

Fluid Concepts and Creative Analogies: A Review

Bruce Burns

■ *Fluid Concepts and Creative Analogies*, Douglas Hofstadter and the Fluid Analogies Research Group, Basic Books, New York, 1995, 518 pp., \$30, ISBN 0-465-05154-5.

There is an apt analogy for the author of this book: Douglas Hofstadter could be considered the Carl Sagan of AI. As Hofstadter points out, analogies are *fluid*, meaning that the analogy between two entities can be drawn differently depending on how these entities are represented. The analogy that is drawn, in turn, can change the representation of the entities being compared. Thus, the analogy between Hofstadter and Sagan can be seen as positive: Both have explained important concepts in their fields to a wide audience and transmitted the excitement of these ideas. Both have inspired a number of people within their fields. Unfortunately, a more negative analogy between Sagan and Hofstadter is possible. Among some astronomers, Sagan's work has not been taken seriously. In a similar fashion, Hofstadter's work has had surprisingly little impact on the field of AI. Some of the reasons for this lack of acceptance in the field are based on irrelevancies. For example, Hofstadter is the only researcher in AI with both a Pulitzer Prize and his own news group on the Internet. However, there are some valid reasons for these reservations. This book illustrates both the negative and the positive analogies.

Fluid Concepts and Creative Analogies is a collection of articles, many already published, by Hofstadter and members of his Fluid Analogies Research Group (here represented by

Dave Chalmers, Daniel Defays, Bob French, Gary McGraw, and Melanie Mitchell). Although Hofstadter is not the first author of all the articles, there is a coherent vision to the whole book, reinforced by the prefaces Hofstadter wrote for each article.

Two major ideas appear to underlie much of Hofstadter's work; although these ideas are related, one is primarily a computational claim, and the other is primarily psychological. The computational claim is that cognition can be modeled using a *parallel terraced scan*, which builds a representation by parallel investigation of many possibilities with different levels of commitment. The psychological claim is that cognition is a form of high-level perception that involves

*Douglas Hofstadter
could be considered the
Carl Sagan of AI.*

a constant interaction between how we represent information and how we process this information. These themes are repeated throughout the book. Those wanting a quick tour should read Chapters 2, 4, and 5. Chapter 2, "The Architecture of JUMBO," gives a good explanation of a parallel terraced scan. Chapter 4, "High-Level Perception, Representation, and Analogy: A Critique of Artificial Intelligence Methodology," explains what high-level perception

is. Chapter 5, "The COPYCAT Project: A Model of Mental Fluidity and Analogy Making," is perhaps the best description of an implementation of Hofstadter's ideas.

The essential idea behind a parallel terraced scan is that initially, the representation of an entity consists of a number of unrelated elements. The representation is then elaborated by many *codelets*, small pieces of code that do specific things. Some codelets recognize potential structures but do not build them; others build structures that have been identified. These structures are often simple, such as the combination of two elements. Because of the parallelism, competing structures can be built; other codelets can destroy a structure that has become inconsistent with the higher-level structure. The structure that emerges is probabilistic because it results from interactions among the actions of many codelets that are executed probabilistically from a *coderrack* of codelets: Every codelet waiting on the coderrack has some nonzero probability of being executed at any time, but once a codelet is executed, it can change the environment so as to change the probability of waiting codelets being executed. Executing a codelet can also lead to other codelets being placed on the coderrack; for example, if one codelet recognizes a potential structure, it can place a codelet on the coderrack to build the structure. As coherent structure is built, the temperature is lowered, making it harder to execute codelets that would tear down the existing structure. Thus, a coherent structure emerges without the need for any top-level control.

A parallel terraced scan is similar in philosophy to other approaches that consider representation as an emergent property of the interaction of many lower-level processes, such as neural networks and genetic algorithms. The link to genetic algorithms seems particularly strong, so much so that it would not be surprising to see some attempt to learn codelets using genetic algorithms. Such learning might address a potential criticism of codelets; they are pre-coded by the designer of a particular

system. However, Hofstadter is less interested in the exploration of the algorithm's properties than in the fact that a parallel terraced scan produces the behavior that is of real interest to him—the probabilistic emergence of structural representations. Hofstadter is attempting to combine the ability of neural networks to detect similarity with the power that symbolic systems get from being able to represent structure.

