

Keywords: intelligent system, knowledge, brittleness, inflexibility, frame problem

Ivo Špigel
OZIR, Zagreb

ABSTRACT Artificial intelligence systems today suffer from problems that are widely acknowledged within the discipline: brittleness, inflexibility, the frame problem and others. These problems are largely due to insufficient methodological foresight in system design. In particular, reduction of a system into components and the explicit representation of knowledge (frames, rules etc.) are misused. Research has begun at OZIR to investigate a different class of systems: integral and implicitly intelligent. This paper explains the hypotheses involved and the direction of research.

INTEGRALNI, IMPLICITNO INTELIGENTNI SUSTAVI Sustavi umjetne inteligencije kakve danas poznamo opterećuju problemi kao što su nefleksibilnost, problem okvira i drugi. Nedovoljna metodološka analiza u pristupu velikim dijelom je uzrok ovakvom stanju. Konkretno, rastavljanje sustava na dijelove i eksplicitno predstavljanje znanja pogrešno se ili nepotrebno primjenjuju. U OZIR-u je započeto istraživanje nove klase sustava: integralnih, implicitno inteligentnih. U ovom radu iznesene su hipoteze na kojima se istraživanje zasniva te osnovni koraci istraživanja.

I

Artificial intelligence systems as we know them today suffer from various widely recognized problems and limitations. Some of these problems are due to lack of methodological insight on the part of researchers uncritically rooted in the rationalistic tradition. This paper is the announcement of a research project which has begun at OZIR. The aim of this project is the introduction of a new class of machines which should overcome some of the problems limiting current implementations. These systems, of course, will suffer from problems of their own. Hopefully, a shift in problem focus will mean a little progress for the field.

Within the paper, I will focus on some fundamental problems of artificial intelligence that are:

- a) omnipresent throughout the discipline
- b) important enough to be stifling scientific progress

Solutions to these problems will be proposed in the paper and explored in implementations.

The structure of the paper is as follows. In the first part, the problems are identified. An explanation of their origin and importance is provided. In the central part of the paper, a new approach is advocated. Finally, research motivations and some inevitable social issues are considered.

Before we continue, a brief summary of the problems and the proposed solutions:

- a) AI today studies and designs separate parts of intelligent systems. As these parts are in fact inseparable, this decomposition leads to inherent limitations. The solution is to study and design integrated systems.
- b) Knowledge is represented explicitly. What needs to be represented is not knowledge, but the world of an intelligent subject, and this representation should be implicit.

A brief remark. In this discussion, I consider intelligence to be exhibited by animal life in general. I also consider human intelligence

to be essentially of a higher order than that of other animals. I will usually refer to general intelligence as »intelligence«, and explicitly denote »human intelligence«, except where the context makes the difference clear.

II

The Reductionist Eastern Sin

AI is a young discipline. Not many people within the field have engaged in an analysis of the methodological structure at hand, Winograd & Flores being a notable exception {WF}. The views expressed here are strongly influenced by theirs, as well as by those of Doug Hofstadter {DH}.

The basic approach taken in AI research faithfully follows the rules of Western rationalistic tradition, above all the principle of *divide et impera*. »Let us model our devices after man, and since this is a very complicated model, let us look at some of the parts before assembling the whole.« Some of the parts, however, when assembled, give only a sum of the parts, and this is not what an intelligent system is about.

We have today a variety and richness of subfields and techniques. Computer vision, speech recognition and others study perception. Robot manipulators and legged robots study the limbs and motion – the motoric system. Expert systems, cognitive science and many more study the cognitive domain. It has become obvious by now that none of these domains have integration with other domains as their long-term goal. This is not surprising.

An intelligent system must be conceived as an integrated system.

What exactly does this mean? Well, it amounts to a statement that reductionism and holism must be well balanced in the methodological structure of our discipline. Let me explain further.

In the present situation, reductionism is misused. The model is reduced in *structure*, while retaining *scale*. We need to reduce the scale,

while holistically retaining the structure.

Look at the way children do things. There is always a lot to be learned from them. When a child models a man or a woman, it does not practice on an arm, or a head, or a leg for years, before moving on to the next part. The child »constructs« the whole person, although in a simplified, rudimentary way. The same holds for toy automobiles, or castles in the sand. Children do it that way because it is the *natural* thing to do.

Let us suppose we still wanted to model part by part, function by function. Let's start with the motoric system. How do we walk? We walk with legs. Or do we? Can you imagine one leg walking by itself? Or does it take two legs to walk? Two legs — and no body? Our whole body walks. Our hands walk. Our shoulders, ears and nose walk. And especially our eyes — they do a lot of the walking.

