Folkscience: coarse interpretations of a complex reality

Frank C. Keil

Yale University, New Haven, CT 06520 8249, USA

The rise of appeals to intuitive theories in many areas of cognitive science must cope with a powerful fact. People understand the workings of the world around them in far less detail than they think. This illusion of knowledge depth has been uncovered in a series of recent studies and is caused by several distinctive properties of explanatory understanding not found in other forms of knowledge. Other experimental work has shown that people do have skeletal frameworks of expectations that constrain richer ad hoc theory construction on the fly. These frameworks are supplemented by an ability to evaluate and rely on the division of cognitive labour in one's culture, an ability shown to be present even in young children.

As the formal sciences advance, they illustrate in ever more detail the causal complexity of natural and artificial systems. A seemingly simple process like the beating of a heart requires an intricate set of causal interactions to function correctly. The ways in which DNA codes for morphological structures in development are not understood despite thousands of papers discovering causal processes that are necessary parts of the developmental story [1]. The same example holds for countless other cases in physics, chemistry, and cognitive science. With artificial systems as well, complexity continues to accelerate. The software that runs routine operations on most of the world's personal computers now has more than 50 million lines of code, organized in a wickedly intricate set of functional components; and even the most seasoned automobile mechanics must now rely on computer aided diagnostic equipment to uncover problems in black boxes that they no longer understand themselves.

This extraordinary complexity poses a critical question for the cognitive science of science, and for the intuitive understandings known as folkscience that we all use in our daily lives. As we cannot possibly grasp all the details of causal patterns in the world around us, at what level of granularity do we code the details and how do we use that level effectively? Although the same problem exists for the formal sciences, which somehow make forward progress while still being glaringly incomplete in key respects [2], it is especially acute for folksciences that we all use every day.

Recent studies are shedding light on the coarseness of our understandings in both formal and intuitive science, and the results are surprising. Our sense of how the world works is often vastly cruder than we think. Explanatory forms of understanding in particular are far different from what our first introspections suggest. A closer look at the formal sciences also suggests that reality falls short of popular images of what individual scientists know and do. This mismatch between what we think we know and what we really know creates a challenge for many empirical demonstrations in cognitive science over the past two decades. In several areas of high-level cognition, there seems to be compelling evidence for the influences of intuitive theories [3,4]. Intuitive theories have become a mainstay of many arguments in cognitive science. But, if we look closer at the theories, are they rich enough in detail to support all that they are claimed to do?

Out of our depth

In the philosophy of science it has become evident that scientific explanations are often much shallower and less complete than they might seem to the outsider [5]. It was still a surprise, however, to learn that scientists proceed in ways that are often highly intuitive. One study examined the daily thought processes of members of a leading laboratory of molecular biology and documented how they made a major breakthrough [6]. The researchers did not proceed by carefully planning out a long series of studies that systematically marched through a tree of hypotheses and findings. Instead, they were highly opportunistic, relying on serendipity and, more importantly, on an intuitive sense of when they were heading towards interesting new insights. The notion of a purely analytic process marching forth through a decision tree and eliminating logical alternatives is largely a myth.

As outsiders we are often surprised at how interdependent scientists are on the expertise of others that have come before them and work elsewhere [7]; but even with such dependencies, most individual scientists usually do know quite deep causal patterns in local domains. How much less do laypeople know and how well do they know their own limitations? In fact, laypeople have surprisingly shallow understandings that are masked by an 'illusion of explanatory depth', in which they think they understand the world in far more detail than they really do. In a series of studies people initially judged their explanatory understanding of devices and natural phenomena as being quite detailed, only to be greatly surprised later at their own ignorance [8] (see Fig. 1; Box 1).

Corresponding author: Frank C. Keil (frank.keil@yale.edu).

