

Modeling Cognitive Flexibility of Super Experts in Radiological Diagnosis

Eric Raufaste (raufaste@univ-tlse2.fr)

Laboratoire Travail et Cognition; Université de Toulouse le Mirail
31058 Toulouse Cedex; France

Abstract

The paper presents theoretical propositions for modeling the expert radiologist. The propositions are twofold. First, a *basic model* is given to complement a recent connectionist symbolic framework (Raufaste, Eyrolle, & Mariné, 1998). Empirical data have showed dissociation between two kinds of experts ("basic" and "super") with regard to cognitive flexibility. The difference is conceived as a kind of perseveration in basic experts. Hence, the basic model was combined with a Supervisory Attentional System (Norman & Shallice, 1986) into an "extended model". An analysis of cognitive activity is then presented within this framework, along with a new theoretical explanation of cognitive flexibility.

Modeling the expert radiologist

The Need for a Connectionist-Symbolic Approach

A well-documented ability in expert physicians is early selection of pertinent diagnostic hypotheses (Elstein, Schulman & Sprafka, 1978). In radiology, perceptual processes have a dramatic importance (Lesgold et al., 1981). However, a search time study suggested the existence of two components, the earlier rapid, and the latter slow (Christensen et al., 1981). It has been proposed that a "visual concept" shapes perception (Kundel & Nodine, 1983). Hence, the first component might be more plausibly described by a connectionist approach. But medical diagnosis is a complex task that requires a lot of deliberate reasoning, so it seems also necessary to have a symbolic layer. Empirically, Lesgold et al. (1988) found a nonmonotonic performance curve on some films: Novices (1- and 2-year residents in radiology) sometimes performed better than intermediates (3- and 4-year residents in radiology). Experts performed the best. To explain this nonmonotonicity, the authors proposed a three-stage framework. In the first stage, novices would acquire basic subsymbolic abilities such as recognizing the normal anatomy on the film. On a second stage, intermediates would develop "cognitive" (i.e., symbolic) abilities. For some cases, however, the cognitive processing would conflict with previously developed perceptual processing, resulting in a decreased performance. In experts, cognitive processing would have reached its plain development so that conflict no longer spoils performance. Thus, a connectionist-symbolic approach (Holyoak, 1991) seemed to be appropriate for modeling diagnosis.

Basic and super experts

In their experiments, Lesgold et al. used expert radiologists who were recognized as "outstanding" by peers. But only a few radiologists become outstanding. Moreover, such radiologists often have professional attributions that are substantially different from "normal" radiologists. These attributions induce more "deliberate practice" (Ericsson, Krampe, & Tesh-Römer, 1993) and more symbolic reasoning. In a recent study of expertise in radiological diagnosis, Raufaste, Eyrolle, and Mariné (1998) called "basic experts" the common radiology practitioners, and "super experts" the outstanding radiologists. The distinction appeared to be empirically fruitful and allowed us to reinterpret classical results by Lesgold et al. (1988). Raufaste et al. tested a framework that was initially devised to account for both subsymbolic and symbolic aspects of medical reasoning. Since subsymbolic and symbolic processes are integrated, they should not generate a conflict. Hence, there should not be nonmonotonicity. We found that performance curves on typical features was monotonically increasing from novices to basic experts. In contrast, performance curves on atypical features was monotonically decreasing from novices to basic experts. Such a result was in accord with our framework. However, super experts exhibited several features that could not stem from our model. In particular, they always had a better performance than the other groups, *even on atypical features*. This replicated the apparent non-monotonicity in the results of Lesgold et al. (1988) and showed a dramatic distinction between basic- and super-expertise: the former is accompanied with growing dependence on automatic processes whereas the latter allows a better independence from automatisms. A key point here is the fact that super experts were better than basic experts at detecting an inconspicuous feature *that could not be awaited* from the hypotheses associated with salient features. Such an effect could not be attributed to better perceptual abilities but rather to a better cognitive flexibility in super experts.

The present paper proposes an extension of the initial framework in order to model cognitive flexibility in super experts. After a new formulation of the model, a new level is added in the model. An original analysis of cognitive activity is conducted and explanatory mechanisms are proposed.

The basic model: a new formulation

The initial model had to account for previously known results about general expertise (e.g., Lesgold et al, 1988) and for what we called *pertinence generation*, that is the acquisition through experience –and the use, of the ability to select rapidly pertinent hypotheses. We present here a refined version of the model: The *basic model*. The basic model is grounded on the concept of Long-Term Working-Memory, that is, W-M is viewed as a more activated part of LTM (Anderson, 1983, Ericsson & Kintsch, 1995).

