

CASYS'97 – International Conference on Computing Anticipatory Systems Liège, August 11–15, 1997. Published in D.M.Dubois (Ed.) *Computing anticipatory systems* (38–47), Woodbury, NY: American Institute of Physics, 1998.

208

Anticipation in the Constructivist Theory of Cognition

I am not a computer scientist and I do not speak the languages of Quantum Computation, Hyperincursion, or Cellular Automata. But a couple of weeks ago I read some of the prose sections of Robert Rosen's *Anticipatory Systems*, and it gave me the hope that the "anticipation" referred to in the title of this conference would not be altogether different from the anticipations we depend on in managing and planning our ordinary lives as human beings. It is a topic I have been concerned with throughout the many years that I battled against the mindless excesses of the behaviorist doctrine in psychology. I was involved with languages and conceptual semantics, and I was among those outsiders who thought that, when people speak, they mostly have a purpose and are concerned with the effect their words will have. This view was generally considered to be unscientific, and I am therefore very happy that "anticipatory systems" have now become a subject for open discussion among "hard" scientists.

Since most of you are probably deeply immersed in specialized research of a more or less technical nature, you may not have had occasion to consider that anticipation would have to be a fundamental building block of any theory of psychology that merits to be called *cognitive*. And you may not have had reason to wonder why the discipline that is now called 'Cognitive Science' still has not moved very far from the input/output, stimulus/response paradigm. But although everyone now agrees that intelligence and intelligent behavior are the business of a MIND, few are ready to concede much autonomy to that rather indefinite entity.

I do not intend to bore you with a survey of a situation that, to me, seems quite dismal. Instead, I shall focus on the one theory that, in spite of all sorts of shortcomings, is in my view the most promising basis for further development, and consequently the most interesting for people involved in the construction of autonomous models.

The theory I am going to talk about is the one Jean Piaget called Genetic Epistemology. The name was not chosen at random. He wanted to make clear that he intended to analyze knowledge as it developed in the growing human mind, and not, as philosophers usually have done, as something that exists in its own right,

independent of the human knower. The name should have warned psychologists that Piaget's theory was not merely a theory of cognitive development, but also constituted a radically different approach to the problems of knowledge. However, especially in the English-speaking world, Piaget was mostly considered as a child psychologist, and his readers disregarded his break with traditional Western epistemology.

This neglect was unfortunate. Piaget actually provided a theoretical model of the human cognitive activity that is more complete and at the same time contains fewer metaphysical assumptions than those of most philosophers.

Towards the end of his working life, which lasted over more than six decades, he said:

The search for the mechanisms of biological adaptation and the analysis of that higher form of adaptation which is scientific thought [and its] epistemological interpretation has always been my central aim. (Piaget, in Gruber & Vonèche, 1977; p.xii)

The term 'adaptation' is the salient point. In many of his writings (he published over 80 books and several hundred articles) he reiterates that what we call knowledge cannot be a representation of an observer-independent reality. And every now and then, as in the passage I quoted, he says that the human activity of knowing is the highest form of adaptation. But he rarely put the two statements together – and this may have made it easier for both his followers and his critics to ignore the revolutionary conceptual change his theory was demanding.

If you consider that in the context of the Darwinian theory of evolution, "to be adapted" means *to survive by avoiding constraints*, it becomes clear that, for Piaget, "to know" does not involve acquiring a picture of the world around us. Instead, it concerns the discovery of paths of action and of thought that are open to us, paths that are *viable* in the face of experience.

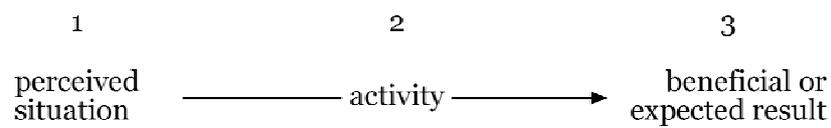
The passage I quoted also indicates that there is more than one level of adaptation. On the sensorimotor level of perception and bodily action, it is avoidance of physical perturbation and the possibility of survival that matter. On the level of thought we are concerned with concepts, their connections, with theories and explanations. All these are only indirectly linked to the practice of living. On this higher level, viability is determined by the attainment of goals and the elimination of conceptual contradictions.