The central idea behind Hofstadter's work is that cognition is high-level perception, an idea captured in the slogan *cognition is recognition*. *Low-level perception* is defined as the processing of information from the sensory modalities, and *high-level*

with especially tough problems" (p. 187); rather, it provides a powerful way of fleshing out the representation of any given situation. Further, the flexibility of analogies (they allow concepts to "slip" to related concepts) explains the fluidity of concepts; Hofstadter therefore uses analogies to account for creativity.

The line between high-level and low-level perception is fuzzy, but the implication Hofstadter draws for AI research is clear: Too often AI has focused on the processing of rigid representations and has ignored how these representations are formed and how their formation interacts with their processing. (Hofstadter exempts machine vision and language pro-

competing programs have addressed real-world problems in science and politics. Hofstadter and his colleagues argue with some force that this criticism is misguided because there is so much information available in the world that a fully developed model of high-level perception is not possible at this stage. He contends that the relevant issues can be addressed in a restricted domain. To do so might seem less impressive than solving real-world problems, but programs that address such problems actually restrict their domains, making them tractable. Further, the part played by the stripped-down nature of real domains in a program's success has not always been acknowledged. Hofstadter has a worthwhile point here; AI research has ignored the advantages of carefully constructed microdomains, but it has allowed the use of vaguely constructed domains that are also restricted. The use of microdomains is not without problems, however, because they, too, can contain hidden assumptions. It is not clear to what extent the programs presented in Hofstadter's book succeed because they avoid representational issues that are relegated to low-level perception.

Another possible reason why Hofstadter's impact has been limited can be traced to his method of engagement with those he disagrees with. His form of argument is often based on analogies and rhetorical questions rather than the analysis of what a program aims to achieve and how well it achieves its goals. The facility of his writing sometimes makes it too easy for him to stray in the direction of ridicule. Thus, in Preface 4, The Ineradicable ELIZA Effect and Its Dangers, Hofstadter claims that the apparent successes of a number of AI programs are analogous to that of Weizenbaum's (1976) ELIZA program, which appeared to be interacting with a person but was really only giving back canned responses. The danger of a program's success being based on its simply giving back what the programmer built into it is well known, and there is probably no program that does not partly succeed because of what is built into it. How-



Hofstadter makes many compelling and evocative points in this book. Why, then, has his work had limited impact on the field of AI?

perception is extracting the meaning from the raw material of low-level perception by accessing concepts and making sense of information at a conceptual level. Such an approach to cognition is not new; it was the central idea of the Gestalt psychologists, and Hofstadter traces its roots to Kant. However, modeling cognition as a form of perception contrasts with the more dominant approach in AI, which assumes that cognition is a process of manipulating symbols. In Hofstadter's view, cognition is instead modeled as an interaction between the building of representations and the manipulation of representations. Hofstadter's work has come to focus on analogy: He claims we are constantly perceiving situations in terms of what we already know. Analogy, defined this broadly, is not just a "heavy weapon wheeled out now and then to deal

with especially tough problems" (p. 187); rather, it provides a powerful way of fleshing out the representation of any given situation. Further, the flexibility of analogies (they allow concepts to "slip" to related concepts) explains the fluidity of concepts; Hofstadter therefore uses analogies to account for creativity. The line between high-level and low-level perception is fuzzy, but the implication Hofstadter draws for AI research is clear: Too often AI has focused on the processing of rigid representations and has ignored how these representations are formed and how their formation interacts with their processing. (Hofstadter exempts machine vision and language pro-

cessing from this criticism because they stop short of modeling processes at a conceptual level.) It is not enough to leave representation to a yet-to-be-developed representation module because the interaction between representation and process is crucial. For example, Hofstadter argues that BACON (Langley et al. 1987) misses the point by aiming to show that given the appropriate data, the program can derive Kepler's third law of planetary motion—because the key to Kepler's discovery was identifying the appropriate data. Hofstadter makes many compelling and evocative points in this book. Why, then, has his work had limited impact on the field of AI? One criticism has been that all the programs created by his group have addressed only microdomains: letter-string analogies, table settings, word puzzles, and letter fonts. Meanwhile,

ever, the ELIZA effect is not an all-or-none proposition, and the evaluation of the degree to which a program succeeds for interesting reasons and to what degree it succeeds for uninteresting reasons needs to include an analysis of the program's assumptions and goals. Such an analysis also needs to be applied to Hofstadter's programs: Do they really solve the problems that they criticize others for failing to address? For example, Falkenhainer, Forbus, and Gentner's (1990) SME analogy-mapping program is repeatedly criticized because the symbols it uses are not grounded; for example, it would operate just as well if the symbols *water* and *ice* were replaced with *A* and *B*. This result is unsurprising given that syntax drives the mapping process in their theory. However, the same argument turned on COPYCAT is quickly dismissed because if an astute human observer continually saw the symbol SIGMA invoked in the presence of successor groups and successor relationships, then this observer would quickly figure out that SIGMA stands for the idea of successor. But the symbol grounding problem is not so easily solved. How does the observer know that the concept of successor group is present? Why could not an equally astute observer, on seeing *water* and *ice*, infer that the arbitrary symbol with them means *melts*?

