In G. W. Cottrell (Eds.), *Proceeding of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 207-212). Hillsdale, NJ: Lawrence Erlbaum. (1996).

Functional Roles for the Cognitive Analysis of Diagrams in Problem Solving

Peter C-H. Cheng
ESRC Centre for Research in Development, Instruction and Training
Department of Psychology, University of Nottingham
Nottingham, NG7 2RD, U.K.
peter.cheng@nottingham.ac.uk

Abstract

This paper proposes that a novel form of cognitive analysis for diagrammatic representations is in terms of the functional roles that they can play in problem solving. Functional roles are capacities or features that a diagram may possess, which can support particular forms of reasoning or specific problem solving tasks. A person may exploit several functional roles of a single diagram in one problem. A dozen functional roles have been identified, which can be considered as a framework to bridge the gulf between (i) studies of the properties of diagrams in themselves and (ii) investigations of human reasoning and problem solving with diagrammatic representations. The utility of the framework is demonstrated by examining how the functional roles can explain why certain diagrams facilitate problem solving in thermodynamics. The thermodynamics diagrams are interesting, in themselves, as examples of complex cognitive artefacts that support a variety of sophisticated forms of reasoning.

Introduction

Work on the nature and use of diagrammatic representations in cognitive science and related fields can be roughly divided into two quite general approaches. First, there is empirical and computational modelling work that examines cognitive processes involved in reasoning with diagrams. For instance, Larkin & Simon (1987) demonstrated that diagrams can (sometimes) be more effective than informationally equivalent sentential representations for problem solving. Koedinger and Anderson (1990) have shown that expert knowledge in some domains may use diagrams encoded in perceptual chunks, or diagrammatic configuration schema. Cheng (Cheng & Simon, 1995; Cheng, in press) has suggested that diagrams may have had a significant role in some important historical episodes of scientific discoveries. (See Kulpa, 1994, for further examples.) The second approach involves the theoretical and empirical study of the properties of diagrams in themselves, without special regard to the problem solving contexts in which they may be found. For instance, Bertin (1981) provides a taxonomy of graphical objects and their relations. By asking subjects to rank diagrams on a number of dimensions, Lohse *et al.* (1994) have obtained a taxonomy of visual representations. (See Wickens, 1992, for further examples.)

However, there is a need for work in the middle ground between the two approaches, which attempts to give some understanding of how properties of diagrams are related to the cognitive processes embedded in different diagrammatic representations. Diagrams are not a homogenous class of representations, but have diverse formats and uses. Many studies of diagrams from a cognitive perspective have tended to focus on a particular type of diagram for one kind of problem, as in some of examples mentioned above. As yet, there is no systematic way to judge whether particular findings for one type of diagram are applicable to other diagrammatic formats, without merely repeating the studies with those formats.

For example, Larkin and Simon's (1987) seminal work analysed the computational benefit of diagrammatic versus sentential representations, comparing representations that are informationally equivalent; information in one representation can be directly translated or mapped into the other. However, in real problem domains the available alternative representations are rarely equivalent in this sense. So, an alternative way to analyse aspects of different representations may be useful, serving as a basis for comparisons and to explain differences between diagrams. Given a particular domain and some characterization of diagrams, how can we determine what will be useful and effective diagrams for particular problems? The design of diagrammatic representations is an issue in cognitive science and related fields, such as learning and instruction, program visualization, and human computer interaction.

There are many ways in which complex phenomena or artefacts can be decomposed, with different units of analysis giving insights on different levels. Previous studies of the properties of diagrams have tended to focus on either (i) whole diagrams (e.g., Lohse *et al.*, 1994) or (ii) on diagrammatic elements (lines, angles; e.g., Bertin, 1981). Here, the properties of diagrams will be considered at an intermediate level that is more directly suited to considerations of the forms of information processing that may be done with them.

Functional roles of diagrams will be the units of analysis. They are capacities or features that diagrams may have, which support particular forms of reasoning or specific problem solving tasks, by making relevant information available to reasoners with little symbolic computation and only minor additions or changes (if any) to the diagrams themselves. Different functional roles make different kinds of information readily accessible.