 1364-6613/03/$ - see front matter doi:10.1016/S1364-6613(03)00158-X

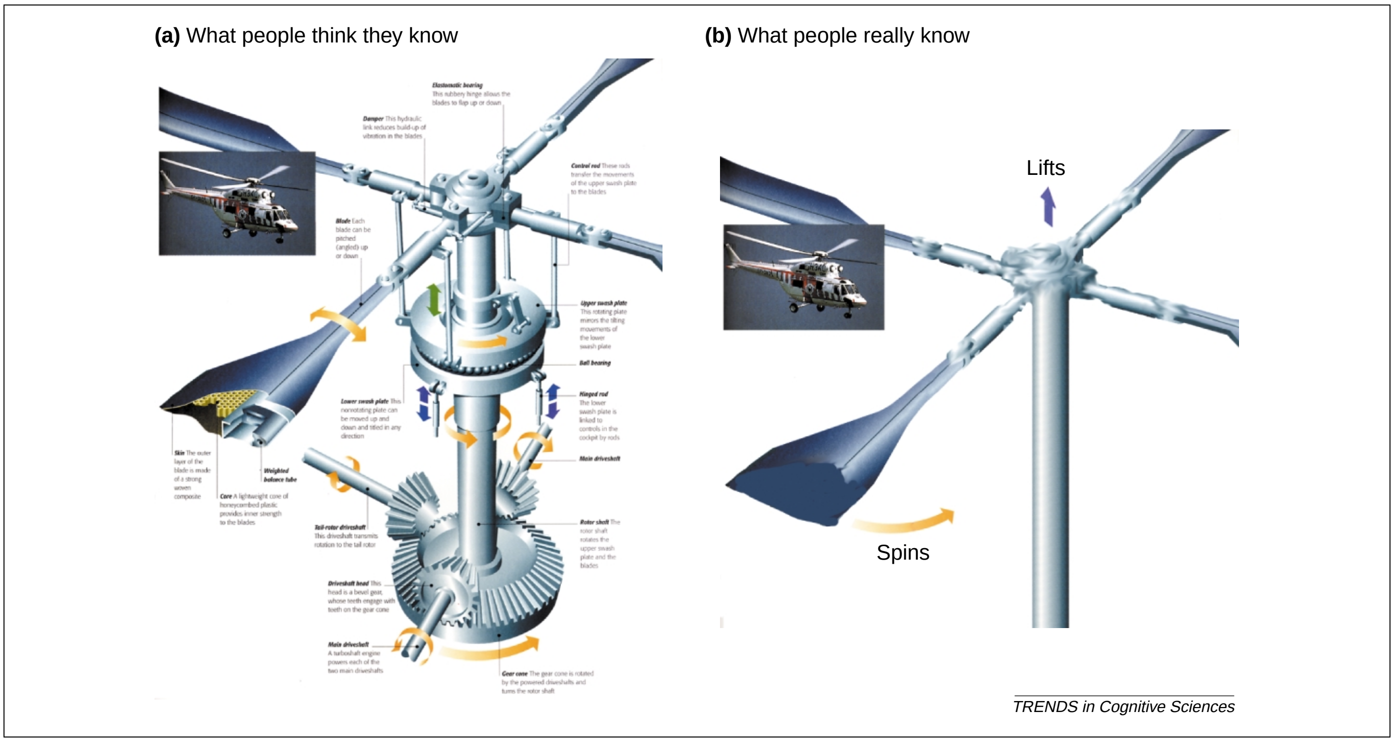

Fig. 1. When people estimate how deeply they understand the workings of various systems, they tend to think they know for more depth of detail than they actually do. When asked how a helicopter works, they seem to think they have knowledge approximating a detailed annotated drawing, but actually have a much coarser understanding corresponding to little more than the sense of a thing with blades that turn and provide lift. This illusion is quite specific to explanatory kinds of knowledge. People estimate the depth of their knowledge of procedures, facts and narratives much more accurately. Adapted with permission from [51].

