A TWO-PROCESS ACCOUNT OF ANALOGICAL CATEGORY LEARNING†

Fintan J. Costello (fintan.costello@ucd.ie)
Cesare Bianchi (cesare.bianchi@ucd.ie)
School of Computer Science and Informatics,
University College Dublin,
Belfield, Dublin 4, Ireland.

ABSTRACT

Structure-mapping theory suggests that a mechanism of explicit comparison is fundamental to the process of analogical learning and development. In this paper we argue that analogical learning does not occur only through explicit comparison; instead, we show that an mechanism based on the priming and activation of the substructures which make up representations can also explain analogical learning, with no comparison taking place. We show that these two mechanisms (explicit comparison and substructure priming) have different computational characteristics and apply at different stages of the learning process. We therefore propose a two-process account of analogical learning, in which substructure priming dominates in early and mid-stage learning but structure-mapping dominates in late-stage learning. We also link these two mechanisms for analogy to current two-process models of learning in general, with the priming mechanism for analogical learning being related to implicit associative learning, while structural alignment is linked to learning via explicit hypothesis testing.

Structure-mapping theory suggests that the process of comparison plays a fundamental role in the learning and development of category rules (rules which identify category members and distinguish members of different categories). Structure-mapping theorists give two separate roles for the process of comparison during learning. First, they propose that as a learning mechanism, the process of comparison between representations facilitates the identification of structural commonalities and the abstraction of category rules. Second, they propose that comparison facilitates the application of abstract knowledge (general category rules) to new instances (see e.g. Doumas, Hummel & Sandhofer, 2008; Fisher & Sloutsky, 2005; Gentner, 2003; Gentner & Medina, 1998; Gentner & Namy, 1999; Namy & Gentner, 2002; Namy, Smith, & Gershkoff-Stowe, 1997, on the structural comparison process in learning and development).

In this paper we argue that the identification of structural commonalities during learning does not occur only through the comparison and structure-mapping of different representations. Instead, we show that a simple account, based on the priming and activation of the substructures which make up representations, can also explain the production of structural commonalities during learning, with no comparison taking place. In this priming account, relational substructures which a learner has successfully used to understand and form rules for one domain are likely to be re-used by the learner to understand and form rules for a second domain. This re-use of relational substructures will produce structural commonalities between the learner's representations of both domains and therefore cause the learner to see an analogical relationship between both domains.

The idea that relational priming plays a role in analogy is not new (see, e,g. Leech,

† This research was supported by a grant under the FP6 NEST Programme of the European Commission (ANALOGY: Humans the Analogy-Making Species: STREP Contr. No 029088).

Mareschal, &. Cooper, 2008). The main novelty in this paper is an analysis of the relationship between structure-mapping comparison and substructure priming as mechanisms for analogical learning. We show that these mechanisms have different computational characteristics and therefore apply in different phases of learning. These two mechanisms for analogical learning also have suggestive parallels in current two-process accounts of category learning in general (see e.g. Maddox & Ing, 2005; Zeithamova, & Maddox, 2005).

The organization of our paper is as follows. In the first section of the paper we present two computer models of analogical learning, one where analogical learning is based on structure-mapping comparison (Kuehne, Gentner, Forbus & Quinn, 2000; Kuehne, Gentner, Forbus, 2000) and the second where analogical learning occurs via a form of substructure priming with no comparison process (Bianchi & Costello, 2009). Both models have been applied to experimental learning data, providing good levels of fit.

In the second section we assess the performance of these two approaches at different stages of learning. We show that these two mechanisms are in some ways complementary, with comparison not performing effectively early in learning (when there are many alternative category rules to be considered and compared) but providing superior performance later in learning (when it allows the detailed comparison of a small number of remaining category rules). Conversely, substructure priming has its strongest influence early in learning (when it provides an efficient way of distinguishing between the many alternative category rules to be considered at that stage), while priming does not perform as effectively later in learning (when there are only a small number of candidate rules available, all equally primed).

