

A Theory of Reflexive Relational Generalization

Leonidas A. A. Doumas (adoumas@indiana.edu)

Department of Psychological and Brain Science, 1101 E. Tenth Street
Bloomington, IN 47405

Miriam Bassok & Amy Guthormson (mbassok@u.washington.edu)

Department of Psychology, Box 351525
Seattle, WA 98195-1525

John E. Hummel (jehummel@cyrus.psych.uiuc.edu)

Department of Psychology, 603 E. Daniel Street
Champaign, IL 61820

Abstract

We present the beginnings of an account of how representations and processes developed for the purposes of reflective reasoning provide a basis for reflexive reasoning as well. Specifically, we show how the symbolic-connectionist representations that underlie the DORA model (Doumas & Hummel, 2005), and the comparison based routines that DORA exploits in the service of addressing reflective problems, such as analogy making and the discovery of novel relations, can be extended to address reflexive reasoning phenomena. We use the reflexive reasoning routines developed in DORA to simulate findings demonstrating that reflexive processes operate when subjects solve real-world mathematics problems.

Keywords: Reflexive reasoning, reflective reasoning, relational reasoning, representation.

Reflective and Reflexive Reasoning

Relational processing plays a central role in human perception and thought. It permits us to perceive and understand the spatial relations among an object's parts (Hummel, 2000; Hummel & Biederman, 1992; Hummel & Stankewicz, 1996), comprehend arrangements of objects in scenes (see Green & Hummel, 2004, for a review), and comprehend abstract analogies between otherwise very different situations or systems of knowledge (e.g., between the structure of the solar system and the structure of the atom; Gentner, 1983; Gick & Holyoak, 1980, 1983; Holyoak & Thagard, 1995). Relational thinking is powerful because it allows us to generate inferences and generalizations that are constrained by the *roles* that elements play, rather than strictly the properties of the elements themselves. For example, in the analogy between the an atom and a solar system the sun is similar to the nucleus of the atom, not because of its literal features, but because of their shared relations to planets and electrons, respectively. Moreover, given that gravity causes the earth to revolve around the sun, you can infer that some force might also cause electrons to revolve around atoms. This kind of inference is effortful and requires reflective thought (Hummel & Choplin, 2000; Hummel & Holyoak, 2003).

While relational reasoning is *reflective*, in that it is effortful and deliberate, many of the inferences we routinely make are so effortless that we are hardly even aware of the fact that we are making them. For example, if you are told that *Susan went to the movie theater* you might infer that Susan saw a movie. Moreover, you probably assume that Susan is human and a female (as opposed to, say, a male raccoon). This kind of inference is made so automatically that it is often referred to as *reflexive* (e.g., Shastri and Ajjanagadde, 1993).

In the study of human cognition, reflexive and reflective inference have, for the most part, been examined separately. For example, reflexive inference is often studied in the context of reading and story comprehension (e.g., Kintsch & van Dijk, 1978; Shastri & Ajjanagadde, 1993; St. John, 1992; St. John & McClelland, 1990), while reflective inference is emphasized in studies of problem solving and higher level reasoning (e.g., Anderson & Lebiere, 1998; Byrne & Johnson-Laird, 1989; Forbus, Gentner, & Law, 1995; Gentner, 1983, 2003; Holyoak & Thagard, 1989, 1995; Newell, 1990). Consequently, most computational models of reflexive reasoning are not suited to account for reflective processes (e.g., Shastri & Ajjanagadde, 1993; St. John, 1992), and models of reflective reasoning are not suited to account for more reflexive inferences (e.g., Falkenhainer et al., 1989; Forbus et al., 1995; Holyoak & Thagard, 1989).

While reflexive and reflective processes do seem to follow different kinds of computational constraints (Shastri & Ajjanagadde, 1993), in many cases, these two types of processes interact and need to be integrated in the performance of a single task. For example, a mathematical solution that requires a reflective rule-based reasoning often involves retrieval of arithmetic facts (e.g., $3 + 5 = 8$). The activation of such facts is highly reflexive (e.g., LeFevre, Bisanz, & Mrkonjic, 1988; Niedeggen & Rosler, 1999). For this reason, it might be limiting to view reflective and reflexive reasoning as isolated phenomena. After all, both occur within the same cognitive architecture, and both processes should operate on the same set of mental representations (Hummel & Choplin, 2000).

