Models of Cognition, the Contextualisation of Knowledge and Evolutionary Organisational Theory

Edward Lorenz
Centre d'Etudes de l'Emploi
Noisy-le-Grand
Lorenz@alum.mit.edu

Preliminary version: not to be cited without author's permission

Paper prepared for the Nelson and Winter Conference, Aalborg, June 12-15, 2001
Part of the material presented in this paper was presented at the DRUID Winter Conference, Klarskovgaard, 18-20 January 2001 and at the Workshop on Cognition and Evolution in the Theory of the Firm, Jena, September 25-27, 2000. I would like to thank the participants for their comments and suggestions. Thanks also go to Christian Bessy whose remarks have been invaluable in improving the quality of the argument.

Introduction

There now exists a considerable body of evidence pointing to persistent differences in the innovative behaviours of companies located in different national spaces. At the level of sectors and countries, these differences can be seen in the relative stability over time of such indicators as rates of innovation and leads and lags in the adoption of innovations. (Dosi, 1991, 1999; Archibugi and Michie eds., 1997; Pavitt and Patell (1988). The durability of such differences in innovative behaviours suggests that the ability of companies to absorb, generate make use of knowledge is strongly influenced by the social or institutional context in which they operate. What can evolutionary theory tell us about the social embeddedness of organisational knowledge and learning?

Theoretical work on the firm from an evolutionary or knowledge-based perspective[1] has, of course, developed a number of useful concepts for explaining diversity as such among firms. One key idea is that the knowledge underlying organisational routines and capabilities is distributed across organisational members in firm-specific team configurations. (Lippman and Rumelt, 1982; Barney, 1986). Such firm-specificity, combined with the tacit nature of knowledge, makes firm capabilities difficult or impossible to imitate by competitors. Another central idea is that due to cognitive limits on rationality organisational search tends to be local or sharply circumscribed by existing capabilities. This confers a strong path-dependent character to learning processes within organisations, serving to reinforce the existence of diversity among firms over time. (Nelson and Winter, 1982; Teece, Pisano and Shuen, 1997; Dosi, 1991). However, little progress has been made in conceptualising the way national institutional arrangements interact with and shape firm-level processes of knowledge use and development.

The basic thesis of the paper is that theoretical advance in this area will depend on some reconsideration of the underlying models or conceptions of human cognition and reasoning that can be found in the evolutionary literature on the firm. I would argue that a serious inquiry into the socially embedded nature of organisational learning requires a conceptual language for thinking about the following two things. The first is how representations, including the heuristics we apply in problem solving, are shaped by the wider social context

[1] I am using these labels in a broad-brush manner to refer to the resource-based perspective associated, for example, with the work of Barney (1986) and Lippman and Rumelt (1982); the dynamic capabilities perspective associated with such authors as Teece, Pisano and Shuen (1997) and Dosi, Nelson and Winter (2000); and

in which we live. The second is how the context that matters stretches beyond our interactions in particular local settings to include wider societal conventions and institutions. In my view, neither of the two dominant conceptions of human cognition informing the evolutionary literature on the firm are adequate for this research programme.

The first, and by far the dominant one in the economic literature, is the information processing or physical symbol system approach associated with the work of Newell and Simon (1958, 1972). It was translated into the realm of organisational theory through the work of Simon and March (1958) and Cyert and March (1962), and has been most fully developed in the contemporary research programme in computational organisational theory (COT) associated with such authors as Carley (1995) and Masuch (1992). It argues that human cognition and reasoning may be analysed and understood independently of social and historical context.

The second basic influence is the situated knowledge and learning perspective, often associated with the work of such authors as Lave (1988), Suchman (1987) and Scribner, (1984). This conception of human thinking and problem-solving has been extended into the realm of organisational theory in research analysing organisational learning in terms of communities of practice (Brown and Duguid, 1991; Lave and Wenger, 1991; Snyder, 1997; and Wenger, 1998a and 1998b). While it takes the opposite stance that knowledge and learning can only be understood in relation to an actional context, it appears to be firmly committed to the view that local context is determinant. This gives rise to a conception of knowledge as fleeting and necessarily tied to human activity in a particular place. For this reason, it appears to close the door on a serious enquiry into more generalised forms of knowledge that stretch across different local settings.

In order to avoid possible confusion, I should make it clear that I am not suggesting that the only way ideas about the importance of social context for knowledge have entered the evolutionary or knowledge-based literature is via the work of the authors cited above. For example, Nelson and Winter's (1982, pp.78-79) develop the idea that planning and implementing skills is dependent on detailed features of the social context. This is quite similar to the idea developed by Lave (1988) and others that reasoning processes, including

computational approaches associated notably with Carley (1995) and Masuch (1992).

basic cognitive aptitudes, are highly sensitive to the local context in which they exercised. The main influence on Nelson and Winter, though, would appear to be Polayni's (1969) phenomenological account of expertise and skilled performance. Clearly one can see in the organisational literature the impact of multiple and possibly overlapping influences.[2] The situated knowledge and learning perspective, however, neatly summarises a view that in one manner or another has worked its way into the evolutionary literature.

