Philosophia Scientiæ

Travaux d'histoire et de philosophie des sciences

18-3 | 2014

Logic and Philosophy of Science in Nancy (I)

The Principle Based Explanations Are Not Extinct in Cognitive Science: The Case of the Basic Level Effects

Lilia Gurova

Electronic version

URL: <http://journals.openedition.org/philosophiascientiae/1017>

DOI: 10.4000/philosophiascientiae.1017

ISSN: 1775-4283

Publisher

Éditions Kimé

Printed version

Date of publication: 1 October 2014

Number of pages: 203-214

ISBN: 978-2-84174-689-7

ISSN: 1281-2463

Electronic reference

Lilia Gurova, « The Principle Based Explanations Are Not Extinct in Cognitive Science: The Case of the Basic Level Effects », *Philosophia Scientiæ* [Online], 18-3 | 2014, Online since 19 January 2015, connection on 05 November 2020. URL : <http://journals.openedition.org/philosophiascientiae/1017> ; DOI : <https://doi.org/10.4000/philosophiascientiae.1017>

Tous droits réservés

The Principle Based Explanations Are Not Extinct in Cognitive Science: The Case of the Basic Level Effects

Lilia Gurova

New Bulgarian University (Bulgaria)

Résumé : On observe une nouvelle tendance dans la philosophie des sciences cognitives, manifeste dans les écrits de Bechtel *et al.* qui met en avant l'importance des explications mécanistes au détriment du rôle explicatif des principes. Cet article est un plaidoyer pour rétablir l'équilibre. Il met l'accent sur l'effort d'explication des effets du niveau de base, l'une des plus importantes découvertes empiriques dans l'histoire de la recherche en catégorisation. L'analyse de trois différentes périodes de cette histoire révèle que le recours aux principes y a joué un rôle crucial. Cependant, afin de reconnaître pleinement le rôle explicatif des principes, nous devrions nous préparer à admettre que les explications déductives-nomologiques ne sont pas les seuls types d'explications basées sur des principes.

Abstract: There is a tendency in recent philosophy of cognitive science, best seen in the writings of Bechtel *et al.*, to overstate the significance of mechanistic explanations and to neglect the explanatory role of principles. This paper is a plea for restoring the balance. It draws attention to the search for explaining the so-called basic level effects, one of the most important empirical findings in the history of categorization research. The analysis of three different episodes from this history reveals that appeals to principles have played an important role in it. However, in order to fully recognize the explanatory role of principles one should be ready to admit that the deductive-nomological explanations are not the only species of principle-based explanations.

1 Introduction

There is a tendency in recent literature on the philosophy of cognitive science and psychology to overstate the significance of mechanistic explanations in these areas (see [Cummins 2000], [Bechtel & Abrahamsen 2005], [Abrahamsen & Bechtel 2006], [Wright & Bechtel 2007], [Bechtel 2008, 2009, 2010]). One can read, for example, that cognitive science “is, more than anything else, a pursuit of cognitive mechanisms” [Abrahamsen & Bechtel 2006, 159], that “the term mechanism is ubiquitous when psychologists and neuroscientists offer explanations for mental activities” [Bechtel 2008, ix], and that “examination of the explanatory discourse of psychologists reveals a shift in emphasis from laws to mechanisms [...]” [Wright & Bechtel 2007, 31]. Such claims are sometimes accompanied by word counts showing that recourses to “mechanisms” are much more frequent than mentions of “laws” and that the latter are very rare (see e.g., [Bechtel & Wright 2011]). Such word counts, however, are highly biased. The word “law” is neither the only word used in psychological language for designating general explanatory statements, nor the one used most frequently.¹ There is another substitute for “law” which has been broadly used in psychological language but which, for some reasons, was completely neglected by the proponents of “mechanistic science”.² This broadly used term for designating law-like statements is the term “principle”. A simple search in Wiley’s abstracts of publications in the field of cognitive science for the period 1991-2011 shows 196 hits for “mechanism” and 86 hits for “principle”. These numbers reveal that, roughly, for any two uses of “mechanism” there is one use of “principle”, a ratio which suggests that in cognitive science, looking for mechanisms is hardly the only game in town. It is an open question whether all recourses to principles in cognitive science literature are made for explanatory reasons.³ To answer this question, one needs to have a closer look at the real uses of principles in cognitive sciences. However, given the excessive attention to mechanisms and mechanistic explanations, such an aim, if considered at all, would look