In this paper we provide initial ideas about how a representational architecture that can support structured relational (i.e., heavily reflective) reasoning might also provide a basis for more reflexive forms of reasoning. This work follows from initial ideas presented by Hummel and Choplin (2000). We show how the same basic representations and processes that underlie DORA (Discovery Of Relations by Analogy; Doumas & Hummel, 2005; Doumas, Hummel & Sandhofer, submitted), a model built to handle heavily reflective reasoning processes such as relational reasoning and inference, and the discovery and predication of relational concepts, can also provide the basis for an account of reflexive reasoning.

The DORA Model

DORA is a symbolic connectionist network that learns structured representations of relations from unstructured inputs. DORA is an extension of Hummel and Holyoak's (1997, 2003) LISA model of relational reasoning. Like LISA, DORA dynamically binds distributed (i.e., connectionist) representations of relational roles and objects into explicitly relational (i.e., symbolic) structures. The resulting representations enjoy the advantages of both connectionist and traditional symbolic approaches to knowledge representation, while suffering the limitations of neither (see Doumas & Hummel, 2005). DORA's basic representational scheme is adapted from LISA. In DORA, propositions are encoded in LTM by a hierarchy of *structure units* (Figures 1). Predicate and Object (*PO*) units (triangles and large circles in Figure 1) locally code for specific roles and fillers. While LISA must use different types of units to code for roles and their fillers, DORA uses the same types of units to code both roles and fillers and differentiates between roles and fillers via its binding mechanism (see below; though for the purposes of clarity, in figures we use triangles for POs representing roles and circles for POs representing fillers). For example, the proposition *bite* (Fido, Brian) would be represented in part by PO units representing the relational roles *biter* and *bitten*, and the fillers Fido and Brian. POs are connected to *semantic units* (smaller circles in Figure 1) that code their semantic features and represent both objects and relational roles in a distributed fashion. For example, the PO unit representing Fido would be connected to a set of semantic units denoting Fido's features (e.g., "dog", "male", "fierce") and the PO unit representing Brian to a set of semantic units denoting Brian's features (e.g., "cat", "male", "tabby"). Similarly, the *biter* and *bitten* roles would be connected to the semantic units denoting their features. *Role-binding units* (*RBs*; rectangles in Figure 1) bind roles to objects in LTM. *Bite* (Fido, Brian) requires by two RBs, one binding Fido to *biter*, and one binding Brian to *bitten*. At the top of the hierarchy, *proposition* (*P*) units (oval in Figure 1) binding sets of RBs into whole relational propositions. In Figure 1 a P unit binds the RBs representing *biter*+Fido to *bitten*+Brian, thus encoding the relational proposition *bite* (Fido, Brian).

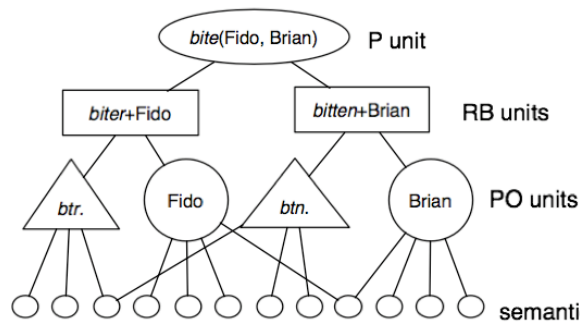


Figure 1. A proposition in DORA. Triangles are used to denote roles and circles to denote objects for clarity. In DORA, the same types of units code both roles and fillers.

In this representation, the long-term binding of roles to their fillers is captured by conjunctive RB and P units. This is sufficient for storage in LTM, however, when a proposition enters WM, its role-filler bindings must also be represented dynamically on the units that maintain role-filler independence (i.e., POs and semantics). In DORA, roles are dynamically bound to their fillers by systematic asynchrony of firing (see also Love, 1999). When a proposition enters working memory (i.e., becomes active), bound objects and roles fire in direct sequence, carrying the binding information in the proximity of firing and the role/filler distinction in the order of firing (e.g., with fillers following the roles to which they are bound). To illustrate, in order to bind Fido to the *biter* role and Brian to the *bitten* role (and so represent *bite* (Fido, Brian)), the units corresponding to the *biter* role fire directly followed by the units corresponding to Fido, then the units for the *bitten* role fire directly followed by the units for Brian. A system that is sensitive to couplets (or pairs) of activation can use this information to represent the bindings of Fido to the *biter* role and Brian to the *bitten* role. As a result, DORA (as opposed to LISA) uses the same pool of semantic units to represent both predicates and objects (Doumas & Hummel, 2005).