The main purpose of this paper, though, is not to attempt a mapping from discussions of knowledge and learning in the evolutionary literature to the two conceptions of human cognition I have identified. Rather, my aims are firstly to bring out in a general way the implications of these two approaches to human cognition for how processes of knowledge use and learning within organisations are conceived. I will point to the reasons why neither approach lends itself to a serious analysis of how these processes are embedded in societal institutions. In doing this, I shall refer to two contemporary research programmes into problems of organisation change and learning that illustrate in a relatively 'pure' manner the underlying models of cognition I have identified: the computational approach on the one hand, and the communities of practice approach on the other.

Following this exercise, I turn to the recent work of Edwin Hutchins (1995) on the cognitive properties of teams and organisations. Although his research is rooted in the situated cognition and learning perspective, in a number of significant respects it overcomes the limitations of this tradition. Firstly, it addresses some of the concerns of those wedded to an information processing conception of cognition by relating the coordinated behaviour of teams to operations on symbolic representations. Secondly, it provides a framework for seeing how the cognitive attributes of teams, including their use of representations, are embedded in a wider social or institutional context.

Organisations as Information Processing Systems

The use computational methods in the modelling or various theoretical schemes is often justified in terms of the intractability of systems of complex equations, or in terms of the opportunities it provides for performing experiments which cannot easily be carried out

[2] For example, researchers on situated cognition in the US have clearly been influenced by the interactionist sociology of Blumer (1969). One of the philosophical influences on symbolic interactionism is the pragmatic philosophy of Dewey and Mead. Further, it is common to draw connections between pragmatism in general, and

within the confines of the laboratory. In the case of computational organisational theory, however, the justification for using these methods has a deeper epistemological basis. Fundamentally it stems from the belief that organisations, much as the human mind and the computer, are complex information processing systems. Correspondingly, it is argued, much can be learned about how organisations use and develop knowledge by studying the characteristics of computer programmes.

Newell and Simon elaborated the information processing conception of human cognition in the 1950s and 1960s. It holds that knowledge consists in representations or collections of abstract symbols which are stored in the mind. Human reasoning and problem-solving behaviours may be characterised and understood in terms of computations performed on the symbols independently of the semantic meanings attached to them. Human intelligence, then, is portrayed as modular in character in the sense that its properties may be understood independently of social context.

It is worth examining is some detail the 1958 article co-authored by Newell, Shaw and Simon, "Elements of a Theory of Human Problem Solving". This piece, which Simon and Newell saw as providing empirical support for their information-processing conception of human thinking, provides a clear statement of the epistemological stance regarding the adequacy of computer programs for studying the properties of real-life problem-solving. They are careful to specify that the identification they draw between computer programs and human problem solving does not go all the way down. As they note, they "are not comparing computer structures with brains, nor electrical relays with synapses" (p. 8). Rather, the argument is that the computer and the human brain constitute two different physical supports or types of hardware for executing the same type of information processes:

> Digital computers… can, by appropriate programming, be induced to execute the same sequences of information processes that humans execute when they are solving problems. Hence, as we shall se, these programs describe both human and machine problem solving at the level of information processes (Newell, Shaw and Simon, [1958] 1989, p. 8)

In a somewhat circular manner, the evidence assembled to prove this proposition was the output of a computer program, the Logic Theorist (LT), which the authors had designed to

---

Mead's work in particular, and the phenomenology of Husserl.

solve simple problems in logic. Thus, their ability to program a computer to solve certain sorts of problems in a way that mimed some of the known stylised facts concerning how humans solve such problems was offered as proof of a theory of human cognitive processes inspired by the operation of a computer.[3]

Simon and Newell later developed a more formal statement of their view in the idea that the basic architecture of human cognition is a physical symbol system. Newell and Simon (1976) defined such a system as follows:

> A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components in another type of entity called expression (or symbol structure)… Besides these structures the system also contains a collection of processes that operate on expressions to produce other expressions: processes of creation, modification, reproduction, and destruction…. A Physical Symbol System has the necessary and sufficient means for general intelligent action (1976, p. 116) [4]

Given the ambition to put cognitive science on a rigorous scientific basis, it was seen as essential that the properties of physical symbol systems, as with any formal systems, could be analysed in a manner untainted by contingent social or cultural affects.[5] As Newell et.al. (1989, p. 107) put it, "Symbol systems are an interior milieu, protected from the external world, in which information processing in the service of the organism can proceed."