1. The proponents of the mechanistic explanatory project have noticed that what psychologists call “effects” corresponds to what is called “empirical laws” in natural sciences. For them, however, effects are not explanatory insofar as they need themselves to be explained [Cummins 2000], [Bechtel 2009, 2010]. I have shown elsewhere that statements describing psychological effects, like empirical laws in physics, can be explanatory, and that they have been used to explain particular phenomena [Gurova 2013]. This paper is focused on a different substitute for the word “law” (the word “principle”) which has been broadly used in psychological language although it has been completely neglected by the representatives of mechanistic science.

2. The term “mechanistic science” was launched by Herschbach & Bechtel as an umbrella term for the attempts “to explain how a mechanism produces a phenomenon by decomposing it into its parts and operations and then recomposing the mechanism to show how parts and operations are organized, such that when the mechanism is situated in an appropriate environment, it generates the phenomenon” [Herschbach & Bechtel 2011, 203].

3. The same note can be made about the recourses to “mechanism”.

peripheral or outdated to most philosophers interested in explanatory practice in cognitive science. The aim of this paper is to take a step towards restoring the balance. It draws attention to three different episodes from the newest history of cognitive science. These are three of the most important attempts to provide an explanation for what is known today as the “basic level effects” in categorization. The analysis reveals that: (1) appeals to principles have been a normal practice rather than an exotic one through the whole history of research on basic level effects; (2) the principle-based explanations which have been advanced are not sub-species of the deductive-nomological explanations; (3) some “design principles” have played a primary role in the construction of mechanistic explanations and in this sense, the latter should be viewed as a variety rather than as a rival of the principle-based explanations. These findings have interesting implications which are discussed in the last part of the paper. Let’s see first, however, how they have been obtained.

2 Episode 1: The first observations suggesting that preferred categories exist and the first attempts to explain them

Roger Brown begins his paper “How shall a thing be called?” [Brown 1958] with a description of the following observation. Although we can use different names for the same object⁴ (e.g., we can name the same coin “a dime”, “money” or “a metal object”, and we can call the same dog “a dog”, “a boxer”, or “an animal”), when talking to children, most adults prefer one of the many possible names for the thing which they are talking about. Moreover, the choices of different adults talking to different children are mainly the same. “How are these choices determined?”—is the question which Roger Brown states as a starting point of his search for explaining the described phenomenon [Brown 1958, 14].

The first explanatory hypothesis which Brown investigates is the belief shared by many adults that when talking to children they tend to use shorter names because of the recognized difficulty for children to pronounce longer and complicated words. It is easy to see that this common belief is untenable. The preferred name “dog” is indeed shorter than (the less preferred) “boxer” and “animal”, but the words “pineapple” and “pomegranate” are longer than the word “fruit”, nevertheless most adults do use them instead of “fruit” when asked by children about the names of these things.

Brown, however, does not give up the common belief immediately. Before dropping it (for the reasons mentioned above), he uses it for a while as if

4. Insofar as “every referent has many names” [Brown 1958, 14].

it were true, i.e., as a principle.⁵ The acceptance of the brevity principle allowed Brown to make use of an important finding reported by Zipf [Zipf 1935]: that “the length of a word (in phonemes or syllables) is inversely related to its frequency in the printed language” [Brown 1958, 4]. The documented correlation between the length of the words and the frequency of their use allowed Brown to generalize the brevity principle into what he called “the frequency-brevity principle” which states that the “choice of a name is usually predictable from either frequency or brevity” of that name [Brown 1958, 14]. After arriving at this new explanatory principle, Brown abandoned its “brevity part” because it “seems not to be the powerful determinant we had imagined” [Brown 1958, 14]. This seemingly surprising move is not illogical: given that the frequency principle is true if the brevity principle is true, the rejection of the brevity principle does not logically entail the rejection of the frequency principle (“non-B” does not follow from “if A then B” and “non-A”).

