MOVING KNOWLEDGE: THE SO–CALLED PROBLEM OF TRANSFER AND HOW TO REFRAME IT

1. INTRODUCTION

It is a truism to observe that all education and training aims at transfer; university-trained nurses are expected to apply their knowledge successfully in the clinical context of the hospital ward; trainee teachers are expected to translate their academic or theoretical knowledge of, say, classroom discipline into the practice of their first classrooms; apprentice car mechanics are expected to be able to apply their technical knowledge to fixing faulty break lines in different models of cars; in general, organizational training is expected to transfer successfully to the work setting (e.g. Baldwin and Ford 1988).

It is an equally common observation that the wished for transfer does not often seem to materialize. The classical view of the problem of transfer is phrased in terms of a learner's inability to use formal knowledge in another context when it would be appropriate to do so. This notion of transfer is grounded in individual accumulations of knowledge (for an overview see Perkins and Salomon 1994).

In this paper I argue that what is commonly termed the problem of transfer is a pseudo problem because its underlying classical view of cognition and information processing is too narrow. Based on symbol processing, it is at odds with what we know from contemporary accounts of the biological brain's architecture and computational capacities, as discussed in recent cognitive science and modelled by artificial neural nets (ANNs) (e.g. Rumelhart and McClelland 1986; McClelland and Rumelhart 1988; Churchland P.S. 1989; Churchland and Sejnowski 1994; Churchland P. M. 1993; Rumelhart 1998). In particular, the theory of cognition where cognition is characterized as distributed, embodied and embedded (Hutchins 1991, 1995, 1996a, 1996b; Clark 1997; Lakoff and Johnson 1999; Hollan *et al.* 2000; Evers and Lakomski 2000), makes it possible to see that what is considered the problem of transfer is simply a feature that characterizes all learning, where learning is defined as *an organism's/agent's adaptive reorganization in a complex system* (Hutchins 1996a, p. 289).

Of particular importance here are new developments that begin to explore *external cognition* – i.e. the many environmental scaffolds humans devise to help their thinking and problem solving. The supposed results of transfer, i.e. the successful "retrieval" or application of "old" knowledge, or successful task performance and problem solving, causally depend on environmental resources and constraints of the situation and thus merit closer examination. The conception of *task environment* and how to determine what are some features of task environments as *activity spaces* that facilitate and constrain problem solving and coordinating activity, are central to this discussion. The paper concludes with some suggestions regarding what topics and issues organizational theorists ought to address to deal with organizational "sticky" (von Hippel 1994) knowledge.

2. SITUATED COGNITIION AND TRANSFER

Organizational theorists have recently been attracted by a research perspective variously known as situated cognition, learning, or situated action because of its interest in the nature and creation of organizational knowledge (e.g. Rogoff and Lave 1984; Resnick, Levine and Teasley 1991; Lave 1988, 1991; Lave and Wenger 1991; Clancey 1993; Salomon 1997). The work of Greeno, Moore and Smith (1993) stands out because they present an explicit model of transfer whose base assumption is that the agent has learnt to participate in a socially constructed domain of situations that includes the situation where transfer can occur. The authors stress that their conception is to be considered from within the context of a general view of cognition where "learning is considered to be essentially situated, an adaptation of a person or group to features of the situation in which learning occurs. Knowledge – perhaps better called *knowing* – is not an invariant property of an individual … knowing is a property that is relative to situations, an ability to interact with things and other people in various ways. (Greeno *et al.* 1993, p. 99). It follows that the question of transfer in their view is the question "to understand how learning to participate in an activity in one situation can influence (positively or negatively) one's ability to participate in another activity in a different situation. The answer must lie in the nature of the situation, in the way that the person learns to interact in one situation, and in the kind of interaction in the second situation that would make the activity there successful." (Greeno *et al.* 1993, p. 100).