This basic understanding of human reasoning and problem solving was then extended to the realm of organisational decision-making by Simon and his colleagues at the Carnegie Institute of Technology in the 1950s and the 1960. Simon set out the basic research agenda in a 1957 address to the Operations Research Association of America.

> Operations research has made large contributions to those management decisions that can be reduced to systematic computing routines. To date, comparable progress has not been made in applying scientific techniques to the judgmental decisions that cannot be so reduced. Research in the past three years into the nature of complex information processes in general, and human judgmental or heuristic thinking processes in particular, is about to change this state of affairs radically. We are now

[3] For a related discussion of Simon's assumption that both the computer and the human brain support similar information processing activity, see Lazaric and Mangolte (1999). These authors draw on Merleau-Ponty's work in phenomenology to argue that progress in the analysis of routines require fuller consideration of the way knowledge may be incorporated in the human body.
[4] Newell et.al (1989, p. 103) caution that the architecture should not be confused with a representation of the external world. The architecture supports such a representation but does not itself provide it.
[5] For a related dicussion of Simon's work

> poised for a great advance that will bring the digital computer and the tools of mathematics and the behavioural sciences to bear on the very core of managerial activity – on the exercise of judgement and intuition; on the processes of making complex decisions (Simon and Newell, [1957] 1982, p. 382).

Simon's 1958 volume with March, *Organizations*, laid out in some detail the view that organisational decision-making, much as human decision making, can be analysed in terms of computer programmes. Many of the elements of the theory of human problem solving developed in his 1958 article with Shaw and Newell reappear in this volume. Thus organisational problem-solving, as with human problem solving, is characterised as involving a certain amount of randomness which nonetheless is kept in check by the use of certain procedures or heuristics (March and Simon, [1958] 1993, pp. 198-201). Similarly, both organisational and human problems solving are described as being hierarchical in character, in the sense that decision making proceeds in stages rather than attacking the whole problem in all its complexity (pp. 211-212).

The 1960s saw a first wave of computational models of organisational behaviour inspired by this research agenda, starting with those published in Cyert and March (1963) and culminating in the famous simulation of the 'garbage can' model of ambiguity in organisational decision-making by Cohen, March and Olsen (1972). Despite the interest which this work created around the issues of bounded rationality and complexity in organisational decision making, it did not generate a sustained programme of research in computational organisational theory. In a recent comment on the sporadic advance of computational modelling, Cohen (1998, p. ix) suggests that part of the explanation may lie in the poor match that existed between the control processes of traditional computing languages (typically FORTRAN), involving iterative loops and conditional branching, and actually occurring organisational coordination mechanisms. This was arguably no where more evident than in his own 'garbage can' model, which as a number of critics have observed is curiously devoid of organisational structure (Padgett, 1980; Carley 1986).

The late 1980s marked the beginning of a new wave of computational models in the field of organisational behaviour. These recent modelling efforts display in various ways the influence of prior developments in the field of artificial intelligence. Firstly, as with the symbolic artificial intelligence community more generally, the information processing or physical symbol system hypothesis is invoked to justify the use of computational methods. Blanning, in his 1992 AI model of organisational decision making, provides an exceptionally clear statement of the position,

> This hypothesis [human organizations are physical symbol systems] derives from the notion that organizational structure and knowledge, both implicit and explicit, impose a degree of formality in information processing and decision making that allows organizations profitably to be viewed as formal systems whose internal workings can be described in symbolic form (1992, p.).[6]

Secondly, the current wave of modelling exercises has employed recent techniques and concepts developed in artificial intelligence. These include the use of genetic algorithms to explore adaptive learning processes, work in distributed artificial intelligence, and object orientation in programming. The latter two of these, as Cohen (1998, p; x) observes, respond to some of the limitations noted above of traditional programming languages.

> Objects provide a highly natural way of implementing model agents who have specialised capabilities and subtle, implicit, networks of interaction. This is a far more congenial framework within which to express intuitions about organisational processes…. A second development is the growing interest in organisational questions shown by researchers in the field of distributed artificial intelligence, who have recognized that distributed computers and programs must deal with many issues that are profoundly similar to those facing human organizations (Cohen, 1998, p. x).

These developments have arguably given new purchase to the idea that a realistic account of knowledge use and decision-making processes within organisations can be provided by studying the properties of computer programmes. They give rise to an ahistorical and asocial conception of organisational learning in terms of algorithmic computations taking place strictly within Newell's 'protected milieu'. In the following section I will turn to an alternative school of thought which starts from exactly the opposite premise, namely that knowledge is inevitably situated relative to an external context.

[6] Blanning's position is by no means idiosyncratic. See, for example, the introduction to the recent edited volume on computer simulation of organisational behaviour by Preitula, Carley and Les Gasser (1998, p. xiii): "Computational organisational theory (COT) is the study of organisations as computational entities. COT researchers view organisations as inherently computational because they are complex information processing entities." Also see Masuch and Lapotin (1989, p. 44), "This hypothesis [the physical symbol system hypothesis] – corroborated by the broader experience of AI models of decision making … – serves as the point of departure for the design of a new, generic model of organizational decision making."