DORA uses comparison-based intersection discovery to isolate and explicitly predicate the shared properties of compared objects and to bind these new predicates to fillers to form bound role-filler pairs. DORA can then learn whole relational representations by joining sets of role-filler pairs (see Doumas & Hummel, 2005; Doumas et al., submitted).

DORA provides an account for a number of empirical phenomena including the discovery of relational representations that support analogical thinking (i.e., representations that are both structure sensitive and semantically rich), children and adult's learning of dimensions and relational representations, the role of comparison and progressive alignment in relation learning, and the shape bias observed in early childhood categorization (see Doumas & Hummel, 2005; Doumas et al., submitted; Hummel & Doumas, 2005). DORA is a model of reflective reasoning, however, as noted by Hummel and Choplin (2000), the representational structure of LISA (and, by extension, DORA) provides an interesting

starting point for an account of reflexive reasoning as well. In both DORA and LISA, propositions are retrieved into WM from LTM via a form of guided pattern matching. During retrieval and comparison, propositions are divided into two mutually exclusive sets: a *driver* and one or more *recipients*. Comparison is controlled by the driver.

As a proposition in the driver becomes active, it generates a systematic pattern of activation on the semantic units. During retrieval, propositions in LTM are allowed to respond to this activation pattern via their shared semantic connections. For example, if the proposition *bites* (Fido, Brian) becomes active in the driver, units encoding *biter* will become active, followed by units encoding Brian, and so forth. As each PO unit becomes active, it activates a subset of the propositions in LTM (those with shared semantics). As propositions in LTM become active in response to patterns of activation imposed by units in the driver, they will themselves feed-back activation to the semantic units. For example, as Fido becomes active in the driver, it might activate other propositions about dogs in LTM (recall that Fido is a dog). These propositions will, in turn, pass activation to any semantic units to which they are connected. The basic idea is to use the feedback from structures in LTM (including both general schemas and specific situations) to reflexively infer additional semantic content of predicates and objects in the driver. As a PO in the driver becomes active and excites a set of propositions in LTM, the semantics that are activated by those LTM representations can be inferred about the PO in the driver.

Consider a minimal case where Fido in the *bite* (Fido, Brian) proposition was connected only to the semantic unit “dog” (i.e., all DORA knows about Fido is that it is a dog). When the *biter* role becomes active it activates a subset of propositions in LTM about biting, and when Fido becomes active, it activates propositions about dogs. The result is a set of active propositions in LTM about biting dogs, which activate a set of semantics connected to biting dogs. These semantics can then be inferred about Fido (i.e., connected to the Fido PO via simple Hebbian learning). The result is a set of features of biting dogs reflexively inferred about an object based on its relational context.

However, in order for this form of reflexive inference to work, the amount of activation that propositions in LTM can pass to semantic units must be limited. The reason is that, left unchecked, spreading activation will simply activate all propositions in LTM.¹ There are a number of ways to limit the spreading activation that results during reflexive inference. One simple way, and one often imposed by the constraints of the task at hand, is to use time. Often we simply do not have the time to allow runaway activation because we must make inferences quickly. Another way to limit the effects of spreading activation is to tier or grade the effect that LTM propositions have on the activation of the

semantics. For example, during reflexive inference the activation of semantics can be graded as a function of when they became active: Semantics that become active earlier have more of an effect than semantics that become active later during inference. In DORA this is accomplished by scaling the activation of semantic units by an inverse exponential function of the iteration they become active (i.e., $scaled(a_i) = a_i e^{-t}$, where a_i is the activation of semantic unit i , $scaled(a_i)$ is the scaled activation of unit i , and t is the time that unit i became active). There is evidence for this type of graded spreading activation from the literature on memory (e.g., Anderson, 1974).

We are not claiming that either of these methods is the only form of limiting activation spread during reflexive inference (the two are, after all, not mutually exclusive). We are simply using these methods as a demonstration that it is that not difficult to avoid the spreading activation problem that arises during reflexive inference, in a network like DORA. What is interesting is that this account of reflexive reasoning arises from the same processes and representations that underlie DORA’s account of explicitly reflective processes like analogy, relation discovery, and relational inference.