In the third section we place our two-process model of analogical learning in a broader context by describing its relationship to the current two-process view of general category learning. This view sees category learning as involving both an implicit, automatic, unconscious, associative learning process and an explicit, voluntary, conscious, hypothesis-testing learning process. We see the priming mechanism for analogical learning as being related to implicit associative learning, while structural alignment and comparison is related to explicit hypothesis testing.

.

PROGRESSIVE ALIGNMENT IN SEQL

SEQL (Kuehne, Gentner, Forbus & Quinn, 2000; Kuehne, Gentner, Forbus, 2000) is a model of category learning based on comparison and progressive structural alignment. During learning of a single category this model maintains a set of *generalizations* (candidate rule structures which may define the category or some subset of that category) and a set of *exemplars* (instances of the category which are not captured by any current generalization) Each generalization is seen as assimilating a set of previously seen exemplars (those exemplars who match the structure of that generalization). As each new exemplar E arrives, it is first compared with each existing generalization using structural alignment (implemented via SME, the Structure-Mapping Engine). Generalisations are sorted according to the number of previously seen exemplars they assimilate, so that E is first compared with the most successful previous generalization (the one that has assimilated most exemplars), then the next most successful generalization, and so on. The first generalization G_i in this sorted list whose structural similarity to E exceeds a preset match threshold, T, is selected as the correct generalization for E, and E is assimilated into G_i by replacing G_i with the structural overlap between G_i and E. If no generalization is sufficiently similar, then E is compared with each stored exemplar E_i. If one of those matches is over threshold, then their overlap becomes a new generalization, and E_i is removed from the stored exemplars. Otherwise, E is added to the set of stored exemplars.

The threshold T determines how conservative this algorithm will be: If T = 1.0, then no abstraction will occur, since only perfect matches would be grouped. If T = 0.0, then any two descriptions could match, leading quickly to an empty description as the concept representation. To ensure that exemplars are not assimilated to vacuous generalizations, Kuehne, Gentner, Forbus & Quinn (2000) suggest that this algorithm requires a high value for T: in their studies they find that values for T in the range 0.85 to 0.98 (roughly, 85% to 98% overlap between structures being compared) give the best fit to human learning performance.

Learning takes place in this model whenever a new exemplar has a high enough analogical overlap with an already successful generalization. When this occurs the generalization is updated by retaining only the shared structure that forms the alignment. Nonoverlapping aspects of a generalisation are thus "worn away" with each new assimilated exemplar, gradually producing a generalized rule for category membership.

EMPIRICAL SUPPORT FOR LEARNING BY PROGRESSIVE ALIGNMENT

Empirical support for the progressive alignment account of learning comes primarily from the fit of the SEQL model to experimental data from studies by Marcus, Vijayan, Rao & Vishton (1999). These studies show that infants as young as seven months can process simple language-like stimuli and build generalizations sufficient to distinguish familiar from unfamiliar patterns in novel test stimuli. In Marcus et al's study, the stimuli were simple ‘sentences’, each consisting of three nonsense consonant-vowel ‘words’ (e.g ., ‘ba’, ‘go’, ‘ka’). All habituation stimuli had a shared grammar, either ABA or ABB. In ABA-type stimuli the first and the third word are the same: e .g, ‘pa-ti-pa’. In ABB-type stimuli the second and the third word are identical: e.g ., ‘le-di-di’. The infants were habituated on 16 such sentences, with three repetitions for each sentence. The infants were then tested on a different set of sentences that consisted of entirely new words. Half of the test stimuli followed the same grammar as in the habituation phase; the other half followed the non-trained grammar. Marcus et al. found that the infants dishabituated significantly more often to sentences in the non-trained pattern than to sentences in the trained pattern. Based on these findings Marcus et al. proposed that infants had learned abstract algebraic rules. Kuehne, Gentner, Forbus & Quinn (2000) applied the SEQL model to this data and found that, unlike other models (particularly connectionist models) the SEQL model of category learning learned the grammar stimuli within the span presented to the infants and did not require supervision to learn successfully.