Theories of Situated Knowledge and the Communities of Practice Approach

The notion of a 'community of practice' grew out of the work on the situated nature of knowledge and learning by John Seely Brown, Jean Lave, Lucy Suchman and their associates at the Palo Alto Institute for Research on Learning in the 1980s. In terms of its intellectual roots, it is probably fair to say sociological theories of situated practice and action were more important influences on their work than psychological theories. For example, the methodological premises of Chicago school symbolic interactionism, associated with Herbert Blumer, have much in common with the view of Brown, Lave and Suchman that knowledge is necessarily tied to an actional context. Thus, Blumer (1969, p. 2) argued that the meanings one attaches to the behaviour of others arises out of social interaction and that these meanings are modified through an interpretative process in course of dealing with others. The reliance on ethnographic methods similarly follows in the tradition of Chicago School qualitative sociology, notably the work of Robert Park and W.I. Thomas. A focus on the everyday or lay methods used by individuals to make sense of what others do and say as tool for exploring the nature of knowledge and representation, especially present in Suchman (1987), draws inspiration from Garfinkel's (1967) ethnomethodology.

While the dominant intellectual roots of the situated learning perspective can be found in sociology, there was a clear concern to draw out the implications for mainstream cognitive science and for theories of knowledge and learning inspired by it. The challenge they posed to the mainstream approach was to contest the view that knowledge consists of representations that are literally stored within the mind, presumably in long-term or semantic memory, and that intelligence and problem-solving capabilities may be understood in terms of formal operations on those representations. The alternative view they elaborated held that knowledge and problem-solving abilities are necessarily situated relative to a local context that includes not only people's interrelations with one another, but also an artefactual context. The way the artefactual context is experienced may differ from one individual to another. For this reason, the structural framework in which people act is necessarily constructed and malleable. (Lave et.al., 1984, pp. 71-72).

In developing this view, researchers at the Institute for Research on Learning for the most part by-passed the question of the internal architecture of human cognition. In particular, they did not seek to challenge the mainstream view by mustering the psychological or neurobiological evidence that might support an alternative description of internal cognitive architecture, in the sense that Newell and Simon used the term.[7] Rather, they challenged the mainstream perspective more indirectly and in terms of providing evidence of observable behaviours that contradict the consequences that should logically follow from the information processing conception of human cognition. For example, Lave (1988) and Scribner (1984), amongst others, have relied on ethnographic methods to demonstrate how basic cognitive aptitudes, such as the ability to carry out simple mathematical operations, are grounded in an individual's contextualised experience. They show that artefacts, such as the arrangement of materials or physical plant layout, may support perceptual inferences that allow individuals to solve problems without having recourse to higher order mathematical forms of reasoning. This in turn serves to challenge the view that experimental laboratory methods might serve to reveal a subject's 'true' cognitive aptitudes.[8]

Scribner's ethnographic study of the problem-solving capabilities of product assemblers or preloaders in a dairy plant can serve to illustrate the methodology. As she observes, the basic tasks of the preloaders do not appear to be especially cognitively rich and certainly they do not involve any esoteric body of knowledge. Based on order forms detailing the products that each wholesale driver has ordered for his day's deliveries, the task of the preloader is to locate and assemble these products in full cases or partially full cases and to transport them to a common assembly area near a moving track that circles the refrigerated warehouse. When all the items are assembled for a driver, they are pulled onto the track and transported to the loading platforms.

Task performance is determined in the first instance by such cognitive aptitudes as the ability to retain in memory the quantities given on the order sheet and the ability to perform simple addition and subtraction operations. In order to assess how the cognitive skills displayed by preloaders take shape in the course of participation in socially organised practices, Scribner designed a series of simulation exercises in which the product assembly solutions of preloaders were compared with those adopted by a sample of tenth-grade students and three novice occupational groups. The results showed that experienced preloaders tended to carry out problem solving in a different way from the inexperienced groups. Office workers, for

[7]See, however, Clancey's (19) work in the area of robotics and artificial intelligence. He seeks to provide empirical evidence of a close coupling between perception, inference and action rather than treating these three processes as occurring in separate and independent times as in the physical symbol system hypothesis.

example, tended to rely heavily on numerical solutions involving counting operations. In contrast, the preloaders adopted 'shortcuts' working directly from their visual perception of the physical displays of products and crates which provided the necessary quantitative information (e.g. automatically knowing that one layer of half-pints gives 16 or that two rows of quarts is 8). As Brown et.al. (1989) have observed, Scribner's "dairy loaders used the configuration of crates they were filling and emptying almost like an elaborate abacus." This resulted in a superior rate of task completion on the part of preloaders, who were able to correctly fill orders without recoding the information gained from visual inspection into the number system.

