] pp. 321ff). But there are also anticipatory processes through which not only *what is* is acknowledged, but also what *might possibly be* is generated (cf. [10] p. 5), i.e., the principle of creative selection. My focus in this study is on particular processes, which can be defined as anticipatory, through which the complementarity between learning and the activity of knowledge production is accomplished. With this focus in mind, I shall refer to experimental evidence. (More scientific reports based on such evidence are available at www.anticipation.info) Evidence from experiments does not in itself replace a broader understanding. What examples can do, especially when associated with data resulting from experiments, is to make us aware of an unusual or unexpected outcome of some process.

Rich Data

Take something as simple as fluctuations in the heartbeat. Data collected in endless physiological observations made doctors aware of such fluctuations. Patients experience them as discomfort. Sometimes we know what they mean. A good diagnostician can distinguish between "normal" irregularities and others related to a condition requiring treatment. E.F. Adolph [16], following Walter B. Cannon [17], deals with self-regulation mechanisms. Change in posture—standing up from a seated position, for example—would cause changes in blood pressure. This is the "physics" of the body: a liquid (blood), pipes (the various blood vessels), a pump (the heart). That this image is an oversimplification, rooted in the rationality of "the human being as machine," will soon become clear. But it serves as a good point of departure. We can understand the mechanism by taking a rubber tube with liquid inside and manipulating it. The pressure of the fluid varies as we change its position from horizontal to vertical. If human beings were subjected to blood pressure variations as often as we change posture, we all would have a terrible time. Dizziness is only one of the symptoms accompanying variations in blood pressure. It appears that a certain reflex, called the *baroreceptor reflex* [18], affects the heart rate via the nervous system in order to maintain the blood pressure within relatively constant and safe limits. How does the heart "know" about a change in position? Obviously, here the analogy to the simple physical system of a machine, with liquid, pipes, and pump ends. The question asked—how does the heart, or whatever controls the heart, know what the body will do—goes beyond the physics of pumping a liquid through pipes (cf. [12] p. 107).

Almost all of us have been educated in the rationalist spirit of a scientific perspective based on the analogy between an organism and a machine. The most

sophisticated adepts of this view will argue that the machine to which the human being can be reduced is probably more complicated than the pump, pipes, and perhaps some controller—but ultimately still a machine. This is the reductionist viewpoint. It maintains that the human being can be reduced to something easier to understand and describe; moreover, that there is some equivalence between the human being and the machine. Reductionism has been practiced along this line of thinking for centuries and, as such, appears to be a valid explanatory model. What animates scientists is the rational thought that what applies to a short time sequence might apply to longer sequences. Beneath this optimistic outlook lie the accepted premises of physical homogeneity, the repetitive nature of physical change (described through the laws of physics), and, of course, reductionist thought.

One of the goals of this text is to show that taking this assertion as a premise can help us satisfactorily explain how humans react to the world in which they live. It also explains, partially, what we can learn in association with experiencing the world as one of action and reaction. However, it cannot explain the example I started out with—why does change in human posture not result in a change in blood pressure. More precisely, it cannot explain the anticipation on whose basis the heart rate compensates for the change in posture *in advance* in order to avoid variations in blood pressure.

The reason for this is that in accepting the reduction of the human being to a machine—no matter how complicated and complex this machine—we simultaneously accept the deterministic premise of the sequential link between cause and effect. And once this is accepted, a time sequence beginning with the cause and continuing to the effect is established. The example of anticipation given above does not fit into this scheme, although it does not preclude causality. The same example suggests that reaction-based learning—learning from something we experienced—is complemented by non-reactive learning, which we should in some ways account for. The simplest example is learning through acknowledging correlations, some of them rather difficult to describe [19] [20].

More examples: Someone approaches with the intention of tickling you. Before the person even touches you, you start laughing. How come? In this case, effect and cause do not seem to follow the deterministic sequence. The effect takes place before the cause (cf. [21], [22]). At work in such situations is the intrinsic dynamics of previous experiences, the learned reflex triggered through association. Speech, a very complicated motoric act (ca. 100 muscles in the face, chest, and throat need to be coordinated [23]) takes place on account of an internal model that "knows" the word before it has been uttered [24]. That we learn part of speaking through imitation is well known. But we also discover speaking as an act of self-definition of the speaking person beyond the learning component: We speak the same language but in so many different ways! Motoric acts not rooted in the cause-and-effect sequence ([25], [26]), i.e., in anticipation of a possible action, are not learned but are rather expressions of choices leading to eventual learning.

In various situations—such as crossing a street or catching a falling object—or when we so often "know" (or do we guess?) what words will follow as someone speaks to us, a joke's punch-line, someone's age, whose footsteps we hear long after we last had contact with the person—something unusual takes place. The event (a car approaches, an object we were unaware of falls, something we surmised would happen, does) to take place in the future—i.e., a *future state*, as it is known in science—affects our immediate behavior—which scientists call *current state*. To be adventurous—a definite means of survival, that is, an evolutionary characteristic—means to be in anticipation of the possible outcome, of the adventure. This means to be in anticipation of what one might eventually learn in the adventure [27], [28].

In regard to all these examples, we can raise the issue of relevance: Is anticipation learned? Can our descriptions of anticipatory processes form a valid model for a theory of learning? Or is learning, as learning theorists would probably be inclined to think, a prerequisite for anticipation? A dogmatic leaning in either direction will actually prevent us from building a coherent explanatory framework, and thus affect our attempts to apply the knowledge gained to practical purposes.

The introductory methodological part advanced, as reference, the notion—by no means unanimously accepted—that to learn means to have an object of learning, an entity subject to our interest (no matter in which form the interest is expressed), and a learner (be this a person, a program, a procedure, etc.). Anticipation is defined as a dynamic characteristic unfolding in a different timeframe, that is, before the object of our interest, and sometimes (as in illusory anticipation) instead of it. This difficult idea of a cause that lies in the future (cf. von Foerster [29]), in the absence of which anticipation becomes meaningless, will be further explained here.

