Instantiated rules and abstract analogy: Not a continuum of similarity

Lee R. Brooks and Samuel D. Hannah

Department of Psychology, McMaster University, Hamilton, Ontario, L8S 4K1, Canada. **brookslr@mcmaster.ca** **hannahsd@mcmaster.ca**

Abstract: We agree that treating rules and similarity as dichotomous opposites is unproductive. However, describing all categorization operations as a continuum of varied similarity process obscures a multidimensional contrast. We describe two processes, instantiated rules and abstract analogy, both of which have aspects of rules and similarity, and question whether they can be compared informatively as points on a continuum.

We agree with Pothos that treating rules and similarity as dichotomous opposites is inappropriate and we strongly endorse his review of the literature supporting a more subtle treatment. However, we are concerned that describing all categorization operations as variations of a similarity process obscures what is essentially a multidimensional comparison. To make our case, we first describe what we think is a common process in categorization.

We have argued (Brooks & Hannah 2000; submitted; Hannah & Brooks, submitted a; submitted b) that identification rules commonly act as a control of attention and learning, not as a complete decision procedure in themselves. Such rules name features whose manifestations in known items are to be learned, and then classify new items by heavily weighting any of those named features whose manifestations are similar to any of those previously experienced. For example, if a medical student is told that the skin disease *lichen planus* sometimes presents with polygonal-shaped papules (medium-sized solid bumps), the student can at least consider that diagnosis when anything that can be called a polygonal bump appears. If a new instance of a polygonal papule were a near twin of a previously seen papule, the student would be confident that the feature bore on the correct diagnosis. However, if the papule were perfectly regular, 10 cm across and bright blue, then that student should rightly be very cautious about deciding that it had anything to do with *lichen planus.* One component of expertise is probably knowledge of the variety of ways a feature can look and still be relevant to the category.

This proposal that rules are applied by attending to the particular appearance of the named features also helps us to understand the fact that most "rules" given either by medical experts or by undergraduates describing everyday concepts are actually only lists of features lacking any specific decision procedure (Brooks & Hannah 2000). A bird, for example, could be described as an animal that flies, sings, and has feathers, without giving any indication of how many of these features it must have or how they should be weighted. However, the terms in the rule have led to learning examples of how bird-like flying, bird-like feathers, and bird-like singing appear. Such learning would allow the learner to immediately reject the word "bird" for an opera singer with a feather boa flying on a plane, even though she has all of the features mentioned in the "rule." That is, no bird had been seen flying in a manner even remotely resembling a jet or singing even a remote approximation to Verdi. The fact that most everyday and expert rules do not have a fixed decision procedure ("best 2 out of 3 features" or "put heaviest weights on these cardinal features") suggests that people normally are using something other than just the number or weighting of features. The additional information normally used, we suggest, is the "goodness" of the feature manifestations, the extent to which they match previously experienced manifestations – information that is not captured in a frequency-based multiple regression. The rule, then, is represented in a much richer fashion than just as a list of terms that rely solely on the resources of the general language for their use. Without some correspondence to general language meanings, the rule would not be useful for initial instruction. However, such general language groundings are insufficient to describe the performance of people with even a modicum of experience. Clearly, the terms in the rule must also have some concept-specific grounding, often to manifestations that are extremely diverse across different members of the same category. The "instantiated rule" process just described applies most obviously to learning with explicit instruction, to double checking or justifying an initial categorization, and to framing an explicit policy.

This idea that the terms in a rule are represented in both an informational (sparse, general language) and an instantiated (perceptual, detailed) form has allowed us to produce and control categorical biasing (Hannah & Brooks, submitted a) and to control the weighting given to familiar manifestations of features (Hannah & Brooks, submitted b). Without representing features in both forms, we were unable to produce important phenomena in concept learning and utilization.

Consistent with Pothos's proposal, these "instantiated rules" are difficult to fit into a strong dichotomy between rule and similarity processes. The matching of a presented feature to the acceptable manifestations of the features for a given category clearly has characteristics that Pothos attributes to similarity. However, the verbal list of features in explicit versions of such instantiated rules have characteristics of rules, including a concentration on small portions of the stimulus and an at least partial failure to track the covariation structure of the domain.