Transfer thus depends on structure in the situation that is primarily socially defined, and that has been included in the person's previous social experience. Of particular importance for their model of transfer are two concepts.. Any activity in which a person is involved is shaped by properties of things and materials in the situation as well as by the characteristics of the person. Support for particular activities, as offered by the joint interaction of those things and materials, are what the writers call *affordances*. Transfer of learning from situation 1 to a different situation 2 requires a transformation of the situation (S2) and "and invariant interaction of the agent within the situation" (Greeno *et al.* 1993, p. 102). Or more formally: " Transfer can occur if the structure of the activity is invariant across the transformation from S1 to S2 with respect to important features that make it successful…" (Greeno *et al.* 1993, p. 102). In addition to invariance in regard to specific features of the situation, the agent has to be able to perceive or be attuned to the affordances in order for transfer to occur.

Unwittingly, this characterisation of transfer goes a long way towards making transfer impossible (see Evers and Lakomski 2000, pp. 130 onwards for detailed discussion; also St. Julien 1997). Becoming highly attuned to the very constraints and affordances of S1 would seem to make it near impossible to "take out" of that situation what might be generalizable to S2 (Bereiter 1997). If the specific constraints and affordances of any situation determine what can be learnt or known, and in the absence of further clarification of the sense in which situations might be said to be invariant, transfer is not possible. More importantly, there does not appear to be a knowing/learning agent who does the knowing. Context dependent cognition is a very important insight, but no context or situation has its meaning written on its sleeve but needs to be interpreted by an agent capable of doing so. To put the matter differently, the identification of what constitutes an affordance or constraint in a situation, or what is an invariant feature between situations, causally depends on an agent's seeing it as such. Such identification in turn depends on the agent's prior learning and, most importantly, on his or her cognitive capacities to be able to make the requisite determinations in the first place.

Writing about *organizational* knowledge in particular, Brown and Duguid (e.g. 2000) join ranks with those theorists, listed above, who advocate a perspective of knowledge production and distribution that is at odds with the prevailing view of knowledge as an individual's property. Arguing for knowledge creation as a socially shared activity, they rightly emphasize that given the very nature of such locally produced knowledge, attempting to move it between communities of practice (Lave and Wenger 1991), or "hybrid communities", is a problem of considerable difficulty which is not sufficiently well recognized. Reliance on technology that was expected to provide streamlined information flows by better search and retrieval systems has proven less than successful to solve the problem. In their view, trying to move the knowledge that was generated in the community's practices without the practices is moving the know–what without the know–how. Brown and Duguid believe that the real or tacit knowledge is locally produced and thus embedded knowledge and that *transfer*ring such knowledge between communities would require, amongst other things, that we abstract it from its embedded context. They believe this to be possible although difficult to accomplish. How this could be done is left open.

For present purposes, the most interesting point which arises from the work of situated cognition theorists is that while they provide an excellent account of why knowledge does *not* transfer there is not much explanation in terms of what does make transfer possible, as indeed can be observed in everyday life. Although the new emphasis on the situatedness of knowledge and its creation in communities of practice are important insights, they end up occupying a place at the other end of the transmission/transfer model occupied by Herbert Simon and colleagues for whom cognition and knowledge is private and individual property. For situated perspective advocates, knowledge and cognition is "all out in the situation" while it is "all inside the skull" for Simon and colleagues. Both positions are overdrawn in important respects.

The so–called question of transfer is phrased both in terms of agents' inability to transfer formal (explicit/propositional) knowledge (classical view) between situations on one hand and in terms of the specificity of the situation and its contingencies which generates *tacit* knowledge (situated action perspectives) on the other. Both sets of approaches implicitly rely on what is called the *sentential view of information processing* (Churchland P.S. 1980; 1989; Churchland P.M. 1993). This perspective (which goes back to Descartes and underlies Simon's symbol processing view) maintains that human cognition consists of acts of symbol processing in the head which, in turn, are externally representable in texts, electronic messages, and the like. Its opposite is the so-called tacit domain which is not amenable to symbolic representation – hence the descriptor *tacit*. This distinction confuses what knowledge *is with how it can be represented.* What easily transfers between groups and communities is abstract or propositional knowledge, which is abstracted from the contexts in which it was created. Tacit knowledge, since it cannot be expressed in symbolic form, shows in the practices, or human action, and cannot therefore be "transferred" in the standard sense. Situated cognition writers have correctly noted this but without providing an explanation of human cognition that takes into account the agent's biological capacities and that we are *embedded* creatures as well. In the next section, a more detailed account of the theory of (distributed) cognition is provided.