The insistence on the importance of local artefactual arrangements and contextualised experience in work such as Scribner's inevitably gives rise to an image of knowledge as something which is fleeting and necessarily tied to a particular place. There is little scope in such an account for developing a serious account of knowledge that stretches across particular times and places in a way that might underpin the existence of widely shared norms, rules and conventions.[9] To do justice to the situated knowledge research agenda, though, the question of the relation of the global to the local, or of the general to the particular, is not entirely ignored. Some effort is made to carve out a middle ground, as in the following quote from Lave and Wenger's 1991 monograph, *Situated Learning*:

> Generality is often associated with abstract representations, with decontexualisation. But abstract representations are meaningless unless they can be made specific to the situation at hand. Moreover, the formation or acquisition of an abstract principle is itself a specific event in specific circumstances. Knowing a general rule by itself in no way assures that any generality it may carry is enabled in the specific circumstances in which it is relevant (pp. 33-34).

One can see in this quote a certain tension in how the relation between the general and specific are handled. On the one hand, there is a desire to recognise the reality of generally held representations and rules, but there is simultaneously a desire to argue that such things only have meaning in relation to peoples' local contextualised experience. The local level, then, is what ultimately determines the way ideas and principles are formed and take shape. This in turn would appear to discourage any interest in undertaking a serious inquiry into the

[8] See Brown, Collins and Duguid (1989) for an overview of this and related literature.
[9] There is, of course, nothing new in this criticism. It is the standard one that has been levied in one way or another against symbolic interactionism, ethnomethodology and other currents of thought in sociology that ground their analysis in the situated nature of practice. See Fisher and Strauss (1978, pp. 485-86) for this

mechanisms that might serve to transmit and diffuse knowledge beyond the confines of particular places and times. At the end of day, the stance elaborated in the situated knowledge perspective can only lead one to doubt the relevance of such notions as societal norms and conventions for an analysis of human or organisational behaviour.

The notion of a 'community of practice' can be seen as an effort, not entirely successful in my opinion, to overcome this limitation of the situated knowledge perspective. The concept was initially developed more as a contribution to a theory of learning in general than as a contribution to a theory of organisational learning (Brown and Duguid, 1991; Lave and Wenger, 1991). The basic idea is that a community of practice consists of people bound together by informal relations who share a common practice. Around this practice they create shared resources in the forms of routines, language and artefactual meanings. These resources may be transmitted from 'old timers' or masters to 'newcomers' or apprentices via the latter's participation in the practices of the community.

By arguing that shared practices give rise to well-defined communities and by identifying a mechanism for the transmission and reproduction of knowledge, the communities of practice literature offers an image of knowledge that is at once less fleeting and less tied to particular places than that developed in the situated knowledge approach. Communities have boundaries, their members share codes or language and they function according to established routines. By embedding knowledge in this way, it becomes possible to use the notion of communities of practice to address some of the concerns of organisational behaviour and strategy theorists. The extension of the concept into the domain of strategy research is largely due to the work of Etienne Wenger (1998a) and those inspired by his research. For example, Wenger (1998b) has been at some pains to define a community of practice relative to the notions of business unit and teams. Snyder (1998) has recently argued that communities of practice provide a foundation concept for analysing the development and transmission of organisational capabilities or competences.

While the communities of practice perspective distances itself from the highly particularised conception of knowledge developed in the situated knowledge perspective, I would nevertheless argue that it falls short of providing a framework that can account for the more

---

traditional criticism raised against symbolic interactionism.

widely-held forms of knowledge that might underpin national specificity in organisational behaviours. Within the logic of the communities of practice framework, more generally-held knowledge can only be accounted for by the dialogues or discourses that take place between the members of overlapping of contiguous communities. But as Wenger (1998, p. 133) has observed, “the discourses that tie communities of practice into broader constellations do not replace practice.”

This gives rise to a tension in treating the relation of the general to the specific, the global to local, that is not so very different from that observed in the situated knowledge perspective:

> What transpires is that knowing is defined only in the context of specific practices, where it arises out of the combination of a regime of competence and an experience of meaning. Our knowing – even of the most unexceptional kind – is always too big, to rich, too ancient, and too connected for us to be the source of it individually. At the same time, our knowing – even of the most elevated kind – is too engaged, too precise, too tailored, too active, and too experiential for it to be just of a generic size. (Wenger, 1998, p. 141).