The "Ante" in Anticipation

What all the abovementioned examples have in common is the fact that they question the dominant explanatory model to which we are accustomed, and which we automatically apply in our thinking: There is a cause (in the past or present) that leads to an effect (now or later). This sequence describes determinism in its simplest form. It is also a simple description of causality in its simplest form. (For more on causality, see [30].) But the heartbeat mechanism given as an example changes (the effect) in advance of any modification of someone's in posture (the cause) within a complex dynamic system of interlocking variables. No one has yet tickled us (cause), but we start laughing (effect). Learning takes care of some of the abilities needed, anticipation takes care of others.

We do not compute as we speak (cf. Guenther [24]). We ourselves do not compute data as we drive through heavy traffic (cf. Zadeh [31]). The hand catches the falling object before the person realizes that it's coming. Good reflexes, some say, always amazed that the hand seems to have a mind of its own. And again: there is no

computation involved, rather an action corresponding to a possible situation that was not triggered by sensory data. We "guess" so many things (just as we fail to guess others) without realizing how we do it and without being aware that we do it. The outcome precedes or predates the action! The next word in a sentence (spoken or written), the punch-line of a joke, the suspect in a crime novel or movie—in each of these cases there is a realization before what we would be inclined or biased to expect. The logic behind this looks like deductive thinking: from something we know to something else that is its consequence. In some cases we might indeed find that a deductive step is part of the anticipation. Reading a crime novel, not unlike solving a puzzle, involves deduction. However, if a merely deductive process lay behind anticipation, machines would be able to read a crime novel with the same pleasure—a qualifier of high complexity—a human being does, and to produce a summary of the novel. The emphasis is on the production of a summary, a task that any machine is anywhere near accomplishing, although we know by now how to program machines for executing deductive tasks. The ingredients of experiencing pleasure might be several—to take a logical argument—but all of them have to rely on anticipation, since the only source of tension is between what we read vs. what is anticipated. Creativity (cf. [10] p. 146) is intrinsic in the dynamics of the living. Such machines would listen to a concert in which performers and listeners were challenging each other with their respective anticipations. The word guessed in a dialog might be deducted from the flow of succeeding words. At times it might make no logical sense, but rather result from a context in which many factors unrelated to what is uttered are involved. Whether deductions are possible or not, what we deal with is always a temporal sequence from the future (as one possibility among many) to the present. This "before" interval, this *ante* (the Latin prefix that denotes "before") can be and has been repeatedly measured, documented, discussed, and questioned.[1]

The description of the cause-and-effect sequence, often understood as action-reaction, is relatively easy to understand. It corresponds to our intuitive sense of time as this was shaped by the view of the world projected upon us by our own experiences as living creatures: from past through present towards the future. But the opposite sequence description, from the future (goal, state, action) to the present (means, methods, choices) seems to run counter to common sense. Think only of love—is it reaction or is it anticipation? Or is it both? Even making love follows anticipations that need not be spelled out here. Is it of the same nature as logical deduction? If it were, machines endowed with deductive capabilities could fall in love and make love. It is not even fun any more to argue with the machine-obsessed followers of Marvin Minsky [32], who predicted,

In from three to eight years, we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare,

---

[1] One of the older versions of the word "anticipation" is the Latin *antecapere*, a "pre-understanding," which eventually became *anticipare*.

grease a car, play office politics, tell a joke, have a fight. At that point, the machine will begin to educate itself with fantastic speed. In a few months, it will be at genius level, and a few months after that, its power will be incalculable.

Here we experience the fundamental need to overcome the bizarre Cartesian optimism embodied in the cult of the machine and to realize the complementary nature of reaction and anticipation. A self-learning machine ("the machine will begin to educate itself with fantastic speed") is as far from us today as is a good theory of learning that incorporates reactive and anticipatory mechanisms!

As stated at the very beginning, examples do not constitute theories. They are not even hypotheses. Rather, they inform us that among the things we can observe and among the experiments we can, to a certain extent, replicate, there are some which do not fit the pattern of what is already known or accepted. Obviously, before looking for explanations outside the knowledge already acquired and tested, we are inclined to apply what belongs to the shared knowledge of the world. Learning comes immediately to mind.

Maybe a cause does exist—even if we are not yet able to point to it—behind the mechanism controlling the heartbeat in anticipation of a change in posture. Or maybe laughter, the effect of tickling, is after all caused by something else, let us say located in the brain or associated with what is called the mind, or genetic mechanisms. Some scientists even believe that they have found the location of the genetic code leading to adventure. But adventure, which makes learning possible, is also in anticipation of the experience of learning; it is a promise, which might turn out empty.

The realization that "ahead of time" (*ante*) is the legitimate prefix of *anti-cipation* presupposes a different description of time, the implicit dimension of all learning experiences.

To Be is to Act... Possibly to Venture

Neuroscience, as part of the science of the living (biological science), took it upon itself to examine the relation between sensory data—what we smell, taste, see, hear, and experience through our sense of touch (the haptic sense)—and our actions. There are many reasons for this undertaking. One of them is the increased realization by the scientific community of the inadequacy of the deterministic-reductionist approach. Researchers are trying to understand the ever-changing (dynamic) relation between the living organism and the environment. Questions are focused on the connection—if any could be established—between what organisms do and how the action is embodied in a change in their condition. This change is an expression of learning, especially of adaptive (not only reactive) learning. But this is also how we evaluate—after the fact—the knowledge involved (or, as some would say, "successful learning," which results in acquired knowledge).

What eventually becomes clear is that living processes are characterized by their integrated nature: the whole (organism, being) and its parts are related and reciprocally interdependent. To understand learning means to understand how interdependence—of actions, agents, individuals, groups, etc.—is instantiated in what we do. Sensory data analysis (also called perception) cannot be meaningful unless we correlate sensory information to action.

