

Can Organisations Afford Knowledge?

R. ANDERSON & W. SHARROCK

Rank Xerox Cambridge EuroPARC, 59–61, Regent Street, Cambridge, CB2 1AB, UK

(Received 14 November 1991; in final form 8 December 1992)

Abstract. “Affordance Theory” has been widely discussed as a potential resource for the design of interfaces for CSCW and other systems. In this paper, we discuss the extension and adaption of this concept beyond the psychology of perception to the social distribution of a common stock of knowledge. We suggest that a working division of labour as that is known, oriented to, and rendered visible by the management of space and artifacts within a working environment can “afford” knowledge of organisational routines and practices. Learning to see the working division of labour is coming to understand the organisation. The grounds for extending the concept in this way are derived from consideration of an actual example taken from fieldwork. Some implications for the design of CSCW systems are reviewed.

Keywords. Affordance, Cognition, Distributed Cognition, Organisation, Knowledge

1. Introduction

In a recent paper, Aaron Cicourel summarised the basic tenets of distributed cognition as follows:

The idea of socially distributed cognition refers to the fact that participants in collaborative work relationships are likely to vary in the knowledge they possess.....and must therefore engage each other in dialogues that allow them to pool resources and negotiate their differences to accomplish their tasks..... A central assumption is that the knowledge sources must co-operate to solve a problem because no one source has enough information to do the job. (Cicourel 1990: 223).

Building upon the work of Hutchins and his colleagues (Hutchins 1990; Norman and Hutchins 1988), Cicourel goes on to propose

.....organizations can have an important influence on a group’s use of cognitive strategies. (Hutchins’ work) suggests that to create a satisfactory account of differential human performance, researchers must shift their focus from the cognitive properties of individuals to on-line studies of groups in natural settings. (ibid.)

To bring off such a shift in focus involves much more than a tolerant pluralism of one current ways of talking about individual cognition. That is, it involves more than simply treating collective problem solving, say, as the aggregative outcome of the efforts of a number of individuals. By the very fact that a group has come together, resources and strategies are available which simply cannot be invoked by individuals working on their own. It is these emergent features of distributed cognition which have, quite rightly, pre-occupied much of the research thus far carried out. It has been concerned with how groups collectively solve problems (Olson and Atkins 1990; Tjosvold 1990; Nardi and Miller 1990), how they access a common pool of knowledge (Wenger 1990), and how they do or do not transfer joint learning from one setting or situation to another (Lave 1988).

The relevance of these novel approaches to cognition for cscw research should be fairly obvious and indeed have been spelled out at length more than once (Brown and Newman 1985; Olson and Olson 1991). Many of the problem solving arenas which distributed cognition has concerned itself with (e.g. joint work tasks, meetings and the like) are precisely those which have looked promising from the point of view of providing computational tools and support. The difficulty is that unlike the classical domain of single-user systems, we have no good model of distributed cognition to transform into the equivalent of the 'model of the user' enshrined in single-user systems. At least, we have no good model which is generalisable enough and flexible enough to allow for the emergent properties of social decision making and group problem solving which are the crux of distributed cognition. What, it seems, is needed is a thoroughgoing radical approach to socialised cognition.

We will try to develop such an approach, building on Gibson's (1979) analysis of affordance within the psychology of perception. In doing so, we will try to avoid the usual impasse of a forced choice between the individual v. the collective determinants of perception and cognition (i.e. between what is "in the head" v. "what is in the phenomenon", the "inner" v. "outer" dichotomies) by insisting upon a socialised conception. This socialised conception states very simply that we can treat organisational forms and phenomena such as a working division of labour as affording knowledge and as having been designed with this possibility in mind. In this sense, then, our objective is the delineation of the *social construction* of affordance.¹

The immediate (and obvious) application of Gibsonian ideas is to the design of the technologies themselves (Costall 1991, Gaver 1991). How can particular systems and their interfaces be constructed so as to afford their functionality? Our suggestion is that we may be able to make some progress towards solving this problem if we think about alternative cultural artefacts (and technologies *are* cultural artefacts) where the same problem can be said to have been solved. The case we pick up is that of a paperwork system within an office. We will suggest that this system has been designed to afford the working division of labour by which tasks are performed. That is, anyone who is familiar with the culture of the

office can see at a glance where things stand and why.