Why explanatory understanding misleads us

The illusion of having detailed and coherent knowledge occurs primarily for explanatory understanding. By contrast, people's ratings of how well they know facts, or procedures, or narratives is well calibrated and they are not surprised at what they really know [8]. There are distinctive structural properties of folkscience that create especially strong impressions of detail and completeness to knowledge. Such factors include confusion of insights into relations among higher order functional units with relations among lower level subsystems [9,10]. Thus, the insight of knowing how a disk drive and a memory store interact may be confused with an understanding of how disk drives and memories work internally. Second, there is often confusion between (1) being able to decipher information from the environment in real time when a device or phenomena is available for inspection, and (2) with having mentally represented all those causal relations. This confusion is analogous to a recent finding of 'change-blindness-blindness' in visual cognition, where people think they have remembered far more from scenes than they really have [11]. In change-blindness-blindness the ability to re-inspect a scene might be confused with having the information stored in memory [12].

The illusion of understanding has been most extensively documented for our understandings of devices and then, secondarily, for knowledge of some biological organs and some non-living natural phenomena, such as the tides. It is likely to also hold for other complex causal systems, such as those governing human behaviour. The illusion of understanding is different from 'overconfidence' effects, in which people tend to overestimate their cognitive skills or the probability of doing well in a task [13].

If our intuitive theories are only vague sketches of how things are structured in a domain, can these theories adequately explain all the effects in the literature that have been attributed to intuitive theories? Errors in mental health diagnosis [14]; the perception of illusory correlations [15]; and science misconceptions have all been attributed to the power of intuitive theories. Moreover, for

Box 1. The illusion of explanatory depth

Participants are first taught how to use a seven-point scale that rates their knowledge on a range from (7) a full working diagram of a system to (1) only the haziest knowledge of some surface features of a system [8]. They then rate how well they know how various devices or systems work, such as a helicopter, a cylinder lock or a zipper. After rating their understanding of a large set of these items, participants are then asked to explain in as much detail as they can the actual workings of a few systems. After giving each explanation, participants are then asked to re-rate the depth of their initial knowledge. They are then asked to answer a diagnostic question that experts consider central to understanding the system (e.g. how does one pick a cylinder lock?). Next participants are asked to re-rate their knowledge a second time. Finally, they are presented with a concise but thorough expert explanation and are asked to rate their initial knowledge again in light of that expert explanation. The results across several studies show a strong drop in ratings of knowledge after each re-rating, often accompanied by shock and surprise by the participants at their own ignorance [8] (see also Fig. 1 above).

concepts themselves, the fundamental units of thought, it appears that many aspects of concepts' use and acquisition are influenced by intuitive theories. These include patterns of categorization, induction of new properties based on prior ones, and patterns of conceptual change in development [4,16].

Faint frameworks, stances, and modes of construal

The notion of a 'framework theory' [17] has been invoked as one way of trying to understand the presence of theory-like effects without the details [18,19]. Framework theories are meant to provide a conceptual umbrella that gives coherence and sense to all the facts in a domain without filling in any of the details [20,21]. When three-year-old children are said to have a 'theory of mind' that enables them to predict the actions of others, it is thought to be a framework theory [17,22]. In many cases, however, that framework may be little more than one or two connected causal beliefs, or just a few fragments of knowledge [23]. For folk psychology in young children, it might simply be the idea that beliefs cause desires, which cause actions. For folkbiology in young children (see Box 2), the overarching theory may be little more than that a vital force drives an organism to grow, need food, and, in the case of animals, move. The putative frameworks are so faint that they do not seem sufficient to explain how children and adults make sense of the world around them. Indeed, calling this knowledge a folk 'science' seems to be almost a misnomer.

Beyond framework theories, however, people track causal patterns that are distinctive to broad domains such as biology and psychology without weaving these patterns together in a law-like theory. That is, people adopt 'modes of construal', in which phenomena in a domain are assumed to correspond to certain causal patterns, even if there is no detailed sense of how those patterns are mechanistically implemented [24]. Modes of construal are similar to conceptual stances, such as an 'intentional stance' towards thinking agents, a 'teleological stance' towards living kinds and artefacts, and a 'physical stance' towards much of the inanimate world [25,26].

