UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Latent Problem Solving Analysis as an explanation of expertise effects in a complex, dynamic task

Permalink

https://escholarship.org/uc/item/31w39621

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 25(25)

ISSN

1069-7977

Authors

Queseda, Jose
Kintsch, Walter
Gomez, Emilio

Publication Date

2003

Peer reviewed

Latent Problem Solving Analysis as an explanation of expertise effects in a complex, dynamic task

José Quesada, Walter Kintsch ([quesadaj, wkintsch]@psych.colorado.edu)
Institute of Cognitive Science, University of Colorado, Boulder
Boulder, CO 80309-0344 USA

Emilio Gomez (egomez@ugr.es)
Department of Experimental Psychology, University of Granada
Campus Cartuja, S/N, Granada, Spain

Abstract

Latent Problem Solving Analysis (LPSA) is a theory of knowledge representation in complex problem solving that argues that problem spaces can be represented as multidimensional spaces and expertise is the construction of those spaces from immense amounts of experience. The model was applied using a dataset from a longitudinal experiment on control of thermodynamic systems. When the system is trained with expert-level amounts of experience (3 years), it can predict the end of a trial using the first three quarters with an accuracy of .9. If the system is prepared to mimic a novice (6 months) the prediction accuracy falls to .2. If the system is trained with 3 years of practice in an environment with no constraints, performance is similar to the novice baseline.

Introduction

In this paper, we introduce a computational theory of representation in experienced problem solving that we call Latent Problem Solving Analysis (LPSA). It is specially suited to model performance in complex, dynamic tasks such as control of dynamic systems. Complex tasks have always been thought to involve high level processes, such as mental models and reasoning. We would like to show in the next sections that, although the conscious, effortful reasoning path is certainly available, people can also use a similarity-based way of action that can give good results in certain situations. LPSA proposes that what people do in some situations that have previously been considered problem solving can be considered memory retrieval and pattern matching.

LPSA is a spatial theory of representation, inheriting the assumptions and concepts of Shepard (1987). That is, the proximal stimulus is supposed to be represented as a point in a multidimensional space, where all other past experiences can be represented as well. The space is created to capture the similarities between objects. Thus, two objects that are similar tend to occupy close areas in the mental space.

LPSA is inspired by Latent Semantic Analysis (LSA), a theory of representation that explains how semantics can be learned from large amounts of experience. LSA has been applied to understand language comprehension phenomena, and has been used extensively in contemporary cognitive science. LPSA needs a corpus of experience, and does not propose mechanisms to act when there is no experience. We need to assume that there are two modes of reasoning, one for situations in which we know very little and another for those situations where we already have a knowledge base. We will review current approaches to expertise and how LPSA relates to them, and present data on how LPSA models prediction in the complex thermodynamic task DURESS (Vicente, 1991).

Expertise Theories

The most popular expertise theories are Long Term Working Memory (LTWM), Elementary Perceiver and Memorizer (EPAM), and Constraint Attunement Hypothesis (CAH). We will briefly describe them in the next sections.

Long Term Working Memory (LTWM). The LTWM theory claims that working memory has two different components: a short-term working memory (STWM), which is available under any condition, but of very limited capacity, and a long-term memory (LTWM), that is available only on the domain where one is an expert, but provides unlimited capacity. STWM accounts for working memory in unfamiliar activities but does not appear to provide sufficient storage capacity for working memory in skilled complex activities. LTWM is acquired in particular domains to meet specific demands imposed by a given activity on storage and retrieval. LTWM is task specific. Intense practice in a domain creates retrieval structures: associations between the current context and some parts of LTM that can be retrieved almost immediately without effort. That is, the retrieval is fast and automatic without requiring voluntary resources as in intentional memory search: the results 'pop out' in memory. The contents of working memory act as the center of a focus that activates other contexts from LTM that are related to them thanks to the retrieval structures.