There is, though, a second strand which we want to emphasise right at the start. If we design CSCW tools for environments such as offices like the one we describe without first paying attention to this aspect of the informal "tacit" working knowledge potential "users" have, we will have disregarded a potentially valuable design resource, but also we run the risk of repeating history by creating much the same clumsy and difficult to use systems as those which the original impetus to research cscw was itself a reaction to.

In *Working for Profit* (Anderson et al. 1989) we set out the sociological frame of reference to the study from which the materials we present here were drawn. This is not the place (and neither do we have space) to rehearse those arguments again. Suffice it to say that the general tenor of sociological analyses of business enterprises and other organisations, even when they have been ethnographically inspired as ours was, has been structural in form. That is, the concern has been with explicating the institutional contingencies which shape, at the structural level, organisational activities (Young 1989, Van Maanen 1986, Van Maanen and Barley 1985). Our focus is somewhat different. We are concerned with depicting the internal configuration of those contingencies: what the social structure looks like from within. Our interest is in the local, day by day, quotidian organisational facts of life and the ways in which they surface in and are oriented to as part of the routine courses of action which define any organisation. As we say, the *sociological* motivations for proceeding in this way are not exactly pertinent here. However, there is one argument which we would like to set out. Designers of technological systems, especially designers of CSCW systems, have come to realise that their technologies are social as well as technological, in the sense that they are deployed and used in social settings and defined by social constructs. In many ways, it is this realisation which makes CSCW designers so distinctive. If ignored, this feature could be a formidable barrier to successful technology transfer. On the other hand if, as many hope, it can be both grasped and turned to advantage, it could be a major resource for successful design. Our description is offered in precisely these terms. Because we focus on the routine, necessarily detailed aspects of the paperwork system at Leisure Time Catering, we hope to offer an account which is at a level of detail sufficient for design decisions to be framed and options ranked. In that sense, we see our work as a bridge between the generalised depictions offered within more conventional sociological accounts of organisations and detailed, formal models of "requirements analysis". A measure of our success will be the extent to which these formal models are adapted to take account of the institutional characteristics of working life in organisations.

The stepping stones we shall put in place are derived from the Gibsonian notions of *affordance* and *information pick-up*. We are well aware that they form central components of his theory of perception. Nonetheless, we believe that, somewhat amended, they can act as useful heuristics, ways of thinking, as we try

to move towards more “ecumenical” collective theories and descriptions of, broadly speaking, thinking. To demonstrate how this might be achieved, we will have to extend and transform Gibson’s concepts, perhaps more than the followers of Gibson would allow. For this reason alone, we do not see what we offer as a Gibsonian account of cognition. It is more an opportunistic appropriation of his ideas.

2. Affording things

J.J. Gibson’s notion of “affordance” is most usually seen as a counterweight to certain (computational) theories of perception. The general aim is to bolster an account of “direct perception” against those which place emphasis upon information processing, representation and interpretation. The object on the floor is a box simply by virtue of its size, shape, and so forth. We see it is a box, or perhaps more accurately we see it is something into which we can put things or on which we can stand things, simply because that is what it is. We do not want here to explicate Gibsonian theories of perception as *theories of perception*. Neither do we want to take sides in the debate over which general approach is preferable as the basis of a psychology of perception.² All we do want to show is that the point which Gibson is making is generalisable beyond perception to other cognitive processes, and beyond individuals to social organisations of various kinds. In that sense, we want to appropriate the move which Gibson makes with regard to perception and apply it to knowledge instead.

The point is really quite a simple one. We always perceive, see, hear *something*.³ One sees a chair; hears a voice; touches a surface. But, in that it is intentional, perception involves us, the perceiving agents, in a relationship with our environment as an environment of perceived objects. Hence Gibsonian theory is ecological. We like to think of this relationship as one of “structuration”.⁴ We do not just see, hear, feel things. We see, hear and feel them in a context, against a background of other things and actions. What is involved, for the social scientist, in structuration is the placing of sensations and perception in the context of patterns of social activities. We feel the surface of the cup when we pick it up to drink from it; we look for our keys when we are going out for dinner.⁵ So, our perception is embedded in courses of action which are themselves patterned, learned and shared. Objects, processes and actions are set against one another in that they emerge from or retreat to a background. The environment is structured so that feeling, hearing and seeing is possible. But, of course, this structure is not unique to us. It is available in patterned or institutionalised ways; ways, that is, which we share with others. In that sense, the possibility of perception can be said to be culturally provided. We encounter the world of objects, as our social world, the one in which we are at home, which we have learned and into which we socialise our children and other “strangers”.⁶

