

SAND2004-4567
Printed September 2004

**Dynamic Context Discrimination:
Psychological Evidence for the
Sandia Cognitive Framework**

Ann Speed
Reliability Assessment and Human Factors

Table of Contents

1. Introduction.....	3
2. Dynamic Context Discrimination: An overview	4
3. Dynamic Context Discrimination as a core cognitive function: A review of the converging evidence	6
3.1 The neurophysiology of sensation and perception.....	6
3.2 Higher-order cognition: Memory and Analogy as the “engines” for DCD	7
3.3 Other examples of DCD.....	12
4. Dynamic Context Discrimination: The computational implementation.....	20
4.1 Psychological inspiration for the model	20
4.2 Knowledge representation in the model: Physiological inspiration	21
4.3 Model Components	22
4.4 Information processing in the Sandia framework.....	23
4.5 Appearance of specific psychological functions in the model.....	25
5. General Discussion	25
6. References.....	28
Distribution	36

1. Introduction

Human behavior is a function of an iterative interaction between the stimulus environment and past experience¹. It is not simply a matter of the current stimulus environment activating the appropriate experience or rule from memory (e.g., if it is dark and I hear a strange noise outside, then I turn on the outside lights and investigate). Rather, it is a dynamic process that takes into account not only things one would generally do in a given situation, but things that have recently become known (e.g., there have recently been coyotes seen in the area and one is known to be rabid), as well as other immediate environmental characteristics (e.g., it is snowing outside, I know my dog is outside, I know the police are already outside, etc.). All of these factors combine to inform me of the most appropriate behavior for the situation. If it were the case that humans had a rule for every possible contingency, the amount of storage that would be required to enable us to fluidly deal with most situations we encounter would rapidly become biologically untenable. We can all deal with contingencies like the one above with fairly little effort, but if it isn't based on rules, what is it based on?

The assertion of the Cognitive Systems program at Sandia for the past 5 years is that at the heart of this ability to effectively navigate the world is an ability to discriminate between different contexts (i.e., Dynamic Context Discrimination, or DCD. While this assertion in and of itself might not seem earthshaking, it is compelling that this ability and its components show up in a wide variety of paradigms across different subdisciplines in psychology.

We begin by outlining, at a high functional level, the basic ideas of DCD. We then provide evidence from several different literatures and paradigms that support our assertion that DCD is a core aspect of cognitive functioning. Finally, we discuss DCD and the computational model that we have developed as an instantiation of DCD in more detail.

Before commencing with our overview of DCD, we should note that DCD is not necessarily a theory in the classic sense. Rather, it is a description of cognitive functioning that seeks to unify highly similar findings across a wide variety of literatures. Further, we believe that such convergence warrants a central place in efforts to computationally emulate human cognition. That is, DCD is a general principle of cognition.

It is also important to note that while we are drawing parallels across many literatures, these are functional parallels and are not necessarily structural ones. That is, we are not saying that the same neural pathways are involved in these phenomena. We are only saying that the different neural pathways that are responsible for the appearance of these various phenomena follow the same functional rules – the mechanisms are the same even if the physical parts are distinct. Furthermore, DCD is not a causal mechanism – it is an

¹ Personality undoubtedly plays a role in behavior, however, including individual differences in behavior that are due to personality structure is beyond the scope of the current discussion.

emergent property of the way the brain is constructed. DCD is the result of neurophysiology (cf. John, 2002, 2003).

Finally, it is important to note that we are not proposing a generic learning mechanism such that one biological algorithm can account for all situation interpretation. Rather, we are pointing out that there are strikingly similar empirical results across a wide variety of disciplines that can be understood, in part, by similar cognitive processes. It is entirely possible, even assumed in some cases (i.e., primary language acquisition) that these more generic cognitive processes are complemented and constrained by various limits which may or may not be biological in nature (cf. Bates & Elman, 1996; Elman, in press).

2. Dynamic Context Discrimination: An overview

In forming our notions about how a computational model of cognition should look and act, we drew on evidence from a variety of literatures in psychology and cognitive neuropsychology. We wanted to ensure that the model we built honored what science knows about how the brain represents and processes information. Basically, the result of this research led us to DCD which states, at its most basic, that human behavior is an iterative, dynamic process that is supported by “higher order” cognitive functions such as memory and analogy. We see evidence for DCD not only in the response to sensory information, but we also see DCD operating in problem solving, text comprehension, language comprehension, and learning. Figure 1 illustrates DCD at a coarse-grained functional level.

To summarize Figure 1, as information comes into the brain from the stimulus environment, the way it is interpreted is affected largely by situations that presented similar stimulus patterns. This similarity is determined via analogy, although we do not subscribe to the notion that there is a separate top-down analogy engine in the brain². Rather, this is a bottom-up process of pattern recognition that originates in the way sensory information is encoded in the sensory cortices (see discussion below on the neurophysiology of sensation) and that is mediated by the already-recognized ongoing context. In this way, the same set of stimuli encountered at two different times can elicit somewhat different interpretations/responses based on the larger context in which they

² Note that this process is not necessarily, or even usually, an explicit or effortful one. (Chi, Feltovich, & Glaser, 1981; Dunbar, 2001; Gick & Holyoak, 1980, 1983; Holyoak & Thagard, 1995; Klein, Calderwood, & Clinton-Cirocco, 1985; Novick, 1988; cf. Holyoak & Hummel, 2001; Roediger, 2000; Ross, 1984). Shastri (1999) calls this *reflexive reasoning*, and we believe this is the same phenomenon Klien and his colleagues observed with expert firefighters’ time-critical decision-making on the fireground (Klein, Calderwood, & Clinton-Cirocco, 1985). In Klein’s RPD, there is rapid retrieval and mapping – and the retrieval is based on rapid recognition of systematic relationships based on schema.

In an interesting laboratory-based experiment, Ratcliff and McKoon (1989) found that retrieval of information based solely on surface features occurs within about 500 ms of target presentation, whereas retrieval of information based on structural features begins at around 700 ms. Clearly, there is a distinction between retrieval based on surface and structural features (the latter really being analogy), but whether or not conscious, effortful mapping occurs in that 200 ms lag time is another issue altogether.

For additional work on this topic, see the literatures on implicit memory and implicit learning.

occur (Capaldi & Neath, 1995; Olafson & Ferraro, 2001; Spellman & Holyoak, 1996; Tulving, 1983).

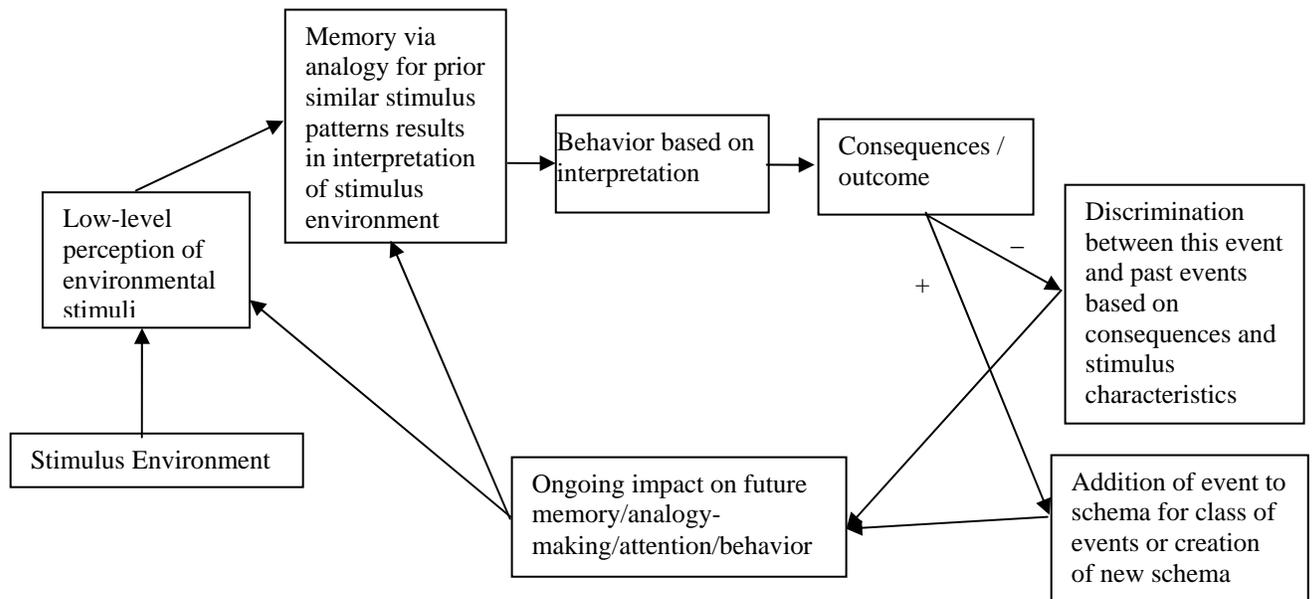


Figure 1: Functional overview of Dynamic Context Discrimination

It is very important to note that there is never a time when the only information coming into the cognitive system is from the external sensory environment. It is also important to note that there is rarely a time when there is only one context operating (e.g., I am hungry *and* I am at work *and* it is 10 am *and* I am typing a paper). These already recognized situations impact the interpretation of new stimulus information coming into the cognitive system, whether that stimulation is external (e.g., it is warm in this room) or internal (I am tired). In part, it is the top-down influences of these already-recognized contexts that make the model dynamical in nature.

Regardless, once an interpretation of the current stimulus environment is formulated, a response follows. This response is based on memory, via analogy, for prior similar events. In general, if the consequences of the response are as expected, this episode contributes to the strength of the situation class that is defined by the present pattern of environmental stimuli. In general, if the consequences of the response are *not* as expected, some distinction is made between the assumed situation class and the current environment. At this point, the organism may then make a finer-grained distinction between two seemingly similar stimulus patterns. Of course, this requires that the organism is sensitive to not only the consequences of its behavior, but that it is also sensitive to the differences in stimulus environments that led to those differential consequences (see the literature on *Einstellung*, or mental set for an example of what happens in the absence of this sensitivity e.g., Luchins & Luchins, 1950; see also literature on negative transfer, e.g., Gick & Holyoak, 1987). The process of either differentiating similar patterns of stimuli or abstracting generalities from non-identical

patterns is an important dichotomy in DCD, as these functions are two sides of the DCD coin. This point will be explained in more detail the following literature review.

3. Dynamic Context Discrimination as a core cognitive function: A review of the converging evidence

In the following discussion, it is important to keep in mind that DCD is an iterative, dynamic process and that events in the past are constantly influencing the way that incoming information is interpreted. Therefore, while the following discussion is presented as though DCD had formal, cleanly separable stages, we recognize that is not necessarily the case.

Regardless, the evidence we present below gives an overview of the kinds of data used to support the architecture of the model that is presented in a later section. These data inform not only the high-level functioning, but the low-level knowledge representation and how these interact with one another. The goal behind building this computational representation of human cognition was to create a model that is behaviorally, psychologically plausible and physiologically inspired. We use data from cognitive neuroscience regarding information representation and processing in the brain as our guide. Then we use the psychological literature to ensure the model is psychologically and behaviorally plausible.