To understand Piaget's theoretical scaffolding it is indeed indispensable to remember that he began as a biologist. He knew full well that the biological, *phylogenetic* adaptation of organisms to their environment was not an activity carried out by individuals or by a species. It was the result of natural selection, and natural selection does nothing but eliminate those specimens that do not possess the physical properties and the behavioral capabilities that are necessary to survive under the conditions of the present environment. All organisms that are equipped with senses and the ability to remember sensory experiences can, of course, to some extent *individually* increase their chances of survival by practical learning. Traditionally this was considered a separate domain, and it was explained by association. Piaget, however, who focused on human development, connected it to the biological principle of the reflex.

In most textbooks of behavioral biology, reflexes are described as automatic reactions to a stimulus. Piaget took into account two features that are usually not mentioned. The first was that the existence of heritable reflexes could be explained only by the fact that a fixed reaction, acquired through an accidental mutation, produced a result that gave the individuals who had it, an edge in the struggle for survival.

It is important to see that the specific property or capability that constitutes the evolutionary advantage has to be incorporated in the genome *before* the conditions arise relative to which it is considered *adapted*.

Remaining aware of the role of its result, Piaget thought of a reflex, as consisting of three elements:



The addition of 'expectation' sprang from the second observation Piaget had made, namely that most if not all the reflexes manifested by the human infant disappear or are modified during the course of maturation. The 'rooting reflex', for instance, that causes the baby to turn its head and to begin to suck when something touches its cheek, goes into remission soon after nourishment through a nipple is replaced by the use of cups and spoons.

Piaget also found that new 'fixed action patterns' can be developed. Such acquired reflexive behaviors are an integral part of our adult living. Among them are the way we move our feet when we go up or down stairs, the innumerable actions and reactions that have to become automatic if we want to be good at a sport, and, of course, the rituals of greeting an acquaintance and of small talk at a cocktail party. There are also reflexes that may lead to disaster – for example the way we stamp our foot on the brake pedal when an unexpected obstacle appears before us on the road.

An acquired reflex that impressed me much when I was young, was the one developed by the adolescent men and women of societies that prescribed skirts for females and trousers for males. In a sitting position, these women would unconsciously spread their skirt when something was thrown to them, whereas the men would clamp their knees together. (In those days, this was still used in the strictly male monasteries of Greece and Macedonia, in order to detect female intruders. Today, they have presumably thought of another test.)

Anyway, the more sophisticated view of the reflex enabled Piaget to take the tripartite pattern of *perceived situation*, *action*, and *result* as the basis for what he called 'Action Scheme'. It provided a powerful model for a form of practical learning on the sensorimotor level that was the same, in principle, for animals and humans.

Studies of animal behavior had shown that even the most primitive organisms tend to move towards situations that in the past provided agreeable experiences rather than towards those that proved unpleasant or painful. Humberto Maturana has characterized this by saying:

A living system, due to its circular organization, is an inductive system and functions always in a predictive manner: what happened once will occur

again. Its organization (genetic and otherwise) is conservative and repeats only that which works. (Maturana, 1970; p.15–16)

This was not intended to imply that primitive living organisms actually formulate expectations or make predictions. It was a sophisticated observer's way of describing their behavior. The pattern of learning, however, is the same as in Piaget's scheme theory, and once we impute to an organism the capability of reflecting upon its experiences, we can say that the principle of induction arises in its own thinking.

This principle has its logical foundation in what David Hume called the "supposition that the future will resemble the past" (Hume, 1742, Essay 3, Part 2). Having observed that, in past experience, situation **A** was usually followed by the unpleasant situation **B**, an organism that believed that this would be the case also in the future, could now make it its business to avoid situation **A**.

Together with its inverse (when situation **B** is a pleasant one, and **A** therefore leads to pursuit rather than avoidance), this is perhaps the first manifestation of *anticipatory* behavior.