Below, we use DORA’s reflexive inference algorithm to simulate subjects N400 ERP responses. Because the N 400 response occurs between 300-500ms after the onset of the stimulus, there is only a short amount of time during which reflexive inference can occur.

Simulations

Bassok and her colleagues (e.g., Bassok et al., 1998) have demonstrated that people are sensitive to the fit between a mathematical operation and the elements upon which the operation is performed. For example, people are happy to add cars and trucks but refrain from adding cars and mechanics. Similarly, people are much happier dividing cars among mechanics than cars among trucks. Such “semantic alignments” make sense in light of the fact that people frequently apply arithmetic operations to solve real world problems (Bassok et al., 1998). Guthormsen and colleagues (2004) used ERP methodology to test the fluency of such alignments. Subjects watched aligned and misaligned addition or division “applied” problems, such as $8 \text{ roses} + 9 \text{ daises}$, flash across a computer screen in sequence. For example, a subject would see the number 8, followed by the word roses, followed by +, followed by the number 9, followed by the word daises (each number, word, or symbol was presented for 650ms). The subject’s task was to solve the problem and generate a numerical answer with an object label. The ERP recording was locked to the second (target) word.

Guthormsen and colleagues (2004) found that the N400 magnitude was significantly larger for target words that created misaligned problems than for those that created aligned problems. This pattern of brain responses is similar (in its polarity, timing, and scalp distribution) to that observed in language comprehension, when people integrate

¹ In short, as a result of one set of LTM propositions becoming active and then activation their semantics, a new set of LTM propositions (those that shared some semantic overlap with the active propositions) will become active, and so forth.

the meaning of consecutive words in a sentence; it is commonly referred to as the N400 effect (Kutas & Hillyard, 1980). That is, presented with a word that does not fit the mathematical operation (e.g., baskets in: 5 apples + 3 baskets), subjects demonstrated the same neural response that is observed when people encounter words that are difficult to integrate in a sentence (e.g., shoes in “he drove to the *shoes*”). These findings provide strong evidence that people integrate mathematical and conceptual knowledge in “real time,” while reading the problem.

Our goal was to simulate the effects observed by Guthormsen et al. (2004) using DORA’s reflexive inference algorithm. We constructed DORA’s LTM (i.e., the knowledge structures that would drive reflexive inference) to reflect the fact that people learn and use mathematics in the context of solving real-world problems. To this end, we randomly selected 22 addition and 22 division word problems constructed by undergraduates in a different study (Reaume & Bassok, 2005). In that study, students were presented with simple addition or division arithmetic problems (e.g., $2 + 7 = 9$; $12 / 3 = 4$, respectively) and asked to generate corresponding word problems. An example of an addition word problem is: “You have 2 oranges and 7 apples, how many fruit do you have in all?” An example of a division word problem is: “Johnny has 12 puppies and wants to put them in 3 baskets, how many puppies should he put in each basket?” These word problems reflected people’s experience with real world situations in which simple arithmetic operations might be used, and with word problems they encountered in school.

To construct DORA’s LTM, we first listed the objects involved in the mathematical operations (i.e., what was being added/divided). Then, we had two undergraduate research assistants create lists of features describing each of these objects. Each of the 44 word problems (22 addition and 22 division) was input into DORA’s LTM as a single proposition (see Figure 2). Each proposition consisted of the objects involved in the mathematical operation (e.g., *added* (roses, tulips), *divided* (apples, baskets)). Each object was attached to the features the undergraduate coders had used to describe it. Each mathematical operation was tied to semantics describing the operation itself (i.e., “division”, “dividend”, “divisor”).

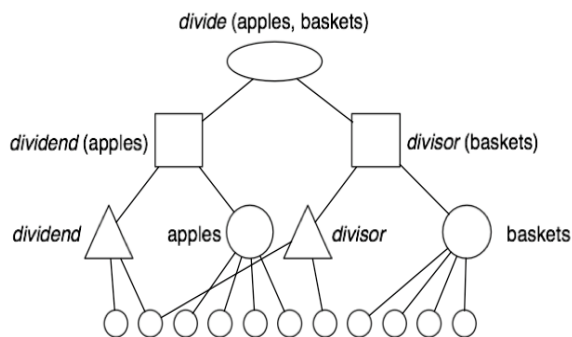


Figure 2. Example of a proposition in DORA’s LTM.