The concept of retrieval structures is inherited from the Skilled Memory Theory (Chase & Ericsson, 1981). A retrieval structure is defined as an abstract, hierarchical knowledge structure used to organize cues used in the

encoding and retrieval of information. LTWM theory proposes that LTWM is generated dynamically by the cues that are present in short term memory. During text comprehension, for example, where the average human adult is an expert, retrieval structures retrieve propositions from LTM and merge them with the ones derived from text.

The evidence for the existence of LTWM comes from two fronts (1) proactive and retroactive interference, and (2) interruption and resumption of performance. If the representations formed during text comprehension are stored in short-term memory, interruption should hinder performance, measured as memory for the text (free and cued recall) and comprehension measures. However, classical experiments (e.g., Glanzer, Dorfman, & Kaplan, 1981) find no detriment in performance at all. The only difference between interrupted and non-interrupted conditions was longer reading times. LTWM permits rapid and reliable reinstantiation of a context after interruption without a decrease in performance.

LTWM has been applied to several domains such as memory for dinner orders, digit recall, chess, and text comprehension, but to date there is no explicit explanations of complex, dynamic tasks. The most promising computational implementations of LTWM retrieval structures have used LSA (see Kintsch, 1998; Kintsch, Patel, & Ericsson, 1999).

EPAM and the chunking theories. EPAM (e.g., Richman, Staszewski, & Simon, 1995) has three main components: an STM, a LTM, and a discrimination net, which allows nodes in LTM to be accessed. Short term memory includes specialized auditory and visual subcomponents, whereas long term memory is divided into declarative and procedural systems. EPAM is the natural evolution of the chunking theory (Chase & Simon, 1973). In EPAM, chunks are extended into templates. Templates are a large chunk, which contains slots (which are variables) that can be filled with concrete values for the current situation that the expert experiences or recreates. The slots might have default values that can reflect the statistically most frequent item that appear in the situation described by the template. Slots are fundamental concepts in EPAM. Within EPAM there are two types of slotted structures: schemas with all slots (generic retrieval structures) and schemas with only a few slots and mostly fixed values, called templates.

The concept of template is intimately bound to the nature of the discrimination net that is assumed as the representational format in EPAM. Slots are created as a function of the number of tests below a node in the discrimination net (e.g., Gobet, 1998). Like the chunking theory, the template theory proposes that expertise is due to '(a) a lager database of chunks indexed by a discrimination net. (b) a large knowledge base, encoded as production and schemas; and (c) a coupling of the (perceptual) chunks in the index to the knowledge base' (Gobet, 1998, p. 127). Like LPSA, EPAM

creates the representations (classification networks) starting from empirical information of similar proportions to what humans accumulate in their experience with the tasks. For example, to mimic DD's (a digit memory expert) behavior, Richman et al. trained the system with exactly the same information the expert had used to reach his expertise level. However, there are no models of continuous, dynamic processes like the one we present in this paper. The main difference is in the representations proposed. LPSA uses a comparatively simple spatial model, whereas EPAM uses discrimination nets, which are elaborated structured representations. The symbolic approach of the discrimination net makes it difficult to apply it to represent domains where variables change continuously, whereas LPSA does not show this problem.

Constraint Attunement Hypothesis (CAH). The LTWM and EPAM theories of expertise are process theories, that is, they try to specify the psychological mechanisms that explain the observable effects. That is, they are theories of 'how'. An alternative view would be to create a product (i.e., input-output) theory of expertise, where the question to answer is 'what' conditions are needed to observe expertise effects. This is the role of the Constrain Attunement Hypothesis (CAH) theory by Vicente and Wang (1998). Contrary to what process theories maintain, CAH does not commit to a particular psychological mechanism to explain the phenomenon of expertise. As a product theory, it aims to address three related issues: '(1) How should one represent the constrains that the environment (i.e., the problem domain) places on expertise? (2) Under what conditions will there be an expertise advantage? (3) What factors determine how large the advantage can be?' (Vicente & Wang, 1998, p. 35).