To be successful, however, both pursuit and avoidance have to be directed by more or less continuous sensory feedback, and this, too, involves a specific form of anticipation. In their seminal 1943 paper, Rosenblueth, Wiener, and Bigelow wrote:

The purpose of voluntary acts is not a matter of arbitrary interpretation but a physiological fact. When we perform a voluntary action, what we select voluntarily is a specific purpose, not a specific movement. (Rosenblueth et al., 1943, p.19)

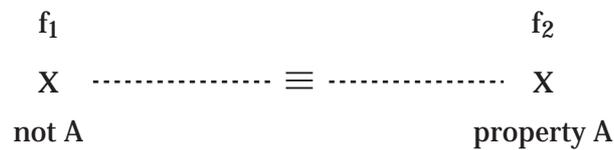
In their discussion of *purposeful behavior*, they used the example of bringing a glass of water to one's mouth in order to drink. The term *negative feedback*, they explained, signifies that "the behavior of an object is controlled by the margin of error at which the object stands at a given time with reference to a relatively specific goal" (ibid.). Such goal-directed behavior, however, has another indispensable component. In order to "control" the margin of error indicated by the feedback – in the given example this would be to reduce the distance that separates the glass from one's mouth – the acting subject must decide to act in a way that will reduce the error. And nothing but inductive inferences from past experience can enable the subject to choose a suitable way of acting.

Let us look at the example more closely. I am thirsty, and there is a glass of water in front of me on the table. From past experience I have learned (by induction and abstraction) that water is a means to quench my thirst. This is the 'voluntary purpose' I have chosen at the moment. In other words, I am anticipating that water will do again what it did in the past. But to achieve my purpose, I have to *drink* the water. There, again, I am relying on past experience, in the sense that I carry out the 'specific movements' which I expect (anticipate) to bring the glass to my lips. It is these movements that are controlled and guided by negative feedback.

When I reflect upon this sequence of decisions and actions, it becomes clear that the notion of causality plays an important role in the event. All my decisions to carry out specific actions are based on the expectation that they will bring about a change towards the desired goal.

The connections between causes and the changes they are supposed to produce as their effects, have forever been the subject of scientific investigation. The concepts of voluntary purpose and that of goal, however, were branded as remnants of teleological superstition and therefore considered inadmissible in the domain of science. The advent of cybernetics and the successful construction of goal-directed mechanisms has demonstrated that the proscription was unwarranted. Yet, today, there are still a good many scientists who have not fully appreciated the theoretical revolution. It therefore seems worth while to provide an analysis of the conceptual situation.

Many years ago, Silvio Ceccato, the first persistent practitioner of Bridgman's operational analysis, devised a graphic method of mapping complex concepts by means of a sequence of frames (I have borrowed the term 'frame' from cinematography; Glasersfeld, 1974). Because no single observation can lead to the conclusion that something has changed, we need a sequence of at least two frames showing something that acquires a difference. Consequently, the mapping has the following form:



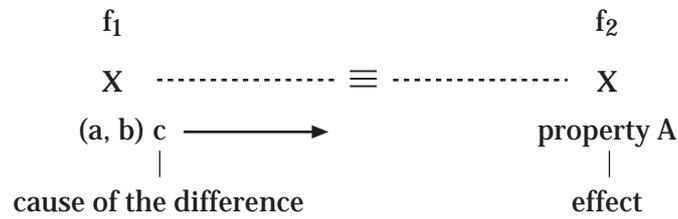
where "X" represents an item that is considered to be the same individual in both frames (indicated by the identity symbol " \equiv "). In short, we maintain an item's individual identity throughout two or more observational frames, and, at the same time, we claim that in the later frame it has gained a property "A" that it did not have in the earlier one (or we claim that it lost a property it had before).

The condition of identity may seem too obvious to mention, but analytically it is important to make it explicit, because of the ambiguity of the expression "the same". In English we say, "This is the same man who asked directions at the airport", and we mean that it is the same individual; but we might also say to a new acquaintance, "Oh, we are driving the same car – I, too, have an old Beetle!" and now we are speaking of two cars. In the first case, we could add, "Look, the man has changed – he's had a haircut!" In the second case, we cannot speak of change although our car is blue, and the other's yellow.