3.1 The neurophysiology of sensation and perception

While our model, in its current state, does not include sensory organs, we do recognize that cognition happens in relation to, and therefore must be partially determined by, a physical presence in the world (cf. Petitot, 1995). We also recognize that the question of how information is represented needs to be addressed at the point at which environmental stimuli are initially registered in the cognitive system. Therefore, in order to develop a model that will eventually interface with synthetic sensory organs; we considered the neurophysiology of sensation in determining our knowledge representation and processing scheme.

This literature indicates that as stimulus information propagates through the sensory cortices, *individual cells* (i.e., simple and complex cells) respond selectively to increasingly complex stimuli. Retinal ganglion cells respond selectively to very specific stimuli (e.g. a dot in a particular place in the visual field), but as information passes through the system, the responding cells are less and less selective about the stimuli to which they respond, essentially creating equivalencies between multiple similar, nonidentical sensory inputs. This is the case for vision (Huebel & Weisel, 1962; 1979), audition (Rauschecker, 2001; Richards & Kidd, 2001) and haptic perception (Rauschecker, 2001), and there is preliminary evidence that this is also true of olfaction (Youngentob, 2001) and taste (Halpern, 2001).

In addition to increasingly complex equivalencies drawn by simple and complex cells, there are neurons in the brain (e.g., in the superior colliculus) that specifically respond to patterns of activity across multiple sensory cortices (Stein, Wallace & Stanford, 2000),

thereby encoding increasingly complex bits of information that include information across multiple senses. Interestingly, the integrative functioning of the neurons in the superior colliculus is dependent on connections to the cortex – once the connections to the parts of the brain that perform higher-level cognitive functions are severed, the neurons in the superior colliculus cannot respond in a way that enables them to integrate information from multiple sensory modalities (Stein, Stanford, Vaughan, & Wallace, 2001).

The important point to take from these findings is that the sensory system is set up to make distinctions between different stimuli as well as to make abstractions across functionally similar stimuli. In addition, it is clear from the existence of necessary two-way connections between the cortex and the sensory pathways that there is iterative processing very early in a human’s perception of the environment. That these processes are apparent so early in the human/environment interface is important for the claims made by DCD that this iterative, dynamic discrimination process is a fundamental process in human cognition.

3.2 Higher-order cognition: Memory and Analogy as the “engines” for DCD

Once the brain has registered sensory information and has perceived meaningful, low-level patterns in the environment, there can be a response to those stimuli. How those stimuli are perceived, and how those perceptions fit into the human’s ongoing understanding of the environment will determine exactly what that response will be. From the cognitive psychological perspective, this process occurs via memory and analogy.

3.2.1 DCD in the absence of an ongoing context

If we could start at the beginning, in the absence of pre-existing contextual interpretation (but not in the absence of world knowledge), understanding the current context would be based solely on bottom-up recognition of the unique meaning of patterns of stimuli currently in the environment³. Consider the following stimuli:

- Large semi-dark room
- Very large screen at front of room
- Rows of large armchairs and sofas facing screen
- Odwalla* juice and *Newman’s Own* pretzel ads showing on the screen

Chances are, even if you had no recollection of choosing a film, driving to the cinema, and buying your ticket, you’d know you were in a movie theater based on experiencing this pattern of environmental stimuli. Furthermore, you’d recognize this as a theater even though most theaters don’t have large armchairs, sofas, and *Odwalla* juice ads. This interpretation in the absence of any prior context happens via memory and analogy (see, e.g., Chase and Simon, 1973; cf. Klimesch, 1996, 1999).

³ We use the term “environment” to include both external stimuli such as sunlight and internal stimuli such as hunger or emotion.

Once a given context is recognized, expectations about appropriate behaviors and about potential upcoming events develop. Once I realize I am in a theater, I would proceed to find a seat and wait for the film to begin. Once the film begins, I know that I should not carry on a conversation with my friends and that I should not stand up in front of the screen for any length of time. In addition I expect the same courtesy from others, I expect that the film will begin within a reasonable timeframe, that there will be previews, and that the picture on the screen will be in focus. If any of these expectations is violated, attention is refocused to the aberrant events or behaviors and a reassessment of the situation and possible responses occur. Likewise, if my behavior falls outside of that expected for moviegoers, I can expect that others in the theater will respond to correct my aberrant behaviors. All of these expectations and behaviors rely on memory for past similar experiences and the consequences of behaviors exhibited in those past episodes.

There are several related hypotheses in the memory literature that account for environmental stimuli acting as cues for recall of prior similar events (e.g., Glenberg & Swanson, 1986; Tulving, 1983). Probably the most well-known of these is Tulving's *encoding specificity hypothesis* which states that the likelihood of recalling a given event, A at some time later than its initial presentation is a function of the similarity of the encoding context (at time 1) and the retrieval context (at time 2; Glenberg, 1984; Glenberg & Swanson, 1986; Neath & Crowder, 1990; Tulving, 1983; Watkins, 1979; 1989; 1990; see also Godden and Baddeley, 1975 and Morris, Bransford & Franks, 1977). The definition of "context" can include sensory stimuli such as odors, emotional state, physiological state, and characteristics of the stimuli themselves, but it does not have to have any intrinsic meaning with regards to the remembered items in order to act as cues - they only do so because of temporal co-occurrence (Capaldi & Neath, 1995). In short, memory is largely a stimulus-controlled phenomenon (Capaldi & Neath, 1995; Nairne, 2002; Tulving, 1983; Watkins, 1989; 1990) and similarity in memory is defined by the correspondence between encoding context at T1 and retrieval context at T2.

However, simply recalling a past experience that had a similar encoding context does not fully explain how we use our past experiences to help us interact with the ongoing world; the research on analogical thinking addresses this problem. Once a relevant memory (a source) is recalled, information known about that source is then applied to the current situation (the target). If the match between the source and the target is an effective one, any conclusions reached about the target based on the source are themselves likely to be effective (positive transfer). However, if the match is a poor one, an ineffective set of conclusions is likely reached (negative transfer; Butterfield & Nelson, 1989, 1991; Catrambone & Holyoak, 1989; Gentner, 1983; Gentner & Markman, 1997; Ghodsain, Bjork, & Benjamin, 1997; Gick & Holyoak, 1980, 1983, 1987; Gick & Patterson, 1992; Holyoak & Thagard, 1995; Kolodner, 1997; Ross, 1996).

Furthermore, the thinker's perception of the current context can have an impact on what is recalled such that as perceptions of the same context change, past events considered to be similar also change (Spellman and Holyoak (1996); see also Capaldi & Neath, 1995; Tulving, 1983). Interestingly, analogy and transfer of training researchers have found that

domain novices tend to recall sources based on surface similarities between the source and the target (i.e., irrelevant cues, or *context alpha* cues; cf. Capaldi & Neath, 1995). Domain experts, on the other hand, tend to recall sources based on structural aspects of problems (i.e., cues relevant to the relational structure of the problem, or *context beta* cues; Chi, Feltovich, & Glaser, 1981; Gentner, 1983; Gentner & Markman, 1997; Gick & Holyoak, 1980; Novick, 1988; Ross, 1996; cf. Capaldi & Neath, 1995; Klein, 1997; Klein, Calderwood, & Clinton-Cirocco, 1985). Therefore, according to the analogy literature, similarity is defined differently for domain novices and domain experts. For novices, similarity is defined by readily observable, surface aspects of the problem while for domain experts, similarity is defined by the less obvious but more functionally relevant features of the problem.

Once the analogous solution procedure is mapped onto and applied to the target, if the application is an effective one, an abstracted version (i.e., schema) of the problem class can then be created or augmented (Gentner & Markman, 1997; see also Gick & Holyoak, 1983 and Watkins & Kerkar, 1985 for additional evidence of schema creation in problem solving and memory, respectively and see Feldman, 2003 for a related idea in the category learning literature). If the application of the source to the target is ineffective, the thinker is then in a position to make a distinction between the source and the target (and their respective problem classes) based on the differences between the two situations (e.g., Butterfield & Nelson, 1989, 1991; Ghodsain, Bjork, & Benjamin, 1997; Gick & Patterson, 1992; Kolodner, 1997. Also see the section below on learning paradigms for additional empirical support). Either way, by relating current events to prior known events, the thinker learns about the world incrementally.

While much of the research on analogy has been done in the laboratory, many researchers recognize that it is, indeed, a core feature of our general cognitive abilities (Forbus, 2001; Hofstadter, 2001; Holyoak, Gentner, & Kokinov, 2001; Holyoak & Hummel, 2001; Holyoak & Thagard, 1995; Kolodner, 1997; cf. Glenberg, 1997). To quote one such researcher:

“If analogy were merely a special variety of something that in itself lies way out on the peripheries, then it would be but an itty-bitty blip in the broad blue sky of cognition. To me, however, analogy is anything but a bitty blip—rather, it’s the very blue that fills the whole sky of cognition—analogy is *everything*, or very nearly so, in my view” (Hofstadter, 2001, p. 499, italics in original).

3.2.2 DCD in an ongoing context

Despite the above discussion, there is almost never a time when we interpret environmental stimuli in the absence of a larger context. Therefore, DCD must be able to deal with the introduction of new stimuli and how stimulus interpretation is affected by larger contextual understanding.

In the presence of an ongoing contextual understanding, the basic process of DCD really doesn’t change – it is still based on pattern recognition supported by memory/analogy. However, the ongoing already-recognized context does introduce some bias into the

interpretation of new incoming stimuli and it biases what is recalled based on those interpretations. In other words, ongoing contextual understanding exerts a top-down influence on both the low-level perceptual processes⁴ in that it exerts influence over attentional processes, as well as higher-level cognitive interpretations (i.e., memory, analogy) of the perceived stimuli.

There are several interesting memory phenomena that illustrate top-down influence on ongoing context interpretation (e.g. Roediger, 1996). For example, recall Brewer and Treyens' (1981) well-known study in which undergraduates had to recall items from a "graduate student's office" in which they were asked to wait prior to the beginning of the "real" experiment. Brewer and Treyens (1981) found that the typicality of items was significantly correlated with their likelihood of recall⁵. In addition, they found that typicality also predicted whether items that were not present in the room would be intruded into subjects' recall protocols. Undergraduates used their prior knowledge of a schematized version of a graduate student's office to help them recall the specific items that were present in a particular instance of such an office – regardless of whether those items were actually present or not.

A related laboratory phenomenon is the false recall phenomenon. In one typical false recall paradigm, subjects learn a list of words, all of which are highly semantically related to a given critical non-presented target word. For example, if the target word is *school*, subjects may learn words like *bus*, *teacher*, *chalkboard*, *homework*, but will never see the word *school*. When asked to recall the list, subjects recall having seen the word *school* on 40% to 55% of trials (as compared to approximately 14% of trials for other English words; Roediger & McDermott, 1995). In addition, when subjects are asked whether they simply *know* the target word was on the initial list or whether they have a conscious recollection of it having been there, they indicate that they have a conscious recollection for the presence of that particular item the vast majority of the time. For example, Payne and his colleagues (Payne, Elie, Blackwell, & Neuschatz, 1996) had a male experimenter and a female experimenter present lists of words to subjects. After recalling words from the lists, subjects were asked whether they actually *remembered* hearing one of the two experimenters read the word or if they just *knew* the word had been on the list. If subjects indicated they remembered an experimenter reading the word, they were asked which of the two had done so. Subjects were willing to say that the critical non-presented item was read by an experimenter 87% of the time, compared to 94% for words that were actually on the list.