Adults and children alike are sensitive to patterns of causation that are distinctive to domains, even though such sensitivities do not translate into laws or deeper understanding. People know, for example, that natural kinds have essences arising from micro-structural properties in ways that artefacts do not [27–29]. Such expectations could arise from a domain-general strategy, in which the first element in local causal chains is seen as the most critical to understanding a domain [30] (see Box 3). Adults, children, and even infants, know that bounded inanimate objects tend to act on each other in ways that require direct contact with nearly immediate consequences whereas intentional agents act on each other at a distance and often with considerable delays [31,32]. By the age of 7 yrs, and possibly earlier, people know that living things have parts with functions but rarely have functions as a whole [26,33].

People also track kinds of properties as causally relevant for domains. For example, adults and children know that in certain domains (e.g. most artefacts) colour is not as important as shape. Thus, colour is judged as largely irrelevant for understanding a completely novel hand tool, whereas shape is considered central. By contrast, colour is seen as more important for living kinds [34]. Even some species of monkeys have similar sensibilities, seeing colour as more important for novel foods than for novel tools [35]. Knowledge of causal relevancy constrains choice among competing explanations and helps guide construction of more detailed explanations, when phenomena or devices are present for inspection, but may contain little information about mechanism (e.g. why this colour is important for this flower) and may not be connected in a systematic framework to other causal patterns that one has noticed. Moreover, this knowledge may be largely implicit. It is now possible to compute what kinds of feature types are seen as causally central to what domains [36] and to show that

Box 2. Causal status and other causal patterns

Unlike detailed theories, people are able to pick up causal patterns in various domains and use those to infer which sorts of properties are most likely to be casually relevant for theory construction. One example is causal status, in which the first element in local causal chains is seen as the most conceptually central to a category [30]. If people are presented with several instances of a new category with a set of four features that occur equally often with those instances; and if those features are also described as part of a causal chain, the first element in the chain will be seen as a stronger determinant of category membership (see Fig. I). Young children are also governed by this causal-status effect [52]. Tracking these kinds of causal relations and the sorts of features that tend to occur first in such chains can provide powerful clues as to what features are most important in different conceptual domains such as artefacts and natural kinds.

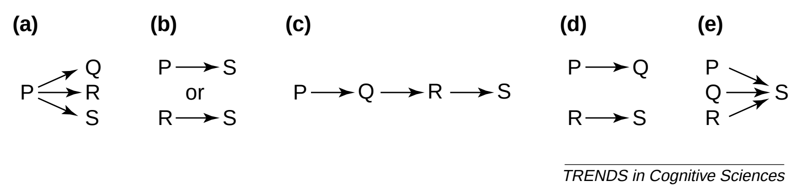

Fig. I. The causal-status effect predicts biases to favour some causal patterns over others as diagnostic of important causal features, most notably the initiating causal factor in a chain. That effect might in turn spring from an essentialist bias for natural kinds in which a single common cause is thought to produce effects. Adults and children alike seem to expect that natural kinds will be explainable by a set of structured relations of the form shown in (a), as opposed to the other possibilities shown in (b–e) [53]. Whereas property P in case (c) would also be predicted to be more central than Q, R or S by causal status, property P in case (a) is especially salient as the initiator of three distinct effects. Adapted with permission from [53].

Box 3. Coarse frameworks in folkbiology

One example of how tracking of causal patterns can constrain theory-construction occurs in the study of folkbiology. Young children certainly do not know the mechanistic details of reproduction, disease transmission, or growth [54,55]; but they do attribute many distinctive patterns to the domain of biology. Living kinds are seen as having a vital force that causes them to grow, reproduce, ingest food and (in the case of animals) move (see Fig. I). This vitalistic causality is understood as different from the mechanical causality of inanimate objects and from the intentional causality of folk psychology [56]. Young children also seem to believe in essences that characterize the nature of living kinds in ways not found for artefacts; and they tend to see these living kind essences as different from essences of non-living natural kinds, such as gold [28]. Finally, they show evidence of feeling that living kinds are best described at the species level and that they are parts of unique taxonomies in ways that other sorts of things are not [57,58].