The CAH theory proposes an important distinction between *intrinsic* and *contrived* tasks. *Intrinsic* tasks are those that are definitive features of the domain of expertise, for example, blindfolded chess, memorizing dinner orders, and memorizing digits. A *contrived* task is one that is not part of the domain of expertise, but designed to fulfill some experimental purposes. For example, chess players just play chess, and remembering chess configurations is not part of the task. This distinction is important because (1) in the expertise literature, *contrived* tasks abound and (2) for some theories such as LTWM the proposed retrieval structures are obtained through a deliberate effort and then will be only relevant on the explanation of *intrinsic* tasks, that is, tasks that are needed to be an expert in the domain, such as memory enhancement in the waiter case. Vicente and Wang consider that most of the tasks used in the literature that studies memory expertise are *contrived*, not intrinsic, and in this sense LTWM and other process theories cannot explain them.

CAH is an ecological theory of expertise in memory recall, inheriting some of the basic ideas from Gibson's (1979) ecological theories of perception. In CAH, the experimenter is after a definition of the goal-relevant constraints in a domain. For example, the concept of *affordance* (what can be done in a particular environment) is reused indirectly and

extrapolated to the domain of memory recall and expertise. However, *affordances* are defined to describe properties of objects, events, and places, and what Vicente and Wang (1998) propose is a description of the whole domain of expertise, so they fall short. Vicente and Wang (1998) proposal needs a mechanism to identify and describe *relations* between the high numbers of components that make up a domain of expertise instead of the components themselves. The solution proposed to study goal-oriented constrains in the environment is the Abstraction Hierarchy (AH). The AH is a hierarchical description of the constrains of the problem domain, but a particular kind of hierarchy. Possible hierarchical descriptions to describe environments are part-whole relations, is-a relations, 'obeys' relations, and mean-ends relations. The definitive characteristic of the AH is that it describes the environment as mean-ends relations, connecting objects within and between levels. Thus, the AH is explicitly goal oriented.

The AH provides a different language for each level of analysis, providing the faculty of abstracting in an out (as in zooming) from the deeper significance of the system goals to the lowest physical levels of description (what physical changes need to be made in order to implement those goals).

The descriptions produced by experimenters using the AH approach are a-posteriori, and there is no guarantee that two experimenters would come up with the same AH when trying to describe the exact same task. AHs have been proposed for complex, dynamic tasks such as DURESS by Vicente. However, our proposal with LPSA is to create a similarity-based set of operations that define the representation of the environment in such a way that the similarity measures can be derived *automatically for any task*. That is, the definition of similarity is bound to experience in the particular environment, so the input for the theory will be the exact same information that humans use when they solve the tasks for the same period of time in which the human was exposed to the environment.

Comparing the theories

Comparing these theories has proven to be a difficult task. Even though some part from the same concepts (for example, EPAM and LTWM share the concept of retrieval structures), and in some cases the same phenomena have been targeted (for example, chess memory), the theories are not well compared in the literature. The reviews that do compare them normally attribute the advantage to the theory that the author of the review proposed (e.g., Gobet, 1998) and that normally originated a retaliation in related articles where the authors of the alternative theories try to amend the criticisms.

In the case of CAH vs. the process theories (LTWM and EPAM) the comparisons are even more difficult because the phenomena of interest are different for the different theories.

The ongoing discussion maintained by Gobet (2000) and Ericsson and Kintsch (2000) seems to be concentrated in two main points: (1) the necessity for slotted schemas.

Eriksson and Kintsch (1995; 2000) predicate that they are not needed, whereas Gobet (1998; 2000) cannot conceive EPAM without them. (2) LTWM proposes a gradual speed-up of encoding in LTWM, but EPAM proposes that there are fixed times for storage, and they are estimated. These two points are not addressed directly in LPSA and will not be commented further. Where LPSA does have a contribution to make is on the definition of retrieval structures, which has been criticized as vague in LTWM, and on the effects of amount of practice on expertise. The CAH assertions about the amount of structure of the environment are explained as well under a computational framework in LPSA.

