



Interactive Effects of Explicit Emergent Structure: A Major Challenge for Cognitive Computational Modeling

Robert M. French,^a Elizabeth Thomas^b

^a*Université de Bourgogne*

^b*Institut National de la Santé et de la Recherche Médicale (INSERM U1093), Université de Bourgogne*

Received 1 August 2013; received in revised form 23 January 2015; accepted 23 January 2015

Abstract

David Marr's (1982) three-level analysis of computational cognition argues for three distinct levels of cognitive information processing—namely, the computational, representational, and implementational levels. But Marr's levels are—and were meant to be—descriptive, rather than interactive and dynamic. For this reason, we suggest that, had Marr been writing today, he might well have gone even farther in his analysis, including the emergence of structure—in particular, explicit structure at the conceptual level—from lower levels, and the effect of explicit emergent structures on the level (or levels) that gave rise to them. The message is that today's cognitive scientists need not only to understand how emergent structures—in particular, explicit emergent structures at the cognitive level—develop but also to understand how they feed back on the sub-structures from which they emerged.

Keywords: Emergence; Interactive emergence; Connectionist models; Active symbols

1. Introduction

David Marr's (1982) three-level analysis of computational cognition argues for three distinct levels of cognitive information processing—namely, the computational, representational, and implementational levels. It was a clear and powerful analysis of the computational metaphor of disembodied mind. Marr was, however, not alone in his tripartite analysis of information-processing systems (McClamrock, 1990). Pylyshyn (1984), one of the leading philosophers of the symbolic artificial intelligence movement at that time wrote about the symbolic, syntactic, and physical levels; in cognitive psychology, the distinction was one of content, form, and medium (Glass, Holyoak, & Santa, 1979). But among cognitive scientists, Marr's three-level analysis was, and remains, the most familiar.

Correspondence should be sent to Robert M. French, LEAD-CNRS UMR 5022, Université de Bourgogne, 21000 Dijon, France. E-mail: robert.french@u-bourgogne.fr

We need to briefly recall the state of cognitive science circa 1980 in order to better situate Marr's famous distinction of levels. At that time the symbolic tradition was the king of the heap. One of the underlying tenets of this tradition was that reality could be encoded as symbols that referred to objects, situations, and actions in the world, and rules consisting of relations between those symbols. While adherents of the symbolic tradition obviously acknowledged the existence of an implementational level, it was considered largely irrelevant to the level at which symbols and rules operated. In other words, a "chair" had a list of properties that could be referred to (and instantiated), whether it was being used to sit on, to reach an object on a high shelf, or to hold back a lion in a circus act. And this list of properties did not depend on any particular "neuronal/implementational" level. What counted for these researchers were the computational and the algorithmic levels. Marr's reintroduction of the neuronal level was, therefore, a significant step and a harbinger of neural network approaches to cognition.

But Marr's levels are—and were meant to be—descriptive, rather than interactive and dynamic. For this reason, we suggest that, had Marr been writing today, he might well have gone even farther in his analysis. He would, we believe, have included a discussion of the following issues (or at least he would have agreed with them): (a) the notion of embodiment, that is, the notion that the computational modeling of intelligence must also take into account the physical specificities of our bodies that allow us to interact with the world around us, (b) the essential blurriness of the boundaries between levels of information processing, (c) the emergence of structure—in particular, explicit structure at the conceptual level—from lower levels, and (d) the effect of explicit emergent structures on the level (or levels) that gave rise to them.

2. Emergence

Emergence has been studied by philosophers since at least the middle of the 19th century (e.g., Mill, 1858) with many significant contributions made in the early decades of the 20th century (e.g., Alexander, 1920; Broad, 1925). After a hiatus from the spotlight, emergence is once again all the rage. Everything from human language to speciation, from economic recessions to political upheavals, from the human eye to the origin of life, and much, much more is described and studied in the framework of emergence. More recently philosophers have discussed emergence in the context of complexity theory and a non-Laplacian universe (e.g., Bedau, 1997).

Our goal in the present article, however, is not to comment on the work of these philosophers, but rather to consider the question of emergence as it applies to Marr's three-level distinction of computational cognition, in general, and to neural-network models of cognition, in particular. Our use of the term "emergence" is close to a recent definition of Goldstein (1999): "the arising of novel and coherent structures, patterns and properties during the process of self-organization in complex systems. Emergent phenomena are conceptualized as occurring on the macro level, in contrast to the micro-level components and processes out of which they arise."