4. THE THEORY OF COGNITION

The theory of cognition studies the organization of cognitive systems where what is considered cognitive "encompass[es] interactions between people and with resources and materials in the environment." (Hollan *et al.* 2000, p. 175) The theory is committed to two important theoretical principles, the first being the boundaries of the unit of analysis for cognition; the second relates to the range of mechanisms that can be considered as taking part in cognitive processes. Adherence to both principles differentiates this theory of cognition from other traditional accounts.

As regards the second principle, unlike traditional approaches of cognition that look for cognitive events inside the agent's skull where s/he is assumed to manipulate the relevant symbols, as prominently argued by Simon and co-workers in the Physical–Symbol–Systems–Hypothesis (for discussion see Evers and Lakomski 2000, Ch.3), the theory of distributed cognition is concerned with a broader category of cognitive events that stretch beyond the individual skull. While the material world provides additional memory for thinking to reduce cognitive overload, it also offers opportunities for reorganizing the distributed cognitive system itself. It is the latter point that is not yet sufficiently appreciated, let alone understood.

Distributed cognition is also *socially* distributed in addition to being *embodied*. Indeed, on the present view, social organization, including organizations and culture, is itself considered a form of cognitive architecture (Clark 1997; Strauss and Quinn 1997; Hutchins 1996a, 1996b). The sense of distribution is thus richer than that which primarily considers distribution across the members of a group. An interesting and complex question in this context is where mind ends and the world begins (see Kirsh 1999). In Kirsh's (1999) view, where we draw the boundary depends on what kind of explanation we want to provide, and its level and focus of analysis. Clark and Chalmers (1998) describe this kind of mutually interacting mind/world relationship with its shifting boundaries by means of the concept of *active externalism*, a linkage that can itself be considered a cognitive *coupled system*.

Human agents rely heavily not simply on sophisticated computing resources but on environmental supports generally, an insight with which we have become familiar: Clark's (1997) example of pen and paper to make long multiplication easier; Hutchins' (1995) description of airline pilots' use of external markers on their controls, and the use of the nautical slide rule in ship navigation; Maglio's *et al.* (1999) experiments with Scrabble players who physically rearrange letter tiles for better word recall, and Kirsh and Maglio's (1992) laboratory experiments with agents playing the interactive video game Tetris all support the claim that agents modify their physical environment to make identification faster, cue recall, and to generate mental images faster than would be possible without such external props. In this manner agents save cognitive and potentially wasteful cognitive labour. It seems quite clear from these examples that there is not just interdependence between internal and external structures but that the external actions agents engage in can be most significant in order to simplify mental computation where speed is important. Such external action can be described as *epistemic action*: "*Epistemic actions* are physical actions people take more to simplify their internal problem–solving processes than to bring themselves closer to an external goal state." (Kirsh and Maglio 1994, p. 514). The latter are called "pragmatic" (see Alterman *et al.* 1998), and it is actions of this kind that have provided the main analytical frame for planning and action theory. However, as Kirsh and Maglio (1994) point out, not all actions are of this kind. Acknowledging *epistemic action* as a category of activity has the consequence that if we continue to consider planning in the traditional mode as state–space search, we must redefine the state–space in which planning occurs.

The conceptualisation of the mind/world relationship, the identification of *epistemic action* as a hitherto overlooked complement to *pragmatic action*, and the complexities of determining whether or where to place the boundary between mind and world, are considerations far removed from the way these relationships have been theorized to date. This is most evident in our traditional understanding of (task) environment that can no longer be seen as unproblematically existing outside our bodies. If cognition, or perhaps cognitive systems, are as "spread" or distributed as indicated in the preceding, task accomplishment or enaction is a far more complex affair than previously assumed.

5. HOW TO DETERMINE A (TASK) ENVIRONMENT

The classical conception of task environment, due to the original and highly influential work of Newell and Simon (1972), kept the mind/world boundary analytically separate, a view no longer supported; by the same token, without their brilliant analytical work this relationship would not have come into focus as it has. In discussing it the different theoretical and philosophical commitments of the two research programs of AI on the one hand and the new cognitive science on the other, become clear.