As with the situated action perspective, the reality of generally held knowledge is not denied in this account. However, it is simultaneously affirmed that knowledge only has meaning for particular individuals relative to their local practice. This would appear to have discouraged any interest in undertaking a serious empirical inquiry into the mechanisms that might serve to transmit and diffuse it beyond the confines of particular practices. In what follows, I will argue that Hutchin’s (1993) work, though rooted in the situated action and communities of practice perspectives, takes an important step towards overcoming these limitations by treating the relation of the general to the specific in a manner that doesn’t simply reduce the general to a backdrop context to be interpreted or positioned relative to based on one’s localised experience and practice.

Culturally Situated Cognition and Ship Navigation

Hutchins' (1995) study is concerned with the work of ship navigators on a US Navy vessel. The navigation team is responsible for determining, charting and logging the ship's position at various time intervals. Correctly establishing a fix or the ship's position is of vital importance to the safety of the ship, especially when entering a harbour and in close proximity to the land. For this reason the degree of error in the fix can be directly read from the result itself: three visual bearings are taken which translate into three lines of position on the ship's chart, thus forming a triangle. The quality of the fix is determined by the size of the triangle formed by the three lines, since if the three lines do not intersect exactly the same point the ship's position is to some degree uncertain.

Navigation work, at least on large vessels such as that examined by Hutchin's, is typically distributed across the members of a team. In the case of the Navy vessel studied by Hutchin's, the computational process of determining the ship's position, the so-called fix cycle, begins with the pelorus operators observing visual bearings of landmarks with a special telescopic sightings device called an alidade. The ship's two alidades, which are located on the port and starboard wings of the vessel outside the bridge, provide angular representations of the ship's position relative to the landmarks and to true north. The printed scales on each alidade's gyrocompass permit to translate these angular representation into digital representations, which are then communicated, either by telephone connection or physically by the pelorus operators, to the ship's navigation recorder, who enters them in the bearing record log. This digitally represented bearing is subsequently communicated to the navigation plotter who represents it as an angle on a one-arm protractor called a hoey. The hoey, once configured, is brought into coordination with the navigation chart, thus reproducing the angle measured between the ship and the landmark in the world as a representation in the space of the chart's latitude and longitude grid.

As this brief sketch indicates, the work of the navigation team can be thought of as a process involving the progressive translation of a visual bearing into a series of angular and digital representations through the use of various physical artefacts in accordance with a well specified algorithm. Hutchins develops a theoretical perspective on this process in terms of the propagation of representational states across a series of representational media. The representational media include both physical objects and tools (e.g. the alidade, the hoey and the chart) and the members of the navigation team. The representations of the position of the ship take different forms in the different media: digital representation in the bearing record log, the memory of that representation in the mind of the navigation bearing recorder, the angular representation in the hoey, etc. Representational states are propagated from one media to another by bringing the states of media into coordination with one another. The tools and other media which are brought into coordination permit difficult tasks to transformed into ones that can be done by the manipulation of simple physical systems or possibly by the mental simulation of the manipulation simple physical system. As Hutchins stresses, “the tools are useful because the cognitive processes required to manipulate them are not the computational processes accomplished by their manipulation” (p. 170-71). This implies that the cognitive properties (e.g. capacity for error detection and problem solving) of the system made of team members in interaction with various artefactual media are not the same as the cognitive properties of the individual members.

Performance in this cognitive system depends on what Hutchin’s, following Norman (1993), refers to as cognitive artefacts, which mediate the realisation of tasks. For example, if a written procedure is used by team members in carrying out the fix cycle, the procedure is a mediating cognitive artefact which is external. However, the procedure may be memorised and exist as an explicit representation within the media of an individual’s internal memory. In this case, the mediating cognitive artefact is internal rather than external. It is in the same class of phenomena as our knowledge and internal use of the order of the letters of the alphabet as a device for remembering the order of a sequence of operations to perform. As Hutchins observes, “The performance guided by the memory of the words in such a procedure is still a mediated task performance, but the mediating structure is now internal rather than external” (p. 303). Moreover, internal artefacts of this sort are capable, much as external tools and devices, of transforming the cognitive processes required to carry out a task.

As Bessy (2001) has observed, in applying this framework to ship navigation Hutchins moves between three distinct levels of analysis, which correspond to three different temporalities. I would argue that the first two of these three levels correspond to the concerns of the situated knowledge and communities of practice perspectives respectively. The third level goes beyond the main concerns of these literatures in showing how practice is embedded in wider forms of culturally transmitted knowledge.

The basic argument is summarised in the closing pages of Hutchins' book. The first level of analysis corresponds to any moment in navigation practice. The accent here is on the way meanings and interpretations are constructed in the process of interaction and task completion. This aspect comes out in Hutchin's striking account (Ch. 8) of the adaptive response of the Navy vessel's navigation team to an emergency situation occasioned by the failure of the ship's propulsion system while undertaking entry to the San Diego harbour. The loss of propulsion led to various electrical failures resulting in a loss of power to the ship's gyrocompass. The emergency was compounded by the loss of back-up power and by the fact that the reserve gyrocompass had been taken out of service. The team was thus dependent on visual readings from a magnetic compass which are subject to compass specific error as well as to error associated with the location specific deviation between true north and magnetic north. The team had no prior experience with coordinating the fix cycle in such circumstances.