Bedau (1997) further elaborates on two “useful hallmarks” of emergence. Emergent phenomena, he writes, are (a) somehow constituted by, and generated from, underlying processes and (b) are somehow autonomous from underlying processes. Both of these properties are very much in line with the approach taken in the present paper. The emergence in a cognitive system of rules—for example, the “vowel-consonant-e” pronunciation rule and the “if *beak* then *bird*” rule that will be discussed later in this paper—provide examples of the latter criterion, although in our case perhaps “semi-autonomous” should replace “autonomous” in Bedau’s definition.

Curiously, the notion of emergence—for example, how the algorithmic/representational level might emerge from the neuronal level—is largely absent from Marr’s three-level analysis. Marr’s vision, at least for what he called “Type 1” systems, is largely a top-down one that starts at the level of abstract concepts and considers what representations or algorithms are needed to handle those concepts, and then explores how these particular representations and algorithms might be implemented in hardware. In short, the flow of Marr’s analysis is, for the most part, from high-level abstractions to hardware.

3. Emergence and neural network models of cognition

When neural network models took cognitive science by storm in the mid-1980s, there was a tendency on the part of some (excessively enthusiastic) researchers and philosophers of connectionist cognition (e.g., Churchland, 1995) to switch the focus from the computational level almost entirely to the implementational level. Although most proponents of the connectionist movement were not nearly this extreme, the underpinning of this movement was that higher levels of organization arose from lower levels of organization and that the explanation of the former resided in the latter. The point on which there was almost universal agreement was that explicit rules of high-level behavior (e.g., “To form the past tense of a regular verb, add the suffix -ed to the verb”) were unnecessary in the sense that there was no need to hard-wire them into the system.

But however successful this approach proved to be in many cases, it frequently provided little or no insight into the underlying processes driving cognition, beyond the bland (non-)explanation that “changes in the weights between nodes in the network produced the outcome.” Ultimately, to *understand* cognitive processing, we need to be able to observe the many intermediate levels in which the patterns derived from raw neural firings are consolidated, organized, augmented, compressed, and correlated with other patterns. In other words, we need to have access to these intermediate levels and to their effects on other levels of processing.

The point is that neither a decomposition from the computational level to the implementational level nor a purely emergent approach in the opposite direction tells us that much about how cognition actually works. This leads us to the main point of the present article: cognitive scientists need not only to understand how emergent structures develop—in particular, explicit emergent structures at the cognitive level—but also to understand how they feed back on the sub-structures from which they emerged.

4. Emergence is interesting; interactive emergence is *really* interesting

A stalagmite emerges from the floor of a cave, the product of untold millions of drops of water falling on it, each one leaving behind a tiny deposit of calcium. But once it has emerged, it just stands there, looking beautiful but influencing nothing. This is not, by and large, the kind of emergence that we are interested in. Rather, we are interested in what might be called “interactive emergence”—that is, the study of emergent structures that influence, among other things, the lower levels from which they emerged—and we believe creating computational cognitive models that embody its principles will constitute one of the major challenges of 21st-century cognitive science.

During the course of learning, we acquire, implicitly or explicitly, certain cognitive and subcognitive structures that help us improve the efficiency of future data processing. By “structures” we mean patterns of organization at the levels of neurons, representations, and concepts that make us better adapted to our environment. At the highest level, these structures include stereotypes of all kinds, mathematical constructs, grammatical rules, and associative regularities and, at the lowest level, they are clusters of neurons within which neural firings are facilitated and between which they are inhibited. These structures invariably will have a significant influence on our perception of new data. Once we have acquired, say, a cognitive-level stereotype, such as “New Yorkers are unfriendly,” our perception of the actions of New Yorkers, even otherwise innocent actions that would not even be noticed if we were observing people in Maine, will be colored by this stereotype. And descending to the cortical level, as cortical maps emerge, they, too, can have a radical effect on the perception and organization of new incoming data.