Hofstadter's criticisms raise the whole issue of how AI programs should be evaluated, the topic of Preface 5: The Knotty Problem of Evaluating Research. In this preface, Hofstadter considers and then dismisses a number of AI models, often explicitly on the grounds that they fail to agree with his intuitions (for example, the reference in the index to SOAR for this section is "SOAR program...ambitiousness yet boringness of" [p. 515]). Hofstadter's point is that ultimately, all AI research is based on intuitions that cannot be defended. However, he firmly believes that AI is a field searching for its foundations; so, it is not surprising that any project is based 90 percent on intuition (he acknowledges that it is dangerous to make this point in print). Hofstadter has a

point here because AI models rely heavily on assumptions that are difficult to test, and thus, they are extremely hard to compare and validate. These difficulties lead to the common occurrence of rival researchers talking past each other. Nonetheless, I cannot agree with his response, which appears tantamount to withdrawal from the argument. Even if only 10 percent is not intuition, then this 10 percent should not be ignored because it might be the basis for a future foundation. Further, if you want people to build on your work, you must convince them that your intuitions are correct; in which case, you are required to do more than just tell them that you are correct and make eloquent analogies.

If Hofstadter truly believes that AI models cannot be defended rationally, then it might explain why his analysis of other models can appear shallow, but by doing so, he risks irrelevance. This outcome would be extremely unfortunate because I believe his basic ideas might be on the right track. However, figuring out what exactly his ideas mean and defending them will take more than Hofstadter telling us his intuitions. I recommend reading this book because it is eloquently written and provocative. However, at the moment, Hofstadter makes his challenge too easy to ignore for those who should take the most notice of it.

References

- Falkenhainer, B.; Forbus, K. D.; and Gentner, D. 1990. The Structure Mapping Engine. *Artificial Intelligence* 41: 1-63.
- Langley, P.; Simon, H. A.; Bradshaw, G. L.; Zytkow, J. M. 1987. *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, Mass.: MIT Press.

Bruce D. Burns is a postdoctoral fellow in the Department of Psychology at the University of California at Los Angeles. He has done both experimental and computational modeling work on analogy.

Computation & Intelligence

Edited by George F. Luger

Published by the AAAI Press /
The MIT Press
700 pp., \$35.00 paperback
ISBN 0-262-62101-0

This comprehensive collection of 29 readings covers artificial intelligence from its historical roots to current research directions and practice. With its helpful critique of the selections, extensive bibliography, and clear presentation of the material *Computation and Intelligence* will be a useful adjunct to any course in AI as well as a handy reference for professionals in the field.

The book is divided into five parts. The first part contains papers that present or discuss foundational ideas linking computation and intelligence, typified by A. M. Turing's "Computing Machinery and Intelligence." The second part, Knowledge Representation, presents a sampling of the numerous representational schemes by Newell, Minsky, Collins and Quillian, Winograd, Schank, Hayes, Holland, McClelland, Rumelhart, Hinton, and Brooks. The third part, Weak Method Problem Solving, focuses on the research and design of syntax based problem solvers, including the most famous of these, the Logic Theorist and GPS. The fourth part, Reasoning in Complex and Dynamic Environments, presents a broad spectrum of the AI communities' research in knowledge-intensive problem solving, from McCarthy's early design of systems with "common sense" to model based reasoning. The two concluding selections, by Marvin Minsky and by Herbert Simon, respectively, present the recent thoughts of two of AI's pioneers who revisit the concepts and controversies that have developed during the evolution of the tools and techniques that make up the current practice of artificial intelligence.

To order, call toll free 1-800-356-0343 (US & Canada) or (617) 625-8569, MasterCard & VISA accepted. Prices will be higher outside the US and are subject to change without notice. Visit the AAAI Press website! <http://www.aaai.org/Publications/Press/press.html>