Finally, top-down influences on the DCD process are well illustrated in the eyewitness testimony research made famous by Elizabeth Loftus and her colleagues (Loftus, 1975; Loftus & Palmer, 1974; Loftus & Zanni, 1975). In one typical study, subjects are shown a series of slides depicting a car crash. Half of the subjects are asked how fast they thought the car was going when it "smashed" into the other car, while the other half were asked

⁴ We are making the classic psychological distinction between sensation and perception in this model. That is, sensations are raw sensory data that come into the cognitive system whereas perceptions are the beginnings of interpretations of those sensations by the brain.

⁵ Typicality was measured for all items used in a separate study.

how fast they thought the car was going when it “ran into” the other car. Subjects in the “smashed” condition reported seeing broken significantly more often than those in the “ran into” condition, even though there was no broken glass in the scene. This particular set of experiments demonstrates that DCD operates not only for ongoing experiences, but also for recall of prior experiences.

The above studies emphasize that memory is *reconstructive* rather than *reproductive* (Greene, 1992; Roediger & McDermott, 1995; Talarico & Rubin, 2003). We don’t remember things verbatim – we recall the gist of events and fill in the holes based on our expectations for what most likely happened. These expectations are dictated both by our general knowledge of the world as well as by the current stimulus environment. Things not recalled about an event in one context are recalled in a different context, and newly acquired knowledge can impact the way that prior events are conceptualized (Roediger, 2000; Spellman & Holyoak, 1996; Tulving, 1983; cf. Zwaan & Radvansky, 1998).

There is a corollary to the encoding specificity principle that is of interest to the DCD in that it explains why schema creation, which is essential to DCD (and human behavior in general) might be a byproduct of the inherently reconstructive (rather than reproductive) nature of memory - the *cue overload hypothesis*. This hypothesis states that, “...the efficiency of a functional retrieval cue in effecting recall of an item declines as the number of items it subsumes increases” (Watkins & Watkins, 1975, p. 443; see also Watkins, 1979). That is, the more times you experience a given type of episode, the less likely you are to recall details of any given instance of that class of episodes (cf. Watkins & Kerkar, 1985; see also Feldman, 2003 for a related idea in the concept learning literature). That is, as you experience a given type of situation more and more often, you develop an abstracted version of that situation, or a *schema*.

In the confines of the memory literature, the cue overload hypothesis explains a variety of phenomenon including:

- buildup and release from proactive inhibition (Capaldi & Neath, 1995)
- delayed match to sample tasks (Capaldi & Neath, 1995)
- the von Restorff effect (i.e., *distinctiveness*; Cimbalo, Capria, Neider, & Wilkins, 1977; Huang & Wille, 1979; Schmidt, 1985, 1991; Wallace, 1965; Watkins 1979).

The cue overload hypothesis is important to DCD because it is the flip-side of the encoding specificity coin. That is, while the stimulus environment serves to limit the potential “search space” for recallable items (cf. the *indexing problem* Kolodner, 1997), the cue overload hypothesis explains observed flexibility for recall such that the stimulus environment does not have to be an exact match in order for a given item or event to be recalled because equivalencies form between multiple non-identical episodes. Furthermore, the formulation of the cue overload hypothesis in the memory literature (as opposed to the analogy literature) provides converging evidence of a compelling mechanism by which schema arise as a result of the normal functioning of memory. If schema creation were always a conscious, effortful process (cf. Gick & Holyoak, 1983), it would be very difficult for humans to interact with the environment – we would either

not be able to apply information from our past experience unless there was a perfect match between current and past contexts, or we would spend so much time explicitly creating schema that we would have no time to do anything else⁶.

3.3 Other examples of DCD

3.3.1 Learning paradigms in the radical behavioral tradition.

Despite its fundamental opposition to studying “mental events,” paradigms from the radical behavioral tradition actually contribute evidence to our argument that DCD is a core human cognitive process. According to behavioral theories, behavior evolves over time as a function of consequences of those behaviors. If a behavior has desirable consequences, the person will be more likely to engage in that behavior in similar circumstances in the future (i.e., the behavior is reinforced). If consequences are aversive, the person is less likely to engage in that behavior in similar circumstances (i.e., the behavior is punished)⁷.

Similarly, when a person attempts to solve a problem and that solution fails, that person will be less likely to attempt the same sort of solution in a similar situation in the future. Likewise, if the solution succeeds in the first place, that solution will be more likely to be used in similar situations in the future.

3.3.1.1 *Stimulus generalization and discrimination*

The question then becomes what constitutes a “similar” situation? Behaviorists define situations according to environmental stimuli that are present or absent at a given time. Examining situation similarity involves measuring response rates in both humans and animals as the stimulus set is manipulated parametrically. The level of *stimulus generalization* is plotted according to the parameters used to modify the stimulus set, creating a *stimulus generalization gradient* that shows the change in response rate as a function of the parametric modification of the stimulus (Donahoe & Palmer, 1994).

⁶ The encoding specificity principle and the schema-creation aspect of cue overload can be invoked to explain why analogies are so difficult to elicit in the laboratory. The finding is that if students (who are invariably domain novices with regard to the types of problems used in analogy research) are given multiple examples of the same class of problem and are then asked to explicitly identify how those examples are similar (that is, they are asked to consciously create a schema for that problem class), they have no difficulty making the analogy to a new example (Catrambone & Holyoak, 1989; Gick & Holyoak, 1983). However, in the absence of this explicit schema creation exercise, and in the absence of explicit hints, students fail to make the analogy the majority of the time (Gick & Holyoak, 1980; 1983). There are a couple of reasons for this. First, they have not had enough experience with this class of problems to recognize the structural features (i.e., the relational aspects of the problem), so they instead suggest solutions from other situations with which they are more familiar – those that share surface features with the current problem. Second, because novices tend to pay more attention to surface features rather than structural features of problems, the T1 context tends to comprise mostly surface features. It is only after they have had sufficient experience with the problem class that the T1 encoding comprises mostly structural features (cf. Butterfield & Nelson, 1989, 1991; see also Ross, 1996), allowing ready transfer to T2 problems that contain the same structural features.

⁷ It is important to note that reinforcers and punishers are highly individualized. Ultimately, whether a consequence is a reinforcer or punisher is defined by its impact on subsequent behavior.

One classic example of stimulus generalization occurs in pigeons. In this research, a pigeon is trained to peck a button below a blue light – a behavior that is reinforced with a food pellet. As the color of the light is changed from blue to green and eventually towards green-yellow, response rates decline parametrically the further away the color gets from pure blue. The same decline happens as the color moves through indigo and into red (see Figure 2). The important issue then becomes what stimuli are considered sufficiently similar to elicit similar responses? Behaviorists appeal to the concept of stimulus classes.

3.3.1.2 Stimulus classes

A stimulus class is a group of similar stimuli that act as indicators that a particular context is in operation, indicating that certain behaviors will result in certain consequences. A stimulus class can also be thought of as a problem class that is based on structurally similar features (Capaldi & Neath, 1995). There are two different types of stimulus classes: *discriminative*, and *functional*. Discriminative stimuli are stimuli that are physically similar (e.g. the blue light in the pigeon example above) that indicate that certain behaviors will result in certain consequences. Functional classes, on the other hand, are not physically similar to one another, but they still elicit similar behaviors. For example, because toys all support “playing behavior” anything considered a toy would belong to a functional class.

Stimulus classes are learned over time through experience and through *differential reinforcement*. With differential reinforcement you learn, by way of the consequences of your behaviors, which contexts will result in which consequences given a certain set of behaviors. For example, some of us may have learned the hard way to actually look for a “Please wait to be seated” sign and/or a hostess station when we walk into a restaurant. Either we stood around waiting to be seated when it was not appropriate for that restaurant, or we seated ourselves, only to be ousted. Or both. The presence of the sign and/or station is part of a functional stimulus class that elicits waiting-to-be-seated-behavior that has been differentially reinforced when members of that class are present (Donahoe & Palmer, 1994). Furthermore, there is evidence that these kinds of stimulus classes form without any conscious effort on the part of the learner and that, in some circumstances, effortful processing is actually detrimental to formation of such classes (Brooks, 1978, cited in Donahoe & Palmer, 1994).

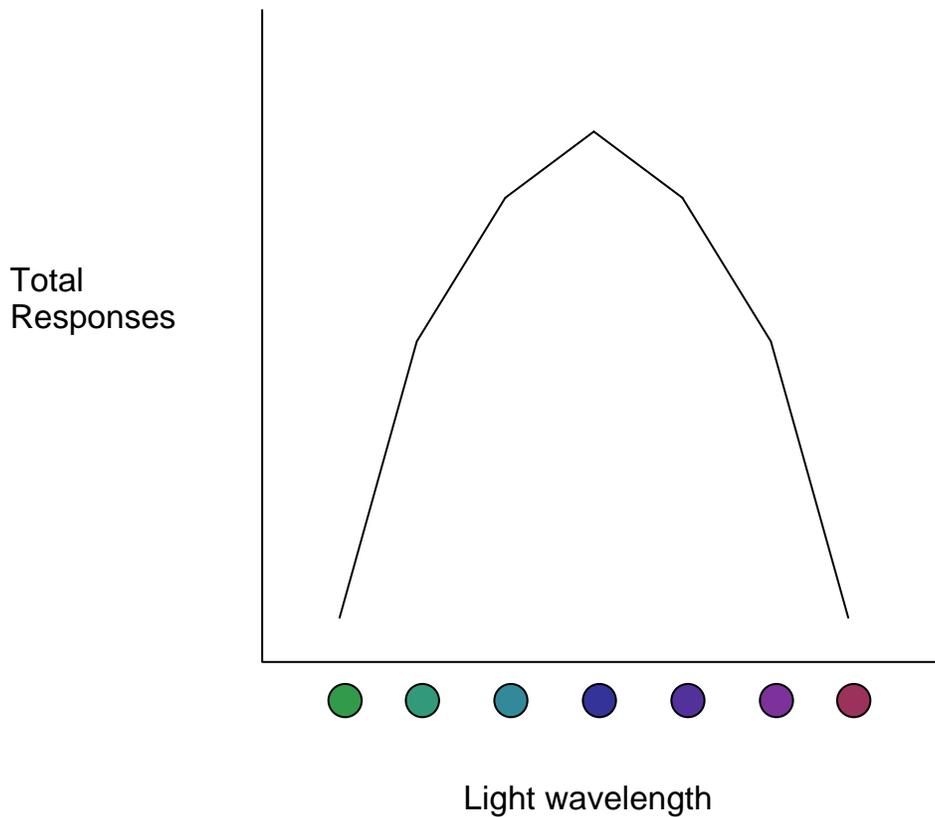


Figure 2. A conceptual example of a generalization gradient.