LTWM claims that the magnitude of expertise effects is related to the level of attained skill and to the amount of relevant prior experience. CAH argues that this claim is incomplete. Expertise effects in memory recall are also determined by the amount of structure in the domain (and by active attunement to that structure).: CAH 'predict[s] … a memory expertise advantage in cases in which experts are attuned to the goal-relevant constraints in the material to be recalled and that the more constraints available, the greater the expertise advantage can be' (Vicente & Wang, 1998, p. 33). A theory that could explain both these assertions (amount of experience and structure of the environment) would be welcomed. LPSA is sensitive both to 'relevant previous practice' and to 'amount of structure in the domain', as we will show in the next sections.

The LPSA proposal for an expertise theory

One of our interests is to show that the abstraction hierarchy, the main innovation and contribution in Vicente and Wang theory (Vicente, 2000 ; Vicente & Wang, 1998) falls short in meeting the requirements for a theory of the environment in its actual form. A good theory that attempts to model the environment should be consistent and effective in different domains. The units and operators proposed should be the same for different environments, even though the basic structures can be very different. We agree with Vicente that it is important that a single theory can model different environments without changes in its basic assumptions. However, when CAH is used for modeling different environments "the details of such models usually differ tremendously from one domain to the next because the relevant cues and their ecological validities can change dramatically (…)" (Vicente & Wang, 1998, p. 603)

LSA is based on the idea of portraying environments as complex networks of coocurrences, that, given a big enough scale, can be mapped onto a multidimensional space of much lower dimensionality. Thus, it provides the means for modeling different domains in a comparable manner. At the moment of this writing, LSA has been applied to a variety of domains including the followings: understanding of source code (Maletic & Marcus, 2000a), text comprehension (e.g., Kintsch, 1998; Kintsch, 2001), categorization (Laham, 1997), metaphor understanding (Kintsch, 2000) and vocabulary acquisition and semantic priming (Landauer & Dumais, 1997). LPSA has been applied to model human

similarity judgments in problem solving tasks (Quesada, Kintsch, & Gomez, 2002), practice effects (Gonzalez & Quesada, submitted), and expert evaluations of landing technique (Quesada, Kintsch, & Gomez, submitted).

A complex, dynamic task: DURESS II. Manipulating previous knowledge by eliminating it has been a dominant in cognitive science, due in part to the need for random assignment of participants to groups that is an exigency of the experimental method. An alternative and very popular take is the expert-novice approach (e.g., Chase & Simon, 1973), that is, to manipulate previous knowledge by pre - selecting participants, forgetting about random assignment of participants to groups. In a wide and diverse range of contexts, from academic disciplines through to games and sports, comparisons of the performance of novices and experts have established consistent relations between knowledge, task performance and level of expertise.

However, not many researchers have the possibility of *manipulating* the environment for the time necessary to make a person an expert in a domain. Most of the studies in expertise and skill acquisition have to content themselves with analyzing diaries and interviews (i.e., Ericsson, Krampe, & Tesch-Römer, 1993) to estimate a posteriori the amount of deliberate practice that their participants invested.

Important exceptions to the problem posited above are single-subject designs such as Richman, Staszewski and Simon (1995), but it is only possible in very simple environments like digit recall tasks. In that case, the experimenter controls the environment (i.e. the sequence of digits that the memonist is to learn) completely and can manipulate it. The basic idea in this research paradigm is to move complexity to the lab, and manipulate previous knowledge by giving exactly the same amount of practice, enough to show expertise levels of skill, to all participants. To simulate expertise environments in labs, we need tasks more complex than the standard ones: more representative, with a long learning curve, and interesting enough to keep the motivation for a long period of time. An example of this kind of tasks is DURESS (DUal REServoir System). DURESS is a thermal-hydraulic process control simulation that was designed to be representative of Industrial process control systems. It consists of two redundant feedwater streams that can be configured to supply either, both or neither of the two reservoirs. The goals of the game is to keep each of the reservoir temperatures (T1 and T2) at a prescribed level (e.g., 40 C and 20 C, respectively), and to satisfy the current mass (water) output demand (5 liters per second and 7 liters per second, respectively). Thanks to the seminar work of Vicente and collaborators (Christoffersen, Hunter, & Vicente, 1996, 1997, 1998), the equivalent of three years of experience with the system DURESS II is available and we used it in our LPSA simulations.

