Dynamic Memory Revisited

Roger Schank's influential book *Dynamic Memory* described how computers could learn based upon what was known about how people learn. Since that book's publication in 1983, Dr. Schank has turned his focus from artificial intelligence to human intelligence. *Dynamic Memory Revisited* contains the theory of learning presented in the original book, extended to provide principles for teaching and learning. It includes Dr. Schank's important theory of case-based reasoning and assesses the role of stories in human memory. In addition, it covers his ideas on nonconscious learning, indexing, and the cognitive structures that underlie learning by doing.

Dynamic Memory Revisited is crucial reading for all who are concerned with education and school reform. It draws attention to how effective learning takes place, and it provides ideas for developing software that truly helps students learn.

ROGER C. SCHANK is John Evans Professor of Computer Science, Education and Psychology, and Director of the Institute for the Learning Sciences at Northwestern University. He also serves as Chairman and Chief Technology Officer of Cognitive Arts.

Dynamic Memory Revisited

Roger C. Schank
Northwestern University

CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK www.cup.com.ac.uk
40 West 20th Street, New York, NY 10011-4211, USA www.cup.org
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alcarón 13, 28014 Madrid, Spain

First published 1999

Printed in the United States of America

Typeset in Palatino 10/13 in Quark XPress [AG]

A catalog record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data
Schank, Roger C., 1946–
Dynamic Memory Revisited / Roger C. Schank.
p. cm.
Includes bibliographical references.
ISBN 0-521-63302-8 (hardcover). – ISBN 0-521-63398-2 (pbk.)
1. Memory. 2. Learning, Psychology of. I. Schank, Roger C.
1946- Dynamic memory. II. Title
BF.371. S365 1999
153.1 – dc21 99–11983
CIP

ISBN 0 521 63302 8 hardback
ISBN 0 521 63398 2 paperback

Contents

Preface to the Second Edition

The first edition of this book was about artificial intelligence (AI). This second edition is about education. It is hard to see how this could be. Are these two subjects really in any way the same? My answer is that fundamentally they are, but of course, I recognize that such a notion would not necessarily be accepted as gospel truth. The common element is *learning*. Without learning there are neither intelligent machines nor intelligent people.

The subtitle of the original *Dynamic Memory* was "*a theory of reminding and learning in computers and people*." In the late 1970s and early 1980s, I was fascinated by the idea that computers could be as intelligent as people. My assumption was that if we could figure out what intelligence was like in people, then we could get computers to model people. Detail the process sufficiently and – *presto!* – intelligent machines. I no longer hold such views.

Since 1981, not as much has happened in AI as one might have hoped. The goal of building a dynamic memory, a memory that changed over time as a result of its experiences, has proven to be quite difficult to achieve. The major reason for this is really one of content rather than structure. It is not so much that we can't figure out how such a memory might be organized, although this is indeed a difficult problem. Rather, we simply were not able to even begin to acquire the content of a dynamic memory. What does a person know? We can attempt to detail the facts an adult might know, as indeed Doug Lenat has attempted to do, but the fact is that facts aren't all there are. For instance, we may know a lot of facts about cars but, even more important, we have a great many experiences with cars, ones that have informed and modified our view of cars over time such that what we really know about cars includes the smell of the back seat of the family car when we were five, the difficulty of dealing with a date when we were teenagers, and fear about changing a tire on an open road. These ideas are not about cars at all. Each experience we have with a car is

also an experience with a person, or a problem, or a life issue. Each informs the other and contributes to what we know about the world.

Content is the primary issue both in building intelligent machines and in helping to form intelligent people. Workers in each of these disciplines sometimes act as if they would just like the content problem to disappear. AI researchers want learning to come for free somehow, either by statistical methods or by making lots of connections. But, any way you look at it, someone has to get a computer to know what a human knows about using a toaster or playing baseball.

We hope we can educate people by simply presenting content or by avoiding the issue altogether and having students learn principles apart from content. It is the details that matter in both cases, and these details must be acquired, unfortunately, through repeated experience. In a dynamic memory, change is the norm and static facts are meaningless; therefore, it is not just the acquisition of content that is the issue. Whatever you learn changes over time. Dealing with change is an important part of intelligence for both machines and people.

When the first edition of this book appeared, I was attempting to lay out the groundwork for building intelligent machines. I did so by attempting to observe and to codify what I could about the nature of the memory that underlies human intelligence. Today I am concerned less with computers and more with people. And today I am concerned more with what it means to learn than with what it means to be smart. I have come to realize that intelligence without learning is not intelligence at all.

In 1981, I had two small children, a daughter who was nine and a son who was six. They were in school and they were learning, but the learning that mattered most in their lives was not taking place in school. What I learned about learning in working on this book was that school as it was constituted (at least in my children's lives) could not possibly actually contribute to learning, at least not in the way one assumed it was supposed to.

Making students memorize and then seeing if they have successfully done so through tests makes no sense in the context of computer learning. We can easily get a computer to know lots of facts. But, in what sense does it really *know* these facts? It is the utilization of information that constitutes knowing, not the recitation of information. Although schools pay lip service to this idea, AI people cannot. No one cares what a computer knows. We care what it can do. The same is true of students, but schools fail to recognize this in any important way.

When I would ask my children what they had learned in school, they occasionally mentioned the factual instruction they had received that day, but more often they talked about the other kids, their relationship with their teachers, and the worries and concerns about social interaction that were truly the instructional material of school. Their dynamic memories were working, and school was causing them to think, ponder, wonder, and generalize, but not about history or mathematics. I began to understand that what I knew about learning wasn't known by the school system as a whole. Schools were trying to teach things that children had no interest in knowing.