However, to say that this instantiated rule process is part of a continuum of similarity processes, potentially modeled by GCM, seems to be a misleading simplification. Many applications of abstract analogy (e.g., Goldstone et al. 1991; Brooks & Vokey 1991) also have an "in-between" position on this proposed continuum in that the features are selective and abstract (Rules) but the comparison process is neither certain nor limited to a small number of terms (Similarity). However, the formation and generalization of abstract analogies is very different from the instantiated rule process of accumulating a rich store of perceptually specific representations around which to generalize. Presumably, a version of GCM (or some other similarity-based model) could model the categorization of new items, but only after specifying a space based on the appropriate features; that is, only after much of the work of psychological interest had already been accomplished. We further suspect that there are critical differences between the authoritative status of the medical rule and the less specific and changeable structure of a discovered abstract analogy in controlling the decision process. While sympathetic to the value of rejecting a rule/similarity dichotomy, the multiple differences between instantiated rules and abstract analogies are not informatively captured as two middling points on a continuum.

Rules, similarity, and the information-processing blind alley

Francisco Calvo Garzón

Department of Philosophy, University of Murcia, Campus de Espinardo, Murcia 30100, Spain. **fjcalvo@um.es**
http://www.um.es/~logica/paco_calvo/paco_calvo.htm

Abstract: Pothos's revision of rules and similarity in the area of language illustrates the impression that the classicist/connectionist debate is in a blind alley. Under his *continuum* proposal, both hypotheses fall neatly within the information-processing paradigm. In my view, the paradigm shift that dynamic systems theory represents (Spencer & Thelen 2003) should be submitted to critical scrutiny. Specific formalizations of the Rules versus Similarity distinction may not lead to a form of unification under Generalized Context Models or connectionist networks.

Pothos's revision of rules and similarity in the area of language (sect. 7) is perhaps the latest episode in a series of exchanges (e.g., Marcus & Berent 2003; Seidenberg et al. 2003) that serves to illustrate the impression of a growing minority: that the classicist/connectionist debate in cognitive science is in a blind alley. The friends of classical orthodoxy (Marcus et al. 1999) continue to

https://doi.org/10.1017/S0140525X05260016 Published online by Cambridge University Press

search for cognitive abilities that, defying statistical explanation under the *poverty of the stimulus* lens, embarrass their foes. Rule-following skeptics (Calvo & Colunga 2003) rejoin by finding ecological data that (1) can be exploited statistically and (2) allow connectionist networks to remain computationally adequate. Put bluntly, the connectionist's overall strategy is to show that stimuli are not so poor after all! Although things are never black and white, the debate has moved along these lines since the re-emergence of connectionism in the mid 1980s. The debate has been fruitful insofar as contributions have filled in empirical gaps at algorithmic levels of description.

It is noteworthy, however, that under Pothos's *continuum* proposal, both hypotheses, the classical and the connectionist, fall neatly within the information-processing paradigm that has shaped the discipline over the last century. The architecture of cognition has been questioned, but assumptions about its computational underpinnings have remained unchallenged. The past-tense debate (Pinker & Ullman 2002; Ramscar 2002), the systematicity debate (Fodor & Pylyshyn 1988; Hadley 1999), and, in the last five years, the algebra-versus-statistics debate (Calvo & Colunga 2003; Marcus et al. 1999), have benefited partially from the classical-connectionist, *within*-paradigm "battle to win souls." This is, however, a "cognitive decathlon" (expression borrowed from Anderson & Lebiere 2003), in which we might never be able to declare a winner! Perhaps we are stuck in an endless dialectic of positing challenges to connectionism and then trying to account for them statistically, forever and ever. In view of this scenario, we may need to consider turning to questions concerning the role that potential contenders, such as Dynamic Systems Theory (DST) (Thelen & Smith 1994) may play in the future.

Unlike information-processing-based frameworks, DST tries to model and explain the behavior of concrete systems by identifying them with sets of variables that change continually over time. A dynamical system, in this way, can be analyzed in terms of the differential equations that contain the quantitative variables whose interdependencies describe the laws that govern the behavior of the system. DST has proved extremely useful in the physical sciences. Limb movement is a classical example in the literature. Kelso (1995) studied the wagging of index fingers, and a number of properties were successfully described and predicted dynamically. The phenomenon could be explained as a property of a nonlinear dynamical system that achieves self-organization around certain points of instability. Thelen et al. (2001), on the other hand, go exhaustively over the literature on the well-known, but still highly controversial, "A-not-B error," and offer a non-classical explanation of motor control and development in that context. In their view, the A-not-B error can be perfectly explained, with no need to invoke information-processing concepts and operations. Specifically, the dynamical evolution of the coupling of perception, movement, and memory can explain by itself the A-not-B error. Crucially, cognitive activity cannot be accounted for without considering the perceptual and motor apparatus that facilitates in the first place the agent's dealing with the external world.

