Memmings
Herbert Hrachovec

"Provocative insights and controversial views on the future of mankind emerge, as a select group of scientists and artists discuss new approaches towards the formation of bio-social memory." Some such advertisement is going to be attached to the current Ars Electronica project. I will not be able to feign ignorance when I finally see it in print. So let me try to make sense of this interdisciplinary, open dialogue with the verdict not yet in. I will pick three distinctively different approaches to information processing along the vector of time. Meaningful dialogue presupposes initial well-defined inputs from the participants. After attempting to figure out how some of these might look [albeit from the position of an interested layman] a general question will be asked. What, if anything, do those approaches have in common?

Memas

When computer scientists talk about "read only memory" they are referring to some hardware destination where data are permanently stored. This use of "memory" is largely metaphorical. But consider the following definition of the so-called Fibonacci sequence: "a sequence of numbers each of which, after the second, is the sum of the two preceding numbers". An unexceptional paraphrase of this could go as follows. Take any two successive numbers of a sequence and add them together; if the result is the next member of the succession it is a Fibonacci sequence. In order to make sense of this prescription one needs a notion of memory stronger than in the case of ROM and RAM. Extension in time, e.g. simple succession, does not suffice either. The Fibonacci function is sensitive to context. It has to "remember" the two most recent numbers in the sequence. Now, though this manner of speaking is not completely superficial, it is, obviously, highly figurative and can be dispensed with by a suitable algorithm. [This is indicated by the subscript "a" in memas.] Here is one such set of instructions, written in Oberon.
No mention of memory is made here. "BEGIN" and "END" indicate the definition of a process. And there are provisions for looping in case a given number is bigger than 1. If this happens, a certain extension of the process [namely "(fib(n-1)+fib(n-2))"] is run on certain numbers. "Keeping something in mind" is completely substituted by an algorithmic formula.

```
PROCEDURE fib [n:INTEGER] : INTEGER ;
BEGIN
IF n>1 THEN
RETURN [fib[n-1] + [fib[n-2]]
ELSIF n=1
RETURN [1]
ELSIF
RETURN [0]
END
END fib ;
```

Tremendous progress has been made by writing programs that can perform tasks formerly thought to demand old-fashioned human consciousness. But one should carefully delineate the systematic presuppositions allowing this kind of achievment. Granted that "PROCEDURE fib [n:INTEGER] : INTEGER" is a [very primitive] example of computers redefining "memory", the present remarks are far from giving the whole story. Procedures are defined in structured programming, which, in turn, needs compilers, assemblers, appropriate hardware,

electricity and so on. Within circumstances carefully provided by researchers, computer industry and big business, numerous small, effective units of computation, performing memory-like tasks, have been designed. The Memas habitat is sets of instructions meant to run on a computer.

Memos

For an example from the realm of organisms [subscript "o"] look at the immunological system. In addition to "nonspecific immunity", animals — starting from the vertebrate — have developed increasingly complex immunological responses. When antigens invade a given organism it responds by producing antibodies, trying to render the invaders ineffective. Two types of cells are usually distinguished. Effector cells of various kinds act to attack the antigens directly. A second type is sometimes called "memory cells". They are generated at the first impact of the antigen, but do not take part in immunological defense immediately. By duplicating the information of the activated antibodies they increase their population. Since these are long-living cells they enable the organism to react against future attacks of this particular antigen. The working of immunological defense depends on two co-ordinated processes. There are provisions to combat an invader cell immediately, and the organism retains an immunological memory.

Vertebrates don’t remember antigens in any conscious way — neither do humans. So, how does an adequate explanation of "memory" in organisms look? I am, to repeat myself, an amateur, collecting fragmentary knowledge from dictionaries and general introductions into these complicated matters. The course of events seems to go like this. Lymphocytes in general are small, rather featureless cells, derived from stem cells of the bone marrow, acquiring their characteristic components from the lymphoid tissues to which they migrate. The thymus gland and the lymphoid system produce T-lymphocytes and B-lymphocytes respectively.

When coming into contact with antigens those lymphocytes may be stimulated to differentiate into large, active cells that divide several times, resulting in a marked increase in the number of lymphocytes able to react against the given antigen. Stimulation of B-lymphocytes by antigens results both in an increase in the production of circulating antibodies and an increase in the number of the relevant B-lymphocytes themselves. These interactions result in a change of the organism’s overall immunological capacity. When particular antigens are encountered for the second time there are many more antibodies around to fight them off.