3.3.1.3 Stimulus classes, prototypes, and schemas

According to the radical behavioral tradition, prototypes can be thought of as examples of a given stimulus class that exhibit the most typical combination of characteristics that are exhibited by that class (which itself is determined by reinforcement). A prototypical dog has certain characteristics that differentiate it from other animals. Similarly, schemas are abstracted forms of instances or episodes – they contain the key aspects of the class which differentiate that class from others. Both the notions of prototype and schema allow for additional details to be incorporated as they are encountered. Furthermore, both can be thought of as the core properties of a given concept or context class – the structural features, those characteristics that make the concept or class what it is. Interestingly, radical behaviorists maintain that stimulus classes (a.k.a. prototypes, schemas) have several characteristics, three of which are:

- class boundaries are “fuzzy” such that no one feature or characteristic is necessary for the discrimination of the class from other classes (c.f., Chi, Feltovich, & Glaser, 1981; Spellman & Holyoak, 1996)
- final class characteristics are individualized in that each individual organism has a different history with contingencies that affected the creation of the class for that individual (c.f., Chi, Feltovich, & Glaser, 1981; Ross, 1996)
- the “existence” of a class/prototype/schema is an emergent property of the reinforcement history of the organism – in actuality the organism *behaves as though*

the class actually exists, but there is no one place or central agent that controls, stores, or owns the class/prototype/schema (Donahoe & Palmer, 1994; c.f., Kolodner, 1997).

3.3.1.4 Differential reinforcement, problem solving, and DCD

We know that in problem-solving, and in generally interacting with the world, people appeal to their past experiences in order to determine appropriate behavior for current situations. We also know, from the previously reviewed literature that this phenomenon is not always an explicit and effortful process. The method by which this happens can be conceived of as a more complex example of stimulus generalization and discrimination. If the stimuli in the current environment are sufficiently structurally similar to those encountered during Situation A (and more similar to A than to any other experience), then the behaviors exhibited in the current situation will most probably parallel those exhibited in Situation A. If the consequences of exhibiting those behaviors in Similar-To-A situations are aversive, the person will display what behaviorists call “stimulus discrimination” over a series of similar experiences. That is, the person will begin to determine how the current situation class differs from the class Situation A. When those differences are present in future Similar-To-A situations, they will have an impact on which set of behaviors are chosen – specifically, where the person might exhibit Behavior 1 in Situation A, he will avoid exhibiting Behavior 1 in certain Similar-To-A situations because Behavior 1 brings about unpleasant consequences in Similar-To-A (cf. Butterfield & Nelson, 1989, 1991; Gick & Patterson, 1992).

Interestingly, one recommendation in the literature on analogy and transfer of training is to allow students to make and correct their own errors in order to enable them to make more finely tuned discriminations regarding when particular solutions are and are not appropriate (Butterfield & Nelson, 1989, 1991; Ghodsain, et al., 1994). This recommendation, in essence, allows the student to develop a unique reinforcement history for solving problems in the domain of interest.

3.3.1.5 Context discrimination and non-humans

Up to now, we have not explicitly stated that DCD is uniquely human. However, in most areas of cognitive research, this assumption is implicit as humans are the only subjects studied. This assumption is important because it implies that DCD is, evolutionarily speaking, a relatively late-arriving feature of our cognitive systems. Interestingly, there is evidence from the radical behavioral literature that this kind of ability, at least at some level, is not unique to humans, primates, or even to mammals.

Zentall and his colleagues (discussed in Zentall, 2003; see also Capaldi & Neath, 1995) have found evidence for social learning by imitation in non-mammals. For example, when a bird (Bird A) watches another bird (Bird B) receive reinforcement for stepping on a treadle, Bird A’s instances of subsequently stepping on the treadle are much higher if Bird B was reinforced (i.e., given a food pellet for stepping on the treadle). In order for this to happen, Bird A must (at some level) recognize the similarity of itself to Bird B, and it must be able to recognize that it can produce behavior that is similar to Bird B’s behavior. These data are particularly compelling considering the behaviors Zentall and his colleagues studied are *opaque*. The way that stepping on a treadle appears to a quail

when it is watching another bird is vastly different from the way that same behavior looks to the bird when it is performing the behavior itself. This imitation behavior will occur even if the observation period and the observer performance is deferred up to 30 minutes in Japanese quail (Zentall, 2003; see also Herrnstein & Loveland, 1964 for additional non-human examples).

3.3.2 Language comprehension and production

The use of language is a quintessential human capability – to our knowledge, no other organism has as broad and as flexible a communication system as we do. What makes human communication so distinctive is the fact that we can, with little to no difficulty and conscious awareness, produce and comprehend thousands of novel sentences a day. Moreover, we rarely produce sentences that are ungrammatical, but when we do, our audience rarely has difficulty understanding what we intended to say. That is, because of a deep understanding of the mechanics of our native language (or any language in which we are proficient, really), we can combine and recombine individual words and sounds into perfectly understandable, novel strings that impart meaning to the listener. This, in turn, enables the listener to respond appropriately. This ability requires a flexibility that cannot be accounted for by explicit rules. This section will review different aspects of the human ability to use language effectively and flexibly: phenomena that can be interpreted to support the perspective that DCD enables agile use of language. There are a couple of separate problems that listeners must be able to solve if they are to understand spoken language: segmentation of speech into individual words and disambiguation of word meaning.

3.3.2.1 *Speech segmentation*

Listening to an unfamiliar foreign language makes abundantly clear one problem every human being has to solve: how to segment the speech stream into individual sounds and words. This process is inherently dynamic: language, whether written or spoken, is temporal in nature thus the ongoing interpretation of the message being relayed is a function of what has already been communicated and of the incoming stimuli. This is true at the level of the phoneme as well as at the lexical level (Carroll, 1994).

Methods for speech segmentation, as it turns out, seem to rest on particular aspects of the language in question. For example, English is based on differential syllabic stresses, so native speakers of English use the correlations between syllabic stress and word endings to determine word boundaries. French, on the other hand, is based on syllable rhythm; so French speakers use this rhythm to determine word boundaries. Speakers of Japanese use a subsyllabic unit, the *mora*, to determine word boundaries (Otake, Hatano, Cutler, & Mehler, 1993). For example, in one set of experiments, native speakers of French were faster to detect the segment *ba* in the word *ba.llance* than they were in the word *bal.con* (where dots indicate syllable boundaries), but native speakers of English did not show this difference (Vroomen & de Gelder, 1995). English natives, on the other hand, have difficulties when the segment to be identified crosses stress boundaries.

3.3.2.2 *Disambiguation of meaning and lexical access*

Segmenting the speech stream into individual words is not the only problem listeners have to solve. Interpretation of the linguistic message changes over the timecourse of an utterance. Interpretation is determined by several factors including word choice, use of syntax, and the prosodic and intonational characteristics of the stimulus. In addition, conditions under which an utterance is produced have an impact on meaning (Carroll, 1994).

Interestingly, many utterances that are perfectly understandable in the context of a stream of speech are totally unintelligible when isolated. For example, in one compelling experiment, Pollack & Pickett (1964; cited in Carroll, 1994) recorded conversations of women sitting in a soundproof room while waiting for a psychology experiment to begin. When they spliced out individual words from the speech stream, only about 50% were intelligible outside of the context of the conversation, whereas virtually all were intelligible when embedded in the speech stream, demonstrating the powerful influence of overall contextual understanding on ongoing language comprehension.

Not only does the intelligibility of individual words change with the context – so can the interpretation of the meaning of individual words. Take for example the effects of word choice on meaning for two sentences that follow (taken from Spivey-Knowlton & Sedivy, 1995). The stem of the sentence is the same for both: *The Arapahoe Indian shot the cowboy with...* This stem can be completed several different ways, depending on the choice of words that complete the final prepositional phrase. Only once that final phrase is complete will the meaning of the entire sentence become clear.

1. The Arapahoe Indian shot the cowboy with the bow and arrow.
2. The Arapahoe Indian shot the cowboy with the leather vest.

As the sentences unfold, the way that the final prepositional phrase is interpreted changes the interpretation of the entire sentence. Interestingly, sentence 1 can be interpreted in two ways: a bow and arrow was used to shoot the cowboy, or the cowboy with the bow and arrow was the one who was shot. However, people prefer the former interpretation (Spivey-Knowlton & Sedivy, 1995). The fact that people tend to prefer to affiliate the prepositional phrase at the end of the sentence (as in sentence 1) rather than with the noun phrase (as in sentence 2), but that we can quite rapidly reinterpret sentence 2 so that it makes sense demonstrates the flexible and dynamic nature of language comprehension.

In addition to word choice and syntax, prosody and intonation are two cues used by both infants and adults to aid language comprehension. Using prosodic cues, young infants can distinguish between languages that have different prosodic rhythms, but not between languages that have similar rhythms (Cutler, 2001). In non-infants, prosody also contributes to understanding the intent of the speaker, especially in infant-directed speech (Fernald, 1989; Hofstadter, 2001). However, these effects are not limited to infant-directed speech (e.g., say aloud the phrases *LIGHThouse keeper* and *LightHOUSE keeper*, placing emphasis on the words in capital letters; Carroll, 1994; Vroomen & de Gelder, 1995). As Hofstadter asserts, there are “probably” situations and then there are “probab-*lee*” situations (Hofstadter, 2001 p. 507).

Contextual influences on language comprehension are also found in categorization of phonemes (Berko-Gleason, 1993; Menn & Stoel-Gammon, 1993). Speech stimuli with identical physical characteristics are perceived as entirely different sounds depending on the context in which they occur and on the native language of the speaker. For example, /b/ is heard as the letter “b” by native speakers of English and as the letter “v” by native speakers of Spanish.

As with context interpretation, language comprehension is both a bottom-up and top-down process. When top-down influences are strong, the sensory stimuli are more quickly recognized than when top-down influences are weak. This is true for both spoken and written language (Carroll, 1994; Spivey-Knowlton & Sedivy, 1995). This is also true for assigning particular meanings to more abstract words (such as a *container* that holds apples versus a *container* that holds soda) and for later recall of a target word (Carroll, 1994). Overall, language comprehension can be thought of as a class of context discrimination problems – determining the intended meaning of the speaker and formulating an appropriate response relies upon the listener’s ability to attend to and accurately interpret various lexical and prosodic cues in the speech stream. Misunderstandings occur when errors in context discrimination are made.

3.3.2.4 *Non-human language capabilities*

As previously mentioned, one very difficult problem that infants must solve in order to acquire their native language is the ability to segment the continuous speech stream into phrases and words and to learn syntactic structure from the incomplete, or sparse input to which they are exposed. These problems has been used by some theorists to argue for a specific language acquisition-type device because the problem seemed intractable for general learning mechanisms (Bates & Elman, 1996; Pinker, 1991, 1994). However, the *less is more* hypothesis states that it is the limited nature of infant’s cognitive capabilities that enable them to perform this parsing task, rather than some specialized device that atrophies after a certain age (Newport, 1988, 1990; see Elman, 1993 for a related computational model). If there is, indeed, a human-specific language acquisition device, we should not expect to see similarities between the discriminations that human infants can make and those that other non-humans can make. However, Hauser, Newport, & Aslin (2001) have demonstrated that non-human primates are able to make some of the same statistical discriminations in a continuous speech stream that human infants are able to make (Saffran, Aslin, & Newport, 1996; see also Morgan, Meier, & Newport, 1987). Specifically, cotton-top tamarins are able to make distinctions between syllables that commonly occur together and those that don’t, mimicking the human ability to distinguish word boundaries in continuous speech based solely on the statistical regularities of syllable-syllable co-occurrences⁸. By no means is this ability sufficient to enable tamarins to acquire human language, but this ability is a crucial building block.