Expectation, Guessing, Prediction

We become aware of anticipation through the intermediary of some of its aspects: from guessing and expectation to prediction and planning. Each of these various aspects—ultimately a combination of reaction and anticipation—can be pragmatically described. To guess is to choose from what might happen on account of various experiences: one's own, in the case of a similar endeavor or one that is likely to happen; the experience of others; or experience based on unrelated patterns—the so-called "lucky throw" of a coin or dice, for example. The reaction component and the anticipation are combined according to a formula that qualifies some of us as better than others in guessing. Reactions are based on the evaluation of the information pertinent to the situation—different when one visits a casino than when one guesses the correct answer in a multiple-choice test. Anticipations result from the self: From all possible games in the casino, some are more "favorable" at a certain time. In the multiple-choice situation, one infers from the known to the unknown. There is learning in guessing as patterns emerge: the next attempt integrates the observation of related or unrelated information. This associative action is the cognitive ingredient most connected to guessing and, at the same time, to learning. The rest is often statistics at work, combined with *ad hoc* associative schemes pertinent to what is possible. Let us acknowledge that guessing (as well as learning) is reduced to nil in predictable situations. Only surprise justifies the effort even when the result is negative. Recent research of responses of the human frontal cortex to surprising events (cf. Fletcher *et al* [33]) points to the relation to learning that I mentioned above. The dorsolateral prefrontal cortex seems to perform a role in the adjustment of inferential learning. Associative relationships leading to associative learning are based on the action of discriminating the degree (strength) of interrelation.

In comparison to guessing, expectation does not entail choosing ("Heads or tails?"), but rather an evaluation of the outcome of some open-ended process. Look at your child's expression and guess what might happen when he or she will "hang out" with friends. (For instance, facial expression reveals a lot more about coming actions than we actually see.) This evaluation might be difficult, if at all possible, to describe (e.g., "I know what you guys plan to do"). In the act of forming an expectation, the focus changes from the probable to the desired and/or expected. These pertain to what is possible. Several sources of information pertinent to forming an expectation are weighed against each other. What appears most probable out of all that is possible

gets the higher evaluation, especially if its outcome is desirable. If the outcome is judged to be negative, then avoiding it is the basis for action. Again, anticipation—reflected in what is perceived as possible—meets reaction, and information is associated with probable cause. Weather is often expected, not guessed. So are the outcomes of activities that weather might influence, such as agriculture practiced prior to the integration of digital information in agricultural praxis. A cornfield is not equally fertile in every spot. We can learn how to increase production by extracting data (through GPS-based measurements) pertinent to fertility and applying fertilizers or planting more seeds in certain spots. Based on the evaluation of the outcome (see introductory definition) we form new expectations. Events with a certain regularity prompt patterns of expectation: a wife awaits her husband, a child awaits his/her parent, a dog awaits its owner, who usually returns from work at a certain time. We encounter these regular events on many occasions and in many activities. And we try to use them to our benefit.

Prediction results from connecting cause and effect, by associating the data describing their connection. Causality, as the primary, but not exclusive, source of our predictive power is rarely explicit, and even more rarely easy to understand. Accordingly, prediction—explicit or implicit—expresses the limits of our own understanding of whatever we want to predict. In some cases, the prediction is fed back into what we want to predict: how a certain political decision will affect society; how an economic mechanism will affect the market; how technological innovation (let's say multimedia) will affect education. As a result, a self-referential loop is created. The outcome is none other than what we input as prediction, although we are not fully aware of the circularity inherent in the process. The impossibility of disconnecting the observer (subject, in learning) from the observed (the object of learning) is an inherent condition of learning, whether human or machine learning. The constructivist perspective demonstrated the point quite convincingly (cf. Ernest von Glasersfeld [34]).

But there are also predictions driven, to a certain extent, by anticipatory mechanisms. Falling in love at first sight is a prediction almost impossible to make explicit. There is no cause-and-effect connection to establish. The future state (the romantic ideal of a great love, or the calculated outcome of an arranged marriage) affects current states as these succeed each other in a sequence of a time often described as "out of this world." Facial expression as a predictor is yet another example. In very sophisticated studies (cf. Ekman *et al* [35]), it was shown that the "language" of facial expression speaks of facts to happen before they are even initiated—which is anticipation in pure form. For those who "read" the face's expression, i.e., for those who learned the language of facial expression, predictions based on their own anticipation guides their action. (A convincing case is that of a Los Angeles policeman who reads on the face of the criminal holding a gun on him that he will not shoot, leading the officer to spare his life, cf**.** Gladwell [36].) The expectation—criminal pulls out gun and points it at the policeman pursuing him—and the prediction—this person with the particular facial expression, as studied by the interpreter, will not shoot—collide.

Descriptions of the relation between expectation and prediction are informative in respect to the mechanisms on which both are based, and to the various levels at which learning takes place. Expectations pertain to more patterned situations. An acceptable description is that the learner extracts regularities or uses innate knowledge (such as in speech). At times, they are an expression of what in ordinary language is described as stereotype or, in some cases, as wishful thinking. However, when the individuals become involved in the activity of predicting (literally, *to say beforehand,* i.e., before something happens), they not infrequently expect the prediction to actually take place. It is no longer a wish, but rather the human desire, expressed in some action, to succeed.

The attempt to understand change (macro-level, i.e., behavior) without looking at the processes through which it takes place (micro-level processes, cognitive processes) also leaves anticipation out of the picture that this understanding of change eventually provides. Change means modification over time. Predictive efforts are focused on understanding sequences: how one step in time is followed by another. However, these efforts focus on what ultimately anticipatory processes are—a modeling of the entity for which they are an agency, and the execution of the model in faster than real time speed. The limited deterministic perspective, mechanic in nature—i.e., which cause leads to which ensuing effect—affects the understanding of anticipation through a description of predictive mechanisms. Predictions following known methods (such as time series analysis and linear predictors theory) capture the reaction component of human action. (For more information and bibliography, see, respectively, http://www.ubmail.ubalt.edu/~harsham/stat-data/opre330Forecast.htm, and http://www.neurocolt.com/abs/1996/abs96037.html).The anticipatory component is left out most of the time, as a matter of definition and convenience: complexity is difficult to handle. Once a predictive hypothesis—let's say every minute the clock mechanism engages the minute hand—is adopted, it defines the cognitive frame of reference. Should the predicted behavior of the mechanism somehow not take place, the expectation is tested. However, mechanisms, as embodiments of determinism, rarely fail. And when they do, it is always for reasons independent of the mechanism's structure.

Predictions concerning the living are less obliging. It happens at all levels of the living that predictions—what will happen next (immediate or less immediate future)—are either partially correct or not at all. In studying learning and selective attention, Peter Dayan *et al* [37] refer to reward mechanisms in the Kalman filter model (more experience leads to higher certainty). In such cases, expectations turn out to be a measure of how much the deterministic instinct (culture, if you prefer) takes over the more complex model that accounts for both reaction and anticipation in the dynamics of the living.