Schemas. Categories, called schemata, may be represented at two levels: They may be symbols (in the sense of Hinton, 1990) and/or they may be patterns that are distributed within subsymbolic networks. A symbol may or not be associated with a subsymbolic pattern. A subsymbolic pattern may or not be associated with a symbol.

Activating Attention postulate: any symbol in Working Memory is a source of activation. Although restricted to a single symbol, a similar postulate can be found in Collins and Loftus (1975).

Inferences and Reasoning. Inferences in the basic model may occur through two distinct processes. The first is spreading activation from a node to another node. For example, if a radiologist detects a cue, the corresponding visual pattern lends activation to its symbol (e.g., a specific syndrome) and activation can spread towards the symbols of the pathologies that are associated with the syndrome. The second process is activation of a production rule in procedural memory. We define a *focal threshold* as the quantity of activation that a symbol must reach for being consciously processed: A category can be symbolically processed (e.g. verbally reported) only if it is associated with a symbol whose activation is above the focal threshold.

As conscious attention works with limited resources, a plausible mechanism for conflict resolution is a competition based on the level of activation (e.g., “contention scheduling”, Norman & Shallice, 1986). Thus, we add a new

postulate:
Captivating Activation postulate: the most activated symbols tend to obtain the focus of attention.

Those premises entail several interesting consequences: (1) a distributed pattern that was implicitly acquired cannot be symbolically processed until it has been associated with a symbol. (2) When the activation of a category increases, its probability to be symbolically processed also increases. (3) If the symbol of a category is inhibited, its probability of being symbolically processed decreases. (4) Through conscious call, activation may spread from a symbol to an associated subsymbolic pattern. Thus, some “substance” can be given to abstract concepts. Reciprocally, (5) activation may spread to a symbol from a distributed pattern that was activated by environmental stimuli. (6) The conscious representation can be defined as the set of categories having a symbol whose activation is above the focal threshold.

Initial Learning and Effects of Experience. The basic model is essentially a spreading-activation model of memory. Conscious processing takes place at the symbolic level and can initially generate theoretical knowledge by creating nodes (*canonical schemata*) in the network, and links (*canonical links*) between the nodes. In our view, experience has two main effects. The first effect is the classical view of connectionist networks: acquiring new nodes and/or new links; modifying the strengths of the links and the base-levels of the nodes. The second effect is to complement the knowledge base with examples that are encoded in episodic memory. In other words, declarative knowledge acquired through University learning, constitutes a pre-structured network of canonical schemata and canonical links around which further subsymbolic acquisitions will be arranged. Indeed, symbolic reasoning may still create new nodes and links after the medical degree course is over, through further reading, reflection on the results of actions, and so on...

Low accessibility postulate: without reinforcements from experience, the weak accessibility of symbols only allows a deliberate access to canonical schemata.

This postulate explains why sometimes novices were found not to use knowledge which they have (Custers, Boshuizen, & Schmidt, 1996). Until the link between symbols and subsymbolic background becomes enough strengthened by experience, only deliberate reasoning may trigger a schema. For example, a novice detects an abnormal feature on a film but does not know how to evoke pathology from the feature. The pathology, however, may be activated by a procedural rule whose action part provides activation to its symbol. The more the subject encounters simultaneously a context and a symbol, the more the link between both is strengthened and the more further encountering of the context will automatically activate the symbol. Thus, with experience, rule-based reasoning is replaced by spreading activation. From the same context, the expert will be able to activate more symbols with more links between them, that is, to generate a richer representation. Due to lateral inhibition effects, pertinent symbols will receive more activation from the context and, therefore, are more likely to win the competition for attentional resources. Hence, because they are more integrated, expert representations should also be more pertinent. Because the representation is richer, complex productions rules may also be triggered and so complex reasoning also becomes available in experts.

Basically, our results as well as the literature fit the model. However, one result was clearly not in agreement with the basic model, even in its current form.