A person may exploit several functional roles of a single diagram in one problem, or may use alternate ones in different problems. The roles are distinct but are naturally dependent on each other. They should have a familiar feel, even though they may not have been explicitly recognised before. A dozen functional roles have been identified and constitute a framework linking properties of diagrams to cognitive processes, which begins to address the problem of how to design effective diagrams for particular problem domains.

The paper has three main sections. In the first section a dozen roles of diagrams are described. In the following section the framework is used to explain the occurrence of particular diagrams in thermodynamics to illustrate the functional roles and to begin to demonstrate the utility of the framework. The final discussion section considers some implications of the framework.

Functional Roles of Diagrams

The following list of functional roles was compiled by examining many diagrams from a variety of domains (engineering, physical science, social science, and medicine). Sources included: collections of diagrams, text books, instruction manuals, published papers, design plans, and laboratory note books. Some diagrams have multiple functional roles. The types of information processing associated with each diagram was studied, and 12 functional roles were identified (although no claim is made about the completeness of this list).

F1 Showing spatial structure and organization. A functional role of some diagrams is to depict the spatial features of objects and the arrangement of their components, with some fidelity. Such diagrams have close spatial mappings from the shape and location of target objects to the shape and position of symbols representing them. Capturing spatial structure is an important function of engineering and architectural drawings or "blue prints". Such drawings show local and global structure of objects symbolically, rather than pictorially. Details hidden by intervening material may be shown by cut-away sections and symbolic conventions used to hide unnecessary detail or provide extra information (e.g., centre lines).

F2 Capturing physical relations. Diagrams can be used to highlight selected physical relations that are of importance in a target domain, without showing the spatial structure of objects and the spatial relations among their components. In schematic diagrams of electrical circuits, for instance, the inter-connectivity and sequence of the components is shown, but the location of symbols in the diagram may bear little relation to the physical location of the components on the circuit board.

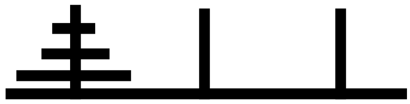

Figure 1: Tower of Hanoi.

F3 Showing physical assembly. A functional role of some diagrams is to show how an object is physically assembled from components; what parts there are and how they go together. This may be achieved by explicitly showing a series of subassemblies or depicting the object as if it had been systematically dismantled. Showing subassemblies is common in engineering "blue prints" and often used in instruction manuals for construction toys (e.g., Lego, Mechano). Where such diagrams provide information about the order of assembly they are also depicting a process (see F8).

F4 Defining and distinguishing variables, terms and components. Some diagrams are used to define or identify components, variables and features pertaining to a target domain. Written labels may be used to name or specify particular components or features. Diagrammatic elements may themselves be special symbols, which have conventional meanings in particular domains, such as the component symbols in electrical circuit diagrams.

F5 Displaying values. A function of some diagrams is to depict values of variables in a manner that facilities qualitative and quantitative reasoning about them, usually in the form of comparisons. Often standard formats or reference systems will be used, such as Cartesian graphs, histograms and pie charts. Problem solvers use their knowledge of the conventions governing the formats when reasoning.

F6 Depicting states. Some diagrams depict the state a system, without special reference to transitions from one state to another. For example, some operating manuals for electronic equipment include schematic diagrams of factory set positions of internal (DIP) switches. By comparing the diagram to the actual switches, we can see whether the equipment is in its default state. Weather charts have the depiction of states as one of their main functions. Figure 1 shows a familiar diagram, a single state of the Tower of Hanoi problem.

F7 Depicting state spaces. A functional role of some diagrams, which logically follows the previous one, is the depiction of state spaces (but not necessarily problem spaces). These diagrams have several components depicting two or more states, with adjacent components normally representing closely related states. The Periodic table of chemical elements is an example. Each element may be considered as an individual state, with its horizontal and vertical position being meaningful in chemical terms. The transition state space for the 3-disc Tower of Hanoi problem is shown in Figure 2. Each node represents one state; at the corners all the discs are on one of the pegs. The lines represent legal transitions between states; one disc moved between pegs.

F8 Encoding temporal sequences and processes.
Diagrams may illustrate the temporal order or flow of a process, by depicting states and the changes to those states. Some of the ways this is done are: (i) placing diagrammatic elements in an ordered sequence; (ii) using arrows to show progress or movement; and, (iii) having contours labelled with time increments.