In most adults, the theory part of folkbiology may also be a coarse framework. It might include notions of adaptation to niches, symbiosis, more elaborated notions about essential properties, finer grained ideas of taxonomies, and more advanced forms of vitalism. Adult biological knowledge might also include a much richer repertoire of facts about individual species and local causal patterns in specific ecologies, such as the reciprocal benefits between some plants and animals. There is considerable variation across cultures in such details, even in the level of analysis that seems most biologically natural [59]. Many nature-deprived urbanites find the level 'bird' most natural whereas traditional peoples of the rural Yucatan find species-level bird categories more natural [60]. These differences, however, should not mask the overall finding that, although most people do not have a richly structured, mechanistic theory of the biological world, they do have a sense of distinctive causal patterns associated with living kinds.

The many details we do know in biology, such as features of sub-features and some local causal relations (e.g. that a particular plant is poisonous to most animals but that a few actually benefit from it), are far different from having a detailed explanatory understanding of nature. As with mechanical devices, most laypeople understand biological mechanisms in far less detail than they think. It is one thing to know a set of local causal relations, but it is quite another to weave them together into a full-scale representation of why properties exist as they do and how they function at the mechanistic level.

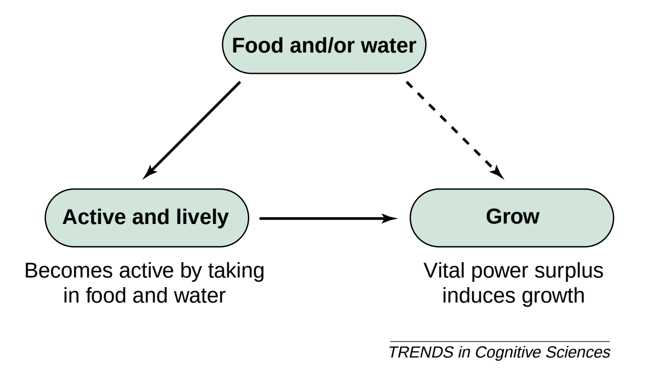

Fig. I. The core set of biological causal relations understood by the young child may be a simple triangle in which food and water are seen as providing a vital power that makes animals active and lively and, when there is a surplus of vital power, induces growth. The direct link between food and water and growth is not as strong in the younger ages. This model might also be applied incorrectly to plants because of a lack of understanding of photosynthesis. Adapted with permission from [56].

such abstract implicit information guides many theory-like effects [14].

How to skate on thin theoretical ice

Theories on the fly

How do we get by with the coarse sketches and gaps in our knowledge? People rarely enter situations with ready-made, detailed theories in mind. Instead, they rapidly decide which domain of causal patterns is relevant and then use their own schematic knowledge of relations and patterns to constrain explanations on the fly. A related phenomenon is ad hoc category construction [37,38], where people create richly structured, but completely novel, categories on the fly, such as things to take out of a burning house. People construct such categories by using sparse causal schemata inferred from goal structures of agents. Similarly, our ad hoc theories derived from situations add details inferred from the situation at hand. Use of causal relational information can be very rapid and is not at the mercy of slow reflective thought. Quite abstract causal patterns guide judgments in the very first moments of categorization [39]. Thus, one advantage of lean causal representations may be rapid deployment.

The construction of richer understandings under situational constraints is related to two other themes in cognitive science known as 'embodied' and 'distributed' cognition [40,41]. Clark offers an elegant analogy with the tuna, which is capable of accelerating faster from a dead stop than the laws of physics and physiology would predict is possible. The tuna achieves this feat by repeatedly flicking its tail and creating a vortex next to its body, which becomes a stored reservoir of kinetic energy. When it needs to accelerate rapidly to catch a prey, it slips into this vortex and sums its own swimming speed with the additional forces in the vortex [42]. Similarly, we supplement our cognitive schemata, by slipping into relational structures in the world that amplify their pattern. An example can be found in physical structures as simple as paths, which are created by other minds and which relieve later path followers of the cognitive load of navigation [43].