It must be emphasized that DST aims not merely at cashing out the axioms of the information-processing paradigm in trendy mathematical terms, but, rather, at articulating a brand new way to understand cognition. In this way, Thelen et al. (2001) speculate as to how higher-level phenomena may be dynamically modeled. A model of mental activity must respect the same principles of nonlinearity, time-dependence, and continuity that are generally invoked in explanations of bodily interactions and neural activity (Freeman 2000). The working hypothesis is that the same mathematical tool kit of differential equations can be put to the use of describing and explaining cognitive activity in general. By contrast, an information-processing agent counts as a computational system insofar as its state-transitions can be accounted for in terms of manipulations on representations.

In my view, the paradigm shift that DST represents should be submitted to critical scrutiny in this context. Its continuous and situated approach may lead us back to the main road by eschewing, rather than revising, the (computationalist) function-approximator approach of connectionism. The question of whether "complex cognitive functions depend on a mixture of statistical and algebraic (rule) mechanisms" (Marcus & Berent 2003), may simply vanish once the continuous interplay that dynamic system theory calls for between brain, body, and environment is modeled in detail. The information-processing paradigm posits sets of internal mechanisms that serve the purpose of information manipulation and storage. Cognitive activity is thus marked by the processing of representational states. Put bluntly, in its classical form, cognition amounts to the manipulation of symbols according to sets of explicit (algebraic) rules. In its connectionist form, cognition would amount to the manipulation of sub-symbols according to sets of implicit (statistical) rules. In my view, no further progress can be made unless the very issue of whether cognition must necessarily be accounted for in information-processing terms is addressed by the scientific community. Dynamic systems theory furnishes us with an appropriate theoretical and mathematical tool kit to make (experimental) predictions in cognitive psychology. According to these predictions, specific formalizations of the Rules versus Similarity distinction may *not* lead to a form of unification under Generalized Context Models (GCM; Nosofsky 1991) or connectionist networks. Traditionally, dynamicism is seen as taking sides with Gibsonian approaches, whereas connectionism pertains to the domain of information processing. Rules and Similarity *in the context of* information processing, I contend, may simply be the wrong framework.

ACKNOWLEDGMENTS

The author was supported by a Ramón y Cajal Research Contract from the Spanish Ministry of Science and Technology during the preparation of this commentary. I thank Esther Thelen for helpful discussions.

Epistemological requirements for a cognitive psychology of real people

John Campion

1 Bembrook Cottage, Woodmans Green, Linch, Liphook, Hants GU30 7NE, United Kingdom. **TJCampion@aol.com**

Abstract: Pothos's analysis is difficult to relate to real human mental processes. He tackles four quite different areas of psychology and adduces evidence from a large number of paradigms. Yet despite this very large scope, he employs a single, simplistic descriptive framework. An epistemological analysis, supported by illustrations from real world decision-making, shows that this steers us away from, rather than towards, an understanding of real human cognitive processes.

I shall focus on the issue of psychological reality as it relates to Pothos's article and illustrate my argument with reference to human decision-making.

Pothos declares that the large array of psychological phenomena that he wishes to address (which runs from human decision-making through to discrimination learning in pigeons) can all be construed as a categorisation process. Further, he states that "nothing is said about how this process of categorization takes place; rather, we are interested in providing a framework for characterizing the categorization as a rules process or an overall similarity one" (sect. 2, para. 4). He also states "there is no conception of what it would mean for Rules to exist, other than as emergent properties of assemblies of neurons . . . [this] is analogous to . . . temperature or pressure, which feel real but actually exist only as emergent properties of assemblies of molecules" (sect. 7.4, last para.).

These views present us with something of an epistemological conundrum in that Pothos seems to adopt a position of naïve realism with regard to molecules and assemblies of neurons (they

https://doi.org/10.1017/S0140525X05260016 Published online by Cambridge University Press