possible to transform similarity processes into rule processes. In contrast to this implication, rule and similarity processes are computationally indistinguishable; it needs some other variable, such as the strategy used (Smith et al. 1998), in order to distinguish rule processes and similarity processes. In order to bolster this notion, let us look at an example from artificial grammar learning (Reber 1967). Participants in an artificial grammar learning experiment may use different processing strategies to classify stimuli: They may, first, explicitly match test stimuli to the rule, taking each property covered by the rule into account. This kind of classification involves systematic processing, needs resources from short-term memory, and includes rule following (see Hahn & Chater 1998, for the distinction between rule-following and rule-describable processes). Second, participants may retrieve items learned during a test and assess their similarity to a given test item in order to classify a stimulus. This often involves less analytical processing than rule application and at best yields rule-describable classifications. In some situations, people may simply assess how easily they were able to process the test stimulus and base their classification on such processing experiences (e.g., Whittlesea & Dorken 1993; cf. Shanks et al. 2002). It is this second kind of artificial grammar learning that got a lot of attention in the last few decades.

Research has found that people use characteristics such as associative strength to classify test stimuli (Knowlton & Squire 1994; 1996; Meulemans & Van der Linden 1997). According to Pothos's distinction between Rule and Similarity processes, this would be a typical similarity process. For each test stimulus, associative strength can be calculated as averaged frequency of bigrams and trigrams in the training stimuli (see Meulemans & Van der Linden 1997, p. 1010). Mean associative chunk strength in their Experiment 1A was about 3.6. Let us assume that participants in an artificial grammar learning experiment without rule knowledge classify stimuli according to associative strength, and we find out that they classify stimuli as being grammatical when associative chunk strength is above a value of 3.6, and as ungrammatical if it is below this value. If we knew that, we could now teach participants the rule and instruct them to apply it for classification of test items: For each test string, calculate first its associative chunk strength and then classify it as grammatical if the result is greater than 3.6, and classify it as ungrammatical if it is smaller than or equal to 3.6. At least intuitively, most researchers probably would classify the latter task as one that involves rule processes. In this case, both rule and similarity have the same number of properties; the difference is that participants with rule knowledge show rule-following classification performance whereas participants without rule knowledge show only rule-describable classification performance. Please note that this rule does not result in 100% classification accuracy; if only 50% of the strings with average chunk strength above 3.6 obey the rule, accuracy will be at 50%. The representations and the computational processes needed in order to classify stimuli are identical in both similarity and rule processes, and hence, the number of properties is identical. Of course, some assumptions are simplified: Participants in an artificial grammar learning experiment may classify stimuli as grammatical if chunk strength is higher than 3.6 on average; some participants may have a lower threshold, others a higher one. Rules, on the other hand, are fixed (see Hahn & Chater 1998). It may even turn out that different participants use different properties in classification according to similarity: Some may look at trigrams in anchor positions, others are assessing global similarity; a rule, on the other hand, is identical for all participants. However, if it were known how individual participants classified stimuli, we would be able to transform these different similarity processes into rule processes with identical computational characteristics.

Taken together, each similarity process can be transferred into a rule process if we knew what the similarity process is based on. However, the mere fact that rule processes can be transferred to similarity processes implies that similarity processes and rule processes cannot a priori be distinguished on the basis of representational characteristics, such as number of properties.

We can distinguish the two kinds of processes in terms of processing modes, however. Whereas people seem to assess similarity processes easily, systematic application of rules soon exceeds working memory capacity. Although there is only one criterion to apply the rule, the algorithm that needs to be applied in our example in order to calculate associative chunk strength is very complicated. It may be easy for the cognitive system to apply this algorithm automatically and in a rule-describable manner, as in similarity processes. In contrast, it is definitely difficult to apply it deliberately and in a rule-following manner, as it happens in rule processes. As a consequence, people can process only few properties if applying rules, but they can process many properties if assessing similarity. This may result in the fact that rule processes normally involve a higher number of properties than do similarity processes.

I have dealt with artificial grammar learning only, but the same logic seems valid for reasoning, categorization, and language. In all cases, similarity processes can be transformed into rule processes, and in all cases, the number of properties would be the same. Again, applying the rule is more effortful than use of similarity processes, although a simple rule may be learned to automaticity, resulting in effortless processing of rule application. In conclusion, the real distinction between rule processes and similarity processes lies in the mode people use to process this information. Computational differences observed between rule processes and similarity processes, such as many properties in similarity processes and few properties in rule processes, are accidental consequences of the different processing modes used in rule and in similarity processes.