A reduced SOS phenomenon in Super Experts

The Satisfaction of Search phenomenon (SOS, Berbaum et al., 1990) is a well-known effect in the literature about observer performance in radiology: an inconspicuous feature (e.g., a lung nodule) has a lower probability to be detected when the film also presents with a more salient unrelated feature. By itself, this phenomenon is clearly

“predicted” by the basic model: (1) the most salient features naturally receive first the focus of attention. (2) Due to spreading activation, the symbols that are strongly associated with this initial context receive more activation than other symbols; (3) Due to the captivating attention postulate, those symbols are reinforced by conscious attention and (4) they become activation sources. From this moment, they control behavior, and rules that relate to those symbols are more likely to fire. In particular, rules for hypothesis testing will lead further exploration of the film, and features that are non-directly relevant to those hypotheses may be missed. As the model predicted, the SOS phenomenon increased with basic expertise, from novices to basic experts. In novices, SOS is reduced by the lower strength of the links between subsymbolic context and symbols. Indeed, other causes for rigidifying effects may be found (e.g., Feltovich, Spiro & Coulson, 1997). On some cases, however, super experts avoided the SOS phenomenon. This is clearly not compatible with the basic model because they should have been even more subject to SOS than basic experts. One might suggest that they might have a better visual ability but an eye-recording study showed that cognitive processes were responsible for the SOS phenomenon (Samuel et al., 1995): most missed nodules were fixated and erroneously categorized as variants of normal. Structural properties of specific knowledge cannot account for the SOS phenomenon because the same diagnosticians do detect the same nodules when no independent salient feature is present. Because neither perceptual abilities nor specific knowledge can be responsible for the better performance we observed in super experts, we need to turn to general mechanisms of control. Thus, we need to extend the basic model.

An Extended Model of Expertise

From a neuropsychological standpoint, it has been argued that an activation-based competition between schemata is not sufficient to model a normal human subject. It can only model a patient with prefrontal lesions (e.g., Shallice & Burgess, 1993). In our attempt to model basic and super expertise, it seems interesting to view the SOS phenomenon as a particular case of perseveration. In such a view, a major difference between basic- and super-expertise is the relative weight of the specific-knowledge base. Basic experts may be modeled by a contention-scheduling process in a knowledge base whereas modeling super experts requires, in addition, the existence of an instance that modulates contention-scheduling. We adopt here, the concept of *Supervisory Attentional System* (SAS, Norman & Shallice, 1986).

Adding a modulation process to the basic model is not sufficient by itself to explain super-expert flexibility: We also need to model cognitive activity. First, we call *mental state* (Smolensky, 1988) a particular pattern of activation in the network. We call *cognitive flow* the sequence of mental states. Now, attentional control has a heavy cognitive cost. Therefore, the SAS is not expected to function actively

without a good reason. We consider that the SAS may be in two states: (1) We call *intervention* the state where the SAS actively modifies the course of contention scheduling. (2) In contrast, we call *survey* the state during which the SAS is not active. We call *natural flow* the cognitive flow when the SAS is in a survey position. Thus, the cognitive flow can be analyzed as a sequence of natural flows which, sometimes, is interrupted by SAS interventions (Figure 1). Natural flow sequences are guided by the procedural schemata associated with the dominant declarative schema (i.e., only the rules that are associated with this schema are activated enough to fire).

Within such an analysis, two questions must be solved: (1) How are the SAS interventions triggered? and (2) How do interventions work? To answer those questions, and in addition to the previous descriptive approach to cognitive activity, we need a general principle that orients the activity:

Principle of coherence maximizing: Cognitive activity is intended to maximize the overall coherence of the cognitive system. At a symbolic level, the principle gives the orientation of cognitive activity. At a subsymbolic level, the principle enables computation in neural nets (for justifications, see Thagard, 1989 and related works).

How are interventions triggered? Because of their importance in daily activity, for both experts and ordinary people, we expect interventions to be a very low-level process. Because cognitive flexibility seems to differ from basic to super experts, and because the two kinds of experts mainly differ by their daily activity, we assume that the low-level procedure can be triggered by procedural schemata which can be learned:

We define a *ruptor* as a schema that operates a categorization on mental states and that activates a low-level procedure of intervention on the cognitive flow. Because of the maximizing coherence principle, ruptors should take as input a kind of information that embodies an estimate of the overall (and/or local) coherence of the network. For example, detecting a deadlock is a good reason for triggering an intervention (See Holyoak, 1991).

With regard to the specific problem of basic and super expertise, we must also justify how interventions explain why super experts are less prone to the SOS phenomenon. The concept of ruptor provides a simple explanation in terms of differential intervention triggering.

— Because ruptors are schemas, they are subject to the contention-scheduling process. This explains why a basic expert becomes more and more dependent on automatism. As the automatism become more efficient, they are more likely to win the competition. With experience they are more and more refined so deadlocks become sparse. Finally, because basic expert daily activity does not include much deliberate practice, ruptors are not systematically trained. Thus, with experience, Basic expertise tends to be more and more in agreement with the basic model because the weight of the SAS is continuously decreasing.