For us, what the concept of “affordance” supplies are the grounds (to put it no more strongly than that) upon which structuration takes place. It shifts the explanatory weight from the individual to the surrounding, socially structured environment. I, the perceiving individual, see this chair, because this object which appresents itself in my environment is structured as an object which affords the potential for sitting on. My environment, notice. Nothing is being said about the environment in general, *sui generis*, independent of the perceiving subject. The point of affordance theory is to locate perception in the relation between the specific individual and the world rather than solely as a process going on “inside the head” of the individual. In so doing, it breaks with what has sometimes been called “the spectator theory” of perception since it no longer conceives of the perceiver as a passive recipient of “sense data” or visual stimulation.

For this theory to be a thorough-going explanatory account, it would have to incorporate some account of the process we have called “structuration.” In line with conventional psychology, Gibson treats this as an issue to be resolved by reference to the physical make-up of objects. There is something about naturally occurring physical objects which predisposes them to be seen as just those objects. The light reflected by objects carries the information enabling us to structure the world in the ways we do. But it is we, as active percipient beings, who do the structuring. Hence, in Gibson’s famous phrase

...an affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of the subjective-objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical. An affordance points both ways, to the environment and to the observer (Gibson 1979: 129).

It is at this point that we wish to make our first important departure from Gibson’s own theory. Since the observer lives in a social world, we want to argue an affordance points out into that world too. The categories we lay upon the world are made available to us through the social institutions within which we live our daily lives. Indeed, often these institutions are defined in terms of the categories with which they are *associated*. When we do our science, we speak and think *scientifically*. When we move around the mundane world, we speak and think *commonsensically*.⁷ Such classification systems are normative, regulative, and hence pre-structure the world for us, but in ways which enable us to make it our own.⁸

The theory of affordance does several things. First it defines the world for an ecological theory of perception. In this theory, the world is stipulated as the-world-given-by-the-perceptual-array. Second, it insists on the dual character of the process of structuration. The meaning of objects is both “in the head” and “in the world”. Third, it denies that ordinary perception is based upon representation

or disambiguation. The qualifying adjective “ordinary” is important. As Coulter and Parsons 1990 argue, “the logical grammar of the verbs of perception do not require perceiving to be coupled with representation”. Or, in John Austin’s (Austin 1962) more trenchant phrase, what we see on the breadboard is bread neither a sign nor a representation of bread. This is why the theory of affordances is an argument for the directness of perception.⁹

In trying, as he does, to avoid having to choose between an “objective” and a “subjective” account of perception, Gibson sometimes ends up veering from making empirical claims about the processes of perception to unsatisfactory metaphysical claims about how the world must be in order to perceive it at all. Here is just one example.

There is much evidence to show that the infant does not begin by first discriminating the qualities of objects and then learning the combinations of qualities that specify them. Phenomenal objects are not built up of qualities. Rather it is the other way around. The affordance of an object is what the infant begins by noticing. The meaning is observed before the substance and the surface, the color and form, are seen as such. An affordance is an invariant combination of variables, and one might guess that it is easier to perceive such an invariant unit than it is to perceive all the variables separately (Gibson 1979: 134).