Rules and similarity processes in artificial grammar and natural second language learning: What is the "default"?

Peter Robinson

Department of English, Aoyama Gakuin University, 4-4-25 Shibuya, Shibuya-ku, Tokyo 150-8366, Japan. **peterr@cl.aoyama.ac.jp**
http://www.cl.aoyama.ac.jp/~peterr/

Abstract: Are rules processes or similarity processes the default for acquisition of grammatical knowledge during natural second language acquisition? Whereas Pothos argues similarity processes are the default in the many areas he reviews, including artificial grammar learning and first language development, I suggest, citing evidence, that in second language acquisition of grammatical morphology "rules processes" may be the default.

Pothos argues that similarity processes in learning and categorization are the "default" because rules processes require many properties of an object to be "suppressed" in coming to a decision about classification. However, Pothos speculates in his conclusion, "even if Similarity is the default, it is possible that some objects might be processed spontaneously in terms of a Rule" (sect. 9), in cases, for example, where a diverse range of objects are grouped together in a category, making similarity judgments difficult to compute. These issues are important to understanding what guides decision making, not just in artificial grammar (AG) learning, and first language (L1) development, but also in natural second language acquisition (SLA). What is the "default" process for learning and categorizing natural second language (L2) stimuli as grammatical or ungrammatical? One possibility, in line with Pothos's speculation above, is that the further the grammatical distance (measured in ways briefly described below) the L1 is from the L2, the more likely a rules process is spontaneously used to (correctly or incorrectly) judge an L2 sentence as grammatical, and that frequency and similarity information available from exposure to the L2 input is less influential, or ignored. Rules processes would therefore be the default basis for classification in these cases, not similarity processes.

https://doi.org/10.1017/S0140525X05420014 Published online by Cambridge University Press

In his review, Pothos summarizes claims about rule versus similarity influences on AG learning. In reporting the results of Vokey and Brooks (1992), Pothos claims, as Vokey and Brooks do, that the judgments were based on "abstract analogy" (sect. 4.3; cf. Brooks & Vokey 1991) and that suppression and rules processes, which Pothos suggests Reber might invoke (cf. Reber 1989), were not involved (sects. 4.3 and 4.4). However, it is an important question how related AG learning and natural L2 learning are, and whether similarity processes, and rules processes (involving suppression) operate in the same way, or differently, on both sets of stimuli. Studies of this issue are rare, and deserving of further empirical research, given the claims Pothos makes in his review. One such study (Robinson 2002; 2005) replicated findings by Knowlton and Squire (1996; to which Pothos refers), with experienced L2 learners. Robinson (2005) found that the high chunk strength, and so similarity of training to AG transfer set items, influenced these learners to incorrectly accept ungrammatical (UG) items but did not influence judgments of grammatical (G) items. In a separate experiment the same learners also learned a new natural L2, Samoan, under incidental training conditions, in which they were required to process sentences of Samoan for meaning, using a rote-learned Samoan vocabulary, with no grammatical instruction about word order or morphology in Samoan. As with the AG stimuli, in the transfer test grammaticality and chunk strength of Samoan test items were controlled for. However, for natural L2 learning, chunk strength, and similarity of training to transfer set items exerted a negative effect not only on correct rejection of UG items (as in the AG experiment), but also on correct acceptance of G items. High chunk strength influenced learners to wrongly accept UG sentences, and wrongly reject G sentences. Why should this be so?

Pothos suggests an answer in his concluding section, "Future directions." Novel UG sentences in that Samoan transfer set "lacked" the necessary morphology present in the input during training (making them ungrammatical), but novel G items included it. The morphology, importantly, in the G items was very different from (distant from) morphology in the subjects' L1 (Japanese), including ergative markers (which Japanese does not have), locative markers preceding, not following, noun complements (in contrast to Japanese head direction), and noun-incorporation and affixation (which Japanese does not have). None of these features characterizes Japanese. Therefore, in the case of G items, the diverse range of morphology involved in categorizing Samoan sentences correctly as grammatical may well have involved spontaneous "rules processes" (as Pothos, sect. 9, speculates) guided by L1 (Japanese) knowledge. Since L1 knowledge was disjunct with the grammatical morphology in G items; this led learners to incorrectly reject them, despite their similarity to the training set. In contrast, the UG items lacking this morphology were less diverse as a category, and high chunk strength therefore influenced learners to incorrectly accept them on the basis of their similarity to the training set.

