

This article was downloaded by: [Yoav Yigael]

On: 18 November 2011, At: 21:27

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



## World Futures

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/gwof20>

## Fundamental Issues in Artificial Intelligence

Yoav Yigael <sup>a</sup>

<sup>a</sup> Gran Shmuel, Israel

Available online: 07 Nov 2011

To cite this article: Yoav Yigael (2011): Fundamental Issues in Artificial Intelligence, World Futures, 67:8, 564-568

To link to this article: <http://dx.doi.org/10.1080/02604027.2011.625742>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

## WORK IN PROGRESS

### FUNDAMENTAL ISSUES IN ARTIFICIAL INTELLIGENCE

YOAV YIGAEI  
*Gran Shmuel, Israel*

Artificial Intelligence (AI) encompasses a wide and diverse field of research, which aims to apply human abilities to the field of technology (computers, robots, etc.). The possibility of achieving such an application is not at all obvious; therefore, there is room to question whether the field of AI has been able to identify and understand the nature of the difficulties involved in the attempt to apply human abilities to machines.

At least at one of the discussion panels held at MIT's "Brains, Minds, and Machines" three-day symposium, the prevailing opinion was that there is something amiss in the current trend of AI research studies.<sup>1</sup> The panel, which included some of the most prominent figures in the fields of AI, cognitive science, and linguistics, was convened in order to discuss the state of AI research today.<sup>2</sup> Most panel members agreed that the last time any significant progress was achieved in the AI field was in the 1980s, and that since then, something had gone wrong along the way. I present here some of the major points raised during the panel discussion, to provide a point of departure for a more comprehensive appraisal of the field of AI.

The panel moderator, Steven Pinker, a renowned cognitive scientist at Harvard University, claimed that current research in the field of AI is focused on narrow issues. He advocated a return to the kind of research that was characteristic of the field in its earlier years, namely, research driven by curiosity rather than by a search for potential applications.

Marvin Minsky, a pioneer in the development of neural networks who early on made significant contributions to the fields of AI and robotics, commented that today's students are busy developing sophisticated machine applications without making the machines any smarter.

Patrick Winston, who has been the director of MIT's AI lab since 1972, agreed with members of the panel that something went wrong during the 1980s. Although some progress has been achieved, it was expected that by now significant strides would have been made. In his view, the focus on mechanisms has led to the investigation of issues that can be explained within this context. As a result, other, more fundamental questions have been pushed aside.

---

Address correspondence to Yoav Yigael, Gran Shmuel, 38810, Israel. E-mail: yoav.yg@gan.org.il

Emilio Bizzi, one of the founders of the Brain Research Institute at MIT, exhibited a more optimistic view regarding our ability to utilize important aspects of human intellect, such as the ability to create generalizations based on a range of experiences, the ability to plan flexible movements in order to avoid obstacles on the way to a specific, predefined goal, for example, when reaching for a pair of glasses.

A well-known linguist, Barbara Partee, claimed that a precondition for what we refer to as *intelligence* is a true understanding of the idea of *significance*. Noam Chomsky, one of the most prominent linguistic scholars to date, noted that the prevailing trend in AI research today is to build systems that imitate parts of human behavior. However, imitation that fails to take into account the *significance* of said behavior cannot—by any means—be considered sufficient.

Nobel laureate Sydney Brenner, one of the prominent investigators of the genetic code, suggested that, not unlike brain research, AI research has been overly focused on describing details on the surface, consequently devoting too little effort and resources to examine the bigger questions that lurk beneath the surface.

I chose to present the points raised by the panel that emphasize aspects that ought to be reconsidered as part of a reassessment of the field of AI. The combined profile of panel members is in itself an indication of the numerous subjects and the complexity of topics that inform the field of AI: technology, cybernetics and information science, cognitive science, linguistics, biology, neurobiology, and so forth.

It should be noted that the opinions expressed by the panelists at this event may not represent the current trends in the field of AI. In fact, the opposing opinion, promoted by the futurist Ray Kurzweil and accepted by many others in the field, is referred to as the *singularity theory*. Supporters of the theory claim that there are no major issues that cannot be resolved; on the contrary, they believe that the current course of development of the field of AI is on the right track. According to Kurzweil, in the not-too-distant future, we will have the capabilities to create machines that can calculate and think in a manner similar to that of humans, and even surpass many aspects of human ability. It would appear that most of the experts in the field of AI concur with Kurzweil's assessment, rather than that presented by the panelists.

I would like to address here some of the assumptions underlying the view proposed by the panelists. The first question I raise was not explicitly discussed by the panelists, yet it is fundamental to the overall issue of AI: Are human capacities in general and human intelligence in particular a mechanism that can be copied and applied to mechanical functions?

Certainly modern biology treats the organism as a type of mechanism.<sup>3</sup> However, it remains to be determined whether the principles underlying the functioning of an organic mechanism are the same as those of an inanimate mechanism. To clarify, this question does not ask whether the same principles of physics govern the organic and the non-organic mechanisms; rather, it asks whether the principles that explain the functioning of an animate mechanism are identical to those that explain the functioning of an inanimate one. To the best of my understanding,

the assumption that guides the current trend of AI research appears to be that the two mechanisms do function—at least in part—on the same set of principles. In other words, those researching the field of AI operate under the assumption that the major issues they face are practical rather than fundamental. Indeed, many of those involved in the field of AI consider the brain to be a type of machine and, therefore, they consider it be possible to copy and apply its capacities to other machines.