Figure 1 : An analysis of the flow of cognitive operations

_ In super experts, daily activity is accompanied with much more explicit reasoning. As researchers, they have to make their reasoning explicit in order to publish, whereas basic experts only have to provide a diagnosis. As teachers of internists, they have to justify their reasoning, to make their implicit inferences explicit. In addition, super experts are often called for diagnosing those cases that were deemed too much difficult by other radiologists. For all those reasons, they are trained to commit high levels of attention in diagnosis. They are more familiar with the fact that automatism may lead to wrong solutions. The factors that favor cognitive flexibility (For a review, see Feltovich et al., 1997) are construed in the extended model as generating and training new raptors. Super expert's SASs are more trained to intervene and their raptors probably have a higher base-level than basic expert raptors. Thus we may expect them to be more independent from automatism than basic experts.

Now, we have a basis for an explanation of the difference between basic- and super-experts with regard to cognitive flexibility. Nevertheless, for the explanation to be complete, we should be able to explain how an intervention may reduce the likelihood of the SOS phenomenon.

How interventions work? The SOS phenomenon can be regarded as a kind of perseveration induced by a positive feedback loop. This loop results from the combination of captivating activation and activating attention postulates. Then, to avoid the SOS, interventions must be able to break the loop. As we stated earlier, interventions are expected to be a low-level process. In our analysis, interventions are short actions from the SAS. After the intervention, new schemata can gain the control over the cognitive flow. Hence, a minimal action of the SAS is to inhibit the current dominating schema (see McCarthy & Warrington, 1990) so that a new sequence of natural flow can begin. More sophisticated explanations based on concepts like Harmony Optimization (Smolensky, 1986) can be found in the theory of stochastic neural nets. However, they are beyond the scope of the present paper.

A Preliminary Test of the Extended Model

A complete test of the extended model is not out of reach, but it requires such methods as eye-movement recording because in order to observe a real SOS phenomenon, one has to ensure that the critical cue was actually seen. We just want to verify the plausibility of the main idea of the model—the SOS phenomenon might be related to a slight form of perseveration, which accompanies basic expertise and can be avoided in super experts by SAS interventions.

If the explanation we proposed is correct, we should be able to find some traces of interventions in verbal protocols. In particular, we should be able to find more interventions in super experts than in basic experts.

With regard to novices and intermediates, a lot of explicit reasoning, and even of deliberate practice (Ericsson et al., 1993) is likely because they have to acquire a vast specific knowledge-base within few years. Moreover, they are trained to learn because before being internists in a specialty, like radiology, they were selected among the best students in general medicine courses. Nevertheless, as their specific knowledge-base grows, they should depend more on automatism and less on general mechanisms such as weak heuristics and SAS interventions. Only the few who will some day become super experts can be expected to maintain a high level of deliberate activity. As a consequence, we can expect a monotonely decreasing curve in the number of interventions from novices to basic experts, and a higher number of interventions in super experts than in basic experts.

The next question is how can we measure interventions? The main consequence of interventions is to change the schema that guides the reasoning process. Therefore, we can trace those changes in verbal protocols. We call *line of thought* the verbal trace of a sequence of natural flow. A sequence of natural flow is not observable, whereas a line of thought can be traced in the verbal protocols. The basic idea behind the test is that interventions change line of thoughts and, therefore, tracing the changes in the line of thoughts

give some indications on the number of interventions. However, not only interventions change lines of thought: The dominating schema can be inhibited because a crucial new information gives a strong positive support to a concurrent hypothesis or a strong negative support to the dominant schema. In other words, we should not count as interventions the changes that can be attributed to a new information arrival. Another important factor in the changes is the use of a systematic strategy of exploration. Radiology residents are taught to explore the films according to a topographical schema that enables them to explore systematically every important zone. Therefore, when a diagnostician uses such a strategy, the dominant schema is not pathology but the topographic exploration schema. Hence we should not count changes in the dominant schema that can be attributed to the use of such a strategy.

In order to give a preliminary testing of these ideas, the verbal protocols that were used in Raufaste et al. (1998) were coded. It should be noted that the new coding was completely independent and different from the original coding. Thus, the number of interventions could be obtained (Figure 2). The test involved 8 novices, 6 intermediates, 4 basic experts, and 4 super experts.