Hutchin's reconstruction of this emergency based on video tape recordings shows how through a trial and error process of mutual adjustment the navigation team managed to put in place a new division of labour and make the necessarily computation to determine the ship's position. As he observes, none of the members appeared to reflect on the process as whole. Rather, team members put constraints on each other's behaviour through presenting each other with partial computations to which they mutually adapted. The analysis brings out in an exemplary manner the idea that the sense an individual attaches to what others do and say in a distributed task is negotiated through on-going interpretative acts. This, of course, is one of the key themes in the situated knowledge literature.

The second level of analysis corresponds rather closely to the view of the communities of practice literature. The internal structure which individual team members bring to their work is the result of relatively long apprenticeships involving active participation in the distributed task of fixing the ships's position. The changes to an individual's internal structure can be thought of as the mental residua of this learning process. These residua allow the individual to coordinate his actions in relation to the external media and the actions of others. Moreover, as Hutchins observes (p. 282), a good deal of the structure that the novice will have to learn in order to perform the various steps is present in the organisational relations among the members of the navigation team. Thus in coming to understand the way the various members of the team depend on each other the novice is simultaneously learning about the nature of the computation being performed and the ways the various parts of it depend on one another.

The third level of analysis operates on a much longer time scale. It concerns the cultural transmission of knowledge that is “crystallised and saved in the physical and conceptual tools of the trade and in the social organisation of work” (p. 374). By physical tools, Hutchin’s is referring to such objects as the adelade, the hoey and the gyrocompass used in the realisation of the computation They constitute possible supports for the cultural transmission of knowledge. By social organisation, he is thinking of the way occupational roles and career structures, such as those described for the navigation team, are the products of an organisation with its own history and traditions, in this case the U.S. Navy. As I observed above, this social organisation conditions the process of acquiring skills and knowledge. By conceptual tools, Hutchin’s is thinking of culturally rooted cognitive artefacts, both internal and external, which influence the way the problem is represented and the nature of the algorithms used in its solution. This aspect of his analysis come out most clearly in the comparisons he makes between the representational assumptions of western and micronesian navigation practice (Ch. 2). The comparison brings out in a striking manner how even such basic premises embodied our representations, such as viewing the vessel in motion relative to fixed objects, rather than seeing the objects in motion relative to a vessel whose position is unchanging, are products of a particular culture and history.

In this third level of analysis Hutchin’s shows how both relatively fleeting processes of adaptation and interpretation and more lengthy processes of apprenticeship are embedded within wider forms of culturally transmitted knowledge. These latter forms of knowledge are stored in multiple ways: in tools, in external representations and in internal representations which provide internal structure. The properties and performance characteristics of the system can only be accounted for by examining the nature of the interactions between these various elements. For this reason, as Bessy (2000) has observed, rather than seeing knowledge as something existing strictly within the minds of individuals, the appropriate imagery is one of culturally situated distributed cognition.

Conclusion

In conclusion it will be useful to examine the particular way in which Hutchin's treats the question of symbols and internal representation. In the situated knowledge and learning literature there has been a reluctance to use the language of symbolic representation, as it would appear to imply a commitment to the mainstream physical symbol system hypothesis. Hutchin's (p. 369) explains, however, that it is important to distinguish the proposition that humans are processors of symbols from the proposition that the architecture of cognition is symbolic. In the mainstream conception, this architecture operates to a large extent at an unconscious level, explaining even such perceptual capabilities that are not open to introspection as face and language recognition. Observing that humans may imagine symbols and operations on symbols, however, as when they mentally carry out place multiplication, in no way provides evidence that the underlying architecture is symbolic.

This aspect of Hutchins' analysis poses a direct challenge to the way those wedded to the physical symbol system hypothesis see the relation between symbols and external reality. Expanding on the idea that internal as well as external symbolic representations serve to mediate performance, Hutchins argues that internal representations exist first in the form of external symbol tokens which are tightly coupled to social relations. It is only in a subsequent stage that these socially constructed representations are internalised and imagined. For example, in Chapter 6 he describes the way a novice and an expert quartermaster organise their activity around a written form. What for the novice is initially a process of carrying out orders becomes in a subsequent time an internalised representation of what to do to carry out a task. Then, with even more experience about the regularities of the world of external symbolic tokens, it becomes possible to imagine symbolic worlds or mental models formed of these tokens and to apply the knowledge gained from this experience to their manipulation (p. 292-93). In this manner Hutchins gives a social account of the origins of internally held symbols and symbol processing, and so contributes to a framework for situating knowledge and reasoning processes within their wider social and cultural contexts.