The empirical support for the frequency principle seemed questionable, too, as Brown himself noticed: the word “pineapple” is not used more frequently than the word “fruit”. Yet again he decided to apply the same seemingly strange logical move: to use for a while the frequency principle as if it were true. Taking the frequency principle for granted raises a new question: what makes some names more frequently used than others? We can assume, Brown states, that

it just happened that way, like driving on the right side of the road in America and on the left in England. The convention is preserved but has no justification outside itself. [Brown 1958, 15]

But he immediately adds that such an assumption is hardly plausible. In order to see that, it suffices to ask whether it could be equally possible to “give coins proper names and introduce people as types”? [Brown 1958, 15]. By taking into account the existence of different naming practices in respect to different objects (people do not follow the same rules when choosing names for coins and people) and among different groups (adults and children often demonstrate different naming preferences) and led by the conjecture that the differences between the naming practices are not the result of arbitrary conventions, Brown arrives at another conclusion: that people tend to name a thing at the level of its “usual utility” [Brown 1958, 16], i.e., to categorize it in the most useful way. This new “utility principle” suggests the following explanation of the fact that parents tend to call a dime “money” when talking to a small child. They do not do so because the word “money” is shorter than the word “dime” (it is not), or because “money” is more frequently used than “dime” (although it is). When talking to their children, parents prefer the word “money” because it is easier to incorporate in the “functional structure of the child’s world” [Brown 1958, 16] which parents anticipate.

5. The default meaning of “principle” assumed in this paper is “a general statement which is assumed to be true”. As we shall see further, the word “principle” allows different uses.

The utility principle, however plausible it may look, does not provide the ultimate explanation for the existence of preferred categorizations in human experience. Taken for granted, the utility principle leads to a new question: what makes particular categorizations in particular circumstances more useful than others? As we shall see, this question played an important role in the following research on categorization.

3 Episode 2: Rosch's principle-based account for the basic level of categorization

About two decades after Brown's first attempt to explain the noticed preferences in naming things for children, Eleanor Rosch and her collaborators showed that similar preferences penetrate the whole cognitive system. Their research was motivated and guided by two assumptions. The first is that categorization, or "the cutting up of the environment into classifications by which non-identical stimuli can be treated as equivalent" is "one of the most basic functions of all organisms" [Rosch, Mervis *et al.* 1976, 382]. The second assumption is that in order to understand the process of categorization, one should "*determine the principles* by which humans divide up the world in the way they do" [Rosch, Mervis *et al.* 1976, 382–383]. Like Brown, Rosch refused to take seriously the possibility of the observed regularities in human categorization to be a result of more or less arbitrary conventions. "Such a view would be reasonable only if the world were entirely unstructured" [Rosch, Mervis *et al.* 1976, 383], but it is not. Some stimulus attributes occur combined with other attributes more frequently than with any of the rest of observed attributes, e.g., wings are more likely to appear together with feathers than with fur. This observation led Rosch to the formulation of the first principle which, according to her, reveals the main determinant of human categorization. The principle states that the world appears to us structured in the sense that some stimuli appear together with other stimuli more often than with the rest of the perceived stimuli. Rosch's principle of perceived world structure, taken together with the consideration that any categorization should facilitate the organism's interaction with the world (an idea that reminds us of Brown's utility principle), makes plausible the conjecture that the most useful category cuts are those which catch the perceived world structure in the best possible way. The categories which satisfy this requirement Rosch called "basic categories" or "basic objects". Respectively, the level in a given taxonomy at which the basic category cuts are made is called "basic level" and its manifestations in different cognitive processes are called "basic level effects". But what determines which categories best correspond to the perceived world structure? A second principle which Rosch called "the principle of cognitive

economy” suggests the answer: the most useful category cuts, or the category cuts that best correspond to the perceived world structure are those that guarantee maximum information with the least cognitive efforts (the least number of categories to deal with). The basic category cuts are also the most discriminative ones: the ratio of the similarity between category members to the similarity between category members and non-members has the highest values for basic categories.