5. The ubiquity of interactive emergence in human cognition

Let us start by considering the image in Fig. 1, an image well known to virtually all cognitive scientists. If you have never seen this image by the photographer R. C. James before, it is very difficult to decipher what it depicts. On the other hand, once you have recognized it as a Dalmatian sniffing the ground near the shadow of a tree, you can no longer *not* see the Dalmatian. In other words, the raw visual input from this image (i.e., the neuronal level) is automatically and correctly chunked every time you encounter it, once you have recognized the Dalmatian (i.e., the cognitive level).

6. Explicit versus implicit emergent structure

The important thing to notice in the above example is the *explicit* nature of the emergent structure formed. We will make a distinction between explicit and implicit emergent structure. What distinguishes explicit from implicit structure is that, in the former there is the emergence of some explicit new structure, something that one was not aware of



Fig. 1. Once you have seen what is represented in this degraded photograph, you can never again *not* see it when you see the image.

before, something that will then have a separate encoding in the system that will be distinct from the structure that gave rise to it. An explicit emergent structure corresponds, roughly, to an explicitly describable (verbalizable) insight about something.

An example from computational cognition that has caused much ink to flow in the last three decades is the question of past-tense formation. Even very young native speakers of English generally form the regular past tense of verbs correctly and do this long before they are taught the explicit rule: “Add -ed to the verb.” The point of Rumelhart and McClelland (1986) was to show that, in fact, there was no need for an explicit rule anywhere in the system to produce regular past-tense rule-following behavior. Regardless of where one stands on this question, something is still left out of the discussion and that is: What are the effects of acquiring the “Add -ed” rule explicitly? What influence does the presence of this explicit rule—which is necessarily encoded in addition to the bottom-up structure that, alone, previously gave rise to the past-tense rule-following behavior—have on future past-tense formation behavior?

7. Interactive emergence with implicitly emergent structures

It is certainly reasonable to say that any model that learns through feedback with its environment is displaying interactive emergence, even if this emergence is not expressed in the form of an explicit (verbalizable) rule.

Edelman (1987) argued for exactly the kind of “linked levels” that we are talking about here. He called it *reentrant connectivity*, which he defined as “Changes in any one level must result in readjustment of all ‘linked’ levels.” In the brain, for example, synaptic changes in response to input will have an effect on future input. So, for example,

Jenkins, Merzenich, Ochs, Allard, and Guic-Robles (1990) looked at how the cortical map is reorganized as a result of learning. In adult mammals there is a separate representation for each finger of the hand in the somatosensory cortex. This separation on the cortical map plays an important role in the capacity to discriminate touch on each individual finger. It was found that if the monkeys were trained to repeatedly touch a wheel with the same finger, the cortical representation for this finger grew bigger. The receptive fields of the neurons in these now expanded cortical zones were also smaller. Having smaller receptive fields is the equivalent of having a finer grained representation for the type of stimulus in question. The bigger cortical map and the smaller receptive field lead to a higher sensitivity for the finger that was repeatedly stimulated. This greater “cortical priority” given to the preferred finger would now change the manner in which stimuli delivered to the fingers is perceived. In short, equivalent stimuli at the fingers would now be perceived with the filter of a priority given to one finger.

Similarly, in artificial networks (e.g., backpropagation networks) that learn to categorize, changes in the weights of the network during learning allow them to correctly classify new input into one of the learned categories, something that they were unable to do before learning. In this way feedback (in this case, from a “teacher” signal that contains the desired response) causes changes in the synaptic weights of the system, which in turn affect how future input to the network is categorized.

It is clear, then, that many current systems, especially connectionist and other neural network systems already implement implicit interactive emergence.

8. Explicit interactive emergence

Let us start with a number of examples, running from the grandiose to the mundane, where the emergence of an explicit high-level concept (a rule, a previously unrecognized pattern, etc.) can have a radical effect on the representational level.

Once Charles Lyell (1830) realized that large-scale geological changes were due to the gradual accumulation of small changes over enormous time spans, he would not have been able to see geological features in the same way as he did before. For example, he would no longer have been able to see rivers in valley bottoms without representing the valleys as having been carved out by the rivers. In short, every geological formation he would have subsequently come across, he would have viewed through his “uniformitarian” lens and would have represented it accordingly.