Predictors reflect our desire to understand how change takes place. They express our practical need to deal with change: However, they leave our own change out of the equation. Actions from thoughts, as Nicolelis [38] calls them, account for the self-

awareness of change. What is learned supports inferences (statistical or possibilistic); uncertainty results as the competitive resources engaged in the inference are overwritten by unrelated factors. Predictions also try to capture the interconnectedness of all elements involved in the dynamics of the observed. And not last by any means, they are an instrument of learning. In this sense, they open access to ways in which we could emulate (or imitate, to use another word) change. It is important to realize that expectations have no direct learning component: one cannot learn explicitly how to expect, even if we accept that there might be structure in the learning process and in the representation. Expectations only occasionally produce knowledge: a series of expectations with a certain pattern of success, or failure for that matter. Predictions, even when only marginally successful, support activities such as forecasting—for short or less than short sequences of change—of modeling, and of inference to the characteristics of the observed dynamic entities. A good prediction says quite a bit about the complexity of the change observed. Even through what it does not capture, it says a lot about complexities we tend to leave aside.

In this respect, predictions regarding the living, although inappropriate for systematically capturing their anticipatory dimension, are a good indicator of what we miss when we ignore anticipation. An example: In focusing only on human beings, predictions of physiological data remain at a primitive stage at best, despite the spectacular progress in technology and in the scientific theory of prediction. Albeit, if we could improve such predictions by accounting for the role of anticipation, we would be in a better position to deal with life-threatening occurrences (strokes, sudden cardiac death, diabetic shock, epileptic seizure, etc. [cf. 12]). This means that we would learn about such occurrences in ways transcending their appearance and probability. Things are not different in the many and varied attempts we undertake in predictions concerning the environment, education, health, and the functioning of markets. Unless and until anticipation is acknowledged and appropriate forms of accounting for it are established, the situation will not change drastically. This becomes even clearer when we look at the very important experiences of forecasting and planning.

Forecasting and Planning

Not fundamentally different from predictions are forecasts. Actually, predictions are a prerequisite of forecasting. The etymology points to a different pragmatics, one that involves randomness (as in casting). Under certain circumstances, predictions can refer to the past: Take a sequence in time—let's say the San Francisco earthquake of 1906—and try to describe the event (after the fact). In order to do so, the data, as registered by many devices (some local, some remote) and the theory are subjected to interpretations. The so-called Heat-Flow Paradox is a good example. If tectonic plates grind against one another, there should be friction and consequently heat. This is the result of learning from other physical phenomena. Along the well-known San Andreas Fault, geologists (and others) have measured (and keep measuring) every conceivable

phenomenon. No heat has been detected. The generalization of the learned in respect to friction proved doubtful at best. Accordingly, in order to maintain the heat dissipation hypothesis as a basis for forecasting, scientists started to consider the composition of the fault. This new learning—extraction of regularities other than those pertaining to friction and heat dissipation—was focused on a different aspect of friction. A strong fault and a weak fault behave differently under stress, and therefore release different quantities of heat. Here is a case in which data is fitted to a hypothesis—heat release resulting from friction. This is an attempt to adapt what was learned to a different context. Therefore adaptive learning and forecast offer a different context for interpretation.

In other cases, as researchers eventually learned, what was measured as "noise" was treated as data. Learning noise patterns is a subject we rarely approach, because we do not yet know of procedures for effectively distinguishing between noise and data. In medicine, where the qualifiers "symptomatic" vs. "non-symptomatic" are applied in order to distinguish between data and noise, this occurs to the detriment of predictive performance.

In general, theories are advanced and tested against the description given in the form of data. Regardless, predictions pertinent to previous change (i.e., descriptions of the change) are not unlike descriptions geared to future change. In regard to the past, one can continue to improve the description (fitting the data to a theory) until some pattern is eventually discerned and false knowledge discarded.

To ascertain that something will happen in advance of the actual occurrence—prediction (the weather will change, it will rain)—and to cast in advance—forecast—(tomorrow it will rain) might at first glance seem more similar than they are. A computer program for predicting weather could process historic data: weather patterns over a long time period. It could associate them with the most recent sequence. And in the end, it could come up with an acceptable global prediction for a season, year, or decade. In contrast, a forecasting model would be local and specific. The prediction based on "measuring" the "physical state" of a person (how the "pump," i.e., heart, and "pipes," i.e., blood vessels, are doing, the state of tissue and bone) can be well expressed in such terms as "clean bill of health" or "worrisome heart symptoms." But it can almost never become a forecast: "351 days later, you will have a heart attack;" or "In one year and seven hours, you will fall and break your jaw."

Predictions do not involve interpretations. They result from whatever explanatory model (expressed or not) is adopted. Forecasts, even when delivered without explanation, are interpretive. They contain an answer to the question behind the forecasted. "The price of oil will change due to...." You can fill in the blank as the situation prompts: cold winter, pipeline failure, war. "Tomorrow at 11:30 am it will rain...." because of whatever brings on rain. "There will be a change in government...." "Your baby will be born in ... hours." A good predictive model can be turned into a machine—something we do quite often, turning into a device the physics

or chemistry behind a good prediction: "If you don't watch the heat under the frying pan, the oil in it will catch fire."

Forecasts are not reducible to the machine structure. They involve data we can harvest outside our own system (the sensorial, in the broadest sense). In addition, they involve data we ourselves generate. The interplay of initial conditions (internal and external dynamics, linearity and non-linearity, to name a few factors), that is, the interplay of reaction and anticipation, is what makes or breaks a forecast.

Our own existence is one of never-ending change. Implicit in this dynamic condition of the living are a) the impossibility of accurate forecasting, and b) the possibility of improving the prediction of physical phenomena, to the extent that we can separate the physical from the living.

Our guesses, expectations, predictions, and forecasts—in other words, our learning in a broad sense, as I defined learning in the methodological introduction—co-affect our actions and affect our pragmatics. Each of them, in a different way, partakes in shaping actions. Their interplay makes up a very difficult array of factors impossible to escape, but even more difficult to account for in detail.