One purpose of this functional role is to provide some kinesthetic sense of the processes being depicted, perhaps as an aid to generating a mental model. The path of a particular problem solver through the problem space of the Tower of Hanoi can be shown by adding arrows or numbering the links in Figure 2.

F9 Abstracting process flow and control. A functional role of some diagrams is to abstractly represent the flow of complex non-linear processes. Such processes may include cycles, iterative loops, contingent branching and parallel tasks. These diagrams use conventional symbols (icons) to represent process stages without depicting the states themselves.

Traditional computer program data flow diagrams have this as one of their main functions. Different processes are named and those of similar type share the same symbol shape; for example, diamonds to depict decisions. The flow of information is shown, but the state of the information is not usually depicted. Gantt charts perform a similar function by naming stages and showing the order and dependency of processes.

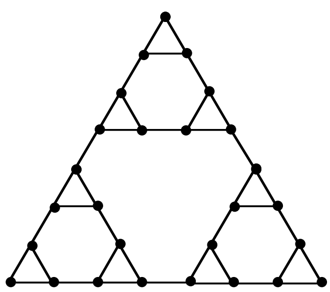

Figure 2 A state space.

F10 Capturing laws. A functional role of some diagrams is to capture a law by the means of its internal structure, such that the diagram does not merely display values of variables, but embodies the law. For example, in physics, the resultant of two forces acting at a point may be found by constructing a parallelogram (two adjacent sides represent the given forces, and the line from the given point to the opposite corner provides the resultant force). Such diagrams capture, and happen to operationalize, vector addition.

Capturing laws is one of the distinguishing characteristics of Law Encoding Diagrams, LEDs, (Cheng, 1994, 1995). LEDs have geometric, spatial or topological constraints that govern their structure, such that the form of the diagram is always consistent with the target laws of a domain.

F11 Doing computations. Computations can to be done directly using the structure of some diagrams. The prime example for numerical calculations are nomograms, which are closely related to slide-rules. Law Encoding Diagrams also allow computations to be made using their structure.

F12 Computation sequencing. Diagrams can help organise, plan and track complex sequences of computations. For example, when doing numerical integration, say by the Simpson's method, a diagram can be used to explain why the calculation takes the form that is does. Tabachneck, Leonardo and Simon (1994) describe how an expert in economics used a graph as a place holder during reasoning and as a summary.

The twelve functional roles have been described and brief examples given. The next section provides a single integrated example covering most of the functional roles.

Diagrams in Thermodynamics

To demonstrate the use of the proposed framework and to further clarify some of the functional roles, this section considers diagrams found in thermodynamics. The domain was chosen because it has complex diagrams that are of interest in their own right, which problem solvers use in sophisticated ways. The diagrams may be considered as cognitive artefacts, which have evolved over the history of the field as effective cognitive tools for problem solving.

Diagrams in thermodynamics texts are mainly of two kinds: (i) component diagrams, which show the structure of particular pieces of equipment or the parts of a plant; and, (ii) property diagrams, which are graphs showing the thermodynamic properties of the fluid in the plant. Figures 3a and 3b are typical examples. An interesting observation regarding these diagrams is their occurrence in complementary pairs when complex heat engine cycles or refrigeration cycles are being considered. In other circumstances they tend to be found alone. Figures 3a and 3b are such a pair for a particular steam power cycle. The diagrams are used to explain the operation of the cycle and to solve problems, such as determining the power output and efficiency of a steam plant, given a few fixed values.

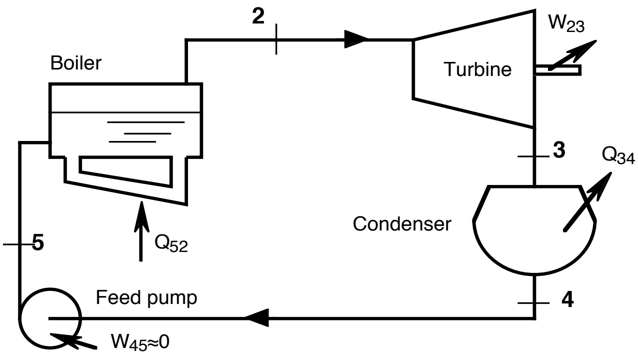

Figure 3a: Steam Plant Component Diagram.