⁸ Interestingly, these same kinds of regularities aid in parsing other input like music, visual patterns, and visuomotor sequences (see Hauser, Newport, & Aslin, 2001 for details).

The compelling thing about these findings is that some of the linguistic components cited as examples of DCD show up in non-humans. These findings provide additional evidence that:

1. DCD and its components are likely not late-arriving in evolutionary terms.
2. Many of the human cognitive processes that we have explained via DCD are not likely to be due to specialized cognitive mechanisms. Rather, they more likely arise from general aspects of the brain, which we hypothesize to be the physiological characteristics of the brain.

This implies that DCD is, indeed a core rather than a peripheral process in human cognition.

3.3.2.5 Summary

Language comprehension and production are two of the most overtly dynamical things humans do. Language is of interest to DCD because it shares many of the same characteristics of DCD with other cognitive phenomena – a compelling finding because language has received somewhat special treatment in the literature. As will be discussed later, language is also of interest to DCD because the process by which we learn our primary language seems to be unique – there is not another obvious example of how we learn to do anything else in quite the way we acquire language. If we can understand how we acquire language, we can gain insight into the way we learn to interact with the world in other, non-linguistic ways.

3.4 Summary

The upshot of the previous sections is this: Because the concept of DCD continues to appear in many guises across many literatures, between which there is often little or no communication, there must be something about DCD that is core to human functioning – especially given the fact that many of the abovementioned subdisciplines often don't agree with one another *philosophically* about what it means to be human and about what is and is not within the purview of scientific inquiry. Yet, evidence of DCD, whatever its various names, continues to appear – across philosophies, across paradigms, across disciplines from low-level perception to high-level cognition.

However, DCD isn't just about passively recognizing patterns in the environment – it is about recognizing general patterns and *regularizing* those patterns: taking multiple specific instances and making an averaged, abstracted, or amalgamated version out of them – it can be thought of as a two-phase process (which are not necessarily correlated with explicit engines or modules in the brain). It is these generalities that then allow us to make inferences about specific instances that seem to match a given generality, thereby enabling us to rapidly deal with novel situations that match, to a greater or lesser degree, things we've encountered in the past.

One interesting question arises from this, then: how does this regularizing process happen? Should that process be considered separately from DCD, or is it an integral part of DCD that can't be separated? As mentioned in the section on memory, one possible explanation is that this regularization phenomenon is a result of the fact that memory is

reconstructive rather than reproductive. That is, we recall the gist of events relatively accurately, then fill in the details with contextually-guided guesses of what they most likely were (Roediger & McDermott, 1995).

4. Dynamic Context Discrimination: The computational implementation

Now that we have presented an overview of converging evidence for DCD, we present the computational implementation of DCD.

The Sandia framework is based on the notion that if we are able to create a computational model of cognition that is both psychologically and physiologically plausible, we can fundamentally transform the way humans and machines interact. It is our contention that cognitive architectures created to date are not as fundamentally plausible as the framework we have been building over the last 5 years. This section introduces the Sandia framework, citing the empirical literature referenced in building the framework.

4.1 Psychological inspiration for the model

Initial inspiration for the Sandia framework came from the literature on naturalistic decision-making (Klein, 1997; Klein, Calderwood, & Clinton-Cirocco, 1985). This notion of decision-making holds that experts don't consider multiple options when faced with a decision – especially in time-critical situations as had long been assumed in traditional models of decision-making (Baron, 1994). Rather, experts are able to quickly assess the nature of the situation based on patterns of cues in the environment. Once a pattern is recognized, a solution procedure that the decision-maker learned through experience in conjunction with that pattern is implemented (Chi, Feltovich, & Glaser, 1981; Klein, 1997; Kolodner, 1997; Ross, 1996). Therefore, decision-making can be characterized as a highly individualized dynamic interaction between environmental stimuli and memory for prior experiences (Kolodner, 1997; Ross, 1996). Interestingly, empirical work in this vein demonstrated that solutions are not always optimal, but they are sufficient the vast majority of the time (Klein, 1997).

If expertise in a domain is conceptualized as knowing and being able to discriminate among a proportionally large number of seemingly similar stimulus patterns, this method of decision-making might be true for those of us not considered experts in a specialized domain. As we move through life, we learn about an increasing number of contexts/patterns/problems and corresponding solutions, or appropriate behaviors. We are then able to discriminate between these contexts with increasing refinement by appealing to what we've encountered in the past. All of us are, in a sense, experts in everyday functioning.

The primary problem with using the NDM literature as the sole psychological foundation for the model is that this literature is primarily based on anecdotal and case study evidence. However, when considered in conjunction with the literatures reviewed earlier, this notion of expert decision making becomes quite compelling – and contributes additional convergence on the notion of DCD as a core psychological process.

4.2 Knowledge representation in the model: Physiological inspiration

In order to implement the DCD account of decision-making (and general human cognitive functioning) in computational form, the issue of knowledge representation had to be addressed⁹. Because we believe that cognitive behavior emerges from neurophysiology, we turned to a neurophysiological theory of representation that originated in research on sensory systems – oscillating systems theory, which has more recently been applied to higher-level cognitive functions (Klimesch, 1996, 1999; Singer 1993). Briefly, oscillating systems theory holds that knowledge representation in the brain is distributed across the cortex, such that complete information about a given concept, item, or event is not encapsulated in a given neuron. Rather, when a stimulus is present in the environment, individual neurons begin to fire in temporal synchrony with one another, and out of temporal synchrony with everything else going on in the brain, forming what is called a *neural assembly*. Membership in a neural assembly is fairly probabilistic – any given individual neuron may or may not fall into temporal synchrony with the other neurons responding to a particular stimulus on a given presentation depending on the current context. However, in general, the same population of neurons does respond with a particular rate of firing that is distinct from the rates of firing of other assemblies (Klimesch, 1996; 1999).

As mentioned in the section on the neurophysiology of sensation, the individual neurons that make up these neural assemblies respond selectively to different types of information including very low-level perception (e.g., horizontal lines in the environment) to higher-level, more abstract concepts that themselves comprise low-level perceptive information. In this way, the information about a particular experience (e.g., eating an apple) that is arriving in the various distinct sensory cortices in the brain (e.g., red and round in the visual cortex, weight and movement in the somatosensory cortex, crunchy sound in the auditory cortex, sweetness in the gustatory cortex) is bound together, via in-phase oscillations of individual neurons across the various sensory cortices, to compose the singular phenomenological experience of eating an apple - the neurons responding to each of the different sensations oscillate as an assembly.

⁹ Knowledge representation has been given a huge amount of press and has been a source of controversy in the cognitive sciences for some time (Bechtel, 1998; Billman, 1999; Dietrich & Markman, 2003; Holyoak & Hummel, 2001; Markman, 2000, 2002; Markman & Dietrich, 2000; Shastri, 1999; Stufflebeam, 1999). At issue is whether the human system contains representations and, if so, what the nature of those representations is - specifically if they are symbolic or distributed in nature. Symbolic systems include production systems but are not limited productions (Anderson, Bothell, Byrne, & Lebiere, 2002; Falkenhainer, Forbus & Gentner, 1989; Holyoak & Hummel, 2000; Holyoak & Thagard, 1995; Keane, Ledgeway, & Duff, 1994; Kokinov, 1994 see also Markman & Dietrich, 2000). Other researchers represent information in a distributed fashion, as in artificial neural networks (McLeod, Plunkett & Rolls, 1998; Plunkett & Marchman, 1993). A third group goes so far as to say that there is no representation in the brain at all (Bechtel, 1998; van Gelder, 1998). However, as will be explained shortly, we adopt a localist representation scheme that assumes distributed representations that underly the current localist implementation (see, e.g., Hummel & Holyoak, 1997 for another example of this strategy).

It is at the level of neural assemblies that the Sandia framework is physiologically based - and this helps keep the model implementation practical. The individual nodes that are implemented in our model conceptually represent assemblies of neurons rather than individual, low-level perceptive and sensory neurons. In this way, despite the fact that the Sandia Framework is currently a localist model, we acknowledge the fact that neurophysiological evidence indicates that knowledge representation is distributed in the brain¹⁰.

4.3 Model Components

This section outlines the basic separable components of the framework. The next section, 4.4 outlines how these components function together.

The basic framework comprises two primary components:

- An associative network made up of concept nodes
- A library comprising knowledge of generic contexts and specific episodes that are defined by unique patterns of activation on the associative network

There are additional components to a full framework, but for purposes of demonstrating the computational implementation of DCD, detailing the basic framework is sufficient.

4.3.1 The associative network

In order for the model to be able to interact with the stimulus environment, there had to be a method by which we represented basic-level information - concepts that would be combined and recombined to define different contexts. However, we recognize that the specific meaning of concepts change with the context in which they occur (e.g., think of the many uses of the word “run”). Therefore, we had to represent information in a “fuzzy” manner. As the model currently stands, this occurs in the associative network. This network comprises nodes, each of which is representative of a concept. These nodes occur as vectors in high-dimensional space (i.e., 100-200 dimensions). In this way, we are able to represent the concept “run” with as many subtly different meanings as we need to by representing those different meanings as slightly different points in that high-dimensional space (cf. Elman, 1995).

Each of these nodes is also representative of a neural assembly, as mentioned above, and therefore each node has the ability to oscillate at a frequency that is determined by a variety of factors including the strength of the stimulus, the current state of the model, decay rates, affective salience of the cues, and the recent state or states of the model. Generally, the oscillation rates of nodes in the associative network fall in the 10-13 Hz upper alpha bandwidth (Klimesch, 1999).

¹⁰ Interestingly, dynamical systems interpretations of cognition have pointed out that explanations at a medium level are lacking in the explanation of human behavior (Petitot, 1995). At the micro level, there are explanations involving neurons, synapses, and neurotransmitters. At the macro level, there exist the more phenomenological and observable behavioral explanations of psychology. However, there is not a strong link between these two explanatory paradigms. This is precisely the level of explanation addressed by the Sandia framework.

Each node has a level of activation associated with it. That level of activation is determined by:

1. the salience of the stimulus in the environment, which determines initial level of activation
2. its semantic relationship with other nodes in the network, the strength of that relationship, and the related nodes' respective levels of activation through *spreading activation* (Anderson, 1983)
3. affective salience such that those cues that are affiliated with strong emotions in given contexts are more highly activated when observed in those contexts
4. the current context – each context carries with it expected cues as well as expected future events that exert *top-down* influence on activation in the associative network
5. a static decay¹¹ rate

Each of these factors interacts with each other factor to dynamically determine a node's individual level of activation at each cycle of the model. Once that level of activation surpasses a predetermined threshold, the node begins to oscillate. Oscillation rate is then determined dynamically by the node's level of activation, given that the level of activation remains above the threshold. Currently, there are no inhibitory processes in the model. Therefore, the only way that a node will cease oscillating is if the decay rate is greater than all of the activation acceleration factors listed above.

4.3.2 The Context/Episodic memory

The Context/Episodic memory is the sum of all of the situations the model knows about – both single instances (i.e. episodes) as well as abstracted versions of events (i.e., schema, contexts). Each situation is defined by a pattern of cues such that as stimuli come in to the model from perceptive mechanisms and activate nodes in the associative network to greater or lesser degrees, the model is able to perform context discrimination dynamically as these activation patterns change over time.