An example from the life of Helen Keller (1905) is particularly salient. Helen Keller, who became completely blind and profoundly deaf at the age of 19 months, came under the tutelage of Anne Sullivan, who tried to find a way to communicate to her that things in the world had names. The 6-year-old girl was unable to make that connection until one day Sullivan held her hand under a stream of water from a pump, over and over signing the word “water” into the palm of her hand. The little girl finally understood the connection (i.e., emergence at the computational level: “everything in the world around me has a name”) and this radically affected her subsequent representations of things in her

environment. On the day of this breakthrough she spent the rest of the day having Sullivan spell the names of objects and actions in her palm, learning 30 words that day.

But these examples, while they clearly illustrate the phenomenon of explicit interactive emergence, tend to obscure its universality. Let us turn to some “mundane” examples from our everyday experience.

In Western Europe old stone churches are everywhere, thousands of them. At the computational level, a “church” is “a building in which religious ceremonies are conducted.” Now, these European churches were built at different times over a period of more than a thousand years. After having visited a number of them throughout Europe, one gradually becomes aware of a number of differences between the oldest of them (Romanesque churches) and those built later (Gothic churches). The former are dark inside with small windows, have square towers with no steeples and semi-circular vaulting; the latter are much brighter inside with big, stained-glass windows, have flying buttresses on the outside and have vaulting that comes to a point in the center. Once you have become explicitly aware of these features, you can never see a European church in the same way again. Just like the degraded-Dalmatian image, once you have seen and explicitly chunked these regularities, you can never return to your original perceptions of European churches. Thus, as for the Dalmatian in Fig. 1, your perception of European churches is forever changed—you cannot look up at the ceiling of the nave of a church and not explicitly notice the shape of its arches, something that initially you were completely unaware of.

Examples from language also abound and are as close at hand as the previous sentence. Sentences in correct English are not supposed to end in prepositions and educated English speakers, without ever thinking about it, generally do not do so. This is, no doubt, the origin of the explicit rule that one must not end a sentence with a preposition. Of course, this explicit rule, once we become aware of it, sometimes causes cognitive conflicts. The hesitation we feel when writing a sentence like the one that concludes the previous paragraph is due to a conflict between the rule (instantiated somewhere in our neural hardware) and our bottom-up mastery of written English. This conflict between our bottom-up English usage and the formal rule can also be the source of humor, as in Winston Churchill’s famous quip, “Ending a sentence with a preposition is something up with which I will not put.” Here, again, is a clear instance of the influence of an explicit cognitive-level rule influencing the lower levels from which our written language, and the rule itself, emerge.

Sometimes there is the realization that a rule can account for what we have previously been doing in a purely bottom-up manner. One of the authors, a native English speaker, was teaching his boy to read English (he lives in France) and noticed what seemed to him to be a regularity in English pronunciation, one that he had never before been aware of (. . .ending a sentence with a preposition, *again!*). The regularity was this: When, at the end of a word, there is a vowel followed by a consonant followed by an “e,” the first vowel is pronounced as a long vowel. Thereafter, when doing reading exercises with his children, whenever he encountered a word ending in a vowel-consonant-e, this rule popped into consciousness, delaying, however slightly, his pronunciation of irregularities with respect to this newfound rule (e.g., words like “glove,” “gone,” “native,” “none,”

“have,” “give,” etc.). Clearly, the emergence of an explicit, chunked rule affected his pronunciation performance—minimally, to be sure—as he taught his children to read English. One might say, “But the rule was present before he ever became explicitly aware of it!” Yes, but *once this explicit emergent rule had been internalized (i.e., encoded), it subsequently had an effect on his perception of vowel-consonant-e words above and beyond that produced by his implicit bottom-up pronunciation capacity.* Further, only after becoming aware of the rule did he explicitly notice that the words “dove” and “live,” in different contexts, have different pronunciations, or that there was anything “odd” about the pronunciation of “native,” “give,” “gone,” “sausage,” or “chocolate.”

So, does the emergent rule have genuine “causal powers” that were not present before the explicit formulation of the rule? It depends. If the presence of the explicit emergent rule affects how subsequent input is processed, then the answer is yes. If not, then no. Do the explicit equations of General Relativity have any causal powers? It is hard to see how one could answer that they do not. Without them no one could have predicted the deflection of light by the sun or black holes. In a similar manner, possessing the explicit vowel-consonant-e rule causes us to hesitate, however slightly, when pronouncing the words “give” or “gone” and, in that sense, is clearly causal, since that tiny hesitation would not have occurred without the presence of the explicit rule.