In distributed cognition, a group of individuals shares a task in ways suggesting that the group as a whole be considered a cognitive organism [41]. If the group achieves a set of cognitive symbioses with mutually supporting roles, then that unit of analysis may be highly illuminating. For example, the members of a navigation crew on a naval ship might individually have incomplete information that would make any one helpless to navigate the ship but that collectively results in efficient navigation (ibid).

Knowing who knows what

The phenomena of embodied and distributed cognition however do not eliminate the need for central cognition in the mind of the individual. Indeed, in some ways they highlight that need. The 'epidemiology of mental representations' in a community poses strong demands on characterizing the domain specific cognitions of individuals and how those cognitions are causally related to their more public products [44]. Consider, for example, how adults and children alike amplify their understandings by relying on the division of cognitive labour that is intrinsic to all cultures [7,45–47]. How do we know whom to seek out when we encounter the limits of our own understanding? If, we want to know more about a phenomenon, such as why water is transparent to, how do we know what experts to seek out for further information? Most western adults would refer to specialists in chemistry or physics, even though they may have never encountered that question before. On what grounds do they assign expertise in that way as opposed, for example, to a water conservation expert or to a manufacturer of swim goggles? A major factor guiding their use of the division of cognitive labour is the ways in which they understand, at a very coarse level, how diverse surface phenomena arise from common underlying causal patterns and systems. They believe not only in distinct essences for various kinds [28] but also in specific patterns of causation for those domains. Thus, relations of light to matter involve action at a distance, in which invisible microstructures have immediate causal consequences, in contrast to the regularities both of social interactions (at a distance, but with delays) and of physical mechanical ones (macroscopic contact and usually without delays). Even preschool children have some sense of these differences and use them to guide judgments about the division of cognitive labour [45].

In short, we overestimate of our own knowledge by underestimating how much information we recover in real time and by underestimating how much we ground our sense of detail in chains of deference to other minds. But those overestimations do not diminish the central importance of our causal gists. These gists make possible the elaboration of details in a constrained way and the ability to know who knows what around us.

Why less might be more

Why should we settle for such limits on the depths of our understanding and why should we be saddled with such illusions of explanatory depth? The answer could lie in the benefits of being shallow. Given that a fully exhaustive understanding in many domains requires an indefinitely deep tracking of causal patterns and regularities, there must be some way in which we know when we have grasped enough to function effectively in everyday life. The problem is analogous to the basic level of categorization, which captures key contrasts among natural categories without getting lost in the details [48,4]. Thus, it is important not only that we be driven to seek out causal patterns that explain surface regularities [49] but also that we know when to stop searching. We need a sense of comprehension that tells us we have what we need. But how does one achieve that sense with incomplete information? In the more formal sciences, we do so when our predictions are confirmed at a sufficiently high rate; but everyday explanatory understandings rarely are used to make explicit predictions about the future [50].

The rush of insight we often get in our attempts to understand the world around us may be at just the right level to provide sufficient constraints to build a much more detailed theory when in a relevant situation and at just the right level to allow us to ground our shallow understandings in much deeper and appropriate areas of expertise in other minds. We may relinquish some accuracy by thinking our understandings of reality correspond to having detailed annotated blueprints in each our minds; but pragmatically that sense of knowing as an individual may be correctly telling us what we really know when we really need to know it and allowing us to use all the resources at our disposal.

Conclusions

The strategy for the future is not to focus on how little we know when think we know more but rather to ask how we are able to develop such efficient ways of tracking the causal structure of the world around us without overloading our computational and storage systems. The intuitive understandings of the world that do exist in the mind of each individual are all the more remarkable for the power and success that they achieve with such compact and efficient means.

Acknowledgements

Preparation of this article and some of the research described herein was supported by NIH grants R01-HD23922 and R37 HD023922-17A1 to Frank Keil. Thanks to Lisa Webb for help in manuscript preparation.