It is very important to note that despite a formal definition of a context as a pattern of activation on the associative network, there does not have to be a perfect match in order for the model to recognize that context. Rather, the model operates on an evidence accumulation process (Coles, et.al., 1995; Kok, 1990; Kounios, et.al., 1994; Wilding, 2000; cf. Koehler, White, & Grondin, 2003; Stein, Wallace & Stanford, 2000), whereby any number of “candidate” contexts can have some evidence, as determined by their similarity to the current pattern of activation on the associative network. However, only the contexts that have sufficient evidence to surpass an empirically-determined threshold can actually be recognized as occurring. The implication of this final statement is that the model can recognize the presence of multiple contexts simultaneously.

4.4 Information processing in the Sandia framework

The method by which the Sandia framework processes information is fairly straightforward. Figure 3 conceptually illustrates the model. At its most basic, the model

¹¹ Decay, in this context, does not refer to the disappearance of memory over time, or forgetting. Rather, decay in this context refers to the rate at which the nodes would “spin down” in the absence of any other stimulation either from the environment or from spreading and top-down activation.

has three components (shown in blue): an associative network comprising nodes that represent domain-relevant concepts, a “library” of known contexts each of which is defined by a particular pattern of activation across the associative network, and the pattern recognition algorithms (which are, in reality, the information populating the library of contexts).

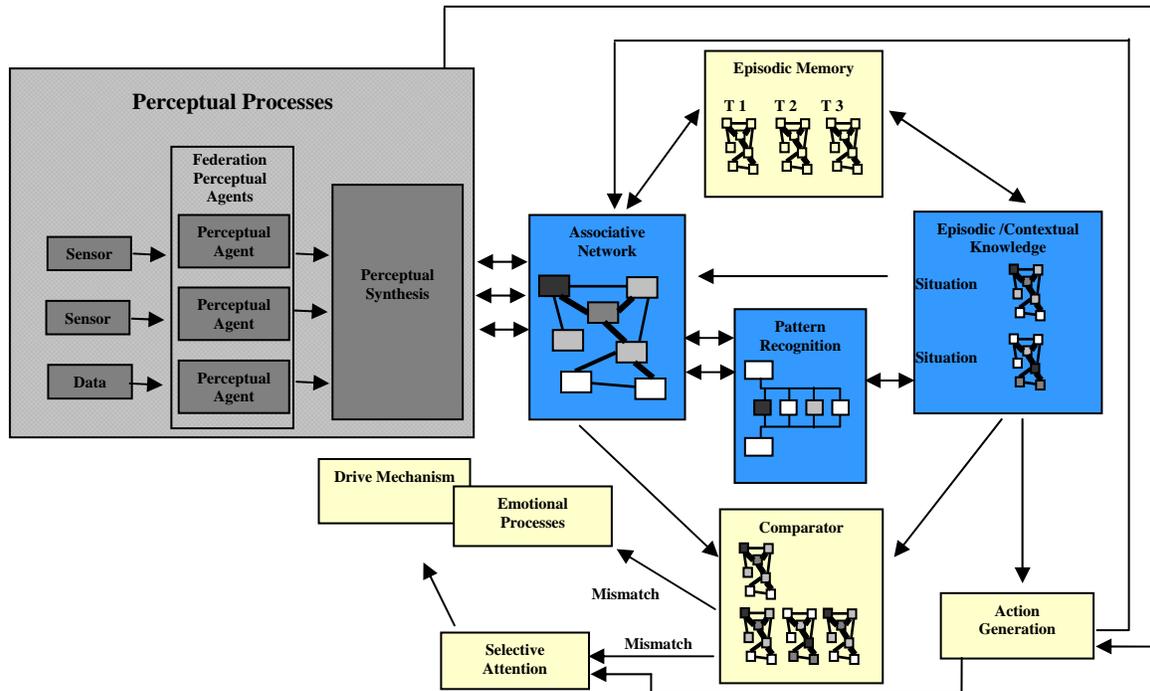


Figure 3. Conceptual diagram of the Sandia Cognitive Framework.

As environmental stimuli are presented to the model (via the perceptive processes which are not currently part of the architecture), nodes in the associative network begin to oscillate at varying rates depending on a variety of factors including current state of the model and salience of the stimulus in the environment. As these concept nodes begin to oscillate, evidence for different contexts in the context library is accumulated through a weighted sum of the activations on the associative net (cf. Koehler, White, & Grondin, 2003; Stein, Wallace & Stanford, 2000). Once a given context, or contexts, accumulates enough evidence to surpass a pre-set threshold, the model recognizes that context as being the current state of affairs. Recognition of a context then implicates an appropriate course of action and defines the current state of the model, including expectations about stimuli that should and should not be present in the stimulus environment, which then impacts the rates of oscillation of nodes in the associative network. In this way, even though situation recognition is a very stimulus-driven process, the framework also accounts for the role of schemata and top-down activation on ongoing cognition and situation interpretation (cf. Brewer & Treyns, 1981; Chi, Feltovich, & Glaser, 1981). That is, depending on the current state of the model, it has expectations about what is more and less likely to happen next which influences perceptions of the environment¹². In

this way, cognition is a dynamical process between perception of environmental stimuli and expectations about what stimuli should be present based on current contextual understanding (van Gelder, 1998)¹³.

In the event that the model encounters a novel situation – that is, a novel pattern of activation occurs across the associative network – the model attempts to relate that pattern to the most similar known context via the evidence accumulation process. However, if the pattern is totally new, that is, if no known context receives enough evidence to surpass the recognition threshold, this new pattern becomes a new episode in the context library. If this new pattern is encountered multiple times, a new schema, or abstracted version of the episode is created.

4.5 Appearance of specific psychological functions in the model

Briefly, there are several psychological phenomena that are captured in this implementation in the absence of explicit modules dedicated to those phenomena.

- Working, short-term and long-term memory are captured in the associate network and the context library simultaneously. The phenomenological experience of working memory is represented as those nodes and contexts that are currently active. Short-term memory is represented via the priming and decay mechanisms in the networks. Long-term memory is represented as those nodes and contexts that are not currently active.
- Analogy-making is captured in part via the evidence accumulation process – specifically, evidence accumulation allows for retrieval from memory of a potential analogue to the current situation.
- Creation of genericized versions of events, or schema, is captured via creating new contexts in the context library and via the modification of those contexts as additional experiences are gathered. This is one very important and fundamental method of learning in both humans and non-humans.
- Awareness of the fact that a single stimulus or pattern of stimuli are novel is captured indirectly via the evidence accumulation process in that if there is not enough evidence for any currently known context to be recognized, that pattern of activation on the semantic space is recorded in context/episodic memory as a unique event.

5. General Discussion

The purpose of this paper was to give a brief overview of the psychological and neurophysiological literature that supports the architecture of the Sandia Cognitive Framework and then to outline the basics of that framework.

The core points of the discussion are:

¹³ We are making a specific distinction between sensation and perception, in the classic psychological sense. Sensation is the method by which our sensory organs (e.g., ears, eyes) detect stimuli in the environment. Perception is the method by which our brains make sense of those sensory inputs (refs here). Sensation is not necessarily influenced directly by cognition (with the exception that attentional shifts may remove a stimulus from the reach of a sensory organ), but interpretation of those sensory inputs, i.e. perception, is definitely open to influence by cognition.

1. The Sandia Framework approaches the issue of system architecture as a reverse engineering problem – if we want to model the human, we have to pay close attention to the human system.
2. Specifically, the behavior of the human system (the behavioral evidence) *is a consequence of* the neurophysiological construction of the system. If the system were physically different, so would be the resulting behaviors.
3. The empirical evidence used to directly inform the architecture of the Sandia Framework is only a small part of the literature available that bolsters the basic idea about the behavior of the model. Dynamic Context Discrimination is a behavioral phenomenon observed in a wide variety of psychological literatures.

In terms of the future of the Framework, one key aspect that has yet to be directly addressed in any implementation is the issue of learning. Up to now, we've talked only briefly about learning and how that might be instantiated in the framework. However, there is one clear distinction between the work that we have done thus far and the way that humans function: how might a beginning cognitive framework learn about it's world *in the same way human infants do?* This is an interesting question, because at present time the methods for framework population we use are clearly different from the ways that children learn about their environment. Our methods of creating an initial model of an individual are not psychologically, physiologically, or behaviorally plausible, but eventually we would like to be able to mimic this aspect of human cognition as well. One candidate model for the development of this functionality is childhood primary language acquisition.

Specifically, a group of researchers has developed theories about the emergence of language as a result of the complex interaction between the stimulus input and the maturational constraints of infants and young children (Elman, 1993, 1995, 1999; Newport, 1988, 1990; see Klimesch, 1999 for a potential electrophysiological correlate of these constraints). Through simulations based on simple recurrent connectionist networks, mechanisms that are dynamical and physiologically plausible (from both neurophysiology and evolutionary perspectives) have been demonstrated that closely mimic human data (McLeod, Plunkett & Rolls, 1998).

Elissa Newport formulated what she called the *less is more* hypothesis to account for the observed critical period for attaining native-like proficiency in a given language. Basically, she hypothesized that primary language acquisition is not due to some explicit device (Bohannon, 1993; cf. Pinker, 1991, 1994), but that it is an emergent consequence of the fact that infants have fewer cognitive resources than do adults and that this limitation changes the nature of the information infants can attend to (Newport, 1988, 1990; see also Elman, 1999 for similar ideas and Klimesch, 1999 for a potential electrophysiological source for these constraints). So, very early, children's working memories can only keep track of very small bits of information. This essentially turns anything but the most simplistic input into noise, allowing infants to detect regularities in these small bits of information that would be lost on adults. Then, as children mature and their cognitive resources become more numerous, they can utilize the information they learned earlier as building blocks for detecting regularities occurring over longer and

longer timeperiods. Characteristics of the speech stream, such as prosody and intonation, provide the earliest regular cues that allow infants to segment the speech stream into individual words and phrases (Cutler, Mehler, Norris & Segui, 1992; Hirsh-Pasek, Nelson, Jusczyk, Cassidy, Druss, Kennedy, 1987; Mehler, Jusczyk, Lambertz, Halsted, Bertoncini, & Amiel-Tison, 1988; Nelson, Hirsh-Pasek, Jusczyk, & Cassidy, 1989), followed by phonological and morphological regularities. Eventually, building upon these basics of language, children are able to implicitly determine parts of speech and the other grammatical aspects of their native tongue.

Using simple recurrent networks, Elman (1993) has demonstrated that when working memory constraints are included in these models, the network not only learns the grammars dictating the construction of simple and complex sentences, they also are able to generalize this knowledge to novel sentences (Elman, 1993). Lewis and Elman (2001) demonstrated that a simple recurrent network can learn complex grammatical constructions using the statistical regularities of input that mimics that heard by children. In addition, once the networks were trained, they exhibited large error rates in predicting upcoming words in novel test sentences when those upcoming words were ungrammatical. Furthermore, the errors made by the networks in this set of experiments are very similar to those made by children.