Newell and Simon (1990, pp. 120–121) locate the philosophical foundations of their approach in Meno's Paradox which Plato solved by means of the theory of recollection. In contemporary form, they suggest, a simpler solution to the paradox can be formulated: "To state a problem is to designate (1) a *test* for a class of symbol structures (solutions of the problem), and (2) a *generator* of symbol structures (potential solutions). To solve a problem is to generate a structure, using (2), that satisfies the test of (1)". (Newell and Simon 1990, p. 120) Symbols are at the core of intelligent action and physical–symbol systems consist of sets of symbols or physical patterns such as "chalk marks on a blackboard", for example (Simon 1996, p. 22). The physical–symbol system is the appropriate theoretical framework for modelling intelligent action, and intelligent action in turn shows up best when humans solve problems. Given this kind of formalization, a problem is well defined "if a test exists, performable by the [physical–symbol] system, that will determine whether an object proposed as a solution is in fact a solution". (Newell and Simon 1972, p. 73). By this means intelligence in problem–solving is exemplified in that symbol structures are generated and progressively modified until the solution is reached. (Newell and Simon 1990, p. 119). There are several important ingredients to problem solving, and the one of most concern here is the nature and role of the environment which Newell and Simon term "task environment". In their view:

> The term *task environment*, as we shall use it, refers to an environment coupled with a goal, problem, or task – the one for which the motivation of the subject is assumed.1 It is the task that defines a point of view about an environment, and that, in fact, allows an environment to be delimited. Also … we shall often distinguish the two aspects of the theory of problem solving as (1) demands of the task environment and (2) psychology of the subject. These shorthand expressions should never seduce

> the reader into thinking that as a psychologist he should be interested only in the psychology of the subject. The two are in fact like figure and ground – although which is which depends on the momentary viewpoint. (Newell and Simon 1972, p. 55)

This formalization of the agent's internal representation of the task environment – the problem space – represents an important abstraction because there is a need to delimit what is the relevant problem space in the face of a potentially limitless search space. Here, the interrelationship between problem space and goal, an agent's representation of the goal to the goal itself, is essential for the task environment *because it allows the task relevant to be sorted from the task irrelevant*. Thus Newell and Simon provide an explicit answer, and sorting mechanism, to what may arguably be the most crucial question in the so called problem of transfer. This is the question to which Greeno *et al.* provide no answer but which they clearly recognize as important in their conceptions of affordances and invariants.

6. ENVIRONMENT AS ACTIVITY SPACE

The move from environment to task environment, and secondly, the strong coupling of agent and environment in terms of agents creating their own problem spaces are important insights. The way Newell and Simon formalize these relationships, however, is too limiting in the face of the real capacities of biological human problem solvers, and thus it constrains the nature and construction of their problem spaces. The reason Newell and Simon's examples work reasonably well – playing chess, scrabble, tic–tac–toe, or solving the Tower of Hanoi problem – is that they are all characterized by restricted problem spaces. In each of them, there are only so many legitimate, i.e. task relevant, moves, and there are no degrees of freedom for improvisation. Engaging in these kinds of problem solving activities is something humans do, and do well, but they are not representative of the everyday problems human agents face, problems whose "problem spaces" are much more diverse, as in the working out of schedules and time tables. This cognitive task is one humans perform routinely, but is one that trips up even the fastest and most sophisticated modern super computer.

Every time human agents are confronted with a problem, they are able to draw on vast stocks of knowledge, acquired in different contexts, for different purposes, at different times, and applied to different tasks. The mind/brain sorts what knowledge was relevant *then* from the knowledge relevant *now*. This ability to sift quickly, and seemingly effortlessly, relevant bits of knowledge from irrelevant bits is an ability that cannot in principle be programmed into a linear computer. Although neurons are much slower than the components in modern computers, operating in milliseconds rather than nanoseconds, brain–style computation, as Rumelhart (1998, pp. 209–210) puts it, is still vastly superior despite the time constraint because there are just so many neurons that operate cooperatively and in parallel fashion rather than in serial steps. A computer program, in contrast, would need to be able to sift a practically limitless search space, and be able to make the right connections out of a potentially limitless repertoire. The computational powers of the most modern computer could not cope with this kind of complexity because it is constrained by the manner of its serial step procedure.