ANALOGICAL LEARNING BY PRIMING

As an alternative to progressive alignment, we propose a mechanism where the identification of structural commonalities during learning does not require comparison and alignment. Instead primed activation of the relational substructures which make up representations can explain the production of structural commonalities during learning, with no comparison taking place. In this type of account, relational substructures which a learner has successfully used to understand and form rules for one domain are likely to be re-used by the learner to understand and form rules for a second domain. This re-use of relational substructures can produce structural commonalities between both domains and therefore cause the learner to see an analogical relationship between both domains.

This type of priming approach to analogy could be instantiated in many different ways. We’ve recently developed a computational model of analogical learning (Bianchi & Costello, 2008) that specifically applies this type of account to the task of category learning; specifically, to situations in which multiple different categories are learned simultaneously by the presentation of a series of exemplars of those different categories. We describe this model briefly here and show it matches participant performance in a multiple-category learning task.

At the heart of our model are two distinct memories, one of which contains the currently successful candidate rules for the categories being learned, the other of which contains a set of failed rules. Both of these memories have a fixed, relatively small size (only a small number of rules can be stored in each memory). When a new exemplar is presented to the model for classification, all currently successful candidate rules that apply to that exemplar are selected from memory; one of those rules is randomly chosen, and the category predicted by that rule is given as the classification response for the presented exemplar. Candidate rules may either be 'full' (identifying a single category as the correct response for a given exemplar) or 'partial' (identifying more than one category as possibly correct). For partial rules, the model selects one of the possibly correct categories at random as its response.

After an exemplar has been presented for classification and the model has given a response, the model is given feedback with the correct answer, from which it can learn. In the learning phase the model records the correctness of the currently successful rules based on the current exemplar's category membership, moves incorrect rules from the 'candidate rule' memory to the 'failed rule' memory, removes rarely used rules from the 'candidate rule' memory while leaving frequently used rules present (priming further use of those rules), and finally creates some new candidate rules to be added to memory.

The creation of new candidate rules is fundamental to the model's mechanism for learning. The model provides four different ways in which new candidate rules can be created. These methods of rule creation reflect different kinds of reasoning. One method for rule creation involves using the current exemplar as a seed. A second method involves randomly generating a new rule. A third method involves creating a rule for one category by randomly modifying an existing rule already in memory for some other category; and a fourth method involves creating a new rule by randomly combining two already-existing rules. The rate of use of the four methods can be set by four different parameters, so the model can have behavior that is more or less anchored in the set of observed exemplars, more or less scientific, and more or less analogical.

The most important feature in this model is that it is able to create partial rules that, if successful, can be further refined through modification. This allows the model to produce analogical rules for two different categories. Consider two categories that have some particular shared structure but are distinguished from each other by alignable differences within that structure. Presented with examples of these categories the model would (gradually and randomly) form a partial rule capturing the shared structure of those two categories (thus distinguishing examples of those categories from examples of other categories), and then (again gradually and randomly) modify that rule to produce two final rules, one for each category. These two final rules would both possess the same shared (analogical) structure common to both categories, but would also contain alignable differences which identify members of those categories separately. The model would thus form an analogical representation of the two categories.

This model makes two general predictions about the course of learning of multiple analogical categories. First, the model predicts that category learning should begin with the formation of 'macrocategories' (pairs of categories that share common features or common analogical structure). Learning should subsequently proceed to the formation of fully-differentiated final categories by the division of those initial macrocategories. Second, the model predicts that the learning of pairs of similar or analogical categories should occur at roughly the same time: once a learner has successfully identified members of one category in such a pair, they will rapidly identify members of the other category. Support for the model comes from experimental confirmation of these predictions. Further support comes from the relatively good fit the model gives to individual participant performance in a multi-category classification experiment. The next section describes this support in more detail.