But how do people recognize the basic categories? A proponent of mechanistic science would insist that the proper answer to this question should point to the mechanism underlying the estimation of the similarity-within to similarity-between ratio. Rosch, however, does not take this path. Her two principles of categorization allow her to arrive at experimentally confirmed hypotheses suggesting that people directly recognize the basic categories: because the members of these categories share a common shape, the subjects easily assign a common image to them. And because the image associated with each basic category is different from the images the subjects create for the other basic categories belonging to the same taxonomy, the basic categories appear as the most easily grasped in perception. This explains why basic categories are the first categories obtained during perception, and why they are the first recognized (and named) by children.

Although Rosch and her collaborators did not ask questions about the possible mechanisms underlying the processes of recognition of basic level categories, other researchers did that. In the following part of this paper we shall discuss what is, to the best of our knowledge, the most successful attempt to provide a mechanistic theoretical account for the existence of privileged categorizations and for the effects these categorizations exert on the whole cognitive system. The following questions will be in the focus of the forthcoming analysis: (1) How does the mechanistic explanatory approach in this case relate to the preceding non-mechanistic explanations? Is it suggested to them as an alternative or does it rather build on what they have stated? (2) Could the proposed mechanistic account be fully analyzed in terms of the “entities” and “activities” that the alleged mechanism consists of, as “the consensus view of mechanisms”⁶ in philosophy of science suggests, or is there something else which the traditional mechanistic analysis tends to ignore? As we shall see, the answers to these questions appear rather surprising for those who have taken the perspective of mechanistic science as it has been understood by most of its recent defenders in philosophy of science.

6. The “consensus view” is described in [Fagan 2012].

4 Episode 3: An attempt for a mechanistic explanation of the basic level effects

In 2004, Timothy Rogers & James McClelland published a book [Rogers & McClelland 2004], the main aim of which, as they stated it, was to lay the foundations of a general theoretical framework for a mechanistic explanation of semantic cognition. This theoretical framework, as they viewed it, should address the questions:

How do we perform semantic tasks...? How do we represent the information that we use as the basis for performing such tasks, and how do we acquire this information. [Rogers & McClelland 2004, 1]

For Rogers & McClelland, “semantic” are all processes of acquiring or using information which “is not available more or less directly from the perceptual input” [Rogers & McClelland 2004, 2]. Given this definition, it’s easy to realize that the processes of categorization (i.e., the processes of cutting the perceived world into different categories and of recognizing an object as belonging or non-belonging to an already known category) are semantic processes. In this sense, the mechanistic theory of semantic cognition which Rogers & McClelland have proposed is at the same time a theory of categorization. Rogers & McClelland admitted that so far the mechanistic approach has not proved successful in this area, but this is because, they claimed, the earlier mechanistic (computational) attempts to account for semantic cognition were based on inappropriate ideas. Among these ideas, to which Rogers & McClelland gave the common name “categorization-based theory” of semantic cognition, are the beliefs that each category, once learned, has a stored local representation in human memory and that all these local category representations are organized in hierarchical (taxonomic) structures. These beliefs, however well they conform to the way we think about categories in our everyday life, form, according to Rogers & McClelland “an incomplete and in some ways paradoxical basis for accounting for the relevant empirical phenomena” [Rogers & McClelland 2004, 5]. Rogers & McClelland’s list of empirical phenomena which, according to them, the traditional “categorization-based” view cannot properly account for, contains phenomena which have been taken to suggest the existence of privileged categories, or levels of categorization in the hierarchical structures of categories. The categorization-based approach to basic categories leads to the assumption that these categories are privileged “entry points” in the stored hierarchical category structures. That means that subjects directly recognize a given object as a member of a certain basic category and only indirectly (by inference) recognize it as a member of the corresponding subordinate and superordinate categories. This view is compatible with most of the observed basic level effects (see [Rogers & McClelland 2004, 22, Table 1–2]), but it is difficult to reconcile with other phenomena. For example, it is taken to be