To simulate a trial in Guthormsen et al.’s (2004) experiment we presented DORA with one of the words in its LTM (i.e., a word it knew) by placing a representation of that word in the driver as a PO unit attached to a single semantic unit (the semantic unit named the object; e.g., if the word was “apple” the PO unit was attached to the semantic “apple”). This corresponded to the first word subjects in Guthormsen et al.’s (2004) experiment saw. We then allowed DORA’s reflexive inference algorithm to run by allowing the PO in the driver to activate propositions in LTM, and allowing these propositions to feedback activation to the semantics (see Figure 3a). For example, if the word was “apple”, a PO attached to the semantic “apple” was activated in the driver, it began to activate propositions in LTM about apples. We then placed a representation of either addition or division in the driver.² Addition was represented by a PO unit attached to the semantics, “addition”, “addend1”, and “addend2”, and division by a PO unit attached to the semantics, “division”, “dividend”, and “divisor”. Of course we are not claiming that these are the “right” semantic primitives of the relations addition and division. Rather, our claim is that relations and their roles are coded by distributed sets of features. The labels we attach to these features are arbitrary and mean nothing to DORA. They are only used to help interpret the model’s behavior.

Again, we allowed DORA’s reflexive inference algorithm to run by allowing the PO in the driver to activate propositions in LTM, and allowing these propositions to feedback activation to the semantics (Figure 3b). For example, when the addition semantics became active they began to activate propositions in DORA’s LTM about addition. As a result of activating the initial word representation (e.g., apple), and the representation of the mathematical operation (e.g., addition) propositions in LTM about performing the given mathematical operation on the object tended to become most active, and thus activate semantics about objects that commonly entered the specific mathematical operation with the specific object. Continuing our example, if apple was activated followed by addition, then propositions about adding apples tended to become active in LTM, which tended to activate the semantic features of objects that were frequently added to apples (namely, “fruit”).

We used the set of semantics that had become active as a measure of what DORA expected to see when the second word appeared. The second word was also a word from DORA’s LTM (i.e., one it already knew). The difference between the semantics that had become active during reflexive inference and the semantics of the second word was used as a measure of DORA’s “surprise” given the

² We did not present DORA with numbers on these trials because it was not our goal to simulate the mathematical reasoning subjects performed. Rather, we were concerned with whether DORA’s reflexive inference algorithm would lead it to expect a certain semantic category given a semantic prime (a word) and a specific mathematical operation (addition or division).




Figure 3. A graphical depiction of a simulation trial in DORA. (a) The first word is activated in the driver. Propositions in LTM become active in response to the active driver unit. Active propositions in LTM activate their semantic units. (b) The mathematical operation becomes active. More propositions in LTM become active in response to the active driver unit. The active propositions in LTM feedback activation to their semantic units. The pattern of activation on the semantic units after (b) was taken as a measure of the properties DORA "expected" to be attached to the second word. Darker = more active units.

second word. The less over-lap between the semantics that had become active by reflexive reasoning and the second word, the greater the “surprise”.

Just like the subjects in Guthormsen et al.’s experiment DORA demonstrated greater surprise for mathematical problem where the elements did not fit with the mathematical operation (e.g., when apples were added to baskets). Just as Guthormsen et al.’s subjects showed a significantly higher N400 response (i.e., between 300 and 500 ms) in response to a word that did not fit with the mathematical context than did subjects who were shown a word that did fit the mathematical context, DORA was more “surprised” by words that did not fit with a given mathematical context, than by words that did. Specifically, when elements did not fit, only 8% of the semantic features attached to the second word were already active upon presentation. However, when the second word did fit with the mathematical operation, 31% of the semantic features attached to the second word were already active upon presentation. In other words, DORA reflexive inference algorithm predicted the words that did fit, but did not predict the words that did not. Exactly as we had hoped, DORA provides an encouraging beginning for understanding how a system built to model reflective processes might be extended to address the problems of reflexive inference.