EMPIRICAL SUPPORT FOR LEARNING BY SUBSTRUCTURE PRIMING

We tested the predictions described above in a category learning task where participants had to simultaneously learn 4 different categories, A, B, C and D. The four categories being learned consisted of two pairs of similar categories (category A was similar to category B, category C was similar to category D). Exemplars of these categories were composed of a number of colored geometric shapes, and were produced by computer according to the category rules for the four different categories. Two category rules (for categories A and B) had a complex structure, based on the quantities of the elements (e.g. same number of yellow and red circles in category A, different number in category B); the other two rules had a simpler structure (presence of a distinctive element - e.g. a blue triangle in category C, a green triangle in D).

Analysis of learning times in the different categories showed that learning of one category was quickly followed by learning of the other similar category, but was completely unrelated to the learning time of the two other, different categories. Analysis of errors showed that even before any category is learned, participants' incorrect responses were significantly more likely to come from the correct macrocategory (i.e. participants answering A when the correct category response was B, participants answering C when the correct response was D). Both of these results correspond to the model's predictions.

To test the model's match with participants behavior in the category-learning task, the model was fitted to each participant's individual responses in the experiment. For each participant a genetic algorithm was used to find the combination of parameter values which allowed the model to best reproduce the specific sequence of answers given by that participant, and the ease with which that participant solved the task (that is, the number of examples the participant needed to identify all four categories). In fitting the model to each participant's data, the model was shown the same examples (in the same order) originally shown to the participant, and Cohen's Kappa Coefficient was computed as a measure of the agreement between the responses produced by the participant and the responses produced by the model. Both the Kappa values achieved and the overall agreement proportions (i.e. the number of examples for which model and participant gave the same response, divided by the total number of examples seen) were significantly more than the random level, for all the participants, thus showing that the model is able to accurately reproduce the human participants' behavior (Figure 1). It was noticeable that the model's fit was better for

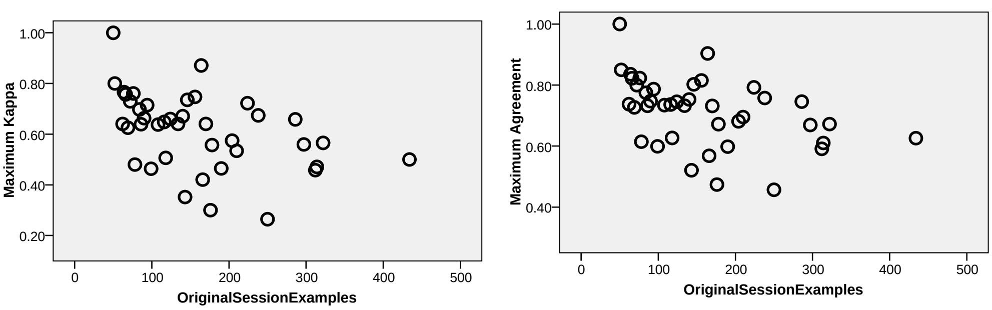

Figure 1. Match between model responses and participant responses for individual participants. Each circle represents the best fit between the model and one individual participant. The left graph shows the maximum values of the Kappa Coefficient obtained in fitting the model for each participant; the right graph shows the maximum agreement proportion for the same fit. Both are given as a function of the number of examples needed to complete the classification task. In most cases the model obtained a close level of agreement with participants' responses.

participants who completed the category learning task relatively quickly: as the right graph in Figure 1 shows, the model produced the same response as the participant at least 60% of the time for all participants who required 100 examples or less to complete the task (expected random agreement is 25%).

COMPUTATIONAL EFFECTIVENESS OF ANALOGICAL LEARNING

In this section we compare the effectiveness of progressive alignment and of substructure priming as category learning mechanisms. We focus on the contrasting performance of these mechanisms at different stages in the learning process. To make this comparison we define a measure of computational effectiveness, or learning rate per computational cost, as follows:

(1) the computational effectiveness of a learning mechanism at a given stage of learning is equal to the probability of learning taking place when a new exemplar is presented at that stage, divided by the computational cost of processing that new exemplar.