Whether imposing these kinds of constraints on a complex enough model could help lead to a computational entity forming as complete an understanding of the world as a human does is a question, fortunately, that still resides somewhat more in science fiction than in science. However, it is important to realize that adult understanding of the world does not simply appear in the brain – it is learned through experience. And the kinds of experiences that contribute to this learning interact with the architecture of the system that is doing the learning, which is biologically determined, to produce a unique individual with a unique perspective on the world. Even if the experiment is never tried, keeping the interaction between environment, biology, and physiological maturation in mind still has the potential to enhance the functioning of a computational modeling effort.

6. References

- Amir, N., Bower, E., Briks, J., & Freshman, M. (2003). Implicit memory for negative and positive social information in individuals with and without social anxiety. Cognition and Emotion, *17*, 567-583.
- Anderson, J.R. (1983). *The Architecture of Cognition*. Lawrence Earlbaum: Mahwah, NJ.
- Anderson, J.R., Bothell, D., Byrne, M.D., Lebiere, C. (2002). An integrated theory of the mind. *Manuscript submitted to Psychological Review*, available at <http://act-r.psy.cmu.edu/people/ja/ja-vita#pubs>.
- Baron, J. (1994). Thinking and Deciding, 2ndEd. New York, NY: Cambridge University Press.
- Bates, E. & Elman, J. (1996). Learning rediscovered. Science, *274*, 1849.
- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist's challenge in cognitive science. Cognitive Science, *22*, 295-318.
- Berko-Gleason, J. (1993). Language development: An overview and a preview. 3rd Ed. In J. Berko-Gleason (Ed.) The Development of Language, 2-37. New York, NY: Macmillan.
- Billman, D. (1999). Representations. In W. Bechtel & G. Graham (Eds) A Companion to Cognitive Science, 649-659. Maldin, MA: Blackwell.
- Blanchette, I. & Richards, A. (2003). Anxiety and the interpretation of ambiguous information: Beyond the emotion-congruent effect. Journal of Experimental Psychology: General, *132*, 294-309.
- Bohannon III, J.N. (1993). Theoretical approaches to language acquisition. In J. Berko-Gleason (Ed.) The Development of Language, 240-297. New York, NY: Macmillan.
- Brewer, W.F. & Treyens, J.C. (1981). Role of schemata in memory for places. Cognitive Psychology, *13*, 207-230.
- Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch and B.B. Lloyd (Eds.) Cognition and Categorization. Hillsdale, NJ: Earlbaum.
- Butterfield, E.C. & Nelson, G.D. (1989). Theory and practice of teaching for transfer. Educational Technology Research and Development, *37*, 5-38.
- Butterfield, E.C. & Nelson, G.D. (1991). Promoting positive transfer of different types. Cognition and Instruction, *8*, 69-102.
- Capaldi, E.J. & Neath, I. (1995). Remembering and forgetting as context discrimination. Learning & Memory, *2*, 107-132.
- Carroll, D.W. (1994). The Psychology of Language, 2nd Ed. pp. 106-209. Pacific Grove, CA: Brooks/Cole Publishing Company.
- Catrambone, R., & Holyoak, K.J. (1989). Overcoming contextual limitations on problem-solving transfer. Journal of Experimental Psychology: Learning, Memory, and Cognition, *15*, 1147-1156.
- Chase, W.G. & Simon, H.A. (1973). Perception in chess. Cognitive Psychology, *4*, 55-81.
- Chi, M.T.H., Feltovich, P.J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. Cognitive Science, *5*, 121-152.

- Cimbalo, R.S., Capria, R.A., Neider, L.L., & Wilkins, M.A.C. (1977). Isolation effect: Overall list facilitation in short-term memory. Acta Psychologica, 41, 419-432.
- Coles, M.G.H., Snid, H.G.O.M, Scheffers, M.K. & Otten, L.J. (1995). Mental chronometry and the study of human information processing. Electrophysiology of mind: Event-related brain potentials and cognition, Oxford: Oxford Science Publications, 86-131.
- Cutler, A. (2001). Prosody, intonation, and processing issues. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 682-683.
- Cutler, A., Mehler, J., Norris, D., & Segui, J. (1992). The monolingual nature of speech segmentation by bilinguals. Cognitive Psychology, 24, 381-410.
- Dietrich, E. & Markman, A.B. (2003). Discrete thoughts: Why cognition must use discrete representations. Mind & Language, 18, 95-119.
- Donahoe, J.W. & Palmer, D.C. (1994). Learning and Complex Behavior. Boston: Allyn & Bacon.
- Dowens, M.G., & Calvo, M.G. (2003). Genuine memory bias versus response bias in anxiety. Cognition and Emotion, 17, 843-857.
- Dunbar, K. (2001). The analogical paradox: Why analogy is so easy in naturalistic settings, yet so difficult in the psychological laboratory. In D.Gentner, K.J. Holyoak, and Kokinov, B.N. (Eds) The Analogical Mind: Perspectives from Cognitive Science. 313-334. Cambridge, MA: MIT Press
- Eliasmith, C. (2002). The myth of the Turing machine: The failings of functionalism and related theses. Journal of Experimental and Theoretical Artificial Intelligence, 14, 1-8.
- Ellwart, T., Rinck, M., & Becker, E.S. (2003). Selective memory and memory deficits in depressed inpatients. Depression and Anxiety, 17, 197-206.
- Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. Cognition, 48, 71-99.
- Elman, J.L. (1995). Language as a dynamical system. In R.F. Port and T. van Gelder (Eds.) Mind as Motion: Explorations in the Dynamics of Cognition, 197-225. Cambridge, MA: MIT Press.
- Elman, J.L. (1999). Origins of language: A conspiracy theory. In B. MacWhinney (Ed.) The Emergence of Language. Hillsdale, NJ: Lawrence Earlbaum Associates.
- Elman, J.L. (in press). Generalization from sparse input. Proceedings of the 38th Annual Meeting of the Chicago Linguistic Society.
- Falkenhainer, B. Forbus, K.D. & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. Artificial Intelligence, 41, 1-63.
- Feldman, J. (2003). The simplicity principle in human concept learning. Current Directions in Psychological Science, 12, 227-232.
- Fernald, A. (1989). Intonation and communicative intent in mothers' speech to infants: Is the melody the message? Child Development, 60, 1497-1510.
- Forbus, K.D. (2001). Exploring analogy in the large. In D.Gentner, K.J. Holyoak, and Kokinov, B.N. (Eds) The Analogical Mind: Perspectives from Cognitive Science. 23-58. Cambridge, MA: MIT Press.
- Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.

- Gentner, D. & Markman, A.B. (1997). Structure mapping in analogy and similarity. American Psychologist, *52*, 45-56.
- Ghodsain, D., Bjork, R.A., & Benjamin, A.S. (1997). Evaluating training *during* training: Obstacles and opportunities. In M.A. Quinones & A. Ehrenstein (Eds.). Training for a Rapidly Changing Workplace: Applications of Psychological Research. 63-88. Washington DC: American Psychological Association.
- 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.
- Gick, M.L., & Holyoak, K.J. (1987). The cognitive basis of knowledge transfer. In S.M. Cormier & J.D. Hagman (Eds.). Transfer of Learning: Contemporary Research and Applications. 9-46. New York, NY: Academic Press.
- Gick, M.L., & Patterson, K. (1992). Do contrasting examples facilitate schema acquisition and analogical transfer? Canadian Journal of Psychology, *46*, 539-550.
- Gilboa-Schechtman, E., Erhard-Weiss, D., & Jeczemien, P. (2002). Interpersonal deficits meet cognitive biases: Memory for facial expressions in depressed and anxious men and women. Psychiatry Research, *113*, 279-293.
- Glenberg, A.M. (1997). What memory is for. Behavioral and Brain Sciences, *20*, 1-55.
- Glenberg, A.M. & Swanson, N.G. (1986). A temporal distinctiveness theory of recency and modality. Journal of Experimental Psychology: Learning, Memory and Cognition, *12*, 3-15.
- Godden, D.R. & Baddeley, A.D. (1975). Context-dependent memory in two natural environments: On land and underwater. British Journal of Psychology, *66*, 325-331.
- Greene, R.L. (1992). Human Memory. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Halpern, B. (1999). Taste. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 826-828.
- Hauser, M.D., Newport, E.L., & Aslin, R.N. (2001). Segmentation of the speech stream in a non-human primate: Statistical learning in cotton-top tamarins. Cognition, *78*, B53-B64.
- Herrnstein, R.J., & Loveland, D.H. (1964). Complex visual concept in the pigeon. Science, *146*, 549-551.
- Hirsh-Pasek, K., Nelson, D.G.K., Jusczyk, P.W., Cassidy, K.W., Druss, B., Kennedy, L. (1987). Clauses are perceptual units for young infants. Cognition, *26*, 269-286.
- Hofstadter, D.R. (2001). Epilogue: Analogy as the core of cognition. In D.Gentner, K.J. Holyoak, and Kokinov, B.N. (Eds) The Analogical Mind: Perspectives from Cognitive Science. 499-538. Cambridge, MA: MIT Press.
- Holyoak, K.J., Gentner, D. & Kokinov, B.N. (2001). Introduction: The place of analogy in cognition. In D.Gentner, K.J. Holyoak, and Kokinov, B.N. (Eds) The Analogical Mind: Perspectives from Cognitive Science. 1-19. Cambridge, MA: MIT Press.

- Holyoak, K.J. & Hummel, J.E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich & A. Markman (Eds) Cognitive Dynamics: Conceptual Change in Humans and Machines, pp. 229-263. Mahwah, NJ: Earlbaum.
- Holyoak, K.J. & Hummel, J.E. (2001). Toward an understanding of analogy within a biological symbol system. In D. Gentner, K.J. Holyoak, and Kokinov, B.N. (Eds) The Analogical Mind: Perspectives from Cognitive Science. 161-195. Cambridge, MA: MIT Press.
- Holyoak, K.J. & Thagard, P. (1995). Mental Leaps. Cambridge, MA: MIT Press.
- Huang, I. & Wille, C. (1979). The von Restorff isolation effect in free recall. The Journal of General Psychology, *101*, 27-34.
- Hummel, J.E. & Holyoak, K.J. (1997). Distributed representations of structure: A theory of analogical access and mapping. Psychological Review, *104*, 427-466.
- Hubel, D.H. & Wiesel, T.N. (1962). Receptive fields, binocular interaction, and functional architecture of the cat's visual cortex. Journal of Physiology, *160*, 106-154.
- Hubel, D.H., & Wiesel, T.N. (1979). Brain mechanisms of vision. Scientific American, September, 150-162.
- John, E.R. (2002). The neurophysics of consciousness. Brain Research Reviews, *39*, 1-28.
- John, E.R. (2003). A theory of consciousness. Current Directions in Psychological Science, *12*, 244-249.
- Keane, M., Ledgeway, T., & Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. Cognitive Science, *18*, 387-438.
- Klein, G. (1997). The Recognition-Primed Decision Model (RPD): Looking back, looking forward. In C.E. Zsombok & G. Klein (Eds.) Naturalistic Decision Making, p. 285-292. Hillsdale, NJ, US : Lawrence Erlbaum Associates, Inc.
- Klein, G., Calderwood, R., & Clinton-Cirocco, A. (1985) Rapid decision-making on the fire ground. (KA-TR-84-41-7). Yellow Springs, OH : Klein Associates Inc. (Prepared under contract MDA903-85-G-0099 for the U.S. Army Research Institute, Alexandria, VA).
- Klimesch, W. (1996). Memory processes, brain oscillations, and EEG synchronization. International Journal of Psychophysiology, *24*, 61-100.
- Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: A review and analysis. Brain Research Reviews, *29*, 169-195.
- Koehler, D.J., White, C.M. & Grondin, R. (2003). An evidential support accumulation model of subjective probability. Cognitive Psychology, *46*, 152-197.
- Kok, A. (1990). Internal and external control: A two-factor model of amplitude change of event-related potentials. Acta Psychologica, *74*, 203-236.
- Kokinov, B.N. (1994). A hybrid model of reasoning by analogy. In K.J. Holyoak and J.A. Barnden (Eds) Advances in connectionist and neural computation theory, vol.2, Analogical Connections, pp. 247-318. Norwood, NJ: Ablex.
- Kolodner, J.L. (1997). Educational implications of analogy: A view from case-based reasoning. American Psychologist, *52*, 57-66.
- Kounios, J., Montgomery, E.C. & Smith R.W. (1994). Semantic memory and the granularity of semantic relations: Evidence from speed-accuracy decomposition. Memory and Cognition, *22*(6), 729-741.