It is in this context also important to note that the biological brain, modelled as a connectionist system, does not store knowledge the way a computer does, in the state of certain units in the system from where it can be retrieved on demand. (See Rumelhart 1998 for a brief account of brain architecture and a discussion of the seven major components of a connectionist system; also Bates and Elman 1993). Knowledge in a connectionist system is in the connections between neurons, more specifically in the strengths, or *weights*, between them. Knowledge representation is non symbolic; knowledge is implicit in the processor itself and is created through tuning of connections while processing is going on. It is neither formulated nor stored in the form of symbols. It also follows from the parallel distributed architecture of the brain that knowledge is not in any one location; it is everywhere – i.e. distributed – in the system.

Does this connectionist account mean that a task environment is in principle as limitless as is the search space discussed earlier? If not, what specifications could be determined that would provide us with a more realistic picture of how agents and their environments operate to solve problems? We have initial indications on how to think about these highly abstract matters, and how to progress our investigations (e.g. Kirsh 1999, 2001; Kirsh and Maglio 1994). The fundamental difference between the AI tradition and the theory of cognition, as would have become clear in the preceding comments, is, as Kirsh (1999, p. 5) puts it so nicely, that "… the real world is a place we inhabit rather than visit. We live here and return to it".

An interesting consequence of *active externalism* is that the environment, unlike Simon's formalization, can be conceptualized more realistically as an *active partner*, "a shifting coalition of resources and constraints, some social, some cultural, some computational…" (Kirsh 1999, pp. 1–2). When all these resources and constraints are brought into an appropriate constellation with each other, we have managed to accomplish the task we set out to accomplish. In other words, *coordination* becomes of critical importance, and especially successful coordination that commonly involves several agents, or tasks would not be achieved nor problems solved. This could not be more different from Simon's agent–as–information–processing–system model that permits no such activity, indeed where coordination is surplus to the requirements of the model *by design*, and where the basic emphasis is on *individual* problem solving. The connectionist account also differs markedly from the situated cognition perspective in so far as it offers a *naturalistic* account of what it means to talk about "cognition in the wild" (title of Hutchins' 1996 book), about individual biological organisms/agents who are implicated in a much larger cognitive system which they in part determine, but which also determines them.

Reiterating the fundamental point that while humans are main actors in the force field of interrelated distributed influences, it is these other influences that co–determine action; coordination happens between all resources and constraints including the human agent. Coordination can thus be thought of as a genuine "team effort" where the composition of the team changes depending on which task is to be accomplished, and the ensemble of constraints and resources available. (See Hutchins 1996a, especially Chs. 7 and 8, Learning in Context and Organizational Learning, which provide some excellent examples of cognitive work as discussed here, including a connectionist account of organizational learning).

Coordination can easily be seen in terms of physical arrangements, ongoing lines of communication, and institutional structures that are readily identifiable as coordinating mechanisms and processes. But at a deeper level, work activities get coordinated by explicit symbolic means, as well as non–symbolic and implicit means as well. The latter might include picking up on the manager's finger tapping, for example, that might help speed up a meeting (itself a mechanism of coordination); examples like these abound in everyday environments where much coordination of activity happens in non–symbolic form and mostly goes unnoticed. It follows from the preceding comments that the mechanics of coordination become of central importance in the explanation and understanding of problem solving or task accomplishment.

7. CONCLUSION

This naturalistic view of cognition developed in the preceding pages draws the radius of human cognition much wider than we are accustomed to doing, given the predominance of the classical, "cognition in the head" view. The theory of cognition describes the cognitive force field in which agents manoeuvre and which can be considered a web of coordination of artefacts, media and processes both internal and external to the agent. In thus expanding the unit of analysis of intelligent or cognitive action in the world what makes for successful learning or adaptation becomes a question of the internal configuration of agents, more precisely their brain states, the reciprocal interactions with the constraints and affordances of the situation, and the mechanisms of coordination pertaining between all of them. On this account, there is no transfer of stocks of invariant knowledge from (A) to (B), there is always "learning on site", which is aided (or hindered) by past experience and mediated importantly by the constraints and affordances of the physical, material, spatial and other features of the situation, be it in a workshop, on a flight deck, or in the bank manager's office.