Mutually reinforcing guesses, expectations, predictions, and forecasts corresponding to a course of events for which we have effective descriptions allow us to proceed successfully in some of our actions. In other cases, they appear to cancel each other out, and thus undermine the action, or negatively affect its outcome. Learning and unlearning (which is different from forgetting) probably need to be approached together. Indeterminacy can be experienced as well. It corresponds to descriptions of events for which we have insufficient information and experience, or lack of knowledge. They can also correspond to events which by their nature seem to be ill defined. We react *and* anticipate. This conjunction defines how effective we can be.

Self-Awareness, Intentionality, and Planning

We human beings are what we do (the pragmatic foundation of identity, cf. [15] pp. 258ff). The only identifier of our actions is their outcome. This is an instantiation of our identity at the same time. The question, "What do you do?" cannot be answered with "I anticipate," followed, or not, by an object, such as "I anticipate that an object will fall," or "I anticipate my wife's arrival," or "I anticipate smelling something that I never experienced before." Anticipation is a characteristic of the living, but not a specific action or activity. We do not undertake anticipation. It is not a specific task. Anticipation is the result of a variety of processes. We are *in* anticipation. As an outcome, anticipation is expressed through consequences: increased performance—an anticipated tennis serve is returned; danger—such as a passing car—is avoided; an opportunity—in the stock market, for instance—is used to advantage. Implicit in the

functioning of the living, anticipatory processes result in the proactive dimension of life. This is where identity originates. Anticipatory processes are defined in contrast to reaction. Characteristic of the deterministic sequence of action-reaction defined in physics, reaction is the expression of our physical nature. Identity is expressed in the unity of the reactive and proactive dimensions of the human being. It appears as a stable expression, but actually defines change. It is the difference between what we seem to be and what we are becoming as our existence unfolds over time. Identity is affected by, but is not the outcome of, learning.

No matter what we do, the doing itself—to which explicit and implicit learning belongs—is what defines our unfolding identity. The outcome is the expression of our physical and intellectual abilities. It also reflects our knowledge and experience. The expression of goals, whether they are specifically spelled out or implicitly assumed, affects the outcome of our actions as well. We identify here the process through which our existence is preserved at the lowest level—not unlike the phototropic mono-cell and progressing all the way up to the human being. But at a certain level of life organization and complexity, the preservation drive assumes new forms through which it is realized. Anticipation is the common denominator. However, the concrete aspect of how it is eventually expressed—i.e., through self-awareness, intentionality, or in the activity we call "planning"—changes as the complexity of the processes through which the living unfolds increases.

Anticipation at the level of preserving existence is unreflected. Facial expression in anticipation of an action is a good example here, too. It seems that facial expression is not defined on a cultural level but is species wide (cf. Ekman [35], Gladwell [36]). It is not a learned expression. We can control our facial expression to an extent, but there is always that one second or less in which control is out of the question. Intentionality is always entangled with awareness—one cannot intend something without awareness, even in vague forms. But this awareness does not automatically make human expressions carry anticipations more than the expression of the rest of the living. The difference is evident on a different level. Humans reach self-awareness; the mind is the subject of knowledge of the mind itself. As such, we eventually recognize that our faces "speak" before we act. They are our forecasts, the majority of these involuntary. Those intent on deciphering them obtain access to some intriguing anticipatory mechanisms, or at least to their expression.

Planning, expressed through policymaking, management, prevention, logistics, and even design, implies the *ante* element—giving advance thought, directing towards something, looking forward, engaging resources (including the self). Moreover, it implies understanding, which resonates with the initial form of the word denoting anticipation: *antecapere*. As such, the activity through which human beings identify themselves as authors of the blueprint of their actions takes place no longer at the object level, but on a meta-level. It is an activity of abstracting future actions from their object. It is also their definition in a cognitive domain not directly associated with sensory input, but rather with what we understand, with knowledge. Plans synthesize predictive abilities, forecasting, and modeling. A plan is the understanding

of what is planned, expressed in goals and means to attain these goals, as well as the time sequence for executing the plan. A plan is a timeline and a script for interactions indexed to the timeline. To what we call *understanding* belong goals, the means, the underlying structure of the endeavor (tasks assumed by one person, by several, the nature of their relation, etc.), a sense of progression in time. As such, they report upon the physical determination of everything we do and of the anticipatory framework. In every plan, from the most primitive to the utmost complex, the goal is associated with the reality for which the plan provides a description (a theory), which is called *configuration space.* If it is a scientific plan, such as the exploration of the moon or the genome project, the plan describes where the "science" actually resides, where those equations we call descriptions are "located." If it is a political plan, or an education plan, the configuration space is made up of the people that the plan intends to engage. Our own description of the people, like the mathematical equations of science, is relative. They "shape" the configuration space, and, within that space, the interactions through which people learn from each other.

The plan also has to describe the time-space in which the goal pursued will eventually be embodied. This is a manifold, towards which the dynamics of our own actions and the actions of those with whom we interact (social context) will move us. Again, in science, this is the landing on the moon, the map of the human gene, a new educational strategy or, in politics, the outcome of equal opportunity policies—to name very few examples. All of these are anticipations projected against the background of our understanding the world as one that unites the physical and the living. Plans spell out the variables to be affected through our actions, and the nature of their interrelationships. Quite often, plans infer from the past (the reactive component) to the future (proactive component). They also project how the future will eventually affect the sequence of ensuing current states. Planning cannot be disassociated from self-regulation, i.e., from the inner dynamics of phenomena and their attractors. These attractors are the states into which they will settle. They are the descriptions of self-organizing processes, their eventual destination, if we can understand it as a dynamic entity, not a static finality. Planning sets the limits within which adaptive processes are allowed. Each plan is in effect an expression of learning in action, the very goal pursued having the nature of an attractor in a dynamic system.