reliably established that at a very early age (before they have learned to speak and to name things) children form very general categories, which after that “progressively differentiate” in the development, allowing children to first learn the names of basic level categories when they begin to talk.⁷ It is difficult to explain the generality of the early categorizations given the assumption that basic categories are perceptually the most salient ones. A second puzzling fact for the privileged entries view of basic categories is the deterioration of concepts in subjects with dementia. Again, it has been confirmed that patients with dementia lose last the abilities to categorize at the most general levels which are superordinate in respect to the basic level. But how is this possible if the superordinate categorizations are available only indirectly (i.e., they are inferred from the basic level categorization), given that the ability to categorize at the basic level has already been lost? Similarly puzzling for the categorization-based approach are the phenomena of interference of typicality and basic-levelness (the atypical members of a given category are preferably named at the subordinate level⁸), and the influence of expertise (experience) on basic level effects.⁹

According to Rogers & McClelland, in order to find a solution to all these puzzles, one should abandon the idea that local category representations are organized in human memory in hierarchical structures which have privileged “entry points”. All basic level effects and the way these effects interfere with typicality, expertise, conceptual progressive differentiation in early childhood, and semantic deterioration in patients with dementia could be explained if one accepted, against the category-based approach, that category representations are distributed. To accept this means to agree that there are no local category representations stored in human memory but rather nodes corresponding to different input stimuli and output categorizations. These nodes are highly interconnected in such a way that each connection between any two nodes has a particular strength corresponding to the frequency with which the events represented by these two nodes appear together. Thus the same node can be part of different category representations which, in this view, are reduced to different patterns of interconnected nodes. Rogers & McClelland’s ideas about distributed category representations are an implementation of the PDP (parallel distributed processing) approach to cognitive modeling which was launched in the 1980s by Hinton & Anderson (see [Hinton & Anderson 1981], [Rumelhart, McClelland *et al.* 1986]). There is not enough space here (and it is not necessary) to explain in details the ideology of PDP and the way Rogers & McClelland built their mechanistic theory of semantic cogni-

7. For discussion on and references to these results see [Rogers & McClelland 2004, 19, 176].

8. E.g., sparrows (typical birds) are preferably called “birds” but penguins (atypical birds) are preferably called “penguins”—see [Jolicoeur, Gluck *et al.* 1984] for discussion.

9. Evidence that experts prefer to name at the subordinate level in their area of expertise has been provided in [Johnson & Mervis 1997].

tion on PDP's ideas. Nor is it necessary to discuss to what extent Rogers & McClelland's claims are justified that they have successfully reconciled phenomena which seemed paradoxical from the perspective of the traditional (localist) categorization-based view. Rogers & McClelland themselves were modest enough to note that their mechanistic theoretical framework of semantic cognition is far from being "a full characterization of the mechanistic basis of semantic knowledge", it's rather an account of a type of the mechanism, further exploration of which might lead to such a full characterization [Rogers & McClelland 2004, x]. Let's suppose that Rogers & McClelland's model of semantic cognition succeeded in explaining what, they claimed, none of the previous theories, mechanistic or not, could explain. Then the following questions arise: (1) How does Rogers & McClelland's explanation of basic level effects relate to the previous ones which were discussed in sections 2 and 3 of this paper? (2) Where does the alleged success of their explanation originate?

Rogers & McClelland pointed explicitly to what their theoretical framework is an alternative to: namely, categorization-based theories of semantic processing which presume that cognitive systems contain explicit category representations structured both hierarchically and around prototypes and which contain privileged "entry-points" corresponding to the basic level categories. Categorization-based theories of semantic processing do build on principles determining what makes the basic categories so special but they are not the only possible instantiation (a processing complement) of these principles. There is evidence that Rogers & McClelland have refuted the categorization-based processing model but not the principles it builds on. Their assumption that the way humans categorize the world is "strongly constrained by the coherent co-variation of ensembles of properties in experience" [Rogers & McClelland 2004, 351], is very close to, if not identical to, Rosch's first principle which states that the perceived world structure (which we derive from the appearance of some stimuli combinations more often than others) is the main determinant of human categorization. Rogers & McClelland also take for granted the main implication of Rosch's two principles which states that basic categories maximize "both *informativeness* and *distinctiveness*" [Rogers & McClelland 2004, 17]. So, in brief, one cannot reasonably claim that Rogers & McClelland's semantic theory is a strong alternative to the previous principle-based explanations of basic level effects which addressed the question of what determines the privileged status of the basic categories. Their theory is rather a processing (mechanistic) complement to these explanations providing the answer to a different question, namely, how the basic level effects appear in the general process of acquiring and using semantic information.