In French, the ambiguity of "*le même*" is analogous, and in German and Italian, although two words would be available to mark the conceptual difference, their use is quite indiscriminate. In fact, the situation in all these languages is worse, because common usage has modified the meaning of "identical" so that it can refer to the similarity or equivalence of *two* objects as well as to the individual identity of *one*.

Without the conception of change there would be no use for the notion of causation. It arises the moment we ask *why* a change has taken place. I have suggested that this question most likely springs from the fact that we attribute a new property (or the loss of a property) to an item which we nevertheless want to consider one and the same individual. This has the appearance of a contradiction, and we are looking for a way to resolve it. Because the frames indicating a change represent sequential observations, such a reason has to be found in the earlier one. We therefore examine what else could have been perceived in frame 1. We may compare the

experience to others we remember, in which X remained unchanged throughout several frames, and try to find something that was present this time, but not present when X did not change. Or we may behave like scientists and replicate the situation of f_1 , adding one by one new elements a,b,c,..., to isolate something we can hold responsible for the change. If we find one, we can map it as follows:



If we really are scientists, we will run all sorts of experiments in order to construct a theoretical model that shows *how* the element "c" effects the change. If we are successful, we will proudly add "c" to the tools we use in attempts to modify the world.

In everyday living, we are not so meticulous. If we find that some element was present two or three times when a given X changed in a desirable way, we are likely to assume that it is the cause, and we will *use* that element in the hope that it will bring about the desired change.

Even if it doesn't, it may take a number of failures to discourage us. If someone provides a metaphysical reason why it *should* work, failures do not seem to matter at all.

I am not sure, but I think it was the literary critic Cyril Connolly who made a startling observation in this regard. Insurance companies, he said, have the most sophisticated questionnaires to assess the risks involved in issuing a life insurance policy to an applicant. The questionnaires are based on meticulous research of mortality statistics. Connolly was struck by the fact that the questionnaires never ask the question "Do you pray?" And he wondered why people continued to pray for survival in all sorts of crisis, when the greatest experts of mortality clearly had found no evidence that it had an effect.

All this involves anticipation. The use of a cause-effect link in order to bring about a change is based on the belief that, since the cause has produced its effect in the past, it will produce it in the future. We project an established experiential connection into the domain of experiences we have not yet had. Hume has explained how we establish such connections: the repeated observation that the two items happened in temporal contiguity led us to infer and formulate a rule that says, if A happens, B will follow. Therefore, if we want B to happen, we try to generate A. In other words, we have a purpose and we act in a way which, we believe, will attain it.

The psychological establishment, which, from the 20s of this century until well into the 70s, was dominated in the United States by the dogma of behaviorism, considered purpose a mentalistic superstition. "Careless references to purpose are still to be found in both physics and biology, but good practice has no place for them," Skinner still wrote in 1971.

Natural scientists in physics, chemistry, and astronomy had found no reason to engage in thoughts about purpose in their disciplines, and they relegated it summarily to the realm of teleology. As Ernest Nagel put it:

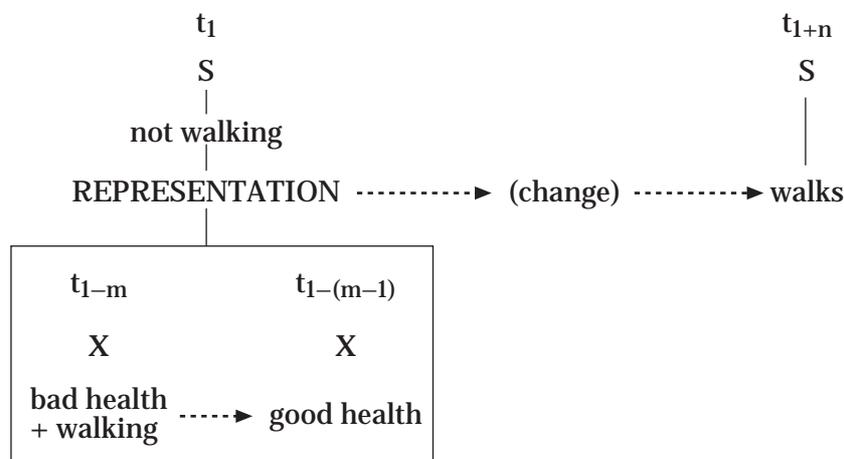
Perhaps the chief reason why most contemporary natural scientists disown teleology, and are disturbed by the use of teleological language in the natural sciences, is that the notion is equated with the belief that future events are active agents in their own realization. Such a belief is undoubtedly a species of superstition. (Nagel, 1965; p.24)