Anticipation and Learning

Learning is an active pursuit of regularities in whatever we do, for instance: meeting a bear, falling in a lake, walking in the dark (to remain with the realm of fear), or smelling something that never existed before (and for which we accordingly have no genetic information or predisposition). But learning is also the realization of time cycles (day and night, seasons), natural cycles (related to our environment), patterns in space, patterns in human interactions, and in the interactions among various organisms. Such regularities appear as independent of our own activity, that is, independent of our self-constitution. They seem to be necessary in nature. For

example, night must necessarily follow day; we cannot avoid this. Accordingly, the anticipation of night (in the form of rhythm changes in the organism), or of seasons (the thickening of animal fur, falling leaves, changes in metabolism), takes the appearance of something unavoidable. A change in the space through which we move exemplifies the same in respect to our space-based experiences. Our pace changes when a hill is followed by a long level pasture, or even when the "tactility" of the road changes from a smooth, dry surface to a rough, uneven surface. The degree of necessity is expressed in the regularities we extracted and became aware of, that is, in our learning. This learning can be further transferred to machines (e.g., cars, to remain in the domain of mobility).

Here a major question has to be posed. If we want to understand how learning is expressed in the form of anticipation, we need to be able to describe what partakes in the "forecast" leading to anticipatory behavior. The hypotheses to entertain are obvious: Anticipation might be associated with conditioning through previous responses—we do not put our hands into fire in order to find out that they can get burned. Anticipation might result from learning experiences. Or, not to be too hastily discarded, even after accounting for Elsasser's [10] storage-free memory model (based on Bergson's [40] considerations of memory), anticipation might be embryonically predefined and influenced by the subsequent dynamics of the organism unfolding in a given environment. Maturing and ageing, with all that these entail in bodily and cognitive changes, that is, in new learning patterns, is the clearest example (cf. Haylick [41]). Here we have a very powerful illustration of the dynamics affecting the relation between anticipation and learning as the body increasingly returns to its "physicality."

Anticipation as Agent

The agency of anticipation is ultimately about supporting the life of an organism. Decreased anticipation corresponds to a decreased living dynamics. As we have just seen, the living returns to the merely physical, which is the substratum of all there is. In some convoluted way, and by focusing on anticipation, we've transcended the question of modality ("How?")—in particular "How does anticipation take place?" It is progressively replaced by the fundamental question, existential in nature, of "Why?"—indeed, "Why does the living individual lose its 'living' characteristics?" By its own condition, learning is driven by this very "Why?" question. To repeat: the answer to "Why?" is the finality called life. We also integrated conditioning—a characteristic of the living very well documented in science (psychology, cognitive science, neuroscience, etc.)—learning, and the predefined characteristics (embryonically and otherwise) into anticipation. The living learns, but some of its anticipatory characteristics are inherited. Better yet, they are implicit in the condition of being alive, or implicit in survival. Adventure, as an inclination, is telling more of survival than of curiosity—unless curiosity itself is related to survival.

It is really difficult to distinguish between what seems to be learned and what is expressed as a new hypothesis. The variability in such activities is the result of a delicate interplay between conditioning through previous responses and learning, from others or from experience. For example, "When the pension funds sell a stock, it is time to sell!"—the participants in the market "guess," "expect," "forecast," or "predict" a future state that becomes the future state of the market through which they identify themselves. In many instances, variability and the anticipation of the sense and magnitude of change can be the result of noise. Rumors have their own dynamics, as does an instable political situation.

Example: A Hybrid Anticipatory Control Mechanism. The "Learning" Automobile

Driving a car involves reaction-based knowledge (acquired through experience, instructions, interaction with other drivers, etc.) and anticipatory skills (sometimes manifested as anticipatory behavior). The hybrid anticipatory control mechanism integrates automatic car supervision and control and human performance. A simple architecture for implementation in the form of intelligent agent technology was suggested by Davidsson *et al* [42].
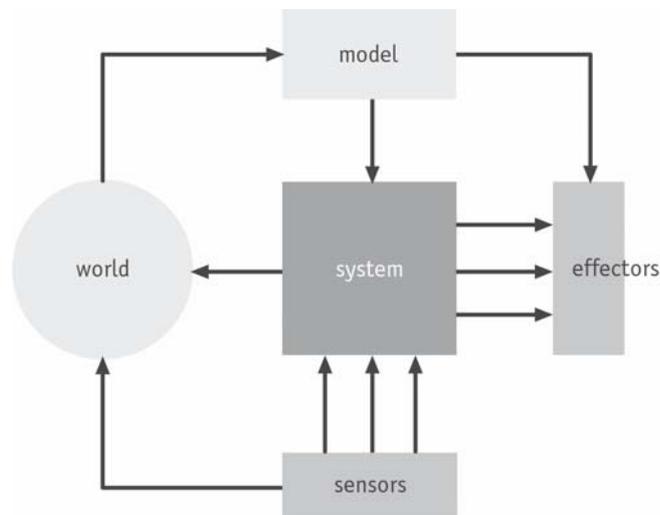
Fig. 1. Anticipation-based architecture for an agent-controlled mechanism.

In this architecture, the model of the car operating in faster than real time (the module described as model) allows predictions to be turned into specific actions. The model affects the system's performance: the prediction of a curve results in lowering speed, for instance.

Here we have two concurrent processes: 1) a reactive process at the object level (the system controlled is the car); b) a predictive process, with anticipatory characteristics, at a meta-level (the model). The integration of the two processes is not trivial. Once we implement an improved architecture of selections corresponding to progressive learning (slow down in a curve, if the road is icy, the slow-down has to be even greater; when driving in heavy traffic on an icy road, deceleration has to be accompanied by a different placement of the car in respect to the median; and so on) , complexity increases. If we introduce competition among models and a reward mechanism, we can augment the adaptive qualities of the integrated control mechanism. In this case, alternatives are generated—some probabilistic (integrating knowledge extracted from statistics), such as "informed guesses," expectations (driving up a hill, one can expect a slower speed), forecasts (after rain, the system can forecast a longer brake-path), and predictions.
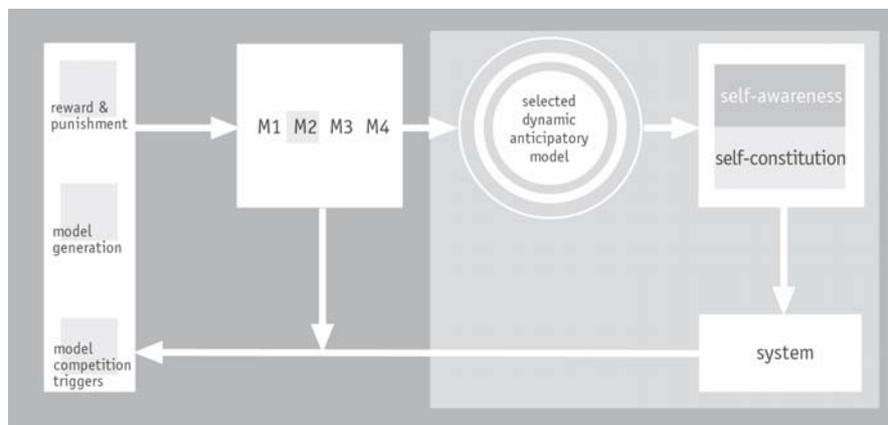
Fig. 2. Anticipative mechanism based on competition among models and the reward mechanism.