Given the previous discussion, what is described as the "problem" of transfer can be considered more productively as a *problem of the mechanics of coordination,* as discussed in the previous pages. If we follow the implications of distributed cognition we cannot fail but notice just how much more complex the agent's embeddedness in and interrelationship with the world is, and that the theoretical and methodological tools we currently have to analyze the agent's actions in the world are still in their infancy. Speculative though much of this work still is, it signals research directions and programs that simply have not yet been recognized in their full significance and potential impact.

What of Brown and Duguid's (2000) specific question of how "sticky" knowledge can be shifted? In the language of the neural net account successful "transfer" becomes a question of whether the appropriate activation patterns came about in the new context, or not. If they did not then it might be concluded that not enough patterns were there to be activated, in other words, the person had not learnt enough in order to "make the connection", or the existing connection strengths were not sufficient to activate the pattern. Activation also depends on the specifics of the context/situation. Here the question of similarity or invariant features between situations becomes important, an issue clearly recognized by Greeno and his colleagues. What it means, though, to speak of invariance or similarity between work settings, from a distributed cognition perspective, and whether there can be invariant structures of different work settings that may be generalized, and may thus lead to cognitively appropriate workplace designs, are questions that cannot be answered at present given the current limitations of our theorizing about such things.

Understanding cognition and structuring contexts better nevertheless are key tasks to undertake if we are to make headway with using organizational knowledge to best effect. It should be clear from the preceding discussion that the concept of transfer, as commonly understood, is not helpful to accomplish this goal. It drops out of the explanatory framework as not supported by the currently best account of knowing and cognition we have. As a consequence, we ought to occupy ourselves with conducting more fine-grained cognitive case studies in work settings; pay more attention to the work space as a cognitive extension of human cognition, including the use of electronic means of communication; undertake work space design studies, and enquire into the requirements of effective and efficient cognitive workflow. The conditions of knowledge creation within highly complex and interconnected systems need to be understood better in order for us to be able to design optimal structures that maximize organizational knowledge and its use.

References

Alterman, R., R. Zito–Wolf, and T. Carpenter (1998). “Pragmatic action”, *Cognitive Science*, 22(1), 53–105.

Baldwin, T. and J. Ford (1988). “Transfer of training: A review in directions for future research”, *Personnel Psychology*, 41, 63–105.

Bates, E.A. and J.L. Elman (1993). “Connectionism and the study of change”, in M. Johnson (ed.) *Brain Development and Cognition: A Reader.* (Oxford: Blackwell).

Bereiter, C. (1997). “Situated cognition and how to overcome it”, in D. Kirshner and J. A. Whitson (eds.) *Situated Cognition.* (Hillsdale, NJ: Lawrence Erlbaum Associates).

Brown, J. S. and J. Duguid (2000). “Organizing knowledge”, in D.E. Smith (ed.) *Knowledge, Groupware and the Internet*,

Churchland, P. S. (1980). “Language, thought, and information processing”, *Nous*, 14, 147–170.

Churchland, P.S. (1989) *Neurophilosophy*. (Cambridge, MA: The MIT Press).

Churchland, P. M. (1993). *A Neurocomputational Perspective*. (Cambridge, MA: The MIT Press).

Churchland, P.S. and T.J. Sejnowski (1994). *The Computational Brain.* (Cambridge, MA: The MIT Press).

Clancey, W. J. (1993). “Situated action: A neuropsychological interpretation, response to Vera and Simon”, *Cognitive Science*, 17, 87-116.

Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again.* (Cambridge, MA: M.I.T. Press).

Clark, A. and D. Chalmers (1998). “The extended mind”, *Analysis*, 58,10–23.

Evers, C.W. and G. Lakomski (2000). *Doing Educational Administration*. (Oxford: Elsevier).

Greeno, J. G., J. L. Moore, and D.R. Smith (1993). “Transfer of situated learning”, in D. K. Detterman and R. J. Sternberg (eds.) *Transfer on Trial*: *Intelligence, Cognition, and Instruction.* (Norwood, NJ: Ablex Publishing).

Hollan, J., E. Hutchins, and D. Kirsh (2000). “Distributed cognition: toward a new foundation for human–computer interaction research”, *ACM Transactions on Computer–Human Interaction*, 7(2),174–196.