In summary, although it may be that a similarity process is the "default" basis for classification in many of the areas Pothos describes, there is some evidence that in learning natural L2 morphology, rules processes may be the default. This evidence is not inconsistent with the claim that some highly frequent, L1-shared cues to form-meaning relationships in the input to L2 learning, such as word order and agency in English, are more susceptible to learning from mere exposure, and similarity-based "strengthening" of such cues, than others, such as the grammatical marking of inflections (e.g., MacWhinney 2001). Neither is it inconsistent with some claims concerning the effects of L2 instruction, that is, that certain aspects of unfamiliar inflectional L2 morphology will benefit from "explicit" instruction, since they activate rules processes during learning (e.g., Robinson 1996), in contrast to inflectional morphology familiar from the L1, which can be learned by similarity processes, following exposure to L2 input alone.

Such issues will be important to address in pursuing further implications of Pothos's proposal for the relationship of AG to natural language learning and classification, and for the issue of whether similarity or rules processes are the default in natural L2 acquisition. With these issues in mind I suggest Pothos should consider adding SLA (related to but different from "learning" and "language," by involving the study of both) to the areas of relevant enquiry into the influence of rules versus similarity in cognitive psychology. Explanations of SLA are characterized by all five types of argument employed in the rules versus similarity debate, summarized in Table 1 (sect. 4.3) – that is, classification, dissociation, and suppression as I have indicated, and also introspective testimony to L2 rule use, or lack thereof; and evidence of differential performance, as well as a priori arguments for why L1 knowledge should exert an influence on classifications of grammaticality in the L2. This makes SLA an important, non-overlapping complement to the other areas of inquiry listed in Table 2 (sect. 8).

Avoiding foolish consistency

Steven Sloman

Department of Cognitive and Linguistic Sciences, Brown University, Providence, RI 02912. **steven_sloman@brown.edu**
http://www.cog.brown.edu/~sloman

Abstract: In most cases, rule-governed relations and similarity relations can indeed be distinguished by the number of relevant features they require. This criterion is not sufficient, however, to explain other properties of the relations that have a more dichotomous character. I focus on the differential drive for consistency by inferential processes that draw on the two types of relations.

In his provocative article, Pothos correctly points out that there must be some coherence among representations to have an effective information processing system. He admirably places this restriction on his own formulation: Whether we call his representational elements rules or similarity or something in between, they must be consistent with one another (i.e., they must not be mutually contradictory with respect to one or another logical calculus). Consistency would seem to be a minimal requirement of a successful system of reasoning.

The problem is that, where people are concerned, contradictions abound. We believe that Linda seems more like a feminist bankteller than a bankteller, and yet we agree that she is more likely to be a bankteller. We are convinced that we should select the "not Q" card in the Wason task (Wason 1966), yet we continue to think the Q card provides more information (and maybe it does, as Oaksford & Chater [1994] point out). We continue to smoke even though we know what a bad idea it is. Divergence in the beliefs held simultaneously by a single person would seem to provide some reason to question the plausibility of a single system in which all reasoning takes place within the same representational medium.

Of course, in principle these contradictions could reflect some general property of a single representational system. Maybe we just don't bother to check for all contradictions, or we live with them, or we just make mistakes. However, the contradictions we observe do not seem to have this homogeneous character. If our more general principles conflict, we try to rectify them. We may not want to believe that smoking is bad for our health, but rejecting such a claim is inconsistent with everything we believe about the integrity of science and our own observations. We are forced to conclude that we should select the P and not the Q card on the basis of a logical argument, even though life would be easier if the conclusion proved our original intuitions correct. Deliberation about which rules to apply at what point in time is what debate often consists of, debate with others as well as with ourselves, because people share a propensity to make our rules consistent. In some cases, it might be rational to keep inconsistent possibilities in mind. For example, one may be considering two inconsistent

https://doi.org/10.1017/S0140525X05420014 Published online by Cambridge University Press