Figure 2: Interventions as a function of expertise

As expected, one can observe a monotonic decrease in the curve: The mean number of interventions is maximal in novices (5.38 , $s_d = 3.16$), it decreases down to 4.33 ($s_d = 2.07$) in intermediates, to 0.75 ($s_d = 0.96$) in basic experts. Moreover, the relation is significant ($F_{(2,17)} = 4.63$; $p = .0135$). Also expected, with 3.75 interventions on average ($s_d = 2.22$), the mean number of interventions is significantly higher in super experts than in basic experts ($t_{(6)} = -2.48$; $p = .024$). The graph presented in Figure 2 is typically of the same kind as the graph one can draw with regard to performance on atypical cases.

When examining the whole curve of planed-triggering, which can be produced from the four groups, we obtain a significant decreasing monotonic relation ($F_{(2,17)} = 5.684$; $p = .006$) ranging from 3.63 ($sd = 1.85$) in novices down to 0.50 ($sd = 0.58$) in super experts. The latter result confirms that the better performance of super experts in the detection of an inconspicuous feature cannot be explained by the use of a more systematic exploration procedure.

Discussion

The model of cognitive flexibility proposed here should be considered as complementary to the prescriptions of Cognitive Flexibility Theory (e.g., Feltovich et al., 1997). Explaining differences between basic and super experts might also have been approached through Rasmussen's model of control (e.g., Olsen & Rasmussen, 1989) which construes cognitive flexibility in terms of the ability to adapt the control mode (skill-, rule-, or knowledge-based) to the specificity of the situation. However, if a radiologist devotes most of his attentional resources to an abstract reflection, the control is knowledge-based with regard to the reflection while, at the same time, there is place for a skill-based control of the film exploration. More generally, automated processes can operate in parallel to attentional processes. Hence, the control is not skill-based only, rule-based only, or knowledge-based only. To the contrary, our model uses two mutually exclusive categories: a state of survey and a state of intervention. In many respects, our model could also be compared to a hybrid model of abduction like UEcho (e.g. Wang, Johnson, & Zhang, 1997). In a recent study (Raufaste & Da Silva-Neves, 1998), basic expert radiologists were found to conform to Possibility Theory (Zadeh, 1978; Dubois & Prade, 1988). UEcho, however, is not compatible with nonstandard approaches, which require two measures of uncertainty. Our model, in contrast, is compatible with these results as well as with results in the literature where subjects tend to conform to bayesian rules of reasoning (Raufaste, Da Silva-Neves, & Mariné, submitted manuscript). Moreover, UEcho combines ECHO and the SOAR architecture whereas our model resembles a hybrid form of the ACT-R architecture and Norman & Shallice' theory of action.

This paper presented a refined version of a model of expertise in radiological diagnosis (Raufaste et al., 1998). The model is now twofold. The "basic model" embodies the previous model as well as two new postulates that relate to attention. In its current version, it is sufficient to model the basic expert. The paper also presented an "extended model". In addition to the basic model, it includes a Supervisory Attentional System (Norman & Shallice, 1986) that accounts for super expert's behavior. The extended model can account for a wide range of results about radiological expertise. Being within the frame of symbolic connectionism, it has the potential to deal with purely perceptual aspects of diagnosis as well as attentional deliberate reasoning. Indeed, much work will now be necessary to test the model. It is expressed in a general form, and might serve for many other domains of expertise. Therefore, it could be useful to any researcher who studies abductive reasoning and/or reflective experts.

Acknowledgments

This research was supported by Grant n° 99N35/0005 of the CNRS – "GIS Sciences de la Cognition".