This biological mechanism was in operation long before someone hit on the idea of describing it as "immunological memory". Its physiological details can be investigated in a piecemeal fashion, not entirely different from the working of an algorithm. And, while the notion of the immunological system "remembering" previous attacks is obviously misleading, the point made in reviewing programming languages applies here as well. Lymphocytes can only fulfill their desired function within the larger framework of their biological hosts. Stripping away picturesque talk we cannot but restrict ourselves to a very specific bio-structural environment outside of which neither proteins, nor their performance of memory-like tasks, makes any sense. Memos are particular functional units to be found in the circulation of body fluids.

Memis

Let us, finally, consider a feature of intelligence, namely the ability to deceive and to discover deception. Describing these phenomena we can use memory-talk directly. This is the origin of such talk. A prisoner wants to escape by wearing civilian clothes and mimicking the

behaviour of some of his guards. He is presupposing that these characteristics are remembered by the rest of the crew and — mistakenly, in this case — attributed to him. We could never be deceived unless we projected some previously acquired knowledge onto the world. Unless, in other words, certain memorized information misled us. And similarly for disillusionment. The prisoner's escape is prevented by someone noticing the difference between the memory-image projected onto the person and his actual appearance or behaviour. The guard notices some detail that does not fit his memory of his impersonated colleague — and the coup is off.

Although I have used an example taken from human conduct, deceptive strategies are obviously to be found throughout living nature. Once upon a time human chauvinism might have regarded this type of behaviour as a prerogative of rational animals, but Artificial Intelligence research has alerted us towards the possibilities of modelling such involved processes by computational interaction. I shall not go into the current debate on the type of computer resources best suited for this task. My point is that a certain procedural logic has to underly every attempt to build the ability to deceive and discover deception into mechanisms and organisms likewise. The following components have to be operative in order for something to be an experience of disillusionment.

[1] Some set of facts must be represented to a cognitive system. [2] The representation has to be stored for future re-application. [3] One such re-application proves successful, i.e. the conditions of such patterns guiding the cognitive system in a changing environment are temporarily met. [4] Some breakdown occurs during step [3]. Successful re-application fails. There is additional evidence, not accounted for in the pattern presently employed. [5] A crucial alternative arises here. The cognitive system can decide [5a] to disregard this disturbance and continue to rely on its acquired pattern or [5b] to check its initial projection, allowing for a possible mismatch between its suggestions and the state of affairs. [6] If [5b] is selected, and a new representation of the environment is attempted, comparison may show that the initial projection was in some crucial way inaccurate. Since there is a very large number of possible divergences between representational patterns, most of them irrelevant to a given epistemic situation, the definition of illusion has to be context-sensitive. It depends on interest-relative crosschecks between cognitive projections.

The conceptual architecture of deception I have sketched here does not determine the kind of system in which it might be implemented. It does, however, presuppose some main tenets of Cognitive Science, most importantly that we may investigate and construct intentional systems. Just as it sometimes proves very effective to describe a collection of meteorological data as "thunderstorm" or "fog", the working of mechanisms and organisms alike can, for various purposes, be described by employing phrases like "it believes, that such and such is the case" or "it assumes that these conditions hold". Now, memory is usually listed amongst such "propositional attitudes". Remembering is the establishment of a relation between some person and representations of former states of affairs. These are tricky issues and I will just point out two details echoing remarks I made in the case of memas and memos. Firstly, there is no explicit mention of "memory" in my functional account of deception. The category seems to dissolve within the intricate adaptive practice of a given [living] system. The second point concerns the framework into which the modules I have been talking about are embedded. Memas reside in sets of instructions for computers, memos need a very specific, orderly production and diffusion of proteins. The memis required environment is intentionality. A system has to be capable of putting a distance between itself and its environment as well as bridging this distance [and control this very process] in order to learn. Learning is just another word for the systematic modification of memory.

Memes

What about memes? Are there any specifications similar to the ones I have been sketching to single them out? The operation of algorithms, biological processes and the logic of certain cognitive attitudes can be investigated in considerable detail without — it seems — any need to refer to memory. A closer look at the programming code fails to reveal any need for this kind of paraphrase, similarly for the other cases. In elaborating what I introduced as cases of science dealing with memory we actually eliminated it from those accounts. Let me sharpen the dilemma somewhat artificially. We have come up with a procedure, written in Oberon, a story about the multiplication of certain cells in certain circumstances and some logical features of system-environment interaction. Taken at face value they have absolutely nothing in common with each other. Formal instructions are not metabolic processes which are not states of mind. Source code does not directly affect lymphocytes and neither of both has any obvious role in comparing representations. When a quack doctor claims that stars or prayers or parapsychological energies determine the growth of a tumor in some defenseless little girl there is a public outcry. Jumping between computer science, biology and cognition does not provoke such reactions. It is widely regarded as innovative and exciting. This is in itself a remarkable phenomenon.