- Lawson, C. & MacLeod, C. (1999). Depression and the interpretation of ambiguity. Behaviour Research and Therapy, 37, 463-474.
- Lewis, J.D. & Elman, J.L. (2001). Learnability and the statistical structure of language: Poverty of stimulus arguments revisited. Proceedings of the 26th Annual Boston University Conference on Language and Development.
- Loftus, E.F. (1975). Leading questions and the eyewitness report. Cognitive Psychology, 4, 560-572.
- Loftus, E.F., & Palmer, J.C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. Journal of Verbal Learning and Verbal Behavior, 13, 585-589.
- Loftus, E. F & Zanni, G. (1975). Eyewitness testimony: The influence of the wording of the question. Bulletin of the Psychonomic Society, 5, 86-88.
- Luchins, A.S., & Luchins, E.H. (1950). New experimental attempts at preventing mechanization in problem solving. The Journal of General Psychology, 42, 279-297.
- Markman, A.B. (2000). In defense of representation. Cognitive Psychology, 40, 138-171.
- Markman, A.B. (2002). Knowledge representation. In D. Medin (Ed.) Stevens' Handbook of Experimental Psychology 3rd Ed. Volume 2: Memory and Cognitive Processes, 165-208. New York, NY: John Wiley & Sons, Inc.
- Markman, A.B. & Dietrich, E. (2000). In defense of representation. Cognitive Psychology, 40, 138-171.
- McLeod, P., Plunkett, K. & Rolls, E.T. (1998). Introduction to Connectionist Modeling of Cognitive Processes. New York, NY: Oxford University Press.
- Mehler, J., Jusczyk, P., Lambertz, G., Halsted, N., Bertoncini, J., & Amiel-Tison, C. (1988). A precursor of language acquisition in young infants. Cognition, 29, 143-178.
- Menn, L., & Stoel-Gammon, C. (1993). Phonological development: Learning sounds and sound patterns. In J. Berko-Gleason (Ed.) The Development of Language, 66-113. New York, NY: Macmillan.
- Morgan, J.L., Meier, R.P., & Newport, E.L. (1987). Structural packaging in the input to language learning: Contributions of prosodic and morphological marking of phrases to the acquisition of language. Cognitive Psychology, 19, 498-550.
- Morris, C.D., Bransford, J.D., & Franks, J.J. (1977). Levels of processing versus transfer appropriate processing. Journal of Verbal Learning and Verbal Behavior, 16, 519-533.
- Nairne, J.S. (2002). Remembering over the short-term: The case against the standard model. Annual Review of Psychology, 53, 53-81.
- Neath, I. & Crowder, R.G. (1990). Schedules of presentation and temporal distinctiveness in human memory. Journal of Experimental Psychology: Learning, Memory, and Cognition, 16, 316-327.
- Nelson, D.G.K., Hirsh-Pasek, K., Jusczyk, P.W., & Cassidy, K.W. (1989). How the prosodic cues in motherese might assist language learning. Journal of Child Language, 16, 55-68.
- Newport, E.L. (1988). Constraints on learning and their role in language acquisition: Studies of the acquisition of American Sign Language. Language Sciences, 10, 147-172.

- Newport, E.L. (1990). Maturation constraints on language learning. Cognitive Science, 14, 11-28.
- Novick, L.R. (1988). Analogical transfer, problem similarity, and expertise. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 510-520.
- Olafson, K.M., & Ferraro, F.R. (2001). Effects of emotional state on lexical decision performance. Brain and Cognition, 45, 15-20.
- Otake, T., Hatano, G., Cutler, A., & Mehler, J. (1993). Mora or syllable? Speech segmentation in Japanese. Journal of Memory and Language, 32, 258-278.
- Payne, D.G., Elie, C.J., Blackwell, J.M., & Neuschatz, J.S. (1996). Memory illusions: Recalling, recognizing, and recollecting events that never occurred. Journal of Memory and Language, 35, 261-285.
- Petitot, J. (1995). Morphodynamics and attractor syntax: Constituency in visual perception and cognitive grammar. In R.F. Port and T. van Gelder (Eds.) Mind as Motion: Explorations in the Dynamics of Cognition, 227-281. Cambridge, MA: MIT Press.
- Pinker, S. (1991). Rules of language. Science, 253, 530-534.
- Pinker, S. (1994). The Language Instinct: How the Mind Creates Language. New York, NY: HarperPerennial.
- Plunkett, K. & Marchman, V. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. Cognition, 48, 21-69.
- Pollack, I. & Pickett, J.M. (1964). Intelligibility of excerpts from fluent speech: auditory vs. structural context. Journal of Verbal Learning & Verbal Behavior, 3, 79-84
- Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: Differences in the time course of retrieval. Cognitive Psychology, 21, 139-155.
- Rauschecker, J.P. (2001). Auditory Physiology. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 52-56.
- Richards, V.M. & Kidd, G.D. Jr. (2001). Audition. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 48-49.
- Ridout, N. Astell, A.J., Reid, I.C., Glen, T., & O'Carroll, R.E. (2003). Memory bias for emotional facial expressions in major depression. Cognition and Emotion, 17, 101-122.
- Roediger, H.L. III (1996). Memory illusions. Journal of Memory and Language, 35, 76-100.
- Roediger, H.L. III (2000). Why retrieval is the key process in understanding human memory. In E. Tulving (Ed.) Memory, Consciousness, and the Brain: The Tallinn Conference, pp. 52-75. Philadelphia, PA: Psychology Press.
- Roediger, H.L. III, & McDermott, K. (1995). Creating false memories: Remembering words not presented in lists. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 803-814.
- Ross, B.H. (1984). Reminders and their effects in learning a cognitive skill. Cognitive Psychology, 16, 371-416.
- Ross, B.H. (1996). Category representations and the effects of interacting with instances. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22, 1249-1265.

- Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month-old infants. Science, *274*, 1926-1928.
- Schmidt, S.R. (1991). Can we have a distinctive theory of memory? Memory and Cognition, *19*, 523-542.
- Schvaneveldt, R.W. (Ed.) (1990). Pathfinder Associative Networks: Studies in Knowledge Organization. Norwood, NJ: Ablex.
- Shastri, L. (1999). Advances in *Shruti* – A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. Applied Intelligence, *11*, 79-108.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. Annual Review of Psychology, *55*, 349-374.
- Skinner, 1974
- Schmidt, S.R. (1985). Encoding and retrieval processes in the memory for conceptually distinctive events. Journal of Experimental Psychology: Learning, Memory, and Cognition, *11*, 565-578.
- Spellman, B.A. & Holyoak, K.J. (1996). Pragmatics in analogical mapping. Cognitive Psychology, *31*, 307-346.
- Spivey-Knowlton, M. & Sedivy, J.C. (1995). Resolving attachment ambiguities with multiple constraints. Cognition, *55*, 227-267.
- Stein, B.E., Stanford, T.R., Vaughan, J.W., & Wallace, M.T. (2001). Multisensory integration. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 574-575.
- Stein, B.E., Wallace, M.T., & Stanford, T.R. (2000). Merging sensory signals in the brain: The development of multisensory integration in the superior colliculus. In M.S. Gazzaniga (Ed.) The New Cognitive Neurosciences, 55-71. Cambridge, MA: MIT Press.
- Stuffelbeam, R.S. (1999). Representation and Computation. In W. Bechtel & G. Graham (Eds) A Companion to Cognitive Science, 636-648. Maldin, MA: Blackwell.
- Talarico, J.M., & Rubin, D.C. (2003). Confidence, not consistency, characterizes flashbulb memories. Psychological Science, *14*, 455-461.
- Tulving, E. (1983). Elements of Episodic Memory, 224-267. Oxford: Clarendon.
- tulving
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. Behavioral and Brain Sciences, *21*, 615-665.
- Vroomen, J. & de Gelder, B. (1995). Metrical segmentation and lexical inhibition in spoken word recognition. Journal of Experimental Psychology: Human Perception and Performance, *21*, 98-108.
- Wallace, W.P. (1965). Review of the historical, empirical, and theoretical status of the von Restorff phenomenon. Psychological Bulletin, *63*, 410-424.
- Watkins, M.J. (1979). Engrams as cuegrams and forgetting as cue overload: A cueing approach to the structure of memory. In Memory Organization and Structure, 347-372. Academic Press.
- Watkins, M.J. (1989). Willful and nonwillful determinants of memory. In H. Roediger III & F. Craik (Eds) Varieties of memory and consciousness: Essays in memory of Endel Tulving, 59-71.
- Watkins, M.J. (1990). Mediationism and the obfuscation of memory. American Psychologist, *45*, 328-335.

Watkins, M.J. & Kerkar, S.P. (1985). Recall of a twice-presented item without recall of either presentation: Generic memory for events. Journal of Memory and Language, 24, 666-678.

Watkins, O.C. & Watkins, M.J. (1975). Buildup of proactive inhibition as a cue-overload effect. Journal of Experimental Psychology: Learning, Memory, and Cognition, 104, 442-452.

Wilding, E.L. (2000). In what way does the parietal ERP old/new effect index recollection. International Journal of Psychophysiology, 35, 81-87.

Youngentob, S. (2001). Smell. In R.A. Wilson & F.C. Keil (Eds) The MIT Encyclopedia of the Cognitive Sciences, 775-776.

Zentall, T.R. (2003). Imitation by animals: How do they do it? Current Directions in Psychological Science, 12, 91-95.

Zwaan, R.A. & Radvansky, G.A. (1998). Situation models in language comprehension and memory. Psychological Bulletin, 123, 162-185.

Distribution

1	MS 0830	Kathleen Diegert (12335)
1	MS 1188	Ann Speed (12335)
1	MS 1004	Patrick Xavier (15221)
1	MS 1170	Russel Skocypec (15310)
1	MS 1188	John Wagner (15311)
1	MS 1188	J. Chris Forsythe (15311)
1	MS 1188	Michael L Bernard (15311)
1	MS 9018	Central Technical Files, 8945-1
2	MS 0899	Technical Library
1	MS 0612	Review & Approval Desk, 9612
1	MS 0188	LDRD office