Method. Complex experimental tasks normally keep a log file containing all the actions and states that every participant has experienced. Since the number of variables is

very high and their relations can be intricate, the analysis are usually beyond the scope of most statistical methods normally employed in experimental psychology, particularly when system states are not in interval scale. As a result, the richness of these log files is underused. However, a clear analogy can be drawn between this particular problem and representational theories of semantics such as LSA: like words, states and actions appear in particular contexts but not in others. Some states and actions are interchangeable, being 'functional synonyms'. Given the right algorithms and sufficient amounts of logged trials, a problem space can be derived in a similar way as semantic spaces are. The underlying idea is that the aggregate of all the action contexts in which a given state does and does not appear provides a set of mutual constraints that largely determines the similarity of meaning of states and sets of states to each other.
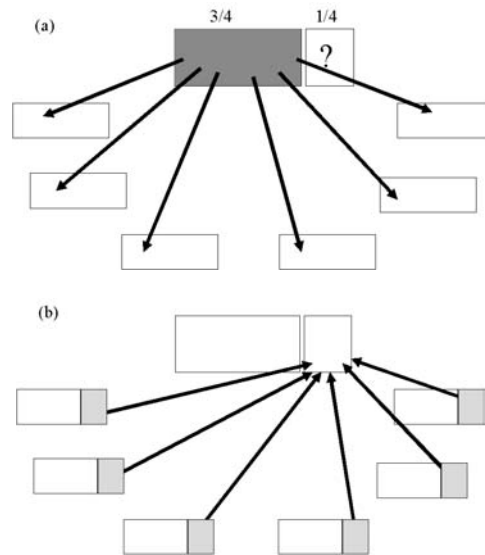

Figure 1: Prediction method employed to estimate the next states in the task in LPSA. Each rectangle represents a trial in DURESS II. (a) The nearest neighbors of the predicting part are retrieved; (b) a composite using the ends of these neighbors serves to predict the target trial's final states.

The AH proposed for DURESS (e.g., Vicente & Wang, 1998) contains four levels: (1) Functional, that describes the purposes: keep the temperature and demand flow rate for each reservoir, (2) Abstract, that describes the system as a function of the laws of conservation of mass and energy, (3) generalized, that uses rates of heat and flow transfer, and (4) physical, that describes the physical position and settings of the components (valves, pumps, and heaters).

A different LPSA corpus was created for each level of the Abstraction Hierarchy. After performing the SVD, the first 100 dimensions were used. Since the goal values are different at each level, the pertinent variables were normalized to goal values in order to make trials more

comparable for Functional purpose and Abstract function levels. All levels were normalized with respect to scale.

To test the 'amount of experience hypothesis', two different corpora, one with 3 years of practice (expert), and another with only 6 months (novice), were created, and the quality of their prediction of future states compared. To test the 'amount of domain structure hypothesis', an additional corpus of 3 years of practice was created, but this time in an environment where the laws of conservation of mass and energy do not exist: the states where randomly assigned to the trials, resulting in a corpus with practically no constraints. This corpus was compared to the 3-years-expert with a constrained environment.

Prediction. Prediction plays a very important role in humans' interaction with the environment. Some scientists argue that many features of cognitive the cognitive system (such as representation, memory, and categorization) can be conceptualized as tools that help to predict the next states of an organism's environment (e.g., Anderson, 1990). The methodology that we used was to test how good of a prediction can be generated using the nearest neighbors of a target slice of performance. For example, in a trial of DURESS, how much of the end can be predicted using the information from the beginning? To do this, we needed to define a cutting point that divided the predicting and predicted parts. The cutting point we defined is the point that leaves ¾ of the trial behind. Such a case is depicted in Figure 1a. Trials in DURESS are represented as rectangles. The shaded area is the part of the trial that is used to predict the remaining part (signaled with a question mark). In LPSA, any passage is a vector, as well as any sub-passage; that is, the shaded (predicting) part and the question mark (predicted) part are both a vector in LSA. Using the predicting vector, we retrieve the nearest N neighbors, depicted as rectangles as well in Figure 1a. In this figure, N=6, that is 6 nearest neighbors are evoked by the first ¾ of the trial. Then, the last quarter of each retrieved neighbor is used to create a composite that predicts the end of the target trial (Figure 1b). The contribution to the composite is weighted by the cosine between the neighbor and the predicting part of the target.