References

1 Gilbert, S.F. (2000) *Developmental Biology*, 6th edn, Sinauer Associates
2 Giere, R.N. (1999) *Science without Laws (Science and its Conceptual Foundation)*, University of Chicago Press
3 Murphy, G.L. and Medin, D.L. (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316
4 Murphy, G.L. (2002) *The Big Book of Concepts*, MIT Press
5 Wilson, R.A. and Keil, F.C. (1998) The shadow and shallows of explanation. *Minds Machines* 8, 137–159
6 Dunbar, K. (1999) The scientist *in vivo*: how scientists think and reason in the laboratory. In *Model-Based Reasoning in Scientific Discovery* (Magnani, L. *et al.*, eds), pp. 85–89, Kluwer Academic/Plenum Press
7 Goldman, A.I. (2002) *Pathways to Knowledge: Private and Public*, Oxford University Press
8 Rozenblit, L.R. and Keil, F.C. (2002) The misunderstood limits of folk science: an illusion of explanatory depth. *Cogn. Sci.* 26, 521–562
9 Simon, H.A. (1981) *The Sciences of the Artificial*, 2nd edn, MIT Press
10 Simon, H.A. (2000) Discovering Explanations. In *Cognition and Explanation* (Wilson, R. and Keil, F., eds), MIT Press
11 Levin, D.T. *et al.* (2000) Change blindness blindness: the metacognitive error of overestimating change-detection ability. *Visual Cogn.* 7, 397–412
12 O'Regan, J.K. and Noe, A. (2001) A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.* 24, 939–1031
13 Fischhoff, B. *et al.* (1977) Knowing with certainty: the appropriateness of extreme confidence. *J. Exp. Psychol. Hum. Percept. Perform.* 3, 552–564
14 Kim, N.S. and Ahn, W.K. (2002) Clinical psychologists' theory-based representations of mental disorders predict their diagnostic reasoning and memory. *J. Exp. Psychol. Gen.* 131, 451–476
15 Medin, D.L. (1989) Concepts and conceptual structure. *Am. Psychol.* 44, 1469–1481
16 Wisniewski, E.J. (2002) Concepts and categorization. In *Steven's Handbook of Experimental Psychology: Memory and Cognitive Processes* (Vol. 2), 3rd edn, (Pashler, H. and Medin, D., eds), pp. 467–531, John Wiley & Sons
17 Wellman, H. (1990) *The Child's Theory of Mind*, MIT Press
18 Vosniadou, S. (1994) Capturing and modeling the process of conceptual change. *Learn. Instruct.* 4, 45–69
19 Carey, S. (1995) On the origins of causal understanding. In *Causal Cognition: a Multi-Disciplinary Approach* (Sperber, D. *et al.*, eds), pp. 268–308, Clarendon Press