References

- Anderson, J.R. (1983). *The Architecture of cognition*, Cambridge Massachusetts, Harvard University Press.
- Anderson, J.R. (1993). *Rules of the mind*. NJ: LEA.
- Berbaum, K.S., Franken, E.A., Dorfman, D.D., Rooholamini, S.A., Kathol, M.H., Barloon, T.J., Behlke, F.M., Sato, Y., Lu, C.H., El-Khoury, G.Y., Flickinger, F.W., & Montgomery, W.J. (1990): Satisfaction of search in diagnostic radiology. *Investigative Radiology*, 25, 133-140.
- Christensen, E.E., Murry R.C., Holland, K., Reynolds, J., Labday, M.J., & Moore, J.G. (1981). The effect of search time on perception. *Radiology*, 138, 361-365.
- Collins, A.M., & Loftus, E.F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 6, 407-428.
- Custers, E.J.F.M., Boshuizen, H.P.A., & Schmidt, H.G. (1996). The influence of medical expertise, case typicality, and illness script component on case processing and disease probability estimates. *Memory & Cognition*, 24, 3, 384-499.
- Darling, S., Della Sala, S., Gray, C., & Trivelli, C. (1998). Putative functions of the prefrontal cortex: Historical perspectives and new Horizons. In G. Mazzone & T.O. Nelson (Eds.), *Metacognition and cognitive neuropsychology: Monitoring and control processes*. Mahwah, NJ: LEA.
- Dubois, D., & Prade, H. (1988). *Possibility Theory*, New York: Plenum Press.
- Elstein, A.S., Shulman, L. S., & Sprafka, S.A. (1978). *Medical Problem Solving: An analysis of clinical reasoning*. Cambridge, Ma: Harvard University Press.
- Ericsson, K.A., Krampe, R.T., & Tesh-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100, 3, 363-406.
- Hinton, G.E. (1990). Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence*, 46, 47-75.
- Holyoak, K.J. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K.A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise*. Cambridge University Press.
- Kundel, H.L., & Nodine, C.F. (1983). A visual concept shapes image perception. *Radiology*, 146, 2, 363-368.
- Lesgold, A.M., Feltovich, P.J., Glaser, R., & Wang, Y. (1981, September). *The acquisition of perceptual diagnostic skill in radiology (Tech. Rep. No. PDS-1)*. University of Pittsburgh, Learning Research and Development Center.
- Lesgold, A.M., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing X-Ray Pictures. in M.T.H. Chi, R. Glaser, & M.J. Farr (Eds.), *The Nature of Expertise*. Hillsdale, NJ: LEA.
- Feltovich, P. J., Spiro R. J., & Coulson, R. L. (1997). Issues of expert flexibility in contexts characterized by complexity and changes. In P.J. Feltovich, K.M. Ford, & R.R. Hoffman. (Eds.), *Expertise in context. Human and machine*. Cambridge, Mass: MIT Press.
- McCarthy, R.A., & Warrington, E.K. (1990). *Cognitive Neuropsychology: A clinical introduction*. San Diego, CA: Academic Press.
- Norman, D.A., & Shallice, T. (1986). Attention to action : Willed and automatic control of behavior. in R.J. Davidson, G.E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation* (vol. 4). New-York: Plenum Press.
- Olsen, S.E., & Rasmussen, J. (1989). The Reflective Expert and the Prenovice: Notes on Skill-, Rule- and Knowledge-base Performance in the Setting of instruction and Training. In L. Bainbridge & S.A. Ruiz-Quintanilla (Eds.), *Developing skills with information technology*. Chichester : Wiley.
- Patel, V.L., & Groen, G.J. (1986). Knowledge base solution strategies in medical reasoning. *Cognitive Science*, 10, 91-116.
- Raufaste, E., & Da Silva Neves, R. (1998). Empirical evaluation of possibility theory in human radiological diagnosis. In H. Prade (Ed.), *Proceedings of the 13th Biennial Conference on Artificial Intelligence, ECAI'98* (pp. 124-128). London: John Wiley & Sons, Ltd.
- Raufaste, E., Da Silva Neves, R.M., & Mariné, C. (submitted manuscript). Consonant Support Theory: A psychological theory of confidence judgements in diagnostic reasoning.
- Raufaste, E., Eyrolle, H., & Mariné, C. (1998). Pertinence generation in radiological diagnosis: Spreading activation and the nature of expertise. *Cognitive Science*, 22, 4, 517-546.
- Samuel, S., Kundel, H.L., Nodine, C.F., & Toto, L.C. (1995). Mechanism of satisfaction of search: eye-position recordings in the reading of chest radiographs. *Radiology*, 194, 3, 895-902.
- Shallice, T., & Burgess, P.W. (1993). Supervisory control of action and thought selection. In A. Baddeley & L. Weiskrantz (Eds.), *Attention: Selection, awareness, and control*. Oxford: Clarendon Press.
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of Harmony Theory. In D.E Rumelhart., J.L. McClelland J.L., & The PDP Research Group. *Parallel Distributed Processing. Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, Mass: MIT Press.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11, 1-23.
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435-502.
- Wang, H., Johnson, T. R., & Zhang, J. (1997). UEcho: A model of uncertainty management in human abductive reasoning. Hybrid Technical Report No. 4 (TR-97/ONR-HYBRID-04). The Ohio State University.
- Zadeh, L.A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1, 3-28.