PROGRESSIVE ALIGNMENT

The progressive alignment process of category learning described above depends on a structure-mapping comparison between each new exemplar and potentially every stored generalization or stored exemplar. The computational complexity or cost of progressive alignment thus depends on two factors: the computational cost of each structure-mapping comparison, and the number of pairwise comparisons that are necessary.

Because structure-mapping is a graph-isomorphism process which must consider a combinatorial number of potential matches between structural elements from the pair of structures being compared to generate an optimal mapping between those structures, it is intuitively an NP-hard problem (Evans, Gedge, Muller, van Rooij, & Wareham, 2008; Veale & Keane, 1997), with a complexity of up to $O(2^N)$, where N is the average number of features and relations in the structures being compared. The average complexity of structure-mapping has been much debated in the literature. Here we focus on the computational cost dictated by the number of pairwise comparisons between a new exemplar and stored exemplars or generalizations.

As described above, in the SEQL model of progressive alignment a new exemplar is compared sequentially to the series of stored generalizations, ordered by number of previously assimilated exemplars. This series of comparisons stops when the new exemplar has a high enough match to a given generalization to be assimilated to that generalization. If the new exemplar does not match any stored generalization, it is then compared to all stored exemplars, and if it doesn't match any of those it is stored as an exemplar itself.

If we assume that a given new exemplar must be compared to M stored representations before a match is found, and if we assume a fixed cost C for each comparison, the cost of this repeated series of comparisons is proportional to M×C. Two factors influence M: the value of the match threshold T, and the number of successful generalizations already found. These factors have different influences at different stages in learning. Consider an early stage of learning, where a number of exemplars have been seen and stored, but no successful generalisations have yet been formed (see Figure 2). If T is high (as required to give a good fit to human learning performance), then the chance of a new exemplar having a match greater than T to one of those stored exemplars will be low. The number of sequential comparisons required to find such a match (or to reach the end of the list of stored exemplars) will thus be large. Early in learning, therefore, the computational cost of the progressive alignment process will rise proportional to the number of exemplars seen so far, multiplied by the cost of structural alignment against each exemplar. Since learning effectiveness is equal to learning rate divided by computational cost, this means that the learning effectiveness of the progressive

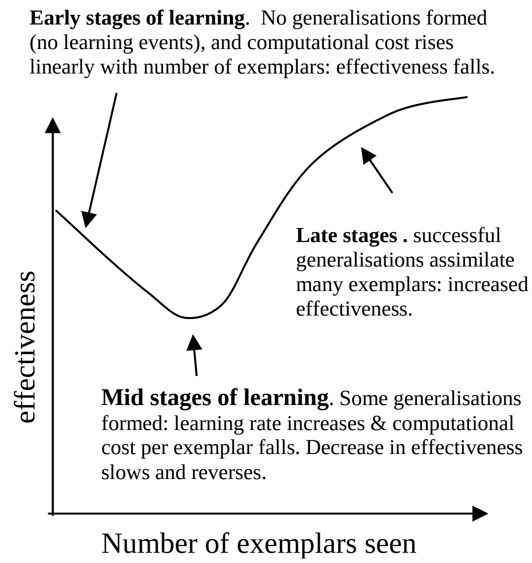

Figure 2. Sketch of relationship between number of exemplars seen and the computational effectiveness of progressive alignment category learning.

alignment process will decline rapidly at this stage of learning.

Next consider an intermediate stage of learning, where a small number of relatively successful generalizations have been formed (generalizations that have already assimilated a certain number of exemplars). The SEQL model compares a new exemplar first against generalizations sorted in order of success, and then against stored exemplars. Since the generalizations available at this stage of learning have already been relatively successful (they have already had a match greater than T to some exemplars), the chances of a new exemplar successfully matching one of these generalizations is increased. Since these successful generalizations are considered first, the number of comparisons, M, required before a successful match is found will be less at this stage of learning than it was at early learning. The learning rate at this stage of learning will thus rise and computational cost per learning event will fall. The effectiveness of the progressive alignment mechanism will thus begin to increase at this stage.