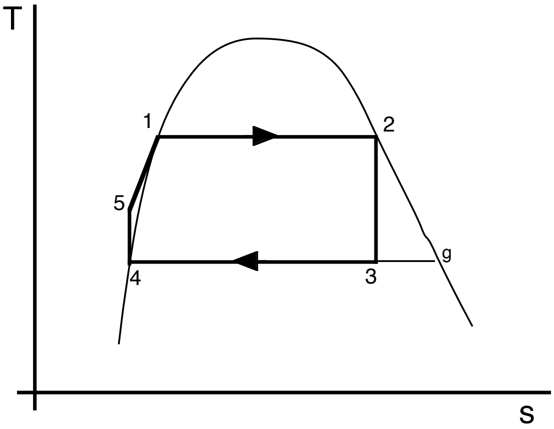

Figure 3b: Steam Cycle Property Diagram.

Rogers and Mayhew (1980) is a popular standard undergraduate text in thermodynamics (a fourth edition was published in 1992). It has three chapters on complex power and refrigeration cycles, in which 23 of the 56 (41%) diagrams are complementary pairs of component and property diagrams. The rest of the book (24 chapters) contains nearly 200 diagrams, but fewer than 10 (<5%) are complementary pairs. Further, observations of experienced instructors, expert in thermodynamics problem solving, showed that they usually draw complementary pairs of diagrams as the first step when solving problems of this kind. Clearly, some benefit is gained by using the complementary pair of diagrams in problem solving on complex thermodynamic cycles.

Explanations of the benefits that problem solvers gain by using these diagrams can be given in terms of existing theories or hypotheses. For example, Larkin and Simon (1987) have shown how diagrams similar to the component diagram can help in the processes of search and recognition during problem solving. Similarly, it is conceivable that experts may have perceptual schemas for different parts of property diagrams, in the form of diagrammatic configuration schemas (Koedinger and Anderson, 1990). However, these theories are not well suited, nor intended, to explain phenomena like the complementary diagram pairs in thermodynamics.

The functional roles framework provides an analysis of a different kind, which is more suited to explaining the diagram pairs phenomena — the two diagrams have different but complementary sets of functional roles. From informal observations of expert problem solvers in thermodynamics working on a range of different problems, and from the study of worked examples from thermodynamics texts, the functional roles possessed by the two kinds of diagrams were identified. Both the component diagram and the property diagram, Figure 3, identify and define things that are important for problem solving (function F4). The plant diagram uses labels and conventional symbols to indicate what items are part of the plant; e.g., the pump is a circle with lines representing its inlet and outlet. Locations of properties of interest are indicated by the diagram; for instance W_{23} is the power output of the turbine and Q_{52} is the heat needed to change the water to steam in the boiler.

The property diagram is a Cartesian graph, with temperature (T) and entropy (s) on the ordinate and abscissa, respectively. The "bell" curve shows the boundary between different phases of water and is known as the *saturation curve*. To its left the water is liquid. To its right the water is a vapour. Under the curve the water is a mixture of liquid and vapour. This is basic background knowledge for problem solvers in this domain. The numbers in the diagram indicate points in the cycle corresponding to the numbered locations in the component diagram, which is the only direct means that problem solvers have of interrelating the two diagrams. There is no '1' in the component diagram, because there is no unique location in the boiler that has properties corresponding to that point in the property diagram. The lines between numbered points in the property diagrams are changes that occur within the components of the plant.

The component diagram shows how the parts of the steam plant are physically connected (F2) (but does not show not true spatial locations, F1). An important function of the property diagram is to show the values of the temperature and entropy around the cycle (F5). This defines the thermodynamic character of cycle and allows comparisons to be made; for example $T_1=T_2$ and $s_2=s_3$. The property diagram may also be considered as a depicting one of many possible states (F6), other states being alternative cycles that have different shapes in the T-s space. The component diagram aids visualization of the process (F8), showing the direction of flow of the fluid around the circuit and the exchanges of heat and work to and from the system. The property diagram also supports visualization, but in terms of the thermodynamic properties of the system and with regard to the physical states of the fluid in the cycle (using the saturation curve).