So, what's the difference between stars affecting tumors and source code affecting lymphocytes? The obvious answer is that in the first case the correlation is entirely unscientific, whereas, having discovered the genetic code, scientists are in a position to establish relevant resemblances to a programming environment, triggering predictions about our future ability to rewrite the biological sequences by means of computers. Much hard work has to go into mutual fine-tuning of the disciplines involved. Scientific claims are empty if they cannot be verified in an overarching, operative framework. The question is: Where does the shared pattern come from? In the shorthand [fake-] terminology of the present essay: How can memas interact with memos and memis? This must be determined within some common, second-level structure. And here we are on familiar terrain. Associations of ideas begin to flow. Both types of modules combine to form a powerful cybernetico-biological entity, the result of fusing computer science and genetic intervention. Incorporating modes of intelligence into the exercise we get an even more encompassing construct, cyberbiology crossbred with self-corrective development of cognitive systems up to and including human culture. In short: the meme.

Any objections? Before committing myself, I want to take a closer look at the argumentative procedure underlying this trans-disciplinary approach. It is suggested that science can establish significant correlations between source code, proteins and intentional attitudes. How is this claim substantiated? In which way does it differ from flashy cyberpunk phantasies? If it is part of the scientific endeavour it has to conform to some set of testable methodological constructions, similar to the ones I started with. Memes, as theoretical constructs, can only derive their respectability by fulfilling precise functions within some perspicuous scientific scheme. I am, again, no expert on such matters. But I will hazard a guess concerning their conceptual role. If they are operative in serious scientific work they will be on a par with memas, memos and memis; there will be no surplus intuitive appeal. The work of a computer-assisted geneticist is not on memes but on recoding chunks of genetic information. And drug research, to take another example, is not on memory either. It investigates physiological enhancements of certain brain activities.

There is presupposition hidden in my way of presenting the case so far. Science has no over-arching notion of memory at its disposal. It consists of a multitude of small- and medium-

scale investigations that may be synthesized by further similar endeavours. All of them are, by definition, confined to a particular set of explorative methods and of a certain subject matter. Any approach to "memory as such" is unfeasible. But how, then, do memes get any content at all? We could, in fact, have raised this question from the outset. Why should we even consider an algorithmic procedure [or inter-cellular transactions etc.] as explanans for something called memory? This question cannot be answered within any of the well-defined procedures claiming to capture [part of] the phenomenon. It touches on the relationship between science and ordinary language. Memory matters to us. It is because of its pervasive importance in human behaviour that scientists are led to search for its physiological, psychological and cultural foundations. Source code and proteins are grouped together because of the overall social interest in finding out about the mechanism of remembering.

All of this opens up a conceptual space for questionable manœvres. I have been reconstructing "memes" as an intermediate abstraction, situated between low-level scientific labour and an extremely general everyday idea. There is an enormous tension here; very different concerns meet at this intersection: politics, the quest for individual fulfillment and the search for scientific knowledge, to name but three. Numerous high level concepts like progress, ecological viability or chaos are embattled in a similar way in modern societies. Those concepts mediate possible interchanges between different traditions and interest groups. ["Evolution", for example, marks a battleground between mainstream science and the religious fundamentalists in the US.] If this is a correct picture a lot depends on how the lines are drawn by each of those mediating concepts. And there are indications that "memes" add to confusion rather than clarify the issues. The term is, on many occasions, used to directly map the layman's understanding of memory to Science. The impression thereby created suggests that there is a sound science of memory, comparable to computer- or genetic science. One main attraction of "memes" is precisely to suggest this analogy to genes. But, as I tried to show, it is unclear what "memory" is mapped onto here. There is, at best, a correlation between an informal notion and a loose cluster of research projects. And the only thing those projects have in common is their being assembled together by using the catchphrase "memes".

Such correlations, I hasten to add, are the sine qua non of innovations in science. Take a phrase like "protein computers". It serves to bring together two threads of thinking that were long considered to be mutually exclusive. After recent advances it is no oxymoron any more. A precise framework to reorganize our conceptual web has been proposed. The following quote is taken from an article on "Protein-Based Computers" [*Scientific American* 3/1995, p.70]: "In addition to facilitating parallel processing, three-dimensional cubes of bacteriorhodopsin provide much more memory space than do two-dimensional optical memories" Robert R. Birge knows what he is talking about — and given the type of analysis proposed above we can begin to make sense of this sentence, even though it freely mixes distinct functional levels. The more general problem, however, remains untackled. There is a tremendous gap between the closely monitored meaning of "memory" in protein computers and everyday intuitions usually associated with the term. My present apprehensions can now be quickly summarized. The invitation to regard ourselves as a society of protein based computers is in danger of falling flat between science and science fiction. "Memes" are a short circuit, not a field of force. A notion whose main purpose seems to be the sudden, irregular discharge of energy.