Let's now turn to the question where the success of Rogers & McClelland's explanation stem from. The core idea underlying their approach is that all "semantic judgments emerge from the sensitivity of a general learning mechanism to coherent co-variation" of perceived properties of things [Rogers & McClelland 2004, 352]. The mechanistic approach to scientific explanation which seems to dominate contemporary philosophy of cognitive science, sug-

gests to look for the “entities” and “activities” which the analyzed mechanism comprises, and for the way these “entities” and “activities” are organized, in order to explain how the studied mechanism works. Rogers & McClelland, however, admitted that at least some of the results they obtained might be reproduced by other processing models describing a different mechanistic structure. But there are principles, they added, “that might be respected by any network that might be proposed as a mechanism for extracting semantic structure from experience” [Rogers & McClelland 2004, 371] and they see as their main achievement the formulation of these principles. One of these principles is the above mentioned principle of “coherent co-variation of properties across items and contexts” [Rogers & McClelland 2004, 350]. Another principle which Rogers & McClelland highly praise is the so-called “convergence principle” which states that the processing and representation of semantic information is to be organized in such a way “that all different sorts of information about all objects in all contexts converge on the same units and connections” [Rogers & McClelland 2004, 360]. For Rogers & McClelland, this principle is important because it suggests a design of the cognitive system which “is likely to be selected for by evolution” [Rogers & McClelland 2004, 371].

In brief, Rogers & McClelland’s mechanistic explanation of basic level effects neither directly refutes the preceding principle-based explanations (it rather builds on them) nor eschews explanatory principles altogether. The mechanistic theory of semantic cognition launched by Rogers & McClelland rests on principles which they view as the most essential part of their theory.

5 Conclusion

The analysis of the three episodes from the recent history of cognitive science reveals the central role of principles in all attempts to explain the phenomena suggesting that a privileged, “basic” level of categorization exists. Even the proposed mechanistic account of the basic level effects (Episode 3) was shown to be essentially determined by a set of “design principles” [Rogers & McClelland 2004]. Rogers & McClelland are not the only ones who make the point that there is more to mechanistic explanations than the analyses in terms of “entities” and “activities” reveal (see also [Chater & Brown 2008]). No doubt, however, when this fact is stressed by two leading representatives of “mechanistic science”, even the most radical mechanist philosophers should take note of it.

Another methodological bias hindering the recognition of the proper explanatory role of principles is the belief that any principle-based explanation is deductive-nomological. Any use of principles as big premises in syllogistic reasoning can hardly be recognized in the cases which were discussed in this paper. The undertaken analysis shows that the advanced principles have either been used as direct explanatory statements (such as Brown’s “frequency

principle”, stating that the preferred names are the ones most frequently used) or as inferential licenses which suggest, support, or make plausible certain explanations without logically implying them.¹⁰ Probably these are not the only possible ways for principles to take part in non-deductive explanatory inferences. New case studies have to be done to eventually confirm this conjecture. The present paper is only intended to persuade those who are interested in the explanatory practice in cognitive science that such a line of research might be rewarding.