9. Computational modeling of explicit interactive emergence

Certain connectionist models, such as the Simple Recurrent Network (SRN, Elman, 1990), can be shown to develop representational clusters at the hidden layer that reflect the organization of the input. But the representational (i.e., hidden-layer) clusters formed in this manner are not active. For an SRN one can argue that implicit structure (i.e., changes in the synaptic weights) emerges, but no explicit “reentrant” chunks of information (Edelman, 1987) are formed and used by the network in subsequent processing. In contrast, the Recursive Auto-Associative Memory (RAAM) model (Pollack, 1990) did chunk input and “reinject” the chunked input into the network to create ever larger chunks of the input data. In a recent model based on the RAAM architecture, French, Addyman, and Mareschal (2011) showed how a recurrent connectionist network exposed to a continuous syllable stream could discover syllable chunks and words in its input. To do this, their model, TRACX, discovers syllable chunks in the input and subsequently uses these chunks as new input to the system. It discovers these chunks by recognizing that it has seen pairs of items together frequently on its input. If the model has seen two items together on input frequently, it concludes that “these two items must be a chunk” and from then on, it replaces these paired items on input with its “chunked” (i.e., hidden-unit) representation of them. This active reuse of the model’s internal representations as input when the input is recognized at chunked is very different from the “mere” emergence of the internal representations in an SRN.

Everyone knows that connectionist system can produce rule-like behavior without recourse to explicit rules (Rumelhart & McClelland, 1986). But this is markedly different than producing not only rule-like behavior but also producing the explicit rule, like the

one with the vowel-consonant-e rule cited above. In other words, neural networks also need to be able to allow *explicit rules to emerge* that will have a direct effect on future processing of input. The emergence and encoding by the system of explicit rules and their effect on the future processing in an unsupervised connectionist network (Kohonen, 1982, 1993) have been studied recently in a model by Cowell and French (2011). Their augmented Kohonen network is designed to be sensitive to patterns in the weights of the network itself. Therefore, when a pattern of weights is spotted indicating that certain features are regularly associated with a particular category (e.g., when “beak” on input and “bird” on output), the network gradually formulates an explicit rule that encodes this frequently observed association (e.g., “if beak, then bird.”). The emergence of this explicit rule in the model can cause conflicts with its purely statistical (bottom-up) categorization abilities based on perceptual features, thereby causing its recognition reaction times to increase (see also Thibaut, Lemaire, & Quadri, 1998; where this type of conflict is studied empirically). Fig. 2 is an example of this kind of incongruity: Our bottom-up perception tells us the animal is a rodent; our top-down rule tells us that it is a bird. (It is, in fact, a bird, but a highly atypical one: a kiwi.)

In short, this model’s incorporation of self-monitoring allows not only for rule-like behavior to emerge in a bottom-up manner à la Rumelhart and McClelland (1986) but also provides a means for an explicit rule associated with this behavior to emerge and be encoded in the system. Thereafter, this emergent rule can interact with the subsequent bottom-up processing of the network. We feel that this work is a small step in the direction of the kind of explicit emergent interaction we are discussing here. However, for the moment, neural network models in which explicit rules emerge and then interact actively with the system’s subsequent processing of input are rare.

There are, of course, a number of non-connectionist computational models whose design is such that chunks are gradually formed from the input data and these chunks



Fig. 2. A kiwi is likely to cause conflict in a categorization task because of the incongruity between its beak, which triggers the rule “if *beak*, then *bird*,” and its visual similarity to a rodent (Cowell & French, 2011).

influence subsequent processing and representation of new input. For example, the “active-symbol” computational models of analogy-making initially developed in D. Hofstadter’s research group in the late 1980s (French, 1995; Hofstadter, 1979; Hofstadter & Fluid Analogies Research Group, 1995; Marshall, 2006; Mitchell, 1993; Mitchell & Hofstadter, 1990) are examples of this type of model in which interactive emergence was a core design principle. In these models there is a constant interaction between representational and computational levels, between bottom-up and top-down pressures. New structures are continually being discovered and influencing subsequent processing of input.