Finally consider a late stage of learning, where a number of highly successful generalizations have been formed (generalizations that have already assimilated a large number of exemplars). Again, since the generalizations available at this stage have already been highly successful (they have already had a match greater than T to a large exemplars), the chances of a new exemplar successfully matching one of these generalizations is high. Since these successful generalizations are considered first, the number of comparisons a successful match is found will be less at this stage of learning than at the earlier stages. Learning rate will be relatively high at this stage, while computational cost per exemplar will be relatively low, resulting in a high degree of effectiveness for the progressive alignment mechanism at this stage.

SUBSTRUCTURE PRIMING

In progressive alignment, the computational cost of processing a newly presented exemplar is a function of the number of exemplars already seen and the degree of learning that has already taken place (because that new exemplar is compared to already formed generalizations and to previously stored exemplars until a match is found). In the substructure priming model of analogical learning, the cost of processing a newly presented exemplar is dictated by a number of model parameters, including the size of the two memories used by the model to hold candidate rules and failed rules, and also the number of times the different ‘rule generation’ methods used by the model are applied for each exemplar. These parameters are fixed across a given run of the model (neither memory size nor number of candidate rules generated can change within a given model run), and so the cost of processing a new exemplar in this model is a constant which does not change as more exemplars are presented.

In assessing the learning effectiveness of substructure priming at different stages of learning, then, we need not consider any changes in computational cost. The only factor that can alter from stage to stage is the model’s

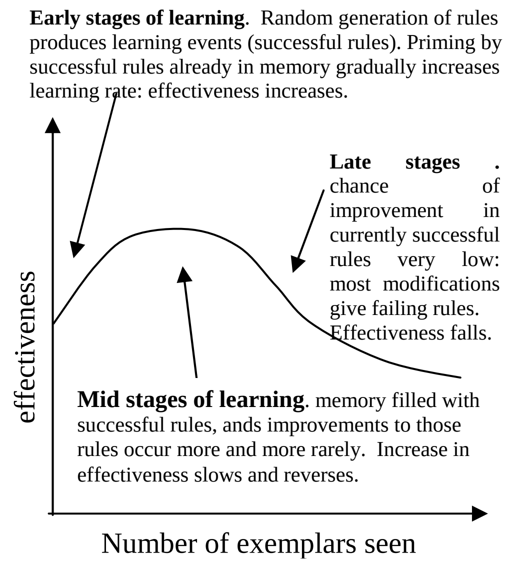

Figure 3. Sketch of relationship between number of exemplars seen and the computational effectiveness of substructure priming category learning.

learning rate. The learning rate, for this model, is equal to the rate at which a rule already in memory is removed and replaced with a more successful rule. Early in learning, such learning events will occur relatively frequently: the initial randomly generated rules stored in memory will be relatively poor and so will have a relatively high chance of being replaced by some improved rule. Further, since the model preferentially stores successful rules and uses those rules to produce new rules, the rate of production of successful rules should increase quickly at this stage of learning. Since computational cost per exemplar is a constant, and during early learning the learning rate increases as more exemplars are presented, learning effectiveness of the model will also increase at this stage of learning (see Figure 3).

As learning proceeds, the model's candidate rule memory will gradually fill with successful rules. At the middle stage of learning, the chance of the model randomly generating a new rule which is more successful than one of those already successful rules will necessarily fall. Learning events, in other words, will become rarer and rarer at this stage. This means that as learning proceeds the initial increase in learning effectiveness of the substructure priming mechanism will slow and gradually reverse. Finally, as the model approaches the later stages of learning, this decline in effectiveness will increase, as the chance of an improved rule being produced at random falls lower and lower.