To believe that the future affects the present is no doubt a superstition, but to declare that purpose and goal-directed action must be discarded because they are teleological notions is no better. It shows an abysmal ignorance of the difference between empirical and metaphysical teleology.

I have suggested elsewhere that Aristotle, who provided the most valuable analysis of the concepts of causation, was well aware of the ambiguity. In his exposition, it becomes clear that what he called 'final' cause, i.e., the embodiment of a *telos* or goal, had two quite distinct applications (Glaserfeld, 1990). On the one hand, he saw the religious metaphysical belief that there was a *telos*, an ultimate, perfect state of the universe that draws the progress of the world we know towards itself. On the other, there was a second notion of the final cause, which he exemplified by saying that people go for walks *for the sake* of their health (Physics, Book II, ch.3, 194b-195a). This was a practical explanatory principle for which there is, indeed, an overwhelming amount of empirical evidence.

In this practical manifestation of finality, no actual future state is involved, but a *mental re-presentation* of a state that has been experienced as the result of a particular action. Even in Aristotle's day, bright people had noticed that those who regularly took some physical exercise such as walking, had a better chance of staying healthy. They had observed this often enough to consider it a reliable rule. Given that they had Olympic games and were interested in the performances of athletes, they probably also had some plausible theory of *why* exercise made one feel better. Consequently, they were confident in believing that going for walks was an *efficient* cause that had the *effect* of maintaining and even improving your health. People who felt that their physical fitness was deteriorating could, therefore, reasonably decide to use walking as a tool to bring about a beneficial change in their condition.

If I use the method of sequential frames to map this conceptual situation and give *temporal* indices to the individual frames, I get the following diagram:



"S" is the person who wants to improve his/her health. Believing that in the past (t_{1-m} to $t_{1-(m-1)}$), other people "X" have caused their health to improve by walking, the person decides to take a walk.

Thus, the cause of S's change of state, from not walking to walking, does not lie in the future. Instead, the operative cause of the person's change of state is a rule that was empirically derived in the past by means of an inductive inference from observations. It is a rule that is based on the notion of *efficient* cause, and the person is anticipating that it will be efficacious also in the future.

There is nothing mysterious or superstitious about this way of proceeding and it certainly is *not* 'unscientific'. In fact, it is no different from what microbiologists do when they place a preparation under the microscope *for the sake* of seeing it enlarged; or from what astronomers do when they carefully program the mechanism of the telescope to track the star they want to observe. But there is nothing particularly *scientific* about these ways of proceeding either. They are analogous to what we as ordinary people do all day long. We turn a door handle and expect the door will open, we put a seed into soil and anticipate a flower, and every time we flip a switch, we anticipate a specific effect. None of these maneuvers works *every* time, but success is frequent enough to maintain our faith in the viability of the respective causal connections.

Yet, we may call all these anticipatory procedures teleological, because they involve goals which, by definition, lie in the future for the individual actor. When I planned to take a particular train from Brussels, my goal was to be in Liège at a future moment. But it was not this future goal that would *cause* me to arrive there. Instead it was the expectation that the train scheduled to go to Liège would do once more what it had done often enough to warrant my anticipation that it would do it again.