So it isn't at all easy to separate the bodily functions into picture-book parts. The way to proceed is to study intelligent systems as integrated, and design them as such.

Time & Space

Steps In Time

The development of intelligence was, and still is, a process — not a sudden event. This is reflected in the structure of intelligence. Given that, we can accelerate this process, but we cannot skip it altogether, if machines are to be more intelligent than they are today. Let us now consider time on two scales of magnitude: the scale of the species/society and the scale of the individual.

Regarding time and intelligence, I consider some facts to be completely obvious and beyond interesting discussion. There was a time when there was no animal life, and thus no intelligence. At some point later in time, animal life had appeared, and with it intelligence. At this point, there were no people, and no human intelligence. Today, there are people, and they embody human intelligence.

On the scale of the individual, there was a time for each of us when we did not exist. At some later point, we exist, exhibiting intelligent behavior.

These are constant points in time where the relationship between time and intelligence is obvious and simple. The periods of time of interest to AI researchers lie between these points. They are the transient periods:

- 1) The period of emerging intelligence. During this period, the process of formation of general intelligence took place.
- 2) The period of the formation of man, relating to the process of the formation of human intelligence.
- 3) The prenatal/postnatal period in the life of a baby, during which each of us goes from splitting cells to saying »mama«. This period extends from roughly 2 months after conception to roughly 24 months after the birth of the child.

If machines are to be intelligent, they must go through the processes bounded by these transient periods.

The analogies are obvious. Intelligent machines correspond to animal life. The best intelligent machines might correspond to people. Each individual machine must go through a process of intelligence formation.

Let's look more closely at what this implies. We shall not begin by modelling man. Study and design shall begin from the simplest forms of intelligent life — maybe worms or insects¹. If and when we succeed in building a true artificial simple animal, we might move on to higher forms. In this way, scientific research can naturally reflect evolutionary processes in nature.

On the other hand, the machines we design must go through individual processes of intelligence formation. This means that we must find out as much as we can about the structure of an intelligent system at the beginning of the third transient period, i.e. what's hard-wired into an animal at the moment it starts learning by itself. We must then reproduce this structure appropriately in the machine and let it develop, if you will, through self-organization (a favorite buzzword lately).

A Spatio-Temporal Outlook on Life

Time and space are forms of perception. There may be others, but we certainly don't know about them and cannot imagine them. Intelligence as we know it exists in a spatio-temporal framework. Try imagining a world without time, or a world without space {IK}. Thus, time and space are essential to intelligent systems in an even more important way than as a crucial factor in their development.

An intelligent system, must therefore have the ability to perceive both time and space. This »mechanism« cannot, of course, be simplistic, in the form of, say, predicate logic: *before(X,Y)*, *after(X,Y)* etc. The system must be organized in such a way that temporal perception is a consequence of the organization, meaning that every activity of the system and every entity within the system is in some way spatio-temporally situated.


III

Integral, Implicitly Intelligent Systems

Some of the reasons for the belief that intelligent systems must be integral to a degree, and their intelligence implicitly defined, have been laid down. I would now like to go into this in some more detail.

Structure of the System

Integral system design does not mean unstructured design. Integral as it is, the system must have a fundamental structure. This structure is defined by three basic functions:



These basic functions are realized by basic subsystems of an intelligent system. It is essential that they be tightly coupled. Analysis of the internal structure of each subsystem and the organisation of the complete system, defined by the relationship between the subsystems, as well as design based on this analysis, is to be a key aspect of the research being proposed in this paper.

It is necessary for intelligent systems to have at least two sources of perception, in order for them to be able to provide feedback for each other². This coincides nicely with the structure given above, even if the basic perceptive function contains only one source of perception itself. Namely, for the motoric function to be intelligent, i.e. to be part of an intelligent system, it must be able to provide feedback information to the two other functions. Given a system with only one »explicit« perceptive subsystem, feedback from the motoric system can be used as another source of perception.

The Body

If anybody needs to be convinced that locomotion is an essential characteristic of intelligent systems, let us consider an intuitive argument: All animals move. No plants move. As has already been said, it is animal life that we consider to be intelligent. If that is not convincing, we can say that for present purposes the necessity of motoric abilities for intelligence is a conjecture, or hypothesis.