This analysis suggests that a substructure priming mechanism for forming analogies will have a stronger impact and be more effective at the beginning of a learning process, when many candidate rule structures are available to a learner and need to be considered during learning. An explicit structural comparison process, by contrast, would be too computationally expensive to be effective until later in learning (when a number of successful generalizations have already been found, limiting the number of comparisons necessary to process a newly presented exemplar (compare Figures 2 and 3). The most computationally effective account of analogical learning, we suggest, is a combination of these two processes, with a gradual transition from substructure priming to structure-mapping comparison as more and more complex rule structures are learned. In this two-process view of analogical learning, the priming mechanism first acts to generate an initial set of relatively successful candidate rules for the categories to be learned: once random rule-generation processes fail to produce a noticeable improvement in these rules an explicit structure-mapping comparison process is used to find and test more structurally complex rules.

RELATING ANALOGY TO TWO-PROCESS MODELS OF LEARNING

One interesting line of research by Todd Maddox and colleagues (see e.g. Maddox, Ashby, Ing & Pickering, 2004; Maddox, Filoteo, Hejl, & Ing,, 2004; Maddox & Ing, 2005; Filoteo, Maddox, Ing, & Song, 2007; Zeithamova & Maddox,, 2006) identifies two distinct mechanisms which can be used in general learning. These mechanisms are usually presented in terms of the identification of a link between stimuli and rewards. One mechanism involves explicit hypothesis generation and testing, and requires the learner to generate hypotheses about what aspects of a stimuli will successfully predict the occurrence of a reward, and to use and test those hypotheses during the course of learning. This mechanism is normally seen as a voluntary, conscious process driven by executive attention and working memory.

A second, quite separate, mechanism involves the formation of implicit associations between stimuli and rewards. This mechanism is normally seen as occurring automatically, based on involuntary perception of stimuli-reward co-occurrence. This associative mechanism is not under conscious control and is assumed not to involve executive attention.

These 'implicit' and 'explicit' processes described by Maddox and his colleagues map well to the explicit structure-mapping comparison process and the implicit substructure priming process for analogical learning. The computationally intensive nature of explicit structure-mapping comparison, and the fact that explicit comparison is a form of directed search, suggests that the structure mapping comparison can be usefully seen as a type of explicit hypothesis-testing for learning. This suggests that structure-mapping comparison during learning is a voluntary, conscious process controlled by executive attention. The random, undirected nature of the substructure priming suggest that this mechanism can be seen as a form of implicit, associative learning, occurring unconsciously and without high memory demands.

Maddox and colleagues have reported a range of substantial differences between explicit and implicit learning mechanisms, in terms, for example, of behavioral dissociation (implicit learning is influenced by the delay between item presentation; explicit learning is not) of neural circuitry (implicit learning is driven by a dopamine-based reward signal in the caudate nucleus , while explicit learning is mediated by a circuit that includes the anterior cingulate and the prefrontal cortex), and of reasoning deficits (with patients with Parkinson's disease able to carry our rule-based explicit learning but not implicit learning, possibly because of the low levels of dopamine associated with that disease). Given our mapping between the two forms of analogical learning and the two implicit and explicit general learning mechanisms we suggest that these same distinctions may apply to structure-mapping comparison and substructure priming.

CONCLUSIONS

Since substructure priming is computationally cheaper than structural comparison, we think priming is more likely to be important when very large amounts of potentially relevant information must be filtered: that is, at the beginning of the learning process. Later in learning however, the explanatory advantages of explicit structural comparison may offset its greater computational cost, especially in situations where relatively complex rule structures are already available for comparison

The overall picture we paint, then, is one where structural commonalities and alignments arise initially via a relatively simple priming process during learning, without any explicit comparison process being activated or engaged. Once this priming mechanism provides a base of relatively complex rule structures, we suggest explicit structural comparison processes may come into play.