References

Barney, J. 1986, "Strategic Factor Markets: Expectations, Luck and Business Strategy", *Management Science*, 32:10, pp. 1231-41.

Bessy, C. 2001, 'Is the Growth of the Firm Limited by tacit Knowledge? The Evolutionary Approach to the Firm Revisited by the Distributed Cognition Perspective, Working Paper, Centre d'Etudes de l'Emploi, Noisy-le-Grand, March.

Blumer, H. 1969, *Symbolic Interactionism: Perpsective and Method*, University of California Press, Berkeley.

Brown, J.S., A. Collins, and P. Duguid, 1989, "Situated Cognition and the Culture of Learning," *Education Researcher*, 18, pp. 32-42.

Brown, J.S. and P. Duguid, 1991, Organisational Learning and Communities of Practice: Toward a unified view of working, learning and innovation' *Organization Science*, Vol. 2, No. 1, pp. 40- 57.

Carley, K. 1986, "Efficiency in a Garbage can: Implications for crisis management," in J. March and R. Weissinger-Baylon (eds.) *Ambiguity and Command*, pp. 165-194.

Carley, K. 1995, Computational and Mathematical Organisational Theory: Perspectives and Directions, *Computational and Mathematical Organisational Theory*, 1, pp. 39-56.

Clancey, W.J. (1997) *Situated Cognition: On Human Knowledge and Computer Representation*, Cambridge, Cambridge University Press.

Dosi, G. 1991, 'Information, Competences and the Firm', Working Paper, University of Rome 'La Sapienza'.

Dosi, G., R. Nelson and S. Winter (eds.) 2000, *The Nature and Dynamics of Organizational Capabilities*, Oxford University Press, Oxford.

Fischer, B. and A. Strauss, 1978, 'Interactionism', Ch. 12 in T. Bottomore and R. Nisbet (eds.), *A History of Sociological Analysis*, Basic Books, New York.

Fox, S., 2000, Communities of Practice, Foucault and Actor-Network Theory', *Journal of Management Studies*, 37:6, pp. 853-867.

Garfinkel, H. 1967, *Studies in Ethnomethodology*, Engelwood Cliffs, N.J.

Hutchins, E. 1995,*Cognition in the Wild*, MIT Press, Cambridge, MA.

Kogut, B. and Zander, U. 1992, "Knowledge of the Firm, Combinative Capabilities, and the replication of technology", *Organization Science*, 3, pp. 383-97.

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

Lave, J. and E. Wenger, 1991, *Situated Learning: Legitimate Peripheral Participation*, Cambridge, Cambridge University Press.

Lazaric, N. and P.-A. Mangolte. 1999. "Routines in Theory and in practice: A Criticism of the cognitive perspective", *Revista de Economia Contemporânea*, No. 5, pp. 7-35.

Lippman, Steven A. and Richard P. Rumelt. 1982. "Uncertain Imitability: An Analysis of Interfirm Differences in Efficiency under Competition". *Bell Journal of Economics* 13, pp. 418-438.

March, J and H. Simon, *Organizations*, [1958] 1993, Blackwell, Cambridge, MA.

Newell, A., J.C. Shaw and H. Simon ([1958] 1989) ‘Elements of a Theory of Human Problem Solving’, Models *of Thought, Vol. II*, Yale University Press, New Haven.

Newell, A. and H. Simon, 1972, *Human Problem Solving*, Prentice Hall.

Newell, A. and H. Simon, 1976, “Computer Science as Empirical Enquiry: Symbol and search,” *Communications of the AMC*, 19 (3).

Norman, D., 1993, *Things that Make Us Smart*,” Addison Wesley.

Padgett, J. 1980, “Managing Garbage Can Hierarchies,” *Administrative Science Quarterly*, 25, pp. 583-604.

Polayni, M., 1969, *Knowing and Being: Essays by Michael Polayni*, Routledge and Kegan Paul, London.

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

Simon, H. 1998, *The Sciences of the Artificial*, Third Edition, 1998, Cambridge, MA. MIT Press.

Snyder, W., 1998, “Communities of Practice: Combining Organizational Learning and Strategy Insights to Create a Bridge to the 21st Century”, Document, Social Capital Group, Cambridge, MA.

Suchman, L. 1987, *Plans and Situated Actions: The Problem of human-machine communication*, Cambridge, Cambridge University Press.

Teece, D., G. Pisano and A. Shuen 1997, “Dynamic Capabilities and Strategic Management”, *Strategic Management Journal*, 18:7, pp. 509-33.

Wenger, E., 1998a, *Communities of Practice: Learning, Meaning and Identity*, Cambridge, Cambridge University Press.

Wenger, E. 1998b, “Communities of Practice; Learning as a Social System,” *Systems Thinker*, June.