Man, as we see, does not walk by legs alone. In fact, no animal does. Birds fly with their whole bodies, and fish and sea mammals swim with theirs. The whole body of an animal, then, is its locomotion system. The mechanical finesses of live bodies seem too complex to be modelled at the present moment. However, we need not copy nature in every detail. Locomotion can be achieved through the simple machinery available to us, as long as one condition is met.

There is more intelligence in the movement of the humblest worm than in the nicest robots of today. This seems quite obvious to me, considering the complexity of motion of either system. The question is: how come? The reason is that the worm, in its wormlike rudimentary way, *learned to move*. This fact is represented in its nervous system by flexibility. A worm will have no trouble at all climbing up any

number of different branches during its life — it just doesn't appear to care whether they are the same or not.

An intelligent system, then, cannot be preprogrammed to move. It must be preprogrammed to *learn to move*, and then it must have a *necessity to move*, forcing it to learn what it has been programmed to learn — to move. Only in this way can truly intelligent motion be designed, making the motoric system an intelligent subsystem of the whole.

The Senses

If the intuitive argument for the necessity of motoric ability for an intelligent system doesn't seem quite convincing, I hope nobody will have to be convinced that perception is a necessary prerequisite for intelligence. Nevertheless, ideas of a »brain in a bottle«, meaning a completely isolated »intelligence« might sound intriguing to some. I should refer those interested to Kant {IK} and Dreyfus {HD}, where the impossibility of such a concept is explained. Dreyfus, of course, takes his argument too far, deducing that AI is impossible in principle, which does not follow from his discussion. For present purposes, we shall assume that perception is essential for intelligence.

Intelligence, on its part, is essential to perception. Artificial intelligence researchers have learned this the hard way, especially in vision research. It has long been realised that brute-force perception is impossible and »knowledge-based« systems are the only answer. Interestingly enough, none of the researchers have concluded what seems rather obvious. You cannot simply »add« intelligence to a perception system. Intelligence is not pepper or paprika. Perception systems can truly function only as subfunctions of an integrated intelligent system.

To give a vivid example, let's focus on vision for a while. How is it that we see so well? Many factors are vital to our ability, and I shall mention only a few.

Prediction. When we open a book, we don't expect a computer to fall out of it. We expect to see letters, numbers, pictures. Two-dimensional symbols on a piece of paper. These symbols exist as such for us, they are a part of our world. Our perceptive system is, therefore, guided and aided by the cognitive system in the process of vision. Automatically, the area of interest within our vision field is determined by this expectation. We don't scan the whole scene stupidly like the vision systems of today, since we know what to look for and where to find it.

Concentration. Do we see everything we see, and do we hear everything we hear? How many times have you been caught reading comics in class, having completely forgotten about the teacher and what she had been saying? »Humberto, repeat what I just said!« »Sorry, teacher, I didn't hear you.« The perceptive and cognitive system work together to focus only on the issue of importance to the whole system, ignoring others and eliminating or reducing unnecessary processing.

An analysis of the »tricks« intelligent systems use to be able to cope with the truckloads of information perceptive subsystems are constantly receiving has not been conducted within this research effort yet, but will obviously be necessary to guide us in design. Once again: there is no intelligence without perception, and no perception without intelligence. The two cannot be separated. Tightly coupled with the motoric system, they constitute an integrated intelligent system.

Cognition

I would have liked to call this section »The Mind« in accordance with previous headings. It would have been misleading, however, strengthening further the old belief that we walk by legs alone, see by eyes alone, and use our brain only when doing mathematics (excuse the exaggeration).


Within the context of this discussion, I prefer using the term »cognitive subsystem« for the subsystem that does what we usually call »reasoning«. In analogy with considering animals intelligent in a general sense of the word, I consider them to be able to »reason«, or, if you will, reason, in a very, very general sense. I am convinced, however, that the reasoning of a frog is far superior to the »reasoning« of any existing expert system, including MYCIN, XCON, Urologist etc. etc.

Far superior, in fact, to any expert system that will ever be built and still deserve the name.

An expert system doesn't know what it's talking about. This is such a notorious fact and has been so nicely illustrated by Doug Lenat that I call it »the GENSYM problem« after his example {DL}. The fact that Lenat is not capable of deriving the consequences of his own insight is sad, but not a central matter at the moment.

Imagine having MYCIN talk to you and use, instead of those nice English words, unique variables generated by the LISP »gensym« function. »The patient should be treated with neocarboanimalis« turns out to be »lkjh llkjh lëkkhjč kj ljkčlhjk uiooz mnnmbqwerty-zuiopš«, and you get a certainty factor of 0.8. You wouldn't be too happy. The machine, however, couldn't care less. Putting it more precisely, the symbols generated would have the same meaning as the English wording — none whatsoever.