10. Conclusion

In this article we have argued that one of the challenges facing 21st-century computational cognitive science will be to gain a better understanding not only of the emergence of new explicit structures over the course of learning but also of how these new structures affect future processing of the system. The so-called connectionist revolution was intent on showing that what appeared to be rule-driven behavior could emerge in the absence of explicit rules from a distributed system of undifferentiated artificial neurons. This segued in many cases into a de facto denial of the existence of explicit rules, which is certainly wrong. Explicit rules, from group stereotypes to grammar heuristics, do emerge from our interaction with the world. These rules then become an integral part of our cognitive system, separate from the purely bottom-up processes that gave rise to them. The first battle of the connectionist revolution—showing that what appears to be explicit rule-following behavior can be produced in a bottom-up manner—has been won. Now, we need to better understand how explicit emergent top-down rules can interact with a system driven by bottom-up processes.

Acknowledgments

This work was supported in part by a joint grant (ORA-10-056) from the French National Research Agency and the ESRC of the United Kingdom.

References

- Alexander, S. (1920). *Space, time, and deity*. London: Macmillan.
- Bedau, M. A. (1997). Weak emergence. In J. Tomberlin, (Ed.), *Philosophical perspectives: Mind, causation, and world, Vol. 11* (pp. 375–399). Malden, MA: Blackwell.
- Broad, C. D. (1925). *The mind and its place in nature*. London: Routledge.
- Churchland, P. (1995). *The engine of reason, the seat of the soul: A philosophical journey into the brain*. Cambridge, MA: MIT Press.
- Cowell, R. A., & French, R. M. (2011). Noise and the emergence of rules in category learning: A connectionist model. *Transactions on Autonomous Mental Development*, 3(3), 194–206.

- Edelman, G. (1987). *Neural Darwinism* (p. 173). New York: Basic Books.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.
- French, R. M. (1995). *The subtlety of sameness*. Cambridge, MA: MIT Press.
- French, R. M., Addyman, C., & Mareschal, D. (2011). TRACX: A recognition-based connectionist framework for sequence segmentation and chunk extraction. *Psychological Review*, 118(4), 614–636.
- Glass, A., Holyoak, K., & Santa, J. (1979). *Cognition*. Reading, MA: Addison-Wesley.
- Goldstein, J. (1999). Emergence as a construct: History and issues. *Emergence: Complexity and Organization*, 1(1), 49–72.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach*. New York: Basic Books.
- Hofstadter, D. R. & Fluid Analogies Research Group. (1995). *Fluid concepts and creative analogies*. New York: Basic Books.
- Jenkins, W. M., Merzenich, M. M., Ochs, M. T., Allard, T., & Guic-Robles, E. (1990). Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation. *Journal of Neurophysiology*, 63, 82–104.
- Keller, H. (1905). Part III from the letters and reports of Anne Mansfield Sullivan. In J. A. Macy (Ed.) *The story of my life* (Letter of April 5, 1887, p. 316). New York: Doubleday, Page & Company.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59–69.
- Kohonen, T. (1993). Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, 6, 895–905.
- Lyell, C. (1830). *Principles of geology*. London: John Murray.
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman & Co.
- Marshall, J. (2006). A self-watching model of analogy-making and perception. *Journal of Experimental and Theoretical Artificial Intelligence*, 18(3), 267–307.
- McClamrock, R. (1990). Marr's three levels: A re-evaluation. *Minds and Machines*, 1, 185–196.
- Mill, J. S. (1858). *A system of logic*. New York: Harper & Brothers.
- Mitchell, M. (1993). *Analogy-making as perception: a computer model*. Cambridge, MA: MIT Press.
- Mitchell, M., & Hofstadter, D. R. (1990). The emergence of understanding in a computer model of concepts and analogy-making. *Physica D*, 42, 322–334.
- Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence*, 46(1), 77–105.
- Pylyshyn, Z. (1984). *Computation and cognition*. Cambridge, MA: MIT Press.
- Rumelhart, D. E., & McClelland, J. L. (1986) On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & The PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*. Volume 2: Psychological and biological models (pp. 216–271). Cambridge, MA: Bradford Books/MIT Press.
- Thibaut, J. P., Lemaire, F., & Quadri, J. (1998). Categorization under the influence. In M. A. Gernsbacher & Sharon J. Derby (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1055–1060). Mahwah, NJ: LEA.