This architecture makes the following available:
- a space of possibilities (in Zadeh's [43] sense) in the model generation module
- conflicting possibilities: for each situation, a certain possibility is better than the others
- a selection and reward mechanism: learning is stimulated by allocating resources to the "winning" mechanism.

The processes with anticipatory, predictive, and forecasting characteristics are described through

Control = function of (past state, current state, future state) system

Knowledge of future states is a matter of possibilistic distributions. The anticipated performance (cf. [44] p. 281) and the actual performance are related. Their difference, together with the reward mechanism, guides the learning component.

Functioning under continuously changing conditions means that control mechanisms will have to reflect the dynamics of the activity. Without learning, this is not possible. If we finally combine the automated part and human performance, expressed in driver behavior features, we arrive at an improved architecture:
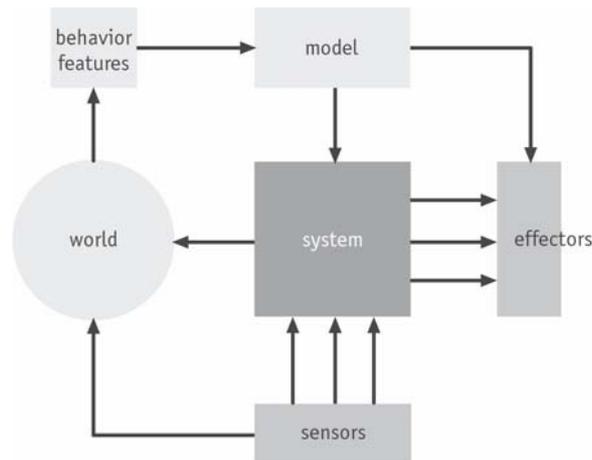
Fig. 3 . Hybrid control mechanism in which driver's behavior features are a source of "learning" for the car

Essential in this new architecture are the indexed behavior features and the mechanism for extracting regularities from them, i.e., the learning. The "learning" car is thus one which combines its own dynamic (modified, evolving knowledge) and that of the driver (anticipation included). This example is only an indication of ways through which implementations of anticipatory characteristics can be achieved. A more detailed analysis guides current research efforts and cooperative projects with industry.

In Lieu of Concluding Remarks

For those who hope that by better understanding anticipation we could improve our prediction and forecast performance, a *caveat:* Anticipatory processes are not reducible to deterministic sequences. Once we accept this, we automatically also accept that although the past remains important, we can no longer act on the assumption that the past contains the present and the future. Moreover, we can no longer rely exclusively on a scientific description according to which the past alone determines the present and the future. We can, however, improve prediction and forecasting by facilitating improved anticipation. This can result from a better knowledge of the possibility space: all that is possible in a given context. This

knowledge extends from better methods for evaluating the weight (significance) of possibilities, for dynamically pruning what is of lesser impact from the possibility space, to learning how to handle the imprecise knowledge that we have, instead of searching for more data and higher precision. This aspect, of the relation between anticipation, human performance, and soft computing, will require our attention beyond the usual inclination to improve the performance of our programs, or to improve the hardware in order to handle more data with higher precision.

Only one more word at this juncture: Planning, ideally a practical activity of integrating reaction and anticipation, can benefit from improved prediction and forecasting. It can benefit more, however, from learning, especially learning associated with anticipatory generalization. Based on this understanding, we have to pursue the more specific levels at which anticipation—the proactive path—and all the physical processes described in relation to it—the reactive path—eventually merge. This is a task of an order of magnitude well beyond a context of understanding for the subject of anticipation.

References

1. Clark, R.W.: Einstein. The Life and Times. Avon, New York (1972)
2. Peirce, Charles S.: CP 2.222. In: Hartshorne, C., Weiss, P. (eds): The Collected Papers of Charles Sanders Peirce, Vol. I-VI. Harvard University Press, Cambridge (1931). Per convention, Peirce is cited by volume number:article number.
3. Rosen, Robert: Anticipatory Systems. Pergamon Press, New York (1985)
4. Rosen, Robert: Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life. Columbia University Press, New York (1991)
5. Rosen, Robert: Essays on Life Itself. Complexity in Ecological Systems. Columbia University Press, New York (2000)
6. Nadin, Mihai: Mind—Anticipation and Chaos. Belser Presse, Zürich Stuttgart (1991), based on lecture: Mind—Intelligence is Process, Ohio State University (1988)
7. Dubois, Daniel, Resconi, G.: Hyperincursivity: A New Mathematical Theory. Presses Universitaires de Liège, Liège (1992)
8. Nadin Mihai: Anticipation: The Cause Lies in the Future, Lars Mueller Publishers, Baden Switzerland (2003)
9. Schultz, Wolfram: Multiple reward signals in the brain. In: Nature Reviews/Neuroscience, vol. 1, December (2001) 199-207