This is the result of formulating knowledge explicitly. Conceptual dependency, frames, scripts, etc. etc. — nothing solves this problem. Just how unaware AI researchers are of this problem can be seen when professors claim that the essence of AI is developing better, more precise formalisms. Neural networks appear to be a very sound step in the right direction, but back to those later. Right now, let's look at the basic functions of the cognitive subsystem:



Learning

A fascinating characteristic of natural learning systems is their non-linearity. A child will need many months to speak her first word. The next one will soon follow, and the speed will increase with the size of the vocabulary. This holds for other domains: motion (walking) for instance. The seeming difficulty, however, of grasping elementary concepts is not a drawback of the system, but a reflection of its strength. The same structural complexity that makes it hard to enter a knowledge domain endows the system with power and flexibility later on.

There is a well known word in many parts of Yugoslavia for uncreative hard-working pupils. They are called »shtreberi« and not very much respected by their bright, lazy colleagues in this country or elsewhere. The fundamental flaw in the knowledge of these kids is a lack of real understanding of the subject matter. In extreme cases, they learn it by heart without having any idea of what they are saying.

No matter how sophisticated machine learning schemes may be today, they are truly ideal »shtreberi«. They have absolutely no understanding of the subject matter and do not relate their »knowledge« to reality.

One of the key reasons for this situation is that AI systems have no real contact with the real world, i.e. they are completely isolated from experience. In a metaphorical way, authors of expert systems talk of their »experience« in the domain and the way they learn from it, but this is of course quite far from truly experiencing the world of an intelligent subject.

The only road to *knowledge* is through *learning* from *experience*, and so these three key issues are inextricably interwoven. This does not mean we are lost in a vicious circle. We must carefully analyse the rudiments of knowledge and the mechanism for learning that are present in existing intelligent systems in the period in which they grow from bunches of splitting cells to animate organisms (albeit prenatal) capable of learning from their experience. It is only this much, or rather an analogy to it, that we shall explicitly program into the system we are designing.

The fact that systems learn through contact and experience in the real world means that we cannot build two identical systems, since they cannot share worlds. »The objective world«, namely, does not exist as such. The world that matters to an intelligent system is its own, unique world, so through learning from this world each system develops unique knowledge.

As the system develops and ages, it loses flexibility. Is this only because the intelligent systems we know are biological? I should think the reason lies in the structure that provides it with learning capabilities in the first place.

Memory

The concept of artificial memory is as old as computers are. One could speak of earlier data storage systems, such as libraries, as memories, but this would be stretching the concept a bit.

The ability of computers to store and retrieve data can be quite confusing when discussing intelligent systems. Apparently, we don't have to worry about enabling the system to remember the way we have to worry about enabling it to learn. This, of course, is a misconception, much in a way electric motors and wheels do not automatically guarantee intelligent motion. We have to purposely make systems remember the hard way, in an intelligent manner. RAM and ROM can only give us the technological foundations.

In biological intelligent systems, a bunch of neurons do not memory make. There are intricate and complex memory mechanisms at hand, and even neurophysiologists do not completely understand them. Since we are mere engineers, it is not our job to analyse memory from a neuroscientist's point of view, but we must at least be in touch with what is known. If we model this well enough, they might benefit from the insight gained in empirical experiments.

Among other aspects, the structure of our memory provides us with the notion of time. The way short-term and long-term memory function and communicate, the way earlier events are »buried deeper« in memory, these and other mechanisms are essential to intelligent behaviour. We shall therefore provide our designs with at least rudimentary analogies of what we know, in hopes of being able to introduce more complex structuring with the passing of time and the growth of our own experience.

Control, Communication

Some views on control have already been explained in the section on motion. Let it suffice for the present to say that there is a certain distinction between control and communication from the position of the intelligent subject that may not be acceptable to all. Namely, communication can be seen as an act of controlling one's communication mechanisms, or subsystems. I consider communication to be important enough to be considered separately.

Apparently, communication is as widespread in the world of intelligent systems as motion. That is, all intelligent systems communicate in some way, and no unintelligent ones do.

We are only beginning to understand the structure of animal communication, and we are constantly pushing the limits of what we see as the potential for animals to communicate in a way more familiar to us.