The distinction we draw between explicit comparison processes for analogical learning and implicit priming based processes may tell us something about analogy as a conscious, executive process of explicit comparison which develops from and builds on an implicit unconscious process of structural priming. It may also tell us something about the relationship between explicit analogy in humans and the origins of analogy via implicit processes in other animals, and indeed about the development of analogy from an implicit to an explicit process in children.

REFERENCES

Bianchi C, Costello F (2008). Analogical Reasoning helps learning of Similar Unknown Concepts: the use of Analogies between Categories in Category Learning. *Proceedings of the 19th Irish Conference on Artificial Intelligence and Cognitive Science.*

Filoteo, J.V., Maddox, W.T., Ing, A.D., Song, D.D. (2007). Characterizing Rule-Based Category Learning Deficits in Patients with Parkinson's Disease. Neuropsychologia. 45, 305-320.

Fisher, A. V., & Sloutsky, V. M. (2005). When induction meets memory: Evidence for gradual transition from similarity-based to category-based induction. *Child Development*, *76*, 583–597.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 95–235). Cambridge, MA: MIT Press.

Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, *14*, 487–513.

Gentner, D & Medina, J (1998). Similarity and the development of rules. *Cognition*, 65, 263-297.

Kuehne, S. E., Forbus, K. D., Gentner, D., & Quinn, B. (2000). SEQL - Category learning as progressive abstraction using structure mapping. *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, 770-775.

Kuehne, S. E., Gentner, D. & Forbus, K. D. (2000). Modeling infant learning via symbolic structural alignment. *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, 286-291.

Leech, R., Mareschal, D. & Cooper, R. (2008) Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences* 31, 357-414.

Maddox, W.T., Ashby, F.G., Ing, A.D., Pickering, A.D. (2004). Disrupting feedback processing interferes with rule-based but not information integration category learning. *Memory & Cognition* 32(4), 582-591.

Maddox, W.T., Filoteo, J.V, Hejl, K.D., Ing, A.D. (2004) Category Number Impacts Rule-Based but not Information-Integration Category Learning: Further Evidence for Dissociable Category Learning Systems. *JEP:LMC*, 30, 227-235.

Maddox, W. T., Ing, A.D. (2005) Delayed feedback disrupts the procedural-learning system but not the hypothesis testing system in perceptual category learning. *JEP:LMC* 31(1), 100-107.

Marcus, G. F., Vijayan, S., Bandi Rao, S. & Vishto, P. M. (1999). Rule-learning in seven-month-old infants. *Science*, 283,77-80.

Namy, L. L., & Gentner, D. (2002). Making a silk purse out of two sow's ears: Young children's use of comparison in category learning. *JEP: General*, *131*, 5–15.

Namy, L. L., Smith, L. B., & Gershkoff-Stowe, L. (1997). Young children's discovery of spatial classification. *Cognitive Development*, *12*, 163–184.

Doumas, L. A. A.,. Hummel, J. E Sandhofer, C. M. (2008) A theory of the discovery and predication of relational concepts Psychological Review 115 (: 1),: 1-43

Sloutsky V.M., Fisher A.V. (2004). Induction and categorization in young children: A similarity-based model. *JEP: General*, *133 (2)*, 166-188

van Rooij, I., Evans, P., Muller, M., Gedge, J., and Wareham, T. (2008) Identifying Sources of Intractability in Cognitive Models: An Illustration using Analogical Structure Mapping. Proceedings of the 30th Annual Meeting of the Cognitive Science Society. Cognitive Science Society.

Veale, T. & M. T. Keane. (1997) The Competence of Sub-Optimal Structure Mapping on Hard Analogies *IJCAI97, the Int. Joint Conference on Artificial Intelligence, Nagoya, Japan, 1997* Nagoya, Japan,

Zeithamova, D., & Maddox, W.T. (2006) Dual-Task Interference in Perceptual Category Learning. *Memory & Cognition*. 34(2), 387-398.