Hutchins, E. (1991). “The social organization of distributed cognition”, in Resnick, L.B., Levine, J.L. and Teasley, S.D. (eds.), *Perspectives on Socially Shared Cognition*, (Washington, DC: American Psychological Association).

Hutchins, E. (1995). “How a cockpit remembers its speeds”, *Cognitive Science*, 19, 265-288.

Hutchins, E. (1996a), *Cognition in the Wild*, (Cambridge, MA: The MIT Press).

Hutchins, E. (1996b). “Organizing work by adaptation”, in Meindl, J.R., Stubbart, C., and Porac, J.F. (eds.), *Cognition Within and Between Organizations*. (Thousand Oaks, CA: Sage).

Kirsh, D. (1999). “Distributed cognition, coordination and environment design”, in *Proceedings of the European Cognitive Science Society*, pp.

Kirsh, D. (2001). “The context of work”, *Human Computer Interaction*, pp.

Kirsh, D and P. Maglio (1994). “On distinguishing epistemic from pragmatic action”, *Cognitive Science*, 18, 513–549.

Kirsh, D and P. Maglio (1992). “Reaction and reflection in Tetris”, in J. Hendler (ed.), *Artificial Intelligence Planning Systems*. (San Mateo, CA: Morgan Kaufman).

Lakoff, G. and M. Johnson (1999). *Philosophy in the Flesh*. (New York, NY: Basic Books).

Lave, J. (1988). *Cognition in Practice.* (Cambridge: Cambridge University Press).

Lave, J. (1991). “Situating learning in communities of practice”, in L. B. Resnick, J. L. Levine and S. D. Teasley (eds.) *Perspectives on Socially Shared Cognition*. (Washington, DC: American Psychological Association).

Lave, J. and E. Wenger (1991). *Situated Learning.* (Cambridge: Cambridge University Press).

Maglio, P.P., T. Matlock, D. Raphaely, B. Chernicky, and D. Kirsh (1999). “Interactive skill in Scrabble”, *Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society*. (Mahwah, NJ: Lawrence Erlbaum).

McClelland, J. L. and D.E. Rumelhart (1988). *Explorations in Parallel Distributed Processing.* (Cambridge, MA: The MIT Press).

Newell, A. and H.A. Simon (1972). *Human Problem Solving.* (Englewood Cliffs, NJ: Prentice Hall).

Newell, A. and H.A. Simon (1990). “Computer science as empirical inquiry”, in M.A. Boden (ed.) *The Philosophy of Artificial Intelligence*. (Oxford: Oxford University Press).

Perkins, D.N. and G. Salomon (1994). “Transfer of learning”, in T. Husén and T. N. Postlethwaite (eds.) *The International Encyclopedia of Education*, Second Edition (Oxford: Pergamon).

Resnick, L. B., J.M. Levine, and S.D. Teasley (1991) (eds.) *Perspectives on Socially Shared Cognition.* (Washington, DC: American Psychological Association).

Rogoff, B. and J. Lave (1984) (eds.) *Everyday Cognition: Its Development in Social Context*. (Cambridge, MA: Harvard University Press).

Rumelhart, D. E. and J.L. McClelland (1986) (eds.) *Parallel Distributed Processing.* Volumes I & II. (Cambridge, MA: The MIT Press).

Rumelhart, D.E. (1998). “The architecture of mind: A connectionist approach”, in P. Thagard (ed.) *Mind Readings.* (Cambridge, MA: The MIT Press).

Salomon, G. (1997). (ed.) *Distributed Cognitions, Psychological and Educational Considerations*. (Cambridge: Cambridge University Press).

Simon, H.A. (1996). *The Sciences of the Artificial.* (Cambridge, MA: The MIT Press). Third edition.

St. Julien, J. (1997). “Explaining learning: the research trajectory of situated cognition and the implications of connectionism” in D. Kirshner and J. A. Whitson (eds.) *Situated Cognition*. (Hillsdale, NJ: Lawrence Erlbaum Associates).

Strauss, C. and N. Quinn (1997). *A Cognitive Theory of Cultural Meaning*. (Cambridge: Cambridge University Press).

von Hippel, E. (1994). “’Sticky information’ and the locus of problem solving: implications for innovation”, *Management Science*, 40(4), 429–439.