Results

The results presented here were calculated using 10 neighbors (striped bars), and the same calculations performed with 10 random neighbors are used as a control group (solid bars). A sample of 100 random trials was selected as target trials, and the results averaged. The predicting accuracy of this method in the 3-years of practice, structured environment' can be observed in Figure 2: the average prediction is .87, which means that our simulated expert can predict the next states of its environment very well indeed. Figure 3 shows that, for the novice simulation, the average prediction is much worse, which is in line with the 'amount of experience' hypothesis. Figure 4 describes the prediction rate (that does not outperform the random control) for the '3 years expert in the unconstrained domain'.
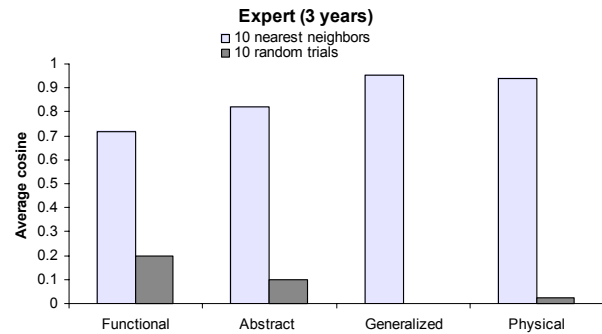

Figure 2: Average cosine between the fourth quarter of a target trial and the fourth quarter of the 10 nearest Neighbors when the three first quarters are used to retrieve the neighbors. The model has been trained with three years of experience.
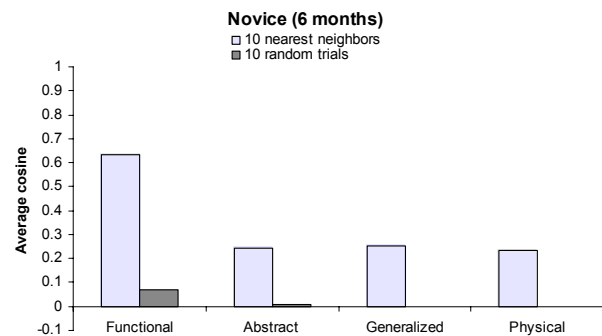

Figure 3: Average cosine between the fourth quarter of a target trial and the fourth quarter of the 10 nearest Neighbors when the three first quarters are used to retrieve the neighbors. The model has been trained with six months of experience.
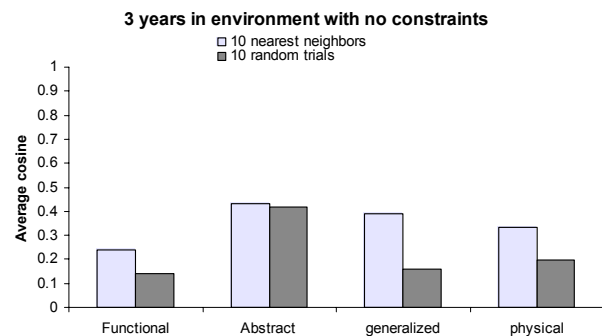

Figure 4: Average cosine between the fourth quarter of a target trial and the fourth quarter of the 10 nearest Neighbors when the three first quarters are used to retrieve the neighbors. Three year of experience in a DURESS simulation with no constraints (random states).

Conclusions

LTWM argues that the magnitude of expertise effects is related to the level of attained skill and to the amount of relevant prior experience. CAH claims that expertise effects in memory recall are also determined by the amount of

structure in the domain (and by active attunement to that structure). LPSA can explain both arguments under the same framework, and proposes a computational model on how the constrains of the environment are internalized and represented. LPSA also extends the area of application of computational expertise theories to complex, dynamic tasks such as DURESS. In doing so, LPSA is a new approach to the expertise and knowledge representation discussions.