10. Elsasser, Walter M.: Reflections on a Theory of Organisms. Johns Hopkins University Press, Baltimore (1998). Originally published as Reflections on a Theory of Organisms. Holism in Biology. Orbis Publishing, Frelighsberg, Quebec (1987)
11. Schall, Jeffrey: Neural Basis of Deciding, Choosing, and Acting. In: Nature Reiews Neuroscience, vol. 2 January (2001) 33-42
12. Rigney, David R., Goldberger, Ary L., Ocasio, Wendell C., Ichimaru, Yuhei, Moody, George B., and Mark, Roger G.: Multi-Channel Physiological Data: Description and Analysis (Data Set B). In: Weigend, Andreas S. and Gershenfeld, Neil A. (eds.): Time Series Prediction: Forecasting the Future and Understanding the Past. Addison Wesley Longman (1994)
13. Ernest, P.: The one and the many. In: Steffe, L. and Gale, J. (eds.): Constructivism in Education. Lawrence Erlbaum Associates, Inc., New Jersey: (1995) 459-486
14. Novak, Joseph and Gowin, D. Bob: Learning How to Learn. Cambridge University Press, Cambridge, UK and New York (1984, 1999)
15. Nadin, Mihai: The Civilization of Illiteracy. Dresden University Press, Dresden (1997)
16. Adolph, Edward F.: Origins of Physiological Relations. Academic Press, New York (1961)
17. Cannon, Walter B.: The Wisdom of the Body (2nd ed.). Peter Smith Publisher, New York (1939)
18. Trzebski, Andrzej: Baroreceptor Reflex Revisited (editorial comment). in: Polish Heart Journal, vol. L, 5, November (1999) 405
19. Nadin, Mihai: Anticipation—A Spooky Computation. In Dubois, D. (ed.): CASYS, International Journal of Computing Anticipatory Systems, Partial Proceedings of CASYS 99, Liege: CHAOS , Vol. 6 (1999) 3-47
20. Meegan, Daniel V., R.N. Aslin, and R.A. Jacobs. Motor timing learned without motor training. In: Nature/Neuroscience, vol. 3, No. 9 September (2000) 860-862
21. Carlsson, Katrina, Petrovic, Skane, S., Petersson, K.M., Ingvar, M.: Tickling Expectations: Neural Processing in Anticipation of a Sensory Stimulus. In: Journal of Cognitive Neuroscience 12, (2000) 691-703
22. Serrien, Deborah J., Mario Wiesendanger. Role of the Cerebellum in Tickling. Anticipatory and Reactive Free Responses. Journal of Cognitive Neuroscience 11, 1999, pp. 672-681.
23. Braun, A.R., H. Guillemen, L. Hosey, M. Varga. The neural organization of discourse, Brain, Vol. 124, No. 10, October 2001, pp. 2028-2044.
24. Guenther, F.H.: Neural modeling of speech production. In: Proceedings of the 4th International Nijmegen Speech Motor Conference, Nijmegen (2001)
25. Max, Ludo: Stuttering and Internal Models for Sensorimotor Control: Deriving Testable Hypotheses for a Theoretical Perspective. in: Maassen, B., Kent, R., Peters, H.F.M, Lieshout, P. van, and Hulstijn eds): Speech Motor Control in Normal and Disorderd Speech. Oxford University Press, Oxford (forthcoming)
26. Gallese, Vittorio: The Inner Sense of Action. Agency and Motor Representations. In: Journal of Consciousness Studies, 7:10 (2000) 23-40
27. Wiesendanger, Mario and Serrien, Deborah J.: Toward a physiological understanding of human dexterity. In: News in Physiological Sciences, vol. 16, no. 5, October (2001) 228-233
28. Wolpert, Daniel M. and Ghahramani, Youben: Computational principles of movement neuroscience. In: Nature/Neuroscience Supplement, vol. 3, November (2000) 1212-1217
29. Foerster, Heinz von: Der Anfang von Himmel und Erde hat keinen Namen (2nd ed.). Döcker Verlag, Vienna (1999)
30. Pearl, Judea: Causality. Models, Reasoning and Inference. Cambridge University Press, Cambridge (2000)

31. Zadeh, Lotfi A.: An Outline of Computational Theory of Perceptions Based on Computing with Words. In: Sinha, N.K., Gupta, M.N., and Zadeh, L.A. (eds.): Soft Computing & Intelligent Systems. Academic Press, New York (2000) 3-22
32. Minsky, Marvin: The Virtual Duck and the Endangered Nightingale. In: Digital Media, June 5, (1995) 68-74
33. Fletcher, P.C., Anderson, J.M., Shanks, D.R., Honey, R., Carpenter, T.A., Donovan, T., Papadakis, N., Bullmore, E.T.: Responses of human frontal cortex to surprising events are predicted by formal associative learning theory. In: Nature Neuroscience, 4:10, (October, 2001) 1043-1048
34. Glasersfeld, Ernest von: Radical Constructivism: A Way of Knowing and Learning. The Falmer Press, London Washington (1995)
35. Ekman, Paul, Rosenberg, E.L. (eds.): What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (FACS). Oxford University Press, New York (1997)
36. Gladwell, Malcolm: The Naked Face. Can experts really read your thoughts? In: The New Yorker (August 5, 2002) 38-49
37. Dayan, Peter, Kakade Sham, and Montague, P. Read: Learning and selective attention. In: Nature Neuroscience, 3 (2000) 1218-1223
38. Nicolelis, Miguel A.: Actions from thoughts. In: Nature 409 (2001) 403-407
39. Ekman, Paul: Universals and cultural differences in facial expressions of emotion. In: Cole, J. (ed.): Nebraska Symposium on Motivation 1971, Vol. 19, University of Nebraska Press, Lincoln (1972) 207-283
40. Bergson, Henri: La matière et la mémoire (1896) (The English edition, cited by Elsasser, was published as Matter and Memory, Allen and Unwin, London)
41. Haylick, Leonard: The future of ageing. In: Nature 408, November (2000) 267-269
42. Davidsson, Paul, Astor, Eric, and Ekdahl, Bertil: A framework for autonomous agents based on the concept of anticipatory systems. In: Cybernetics and Systems. World Scientific, Singapore (1994) 1427-1434
43. Zadeh, Lotfi A.: Fuzzy sets as a basis for a theory of possibility. In: Fuzzy Sets and Systems 1 (1978) 3-28
44. Ragheb, M., Tsoukalas,L.: A coupled probability-possibility method for decision-making in knowledge-based systems. In: S. C-Y. Lu, S.C-Y, Commanduri, R. (eds.): *Knowledge-Based Expert Systems for Manufacturing* New York: ASME, New York (1986)