The issue is not whether kids are interested in topics like math and history. Presented properly, almost everything is of interest to children, and my kids were no exception. But they weren't generally interested in the way these topics were taught in school, namely, in a non-doing learning environment that was devoid of reference to real-world relevance and to their lives.

All my work in AI until 1981 had been about getting knowledge into computers. But, we didn't try to teach history or mathematics, we tried to teach how to function in a restaurant, or how to understand the interactions of foreign diplomats. Of course history and mathematics were valuable to know on occasion in some contexts, but the real issue was the acquisition of the mundane knowledge that people take for granted. And watching my children grow, I could see them acquiring this knowledge. They acquired it even in school sometimes, since they acquired it wherever they were, but such knowledge was simply not the material that counted in school, although it counted in life. Something was radically wrong with what the schools were trying to teach.

In 1989 I moved from Yale to Northwestern. This was not simply a change in venue. The move occurred because of the issues just raised. I felt that we were not likely to build intelligent machines any time soon but that we might help create more intelligent people. The major point of the original *Dynamic Memory* was that humans have complex and constantly changing memories and that if we wanted to build intelligent machines we would have to figure out how each experience that a machine involved itself in would alter its existing memory structures. To accomplish this, we would need to study how people acquired and stored memories and how they constantly learned from new experiences. In short, if we wanted to build machines that learned from experience, we would have to study how people learn from expe-

rience. To make dynamic memories in computers, we needed to know what dynamic memories looked like in people.

This is easier said than done. Real people aren't very willing or able to discuss their memories. We could get a machine to acquire experience in only one of two ways. Either it would have to grow naturally through experience as a child does, or it would have to have knowledge spoon-fed into it. The former approach, though obviously the "right" one, is very impractical: The computer would need to be able to operate in the world, to have needs it had to satisfy, and to find ways (on its own) of satisfying them. The prospect is exciting but daunting. The alternative approach, spoon-feeding all there is to know about the world, is equally daunting, but for different reasons. What do we know? How can we find everything that a given entity knows and put it in a memory that can dynamically reorganize itself with every new experience? That, to some extent, was the question behind the first version of *Dynamic Memory* and it is still discussed here.

Although these subjects are still of great interest – we will always want to know how we acquire knowledge and how we could get a machine to acquire knowledge – they have been supplanted in my mind by a more important, but quite related, issue. In order to get machines to acquire knowledge, I began to wonder about how humans acquire knowledge. As I looked into this issue, I began to wonder about the question of how we could help humans acquire knowledge better.

In the late 1980s, I began to seek funding to help me take what we had learned about learning in the context of the creation of intelligent machines and apply that knowledge in building learning environments that would enhance the learning of people (especially children). What had we learned about learning? In the first *Dynamic Memory*, the main points were that learning depends upon expectation failure and the attempt to explain that failure, and that remindings that come from our store of memories are instrumental in helping us create explanations and are therefore critical to learning. We process new experiences in terms of prior experience, and our memories change as a result. To translate this into a teaching environment, we must set up situations in which students can have their expectations fail and can either be reminded or be instructed about how not to fail next time. To put this another way, learning takes place if, while one is attempting to do something, something else inhibits the doing and causes one to wonder why what one thought would work, didn't.

One key problem addressed in the original version of this book is how the process of abstraction and generalization from experience works. Abstraction and generalization are integral parts of the learning process. In order for an entity to learn, it must be able to generalize from experience. But, more important – and I realize this seems obvious – in order to generalize from one's experience, one must have an experience. The obviousness of this point has not prevented computer scientists from failing to provide proper real-world experiences for their so-called learning machines, nor has it pointed educators toward the idea that learning will require actual experience.

The key idea here is doing. John Dewey and others have noticed that most learning occurs in the context of doing. While I was considering how computers might learn, I came face-to-face with the realization that computers weren't doing much of anything. We were trying to get them to learn about restaurants because they were reading about restaurants. They weren't eating in restaurants, nor were they really likely to be very interested in restaurants. Of course we could remedy this by having them learn to cook or something practical like that, but in fact AI, as a field, never really did such things. AI never concentrated on getting a machine to do something practical and learn about it by doing it.

For that matter, although children learn by doing all the time in their daily lives, they hardly do it in school at all. They learn mathematics or history because they have to, not because they are trying to do something for which such information might be valuable. Thus, both for computers and for people, environments would need to be created that allowed them to do something they wanted to do and allowed learning to take place in the context of that doing. The ideas in *Dynamic Memory* are critical to understanding how learning takes place in a doing environment.

This new version of *Dynamic Memory* is about translating what the original version said about how computers might learn into a more direct version that addresses the issues of enhancing human learning. I am concerned with how to change education. It is necessary to understand how learning works, and because the job of learning is to change memory, we need to understand how memory works. Armed with that understanding we can begin to rethink school.

The work of the last nine years that underlies the rewriting of *Dynamic Memory* was supported in large part by Andersen Consulting. It was a tumultuous relationship, different than the typical federal

grant by quite a bit. Nevertheless, they enabled me to build the Institute for the Learning Sciences (ILS) and this created the venue for building the software that embodies many of the ideas described here. ILS was, at its height, an exciting place filled with many brilliant people. I especially want to thank Alex Kass, Gregg Collins, Ray Bariess, Larry Birnbaum, Kemi Jona, Michael Korkusca, Chris Riesbeck, and a host of others who contributed to building software and playing with ideas about learning. Also, I want to thank Franci Steinmuller, who read early drafts of the manuscript and made copious notes, rewrote sentences, challenged ideas, and occasionally wrote a sentence or two as a suggested improvement. I also want to thank Adam Neaman and Heidi Levin for their help in the preparation of the manuscript of *Dynamic Memory Revisited.*

Roger C. Schank
Palm Beach, Florida
February 1998