Yet, despite reservations such as these, there is something attractive about Gibson’s argument. What it offers is the opportunity to break out of the dichotomy of “inner” and “outer” processes by treating direct perception as embedded in cultural practices, the social institutions we referred to earlier. Becoming members of these institutions means being able to see right off that this is a box, this is a quark, this a limited-over cricket match without the need to reflect or represent the phenomenon to ourselves. What it does not mean, of course, is that we are always correct or that there are no occasions on which we explicitly construct representations of congeries of objects, persons and relations. But when we do, we do so in the light of sets of specific relevances, in very circumscribed contexts.¹⁰ What Gibsonians refer to as the affordance relation – that is the relationship between the perceiving agent and the surrounding environment – can be said on this view, to be *praxeologically structured*, that is organised by virtue of the task in hand.¹¹

The praxeological move, if we may call it that, offers a way around the impasse which straightforward affordance seems to create. The recognisable use-function of the object does not have to be contained in the object in order to be found by the perceiver. That is, use-function does not have to be defined as a property of the object in an environment independent of any use to which it might be put. What something is or what it is for can be treated as constituted and re-constituted in and through any projected course or courses of actions given within a setting. In that sense affordance, as we are now discussing it, provides a pointer to the socio-cul-

tural resources we draw upon on such occasions. In other words, it allows us to begin treating perception as an intersubjective, public, socially organised accomplishment rather than a subjective, private, internal process.

3. Affording knowledge

The central element in the ecological theory of preception is “information pick-up”. By this term, Gibson designates the process of constituting the environment as an environment of objects and relations. Information pick-up is not the passive processing of information given by the ambient array, but the active construction of the-world-for-the-individual-in-a-particular-context. It is the orientation to context or situation which makes it possible to use affordance theory as a way of moving towards a sociology of cognition. It gives us the task of fixing the situation or context within which perception takes place. For social science knowledge is not to be considered apart from the situations and courses of action within which it is deployed. Knowledge and action are conjoined. Thus what social science insists upon is, in the current jargon, *situated cognition* (Lave and Wenger 1992). The connection between knowledge and action is defined in constitutive terms. Patterns of knowledge and patterns of action define each other. Hence knowledge is seen as *social through and through*.

Following Schutz (1967), we will call the institutionally given constellation of considerations which define the context of any activity, that is the horizons of relevances and types of motivation in play, “a finite province of meaning”. The systems of categories we mentioned earlier, are just one component of such provinces. As we made clear, although there are multiple potential alternatives available, the choices between them are never entirely open. The repertoire of actualisable finite province of meaning is regulated within any culture.¹² For our accounts of the world to be intelligible to others they have to be couched in terms of some shared finite province of meaning.¹³ It is this line of thinking which allows us to say that the world is socially organised.¹⁴

The social institutions which contain such frames of reference and schemes of interpretability do not have an existence outside the courses of action which we all engage in. They are constructed and re-constructed in and through the courses of actions in which we engage. The problem set for a sociology of cognition is the description of the processes by which, from within whatever courses of action they are engaged in, social actors frame their activities so that mutual understanding and intelligibility is achieved. Putting it another way, how *do* we know what the relevant social practices are? How *do* we know what we (jointly) are doing?

The extension of the notions of affordance and information pick-up we have advocated is an obvious one. Gibson himself hinted at this in passing although he goes on to say that he doesn’t know what to do about it.

The theory of information pick-up makes a clear-cut separation between perception and fantasy, but it closes the gap between perception and knowledge. The extracting and abstracting of invariants are what happens in both perceiving and knowing. To perceive the environment and to conceive it are different in degree but not in kind. One is continuous with the other. . . (E)xtended or aided modes of apprehension are all cases of information from a stimulus flux. The learner has to hear the speech in order to pick up the message; to see the model, the picture, the writing. but the information is largely independent of the stimulus flux. (Gibson 1979: 258)

Put simply, all that is involved is the suggestion that organisations and other collectivities can be seen as constructions or social artifacts which are designed to provide their members – those who are competent in their ways – with the means to see at a glance, grasp unthinkingly, recognise immediately what is going on around them. It is this extended sense of the notion that we mean when that we talk of organisations affording knowledge. To learn one's way around an organisation, to become at home, socialised, is just to learn how to pick up organisational information from the social environment.

In the next section we examine two examples of ways in which information pick-up may be provided for within organisations. They exhibit the affordance of organisation knowledge in the sense that we have used that term. We then go on to reflect on the implications of this approach for understanding the character of organisations. The examples we have chosen are extremely familiar ones. There is nothing difficult or arcane in them. It is precisely their familiarity which makes them the kind of examples we need, since what we are interested in is teasing out the basis of their ordinariness. It is this which makes them easily learnt, disseminated and adapted and thus extremely powerful vehicles for conveying knowledge. The trivialness of the examples should not mislead us into thinking that the practices they evince are trivial, for they are not. The ability in any context, be it organisational or some other, to "see at a glance" what is going on around us is one of the hallmarks of cultural competence.