Whatever forms communication takes, however, it is always intelligent in the sense that the communicating agent is. Bees do not talk. AI, however, has been trying to make completely unintelligent systems communicate in spite of this. The GENSYM problem in communicating with expert systems has already been explained. Another nice example are speech generation systems, neural network based or not. These systems do not communicate in any substantial sense of the word — they simply transform text from one form into another. This is not what is needed. An intelligent system must communicate for a reason, and in a way that corresponds closely to its structure and complexity. We shall start by building simple systems with rudimentary communication abilities in order to gain a deeper understanding of the processes involved and, more importantly, because the gradual evolution of one system from another is crucial to the structure of higher order intelligent systems.

IV

Where Do We Begin?

What, then, are the implications of what has been said to design and research? The reason for research is an inquest into the nature of intelligence. A set of design principles should follow from insights thus gained, enabling experimental validation of various hypotheses.

The first stage of research involves building a working model of an animal of rudimentary intelligence. This stage breaks up into four steps:

1. Selection of an animal to model. We need to decide upon a specific animal. Our selection criteria are that the animal be as simple as possible, while still displaying rudimentary intelligence. The earthworm is a potential candidate at the moment.
2. Analysis of the structural and functional organization of the animal in terms of the motoric, perceptive and cognitive subsystems described above.
3. Mapping of this organization onto a system lending itself to practical realization based on the computer technology available to us. This, of course, is the crucial step, and amounts to building the core of the model.
4. Specification of a system based on this mapping and physical implementation. The result should basically be a moving robot, which need not necessarily be a physical analogy of the animal (e.g. in terms of a similar locomotion system), but the functional and structural mapping should preserve the basic organization of the original.

Various basic functions and subsystems of our design represent different design problems. One of the fundamental differences is that in some areas — sensing, learning — the problem is development of the function, while in others — motion, memory — there is double trouble because of the need to refrain from the tantalizing possibilities offered by technology. In a Baconian way, we must hang weights on the wings of technology, providing the system perhaps with wheels and motors, but depriving it of the luxury of ready-made control software. The system, like a child, must be forced to learn to do things by itself. Otherwise, it will be the perfect spoiled child — completely unable to cope to a degree that will make it unintelligent.

This fairly short specification implies quite a few design problems, relating to issues already mentioned as unclear in the text. If we can solve these problems even in a rudimentary way, the step from an artificial bug to an artificial frog might be much easier to take, much as the child has the most trouble grasping elementary concepts.

V

Related Issues

Knowledge Representation

The fundamental problem with knowledge representation has already been mentioned in the text, but since this is such an important and favorite child of AI researchers, I would like to say a few more words about it.

Widely accepted concepts usually have an implicit justification that is not necessarily true. The justification for parliamentary government, for instance, is the technological inability of society to enable everybody to directly influence decisions of general importance. This was O.K. until yesterday, but information and telecommunications systems are rapidly challenging the justification.

A similarly implicit, only completely false justification for knowledge representation schemes is that *we must* represent knowledge in some way, since we shurely don't carry people, houses and elephants around in our heads. The only problem with this is that *knowledge* is not what is represented in the cognitive system — it is the world, our world, the unique world of the subject itself. Various knowledge representation schemes are actually representations of their author's understanding of how people's mind's work. This understanding is

not necessarily complete in each particular case.

Storing knowledge explicitly, in the form of frames, rules, logic etc. amounts to creating an illusion that the system understands. Joseph Weizenbaum realised the implications fully a long time ago, and Margaret Boden has extensively commented on the matter {JW, MB}. The systems and schemes, however, proved quite useful and have remained with us to this day.

I don't claim to have any deep insight into the workings of the mind or of the brain. However, I am painfully aware of this, as well as of the fact that I cannot devote the rest of my life to any of the numerous scientific disciplines covering the domain. My colleagues and I must learn enough to be able to communicate creatively with psychologists, neurophysiologists and many others. A collective, interdisciplinary effort is required to solve these difficult problems, and no ad hoc, simplistic solutions will do.

Knowledge must be implicit, and intelligence is an epiphenomenon of the organization of the system.

The Frame Problem

The frame problem is notorious in AI. One of the most frustrating aspects of this problem is the apparent ease with which people handle it. Now, where did this *enfant terrible* originate in the first place?

The roots of the frame problem lie in Eden, in the days before AI committed its eastern sin. When knowledge is represented explicitly through clever ad hoc schemes, the problem can be solved only by devising still cleverer and cleverer and ... counterschemes to tackle it.

When the organization of the system embodies knowledge, everything that is known is distributed within the system as far as is natural to the context. The consequences of any action are then limited by the distribution of the entities involved in the system and naturally affect only the appropriate environment.