Conclusion

One of the slogans of Skinner's behaviorism was: "Behavior is shaped and maintained by its consequences" (1972, p.16). Hence it would be affected by what happens after it. But there is no mysterious time reversal – only a specious description of what actually goes on. And Skinner was wrong to condemn teleology, insofar as he was speaking of intelligent organisms. Their behavior is shaped by the consequences their actions had in the past, and that is precisely what constitutes the 'empirical' teleology I have discussed above. As Maturana said, we "function in a predictive manner: what occurred once will occur again". In other words, we can learn from our experience and abstract regularities and rules regarding what we can expect to follow upon certain acts or events.

The mysterious feature is our ability to reflect on past experiences, to abstract specific regularities from them, and to project these as predictions into the future. This pattern covers an enormous variety of our behaviors. I would suggest that there are three related but slightly different kinds of anticipation involved.

1) Anticipation in the form of implicit expectations that are a prerequisite in many actions. For instance, the preparation and control of the next step when we are walking down stairs in the dark; this does not require the prior abstraction of cause-effect rules, but it does require familiarity with specific correlations of actions in prior experience.

2) Anticipation as the expectation of a specific future event, based on the observation of a present situation; this is a prediction derived from the deliberate abstraction from actions and the consequences they had in past experience.

With regard to these two forms of anticipation, one may say that robots and other artificial mechanisms are today able to simulate them, if not actually, at least theoretically.

But there is a third form: Anticipation of a desired event, situation, or goal, and the attempt to attain it by generating its cause. This, too, is based on the abstraction of regularities from repeated correlations in past experience, and it is not considered 'scientific' without a conceptual model that 'explains' the cause-effect connection. But even where we have such a model, its simulation presents a problem. When we, the human subjects, pursue a goal and attempt to attain it by using regularities we have abstracted from our past experience, we ourselves have chosen the goal because we desired it. In contrast, Deep Blue, the chess program that can beat a human master and has in its repertoire thousands of cause-effect connections, does not know *why* it is playing or why it should be desirable to win.

The first two forms of anticipation I have listed are sufficient to say that, without them, every step we take would be a step into *terra incognita*. We would fall into precipices, crack our heads against walls, and only by accident would we find something to eat. In short, we may conclude that without anticipatory systems, there would be no life to speak of on this planet.

The third form, that involves the choice of goals, has an ominous significance. If we act in the belief that the future consequences of our actions will be similar to what they were in the past, we ought to be careful about the goals we choose. For whenever we attain them, these goal states may have further consequences, and they will be our responsibility. Therefore, if we don't heed the ecological predictions scientists compute today, we may have nothing to anticipate tomorrow.

References

- Aristotle, Physics (Book II). In R. McKeon (Ed.), *Introduction to Aristotle* (translation by W.D.Ross). New York: Random House Modern Library, 1947.
- Glaserfeld, E.von (1990) Teleology and the concepts of causation, *Philosophica*, 46 (2), 17–43.
- Glaserfeld, E.von (1974) "Because" and the concepts of causation, *Semiotica*, 12 (2), 129-144.
- Gruber, H.E. & Vonèche, J.J. (Eds.) (1977) *The essential Piaget*. London: Routledge & Kegan Paul.
- Hume, D. (1742) *Philosophical essays concerning human understanding*. London: Millar. (In 1758, after the 4th edition, the work was called *An enquiry concerning human understanding*.)
- Maturana, H. (1970) *Biology of cognition*. BCL Report No.9.0. Urbana: U.of Illinois.
- Nagel, E. (1965) Types of causal explanation in science, in D.Lerner (Ed.) *Cause and effect*, (11-32). New York: The Free Press.
- Piaget, J. (1937) *La construction du réel chez l'enfant*. Neuchâtel: Delachaux et Niestlé.
- Rosen, R. (1985) *Anticipatory Systems*, Oxford/New York: Pergamon Press.

Rosenblueth, A., Wiener, N., & Bigelow, J. (1943) Behavior, purpose and teleology,
Philosophy of Science, 10,18–24.
Skinner, B.F. (1971) *Beyond freedom and dignity*. New York: Bantam/Vintage.

This paper was downloaded from the Ernst von Glasersfeld Homepage, maintained by Alexander Riegler.



It is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs License.
To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/2.0/> or
send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, CA 94305, USA.

Preprint version of 8 Mar 2006