4. The affordance of organisational knowledge

The following examples are drawn from data collected as part of a previous study (Anderson et. al. 1989). This study concerned the day to day business decision making of an entrepreneurial firm (Leisure Time Catering). As part of the research, some time was spent looking at the financial accounting systems which the Company used. The examples we discuss are all drawn from this data.

To begin with, it might be as well to summarise some relevant features of the accounting practices at LTC before we move to considering particular details.

1. The company receives invoices in a constant flow, but pays them only fortnightly when a series of pre-printed cheques is run off.
2. There are two forms of invoice, Food and Non-food. Non-food invoices are treated in an entirely different way to food invoices. Part of the reason for separating out food invoices is managerial; other reasons are financial and organisational. Given the nature of the business, the vast proportion of invoices are for food (and drink) items, totalling in excess of 400 per month. One person can deal with these efficiently if they specialise in them. Given that food purchases are centralised through main suppliers, the supervisory role can be carried out by this person. She can also check that the appropriate negotiated prices are charged. A cross check of food purchases is also obtained through the outlet's fortnightly returns on which usage, food, liquor and labour costs are calculated. For non-food items, no such cross checks occur, there is no central purchasing policy (by and large) and knowledge of what the appropriate prices are is dispersed among the management. Circulation of the invoices is required for checking purposes alone. However, such circulation also allows supervision of spending at the units – a means by which Director level managers can keep their fingers on what is going on.
3. Any invoice, food or non-food, has a circulation life within the Company, that is, a length of time it takes to process and pay it. The shortest this is likely to be is a week; the longest 6 to 8 weeks. During this life time, it passes through a number of hands and across a number of desks. The invoice contains a record of its own passage displayed as the information entered in the date stamp. (Cf. below). Invoices circulate in bundles which start life as "what was in the post" and accumulate into "a day's worth" after coding by Purchase Ledger. From there on, they can accrete into large sized bundles or heaps, depending on how quickly they are processed in the later stages. The most likely place for this accretion to occur, in the case of non-food invoices at any rate, is in the Chairman of the Company's office, where they can sit for up to a week.
4. Each stage in the process is dependent on the completion of prior stages. It is also scheduled to fit around the "fixed points" of computer input, the weekly wage payments and the fortnightly stocksheet runs. This means that invoice processing is a continuous, fitted-in-where-it-can-be matter for the data processor (Dawn) in the computer room. There is a constant backlog of processed invoices to be typed in.
5. The rationale for having a separate invoice processing function is, of course, twofold. It gives a fair degree of financial control and also allows efficiency of effort by freeing management from the task of checking bills and so forth.
7. As a succession of tasks to be done in a series, the processing of an invoice has itself to be fitted into the daily and weekly routines of those that deal with them. No-one deals just with invoices.

Given the volume involved and the fact that the life cycle of an invoice within the system is or can be quite long and complicated, "tracking" invoices is a major pre-occupation of members of the financial staff. For a variety of reasons, such as the need to respond to queries from suppliers, complaints from units and so on, office staff have to be able to find out reasonably easily where a particular invoice may be in the process, what work has been done on it, what will happen to it next, how long that might take, and so on. Clearly what this involves is a concern for the state of the division of labour with regard to any particular or group of invoices. The division of labour is, of course, a way of structuring the accounting the other relevant work being carried out on the invoices. People who work in the office and who know its procedures well have a number of shared ways of seeing the invoice-relevant division of labour. These they will have learned from one another as they became more and more familiar with "the organisation" of work in the office. Each is used on different occasions for different purposes. We will pick out just two; the office layout as a representation of the work flow and the invoice date stamp as a stratified record.

5. The office as an ecology of activities

Figure 1 displays the floorplan of the main office at LTC. What is immediately obvious is that the major Company relevant distinctions, such as those between the different Divisions are not immediately visible. This is because they are not relevant to the work which the invoice processors carry out. Sue, Amanda, Rosemary and so on are only interested in the type of outlet for which an invoice had been presented in so far as that bears upon the tasks they have in hand.¹⁵