The frame problem is a non-problem if the system is intelligent in an integrated, implicit way.

Neural Networks

The architectural concept of neural networks has begun to answer one of the fundamental problems mentioned in this text — the problem of implicit representation of knowledge. Immediately, results have shown them to be superior to standard techniques in many application areas.

This concept, however, does not address the other fundamental issue — integration. The idea of an integrated system built around a neural network seems promising. Interestingly enough, even the staunchest critics of AI {HD} seem to be sympathetic, or reserved in the worst case, when discussing NN's.

Motivation & Social Aspects

In my experience, it has not usually been the case that scientists analyse their personal and social motivations for doing what they are doing. Again, there is an implicit justification for research, rooted in the Western Judeo-Christian tradition, stating basically that scientific advances automatically benefit all humanity. This has, of course, been questioned strongly in this century, and many people believe today that the major contribution of the space program to society has been the teflon pan. I would almost agree with this view, even if it is a bit extreme.

I believe that a strong personal motivation is present in doing AI research. There is a strange feeling of playing God about this discipline: create something in your own image, something which behaves remarkably like yourself. This motivation has been furthered by ad hoc concepts such as the Turing test, which even define artificial intelligence as the ability of a machine to imitate a human being.

AI is among the few sciences which have the potential of rapidly breeding enormously powerful technology, thus potentially threatening many people, either through weapons or social unrest resulting from industry transformation. This does not mean it shouldn't be investigated. We who are doing it, however, must be intensely aware of the implications and possible consequences of our work. It is our social responsibility to guide and control the results of our research whenever we can, and avoid involvement with projects where

they might be misused.

I agree largely with Weizenbaum {JW} that application domains of AI systems should be carefully selected. Where exactly to draw the line can be a matter of discussion, but it won't do to just blindly stumble into any domain that one considers interesting without giving some thought to the consequences of developments within that domain.

An AI School

AI research as it has been described here requires a different sort of education than any of us get today, in Yugoslavia, in the States or elsewhere. Specialization, so dominant a trend in the decades past, will not suffice any more.

An institution is needed that will provide AI students with a wide knowledge of domains central to the discipline: mathematics, philosophy, computer science, biology, neurophysiology, linguistics, psychology and others. Students would, of course, specialize in one aspect and research subject, but the depth of insight into one particular domain must be partially sacrificed to make way for a breadth of knowledge that is a necessary prerequisite for studying and designing integral, implicitly intelligent systems.

Prologue

There can be no strictly formal theory of artificial intelligence. Human intelligence is creative, thereby having the potential of always going one step beyond any formal definition. This means that we cannot define intelligence in general, because we cannot define its most interesting form: creative intelligence. Since we cannot define one of the central concepts, we cannot ever hope to construct a full, complete system deserving to be called a »theory of artificial intelligence«.

Given this, we at OZIR have stopped worrying about definitive solutions and are trying to do the best we can with what we have. This is all we are attempting by embarking on the study and design of integral, implicitly intelligent systems.

Acknowledgments

I wish to thank all my colleagues at the Department for the Study of Intelligent Systems at OZIR for many hours of inspiring discussion and a few years of AI research, especially Ivan Maršić who began it all. This is the end of the paper but only the beginning of interesting research.

Notes

1. I owe this idea to Ivan Maršić.
2. Vukašin P. Masnikosa, Ivan Maršić. I would certainly like to see stronger biological support for this view, which I nevertheless consider intuitive enough to be convincing.

References

1. {MB} Margaret Boden: *Artificial Intelligence and Natural Man*, second edition, The MIT Press, London, England, 1987.
2. {HD} Hubert Dreyfus: *What Computers Can't Do*, Nolit, Belgrade (in Serbo-Croatian)
3. {DH} Douglas Hofstadter: *Godel, Escher, Bach: An Eternal Golden Braid*, Penguin Books, Harmondsworth, England, 1987.
4. {IK} Immanuel Kant: *A Critique of Pure Reason*, BIGZ, Belgrade, 1976. (in Serbo-Croatian)
5. {DL} Douglas Lenat et al.: *CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks*, AI Magazine, Vol. VI, No. 4, Winter 1986
6. {WF} Terry Winograd, Fernando Flores: *Understanding Computers and Cognition — A New Foundation for Design*, Ablex Publ. Co. Norwood, U.S.A., 1987.
7. {JW} Joseph Weizenbaum: *Computer Power and Human Reason*, Rad, Belgrade (in Serbo-Croatian)