Bibliography

- ABRAHAMSEN, Adele & BECHTEL, William [2006], Phenomena and mechanisms: Putting the symbolic, connectionist, and dynamical systems debate in broader perspective, in: *Contemporary Debates in Cognitive Science*, edited by R. Stainton, Oxford: Blackwell, 159–185.
- BECHTEL, William [2008], *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*, New York: Taylor & Francis.
- [2009], Constructing a philosophy of science of cognitive science, *Topics in Cognitive Science*, 1(3), 548–569, doi:10.1111/j.1756-8765.2009.01039.x.
- [2010], How can philosophy be a true cognitive science discipline?, *Topics in Cognitive Science*, 2(3), 357–366, doi:10.1111/j.1756-8765.2010.01088.x.
- BECHTEL, William & ABRAHAMSEN, Adele [2005], Explanation: A mechanist alternative, *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2), 421–441, doi:10.1016/j.shpsc.2005.03.010.
- BECHTEL, William & WRIGHT, Cory [2011], What is psychological explanation?, in: *Routledge Companion to Philosophy of Psychology*, edited by J. Symons & P. Calvo, London: Routledge, 113–130.
- BROWN, Roger [1958], How shall a thing be called?, *Psychological Review*, 65, 14–21.
- CHATER, Nick & BROWN, Gordon [2008], From universal laws of cognition to specific cognitive models, *Cognitive Science*, 32, 36–67, doi:10.1080/03640210701801941.

10. For example, Rosch’s principle stating that the basic level categories maximize informativeness and distinctiveness makes plausible the assumption that the members of a basic category share a common shape but does not logically imply it. Once confirmed, the property of the basic categories’ members to have a common shape has been explained by the principle which suggested its existence: Why do people prefer to use categories the members of which have the same or similar shape? Because such categories are more informative and more distinctive than the categories which members’ shape significantly varies.

- CUMMINS, Robert [2000], “How does it work?” versus “What are the laws?”: Two conceptions of psychological explanation, in: *Explanation and Cognition*, edited by F. Keil & R. Wilson, Cambridge, MA: The MIT Press, 117–144.
- FAGAN, Melinda [2012], The joint account of mechanistic explanation, *Philosophy of Science*, 79(4), 448–472, doi:10.1086/668006.
- GUROVA, Lilia [2013], Principles versus mechanisms in cognitive science, in: *EPSA11 Perspectives and Foundational Problems in Philosophy of Science*, edited by V. Karakostas & D. Dieks, Dordrecht: Springer, 393–403, doi: 10.1007/978-3-319-01306-0_32.
- HERSCHBACH, Mitchell & BECHTEL, William [2011], Relating Bayes to cognitive mechanisms, *Behavioral and Brain Sciences*, 34(4), 202–203, doi: 10.1017/S0140525X11000318.
- HINTON, Geoffrey & ANDERSON, James [1981], *Parallel Models of Associative Memory*, Hillsdale, NJ: Lawrence Erlbaum.
- JOHNSON, Kathy & MERVIS, Carolyn [1997], Effects of varying levels of expertise on the basic level of categorization, *Journal of Experimental Psychology: General*, 126(3), 248–277, doi:10.1037//0096-3445.126.3.248.
- JOLICOEUR, Pierre, GLUCK, Mark A., & KOSSLYN, Stephen M. [1984], Pictures and names: Making the connection, *Cognitive Psychology*, 16(2), 243–275, doi:10.1016/0010-0285(84)90009-4.
- ROGERS, Timothy & MCCLELLAND, James [2004], *Semantic Cognition. A Parallel Distributed Processing Approach*, Cambridge, MA: The MIT Press.
- ROSCH, Eleanor, MERVIS, Carolyn, GRAY, Wayne D., JOHNSON, David, & BOYES-BRAEM, Penny [1976], Basic objects in natural categories, *Cognitive Psychology*, 8(3), 382–439, doi:10.1016/0010-0285(76)90013-X.
- RUMELHART, David, MCCLELLAND, James, & THE PDP RESEARCH GROUP [1986], *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 1, Cambridge, MA: The MIT Press.
- WRIGHT, Cory & BECHTEL, William [2007], Mechanisms and psychological explanation, in: *Handbook of Philosophy of Science. Philosophy of psychology and cognitive science*, edited by P. Thagard, Amsterdam: Elsevier, 31–79.
- ZIFF, George K. [1935], *The Psycho-biology of Language*, Boston: Houghton Mifflin.