Discussion

Using simple operations that were already in place for the purposes of reflective reasoning, DORA has been able to account for the reflexive reasoning phenomena observed by Guthormsen and colleagues (2004). When DORA’s LTM consists of math word problems generated by undergraduate

students, DORA reflects the same biases for the fit between mathematical problems and real world elements demonstrated by adult reasoners. Just like the subjects in Guthormsen et al.’s experiment, DORA was more “surprised” when encountered mathematical word problems where the elements did not fit naturally with the mathematical operation, then when it encountered word problems where the elements did fit naturally with the mathematical operation. This suggests that the symbolic-connectionist representational structure and the mapping based reflexive inference routines that DORA performs in the service of reflective tasks like analogical mapping, memory retrieval, and relation discovery might provide the beginnings of an account of reflexive inference as well. At the very least DORA, following from Hummel and Holyoak’s (1997, 2003) LISA model provides evidence that the same representational structures and basic processes might underlie and operate in the service of both reflective and reflexive reasoning process.

As a theory of reflexive inference, however, DORA is far from complete. We have demonstrated DORA’s ability to account for simpler reflexive inferences, but it is not clear whether DORA would scale up to account for more complex reflexive inference. For example, as noted in the introduction, if you are told that *Susan went to the movie theater* you might infer she saw a movie. If DORA’s LTM contained a number of propositions about seeing movies at movie theaters it might be able to infer that Susan saw a movie when she went to the theater, but it is not clear how DORA could reflexively infer structured propositions such as *saw* (Sally, movie) solely given feedback to semantic units from LTM. However, DORA does suggest a promising starting point for investigating the requirements

that a representational system must meet in order to account for both reflective and reflexive inference.

References

- Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: LEA.
- Bassok, M, Chase, V. M, & Martin, S. A. (1998). Adding apples and oranges: Alignment of semantic and formal knowledge. *Cognitive Psychology*, 35, 99-134.
- Byrne, R. M. J., & Johnson-Laird, P. N. (1989). Spatial reasoning. *Journal of Memory and Language*, 28, 564-575.
- Doumas, L. A. A., & Hummel, J. E. (2005). A symbolic-connectionist model of relation discovery. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 606-611. Mahwah NJ: LEA.
- Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (submitted). A symbolic-connectionist model of predication and relation discovery.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141-205.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (2003). Why we're so smart. In D. Gentner and S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp.195-235). Cambridge, MA: MIT Press.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Guthormsen, A., Bassok, M., Osterhout, L. (2004, November). ERP Measure of Conceptual Integration Between Mathematical and Semantic Relations. Poster presented at the 45th annual meeting of the Psychonomic Society, Minneapolis, MN.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Holyoak, K. J., & Thagard, P. (1995). *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press.
- Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich and A. Markman (Eds.), *Cognitive Dynamics: Conceptual Change in Humans and Machines*. Hillsdale, NJ: Erlbaum.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, 99, 480-517.
- Hummel, J.E., & Choplin, J.M. (2000). Toward an integrated account of reflexive and reflective reasoning. In L.R. Gleitman & A.K. Joshi (Eds.), *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 232-237). Mahwah, NJ: Erlbaum.
- Hummel, J. E. & Doumas, L. A. A. (2005). Some speculations on the development of the role of visual invariants in shape perception. Paper presented at the biennial meeting of the Society for Research in Child Development.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427-466.
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*.
- Hummel, J. E., & Stankiewicz, B. J. (1996). Categorical relations in shape perception. *Spatial Vision*, 10, 201-236.
- Kintsch, W. & van Dijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review*, 85, 363-394.
- Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, 207, 203-205.
- LeFevre, J-A., Bisanz, J., and Mrkonjic, L. (1988). Cognitive arithmetic: Evidence for obligatory activation of arithmetic facts. *Memory & Cognition*, 16(1), 45-53.
- Love, B. C. (1999). Utilizing time: Asynchronous binding. *Advances in Neural Information Processing Systems*, 11, 38-44.
- Niedeggen, M., & Rosler, F. (1999). N400 Effects Reflect Activation Spread During Retrieval of Arithmetic Facts. *Psychological Science*, 10(3), 271-276.
- Reaume, G. & Bassok, M. (2005). *Apples and Apples of Alignment: The Interaction of Semantic and Formal Knowledge in Simple Arithmetic*. Poster presented at the 7th annual meeting of NOWCAM, Bellingham, Washington.
- Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings. *Behavioral and Brain Sciences*, 16, 417-494.
- St. John, M. F. (1992). The Story Gestalt: A model of knowledge-intensive processes in text comprehension. *Cognitive Science*, 16, 271-302.
- St. John, M. F., & McClelland, J. L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, 46, 217-257.