Acknowledgments

References

Anderson, J. R. (1990). The adaptive character of thought. Hillsdale, New Jersey: Lawrence Erlbaum associates.

Chase, W. G., & Ericsson, K. A. (1981). Skilled memory. In J. R. Anderson (Ed.), Cognitive skills and its acquisition (pp. 141-189). Hillsdale NJ: Erlbaum.

Chase, W. G., & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), Visual information processing. New York: Academic Press.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1996). A longitudinal study of the effects of ecological interface design on skill acquisition. Human Factors, 38, 523-541.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1997). A longitudinal study of the effects of ecological interface design on fault management performance. International Journal of Cognitive Ergonomics, 1, 1-24.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1998). longitudinal study of the impact of ecological interface design on deep knowledge. International Journal of human-Computer Studies, 48.

Ericsson, K. A., & Kintsch, W. (1995). long-term working memory. Psychological Review, 102(2), 211-245.

Ericsson, K. A., & Kintsch, W. (2000). Shortcomings of generic retrieval structures with slots of the type that Gobet (1993) proposed and modeled. British Journal of Psychology, 91, 571-588.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer. (1993). The role of deliberate practice in the acquisition of expert performance. Psychological review, 100(3), 363-406.

Gibson, J. J. (1979). The Ecological Approach to Visual Perception. Boston: Houghton Mifflin.

Glanzer, M., Dorfman, D., & Kaplan, B. (1981). Short-Term Storage in the Processing of Text. Journal of Verbal Learning and Verbal Behavior, 20(6), 656-670.

Gobet, F. (1998). Expert memory: a comparison of four theories. Cognition, 66(2), 115-152.

Gobet, F. (2000). Some shortcomings of long-term working memory. British Journal of Psychology, 91, 551-570.

Gonzalez, C., & Quesada, J. F. (submitted). Learning in a Dynamic Decision Making Task: The Recognition Process. Computational and Mathematical Organization Theory.

Kintsch, W. (1998). comprehension: a paradigm for cognition: Cambridge university press.

Kintsch, W. (2000). Metaphor comprehension: A computational theory. Psychonomic Bulletin and Review, 7, 257-266.

Kintsch, W. (2001). Predication. Cognitive Science, 25, 173-202.

Kintsch, W., Patel, V., & Ericsson, K. A. (1999). the role of long-term working memory in text comprehension. Psychologia, 42, 186-198.

Laham, D. (1997). Latent Semantic Analysis approaches to categorization. In M. G. Shafto & P. Langley (Eds.), Proceedings of the 19th annual meeting of the Cognitive Science Society (pp. 979). Mahwah, NJ:: Erlbaum.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. Psychological Review, 104, 211-240.

Maletic, J. I., & Marcus, A. (2000a). Using Latent Semantic Analysis to Identify Similarities in Source Code to Support Program Understanding, 12th IEEE International Conference on Tools with Artificial Intelligence (ICTAI00). Vancouver, British Columbia.

Quesada, J. F., Kintsch, W., & Gomez, E. (2002). A theory of Complex Problem Solving using Latent Semantic Analysis. In W. D. Gray & C. D. Schunn (Eds.), 24th Annual Conference of the Cognitive Science Society (pp. 750-755). Fairfax, VA.: Lawrence Erlbaum Associates, Mahwah, NJ.

Quesada, J. F., Kintsch, W., & Gomez, E. (submitted). Automatic Landing Technique Assessment using Latent Problem Solving Analysis, 25th Annual Conference of the Cognitive Science Society.

Richman, H., Staszewski, J., & Simon, H. A. (1995). simulation of expert memory using EPAM IV. Psychological Review, 102(2), 305-330.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. Science, 237, 1317-1323.

Vicente, K. J. (1991). Supporting Knowledge-Based Behavior Through Ecological Interface Design. University of Illinois at Urbana-Champaign.

Vicente, K. J. (2000). Revisiting the constraint attunement hypothesis: reply to Ericsson, Patel and Kintsch (2000) and Simon and Gobet (2000). Psychological Review, 107(3), 601-608.

Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. Psychological Review, 105, 33-57.