

The creative process: A computer model of storytelling and creativity

Scott R. Turner

(University of California, Los Angeles)

Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, ix + 298 pp.
Hardbound ISBN 0-8058-1576-7, \$49.95

Reviewed by
Keith Oatley
Ontario Institute for Studies in Education

Within cognitive science and psychology, there has been a good deal of interest recently in the topic of creativity. In this book, Scott Turner of the University of California, Los Angeles, presents a theory of creativity applied to generating small stories.

Turner can be thought of as a member of the third generation of the Schank family: first, of course, there was grandfather Roger Schank, who in the 1970s with Robert Abelson at Yale, embarked on the research project of understanding narrative text using the ideas of goals, plans, and scripts. The attempt was to propose computational models that would accomplish aspects of narrative understanding. The second generation was a talented group of people, including Wendy Lehnert and Robert Wilensky, who did their Ph.D.s at Yale on story understanding. Turner is a member of a third generation, advised by Michael Dyer who also obtained his Ph.D. at Yale and then moved to an academic position at UCLA. Dyer had turned his attention to story generation as well as story understanding.

In the classic mode, Turner took on rather too much for his Ph.D. He wrote a large artificial intelligence program, 17,000 lines of Lisp code, that produces a reasonable output that could, with the suspension of a certain amount of disbelief, pass for the production of a human author. He calls his program "Minstrel." It generates stories of half a page or so about knights and ladies at the court of King Arthur. The program was the core of his Ph.D. thesis, and this is the book of the program.

Here is a sample from one of Minstrel's stories (p. 72):

It was the spring of 1089, and a knight named Lancelot returned to Camelot from elsewhere. Lancelot was hot tempered. Once, Lancelot lost a joust. Because he was hot tempered, Lancelot wanted to destroy his sword. Lancelot struck his sword. His sword was destroyed.

One day, a Lady of the court named Andrea wanted to have some berries. Andrea went to the woods. Andrea had some berries because Andrea picked some berries. At the same time, Lancelot's horse moved Lancelot to the woods. This unexpectedly caused him to be near Andrea. Because Lancelot was near Andrea, Lancelot saw Andrea. Lancelot loved Andrea.

This is about 40% of the story. We human readers are meant to infer that Lancelot fell in love with Andrea. It's not a bad story. In the next paragraph Lancelot is again moved by his horse, unexpectedly, into the woods, and he sees Andrea kissing someone called Frederick. His hot temper is again roused. If you want to know what happens you will have to read the book!

Turner has obtained human opinions about his computer-generated stories—they are judged to be the kind of thing that a 10- to 15-year-old child might write, though without some of the linguistic polish of a human author. Readers of such stories who are familiar with the work of Schank and his followers will hear clankings of the machinery of conceptual dependency in the background. Those familiar with computer games will not be surprised by the world of fights, magic potions, dragons, and so on, that the knights and ladies inhabit. Nonetheless what Turner has accomplished is remarkable, the more so in that he has shown that the creativity of his story generator also works quite well in inventing mechanical gadgets.

Many interesting books have at their center a single idea, and this is true of Turner's book. The idea is this: the creativity that forms the basis of successful story writing, and of some other kinds of problem solving, derives from transformation of the problem at hand followed by search over known examples of solutions near the problem of interest. So search is integrated with adaptation of the problem to plausible previous examples. This kind of creativity is what Johnson-Laird (1993) has called creativity within a genre. Understanding it is far more tractable than understanding how completely new genres in any field of endeavor are formed; nonetheless it presents a considerable challenge, and there is only a tiny number of good theories in the area. What Turner has shown is that a process of adaptation to previous examples can be implemented computationally, and will form the basis of producing text that has many of the properties of acceptable short stories.

Turner has written a program that will generate stories that have a point, that are consistent, that have a twist in them, and that are moderately engaging to a human reader. Though some previous programs have produced pieces of narrative—a character with a goal does this, and then this, and then this—*Minstrel* is the first in which any model of the goals of the author has been implemented. Turner does an excellent job here. He has chosen to cast the writing of a story as problem solving by the author. This allows him nicely to capture the idea of having to satisfy multiple constraints set by the author's goals. In Turner's treatment, four kinds of goals are considered:

Thematic goals: The story must illustrate a theme, in this case provided by the user. The theme of the story fragment given above was 'Hasty impulse regretted'.

Drama goals: In *Minstrel*, four techniques used by human authors have been implemented that give dramatic force to the story: suspense, tragedy, foreshadowing, and characterization.

Consistency goals: The narrative should reflect general understandings of the way the world works.

Presentation goals: Story elements are presented to the reader in particular ways, and in a particular order.

The program starts with an initial goal 'Tell a story', which is then broken down into subgoals that include adopting a theme, illustrating the theme, applying drama goals, and so on. The program starts with a theme, then cycles iteratively through the subgoals. The whole process depends on case-based reasoning. The program has at its disposal in memory a number of previous 'cases'—hand-coded stories in which comparable problems have been solved. What it does is to search over these cases for useful elements, and transform them to solve the problems of the current story. Some people might think this is cheating, but not so. As the author who is (arguably) America's principal author of literary science fiction, Ursula K. LeGuin, said in a recent

interview, “Books come from books”.¹ A process of intelligent searching of what is known, to find examples that can be transformed to suit new contexts that have their own local constraints, is indeed the essence of within-genre creativity.

Like the authors of short stories, the writers of artificial intelligence programs have a number of goals. These include the following: Engineering goals—to write programs that extend human capabilities; theoretical goals—to write programs that are working examples of psychological theories; and domain-specific technical goals—to crystallize useful fragments of knowledge as techniques that can be used directly in human action or in programs. Attaining engineering goals in the writing of enjoyable short stories that people would want to buy is, as Turner says, some way off. The many thousands of human authors struggling to get their short stories into print need not feel anxious; they are not about to be displaced by the computer. Important goals of psychological theory are, however, accomplished by Turner’s program. He has offered an instantiated theory of story generation, the first that takes seriously the central issue of authors’ goals and the structuring possibilities afforded by their multiple constraints. Any new cognitive theory of story generation would do well to start from Turner’s model. As to domain-specific technical goals, again Turner does well. Here are a number of instantiated techniques that provide essential components of story generation. They can be embodied in computer programs or used in the more traditional method of human authors casting words directly onto the page.

Reference

Johnson-Laird, Philip N. (1993). *Human and machine thinking*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Keith Oatley is Professor in the Centre for Applied Cognitive Science at the Ontario Institute for Studies in Education. His research is on the cognitive structure of emotions, and on the relation of literary criticism to cognitive psychology. He has also written a novel, *The Case of Emily V.*, which won the Commonwealth Writers Prize for Best First Novel, 1994. Oatley’s address is: Ontario Institute for Studies in Education, 252 Bloor Street West, Toronto, Ontario, Canada M5S 1V6; e-mail: keith.oatley@cacsmail.oise.on.ca.

¹ International Festival of Authors, Toronto Harbourfront Reading Series, 1994.