20 Thagard, P. (2000) Probabilistic networks and explanatory coherence. *Cogn. Sci. Q.* 1, 91–114
21 Rehder, B. and Ross, B.H. (2001) Abstract coherent categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 27, 1261–1275
22 Gopnik, A. and Meltzoff, A.N. (1997) *Words, Thoughts and Theories*, Bradford Books/MIT Press
23 diSessa, A.A. (1988) Knowledge in pieces. In *Constructivism in the Computer Age* (Pufall, G.F. and Peter, B., eds), pp. 49–70, Erbaum
24 Keil, F. (1995) The growth of understandings of natural kinds. In *Causal Cognition* (Sperber, S. *et al.*, eds), pp. 234–267, Clarendon Press
25 Dennett, D. (1998) The intentional stance MITECS. In *The MIT Encyclopedia of Cognitive Sciences* (Wilson, R. and Keil, F., eds), MIT Press
26 Kelemen, D. (1999) Function, goals, and intention: children's teleological reasoning about objects. *Trends Cogn. Sci.* 3, 461–468
27 Keil, F.C. (1989) *Concepts, Kinds and Cognitive Development*, Bradford Books/MIT Press
28 Gelman, S.A. (2003) *The Essential Child*, Oxford University Press
29 Bloom, P. (1996) Intention, history, and artifact concepts. *Cognition* 60, 1–29
30 Ahn, W. *et al.* (2000) Causal status as a determinant of feature centrality. *Cogn. Psychol.* 41, 361–416
31 Spelke, E.S. *et al.* (1995) Infants' knowledge of object motion and human action. In *Causal Cognition: a Multidisciplinary Debate* (Sperber, D. *et al.*, eds), Oxford University Press
32 Johnson, S. (2000) The recognition of mentalistic agents in infancy. *Trends Cogn. Sci.* 4, 22–28
33 Opfer, J.E. and Gelman, S.A. (2001) Children's and adults' models for predicting teleological action: The development of biology-based models. *Child Dev.* 72, 1367–1381
34 Keil, F.C. (1998) Cognitive science and the origins of thought and knowledge. In *The Handbook of Child Psychology* In *Theoeretical Models of Human Development* (Vol. 1), 5th edn, (Damon, W. and Lerner, R., eds), pp. 341–413, John Wiley & Sons
35 Santos, L.R. *et al.* (2002) Domain-specific knowledge in human children and non-human primates: artifact and food kinds. In *The Cognitive Animal* (Bekoff, M. *et al.*, eds), MIT Press
36 Sloman, S.A. *et al.* (1998) Feature centrality and conceptual coherence. *Cogn. Sci.* 22, 189–228
37 Barsalou, L.W. (1983) Ad hoc categories. *Mem. Cogn.* 11, 211–217
38 Barsalou, L.W. (2002) Being there conceptually: simulating categories in preparation for situated action. In *Representation, Memory, and Development: Essays in Honor of Jean Mandler* (Stein, N.L. *et al.*, eds), pp. 1–16, Erlbaum
39 Luhmann, C.C. *et al.* Theories and similarity: Categorization under speeded condition. *Proc. 24th Annu. Conf. Cogn. Sci. Soc.*, pp. 590–595, Erlbaum (in press)
40 Clark, A. (2001) Reasons, robots and the extended mind. *Mind Lang.* 16, 121–145
41 Hutchins, E.L. (1995) *Cognition in the Wild*, MIT Press
42 Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*, MIT Press
43 Sperber, D. (1999) Conceptual tools for a natural science of society and culture. *Proc. Br. Acad.* 111, 297–317
44 Sperber, D. (1996) *Explaining Culture: a Naturalistic Approach*, Blackwell Publishing
45 Lutz, D.J. and Keil, F.C. (2002) Early understandings of the division of cognitive labor. *Child Dev.* 73, 1073–1084
46 Goldman, A. (2002) *Pathways to Knowledge*, Oxford University Press
47 Kitcher, P. (1993) *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*, Oxford University Press
48 Rosch, E. *et al.* (1976) Basic objects in natural categories. *Cogn. Psychol.* 8, 382–439
49 Macnamara, J. (1999) *Through the Rearview Mirror: Historical Reflections on Psychology*, MIT Press
50 Keil, F.C. *et al.* (1998) Two dogmas of conceptual empircism. *Cognition* 65, 103–135
51 Wright, M., Patel, M. eds (2000) *How Things Work Today*, Crown Publishers, New York
52 Ahn, W. *et al.* (2000) Causal status effects in children's categorization. *Cognition* 76, 35–43
53 Ahn, W. *et al.* (2001) Why essences are essential in the psychology of concepts. *Cognition* 82, 59–69
54 Medin, D.L., Atran, S. eds (1999) *Folkbiology*, MIT Press
55 Simons, D.J. and Keil, F.C. (1995) An abstract to concrete shift in the development of biological thought: the insides story. *Cognition* 56, 129–163
56 Inagaki, K. and Hatano, G. (2002) *Young Children's Thinking about the Biological World*, Psychology Press
57 Atran, S. (1998) Folk biology and the anthropology of science. *Behav. Brain Sci.* 21, 547–609
58 Atran, S. *et al.* (2001) Folkbiology doesn't come from folkpsychology: evidence from Yukatek Maya in cross-cultural perspective. *J. Cogn. Cult.* 1, 3–42
59 Lopez, A. *et al.* (1997) The tree of life: universal and cultural features of folkbiological taxonomies and inductions. *Cogn. Psychol.* 32, 251–295
60 Bailenson, J.N. *et al.* A birds eye view: biological categorization and reasoning within and across cultures. *Cognition* (in press)