Proponents of the assumption that there is no fundamental difference between an animate and an inanimate mechanism should be aware of the following important differential factor: all of the functions of an organism, including all human mental activity, are geared toward the goal of the survival of life (one's own or one's progeny).<sup>4</sup> Each and every organism functions—in fact must function—to sustain life. In contrast, there is no machine that acts to sustain or preserve itself. Neither the existence nor the duration of the existence of a machine depends on how it works or whether it's on or off. A machine is built so that its functions may serve entities other than the machines themselves.

This distinction raises the following question then: Does the difference between the purpose and objective of an organism and those of a machine have implications for the possibility of creating machines that copy or imitate human activities? I will consider this question in light of observations made during the previously mentioned panel discussion.

1. The issue of semantics, which was raised by both Partee and Chomsky, can serve as a point of departure for considering the difference between organic and inorganic mechanisms. Semantics is an important branch of linguistics, which examines the relationships between word and meaning.
2. The act of attributing meaning to actions and things is not unique to human language. Indeed, an organism's every act has meaning, or, more generally stated, every action is an indication of something. Thus, for example, the sensation of hunger is an indication of a deficiency and of physiological changes that signal potential malfunctioning of the body's activities. The amount of attention allocated to the sensation of hunger is an indication of its importance at a given moment. The accumulation of actions directed to obtaining and preparing food is an indication of the efforts invested toward this goal. The aroma and flavor of the food are indications of its quality and freshness, and the satiated sensation that follows eating is an indication that the goal has been reached (activating all of the relevant micro mechanisms that ensure the organism's continuous existence). This is merely one example among the innumerable activation cycles that support the existence of life, in which every component in the cycle indicates the next activity required.

These indications can be categorized according to the particular objective, for example, indications that serve to ensure that a completed action indeed fulfilled the required objective, indications that serve to enhance the functioning and quality of a given action, indications that serve to emphasize the importance of an action in the general scheme, and so on. These indications help the

organism repair itself, and they are all linked—directly or indirectly—to the basic objective of preserving the organism's life. They address a single question: Are the organism's activities and the activities in its environment supporting or threatening its existence? In an organism, actions and activities cannot be separated from their indications. The Even the activities of the simplest, the least sophisticated of organisms, are indications of something else.<sup>5</sup>

To the extent that there are machines that feature structured indications for self assessment and repair, these machines are oriented to assess and repair human objectives, rather than objectives intended for the survival of the machine itself. The indications mean nothing to the machine itself, and the actions that the machine carries out have no significance for the machine itself. Therefore, the issue of copying human abilities and integrating them in machines must address this question: How can a machine be made to behave like a human if it is incapable of perceiving meaning?

3. An issue raised by Minsky regarding the direction of current research was that robots are now more sophisticated and can handle new kinds of actions, yet they have not become any smarter. This raises the question of whether the ability to *compute* and *calculate* can be equated with the ability to *think*. To date, enhanced performance in terms of computation and calculation has not resulted in enhanced thinking abilities. It is safe to assume that human thought is a complex process, of which computing is but one aspect. As an example, it is difficult to imagine how the following aspects of thought might be computerized.
  - a. The ability to generalize from accumulated experience (Bizzi's claim): to perform this act, a machine must have an indication of the relative value that the experience holds for the machine itself. How would it be possible to generalize from experiences that hold no meaning for the agent of the experience? If they are meaningless, how can they be assessed?
  - b. Thought processes, such as exemplified in conceptualization, are based on a logic that cannot be computed using the currently available technologies. In conceptualizing there is a shift from a micro to a macro perspective, an act that combines selection and classification of phenomena and the creation of groupings based on a shared principle. The principles for sorting serve to establish order, which is important for the proper functioning of human society in general, rather than for the individuals that comprise it. These acts of thinking pertain to the logical category referred to as the process of *induction* (which will not be explicated here). For the purposes of the current discussion, it is important to note that for this type of sorting to occur, the process itself should have meaning for the sorter, that is, the machine, or group of machines (if it makes sense to speak of such an entity).
  - c. The claim introduced by Brenner during the aforementioned panel discussion, namely, that research has not devoted sufficient efforts to investigate beneath the surface, can be considered from several aspects. In complex systems, levels of organization constitute a multi-dimensional issue; here I refer to only one aspect. AI systems are supposedly (self) consistent and recursive, and therefore the famous theorem proven by Kurt Gödel is

applicable to AI systems. In general terms, the theorem proves that recursive systems are limited in their ability to relate to themselves. Human thought is able to overcome such limitations (by means of various levels of awareness and consciousness). In order to imitate the performance and apply the abilities of human thought to machines, this limitation of the machine must be addressed and overcome.

These are but a few of the questions concerning the application of human abilities in the field of technology; indeed, there are many more to consider. We may hope that ways will be found to examine and discuss the underlying assumptions that guide current AI research, so that future research may produce significant breakthroughs in this field. As in any other field, so too in AI, identifying the problems is the first step toward their solution.

### NOTES

1. The symposium took place in early May 2011.
2. The current article is based on the minutes of the panel discussion, reported in MIT's *Technology Review*, published May 4, 2011.
3. *Mechanism* is defined as any system that functions according to the known laws of nature and does not require a power unknown to us. For our purposes, the definition of *mechanism* need not address the issue of whether it is animate or not or whether it is natural or artificial.
4. For a discussion on differentiating the phenomenon of Life from other natural phenomena and the implications of this distinction for concepts of evolution of the species (including the human species), see: Y. Yigael, Life: Definition, Origin, and Evolution, *World Futures: Journal of General Evolution* 66, no 8 (2010): 573–596.
5. In the human organism, there is no guarantee that the indications themselves are functioning properly. The issue of malfunctions in indications is complex and beyond the scope of the current article.