The development of temperature and entropy graphs in thermodynamics is significant, because they captures some useful laws in diagrammatic form (F10). (Pressure-volume graphs are more easily conceptualized but they are less useful in problem solving.) In T-s graphs, the area directly under any curve (to the s axis) represents a quantity of heat; for example, the area under $line_{34}$ is the heat lost in the condenser; that is Q_{34} in the component diagram. Similarly, the size of the area enclosed by a loop or cycle, such as $loop_{12345}$ in Figure 3b, indicates the net amount of heat received by the system, and hence from the first law of thermodynamics, the area represents the net amount of mechanical energy produced by the system. The direction of cycle is significant; clockwise means the system produces useful power, and anti-clockwise means that it consumes energy (refrigerators). Because the property graph captures these important relations in a simple diagrammatic fashion, the graph is a useful model for problem solving. For example, a problem solver can visualize how to increase the power output of the cycle by increasing the enclosed area, without changing its shape. This may be done by increasing both T_1 and T_2 or reducing both T_3 and T_4, but not by independently increasing s_3 (as it would no longer equal s_2).

Further, it is possible to do calculations with the property graph (F11). When calculating the output conditions of the turbine, point 3, it is necessary to know how much of the fluid is vapour and how much is liquid, the dryness fraction. This can be found from the lengths of $line_{43}$ and $line_{4g}$. The ratio of their respective lengths equals the dryness fraction (because of its definition with respect to changes in entropy). At point 4 the fraction is zero (all liquid) and at 3 the fraction is approximately 0.8 (largely vapour). The property diagram allows perceptual inferences to be made about the equality of many of the variables, because the cycle is represented by several vertical and horizontal lines. For example, given T_2 the value of s_2 can be found from standard steam tables, because point 2 is on the saturation line. Then by inspection we see that s_3 equals s_2. Knowing the dryness fraction, the temperature T_4 can be found by calculation and hence T_3 is known. This sequence of inferences also illustrates how the property diagram can be used to help plan and execute a series of computations (F12).

In this example, the component diagram and property diagram shared few functional roles, and when they did they involved different sets of information. Thus, a possible explanation of why both diagrams are needed for effective problem solving, on complex thermodynamic cycles, is that between them they support a broad range of the kinds of problems solving tasks that are required.

Discussion

The concept of functional roles and the identifying of actual functions shows that there may be a useful level of analysis of the properties of diagrams that falls between general characterisations of whole diagrams and analyses of the properties of elementary diagrammatic components. Different combinations of functional roles will determine the overall character of a diagram, by making different kinds of information more readily accessible; as seen in the two thermodynamic diagrams. This final section of the paper considers issues raised by the functional roles framework, beginning with the more specific questions.

The next stage in this research will be to apply the framework to an interesting class of diagrams, Law Encoding Diagrams, which appear to be effective representations for some forms of problem solving and for learning (Cheng, 1994, 1995, 1996, in press). A possible explanation of their advantage is that they manage to combine many functional roles in a single diagram.

It will be interesting to study whether experts exploit more of the available functional roles of diagrams than novices. Larkin and Simon (1987) consider a pulley system problem, in which the diagram of the physical arrangement of the system provides useful locational information for problem solving. This is appears to be a factor that underpins both the functional roles of depicting physical relations, F2, and sequencing computations, F12. Although Larkin and Simon's computation models do not deal with learning, it is the case that as problem solvers become more experienced on pulley problems, they can solve them by inspecting the diagram without recourse to written calculations. This might be explained by hypothesizing that they have learned that diagrams of pulley sub-assemblies capture a specific version of the lever law, perhaps in the form of a diagrammatic configuration schema (Koedinger & Anderson, 1990). This more expert-like use of the diagram could be seen in terms of exploiting the functional role of capturing laws by means of the structure of the diagram, F10. Detailed study of expert versus novice problem solving in thermodynamics would be a suitable place to begin the examination of this issue.

Wickens (1992) effectively demonstrates the need for compatibility between visual displays of information with problem solvers' mental models; an important issue for human computer interaction. The framework of functional roles may be considered as providing a further compatibility dimension for the design of such displays. A diagram or display has greater compatibility, in terms of the framework, when more of the problem solving tasks are directly supported by functional roles of the diagram(s) being used. In the thermodynamics example, if either diagram were absent, it is likely that the solutions to problems would be harder, because there is less information available to constrain problem solvers mental models of processes that are happening in the plant.

The effectiveness of a diagram for a particular functional role will, in part, depend on the way information is embodied by particular forms and combinations of elementary diagrammatic components. Cleveland and McGill (1985) have studied the perceptual processes involved in qualitative and quantitative judgements using traditional graphical representations. This may be considered as an analysis of the cognitive processes needed to exploit the displaying values functional role, F5. Similar studies could be conducted for the other functional roles; for example, what cognitive processes that are implicated in the use of a diagram that captures a law in its internal structure, (F10)?

The emphasis of the functional roles framework is on the kinds of information are available for reasoning and problem solving tasks, rather than the way in which information is structured and processed. The kinds of questions that this approach aims to address concern the types of problems to which particular diagrams are well suited. This is different, but complementary, to the previous work, described in the introduction, that considers the cognitive processes found in reasoning with diagrams, without special attention to the semantic content of the information being processed. The motivation, in part, for the present approach is to provide a principled method for the selection and design of effective diagrams for particular problems. By focusing on the functional roles of diagrams, which directly relate to the kinds of information that are easily accessible in diagrams, the framework may provide a link between the task requirements of problems and the types of activity supported by different diagrams. Work is progressing on the specification of a methodology for the selection and design of diagrams using the framework.

Investigation of the framework in relation to representations in general is also currently being pursued. The concept of functional roles is applicable to other graphical representations, such as tables, as well as to character string notations, such as algebra and written English. The central issue in this regard is to devise principled ways to identify and differentiate different the kinds of information that may be available in a representation. The work is concentrating on formal approaches to the classification of types of information, with the examination of ontologies of problem solving tasks and methods.

Acknowledgements

This research was supported by the U.K. Economic and Social Research Council. Thanks must go to members of the Centre for Research in Development Instruction and Training for their help in this work, with special thanks to David Wood, Fernand Gorbet and Shaaron Ainsworth.

References

Bertin, J. (1981). *Graphics and Graphic Information Processing*. Berlin: Walter de Gruyter.

Cheng, P. C.-H. (1994). An empirical investigation of law encoding diagrams for instruction. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. (pp. 171-176). Hillsdale, NJ: Lawrence Erlbaum Associates.

Cheng, P. C.-H. (1995). Law encoding diagrams for instructional systems. *Journal of Artificial Intelligence in Education*, 7(1).

Cheng, P. C.-H. (1996). Learning Qualitative Relations in Physics with Law Encoding Diagrams. In *Proceeding of the Eighteenth Annual Conference of the Cognitive Science Society* Hillsdale, NJ: Lawrence Erlbaum.

Cheng, P. C.-H. (in press). Scientific discovery with law encoding diagrams. *Creativity Research Journal.*

Cheng, P. C.-H., & Simon, H. A. (1995). Scientific Discovery and Creative Reasoning with Diagrams. In S. Smith, T. Ward, & R. Finke (Eds.), *The Creative Cognition Approach* (pp. 205-228). Cambridge, MA: MIT Press.

Cleveland, W. S., & McGill, R. (1985). Graphical perception and graphical methods for analysing scientific data. *Science*, 229, 828-833.

Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14, 511-550.

Kulpa, Z. (1994). Diagrammatic representations and reasoning. *Machine Graphics and Vision*, 3(1/2), 77-103.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.

Lohse, G., Biolsi, K., Walker, N., & Rueter, H. (1994). A classification of visual representations. *Communications of the ACM*, 37(12), 36-49.

Rogers, G. F. C., & Mayhew, Y. R. (1980). *Engineering Thermodynamics* (3rd ed.). New York, NY: Longman.

Tabachneck, H. J. M., Leonardo, A. M., & Simon, H. A. (1994). How does an expert use a graph? A model of visual and verbal inferencing in economics. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. (pp. 842-847). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wickens, C. D. (1992). *Engineering Psychology and Human Performance* (2nd ed.). New York, NY: Harper Collins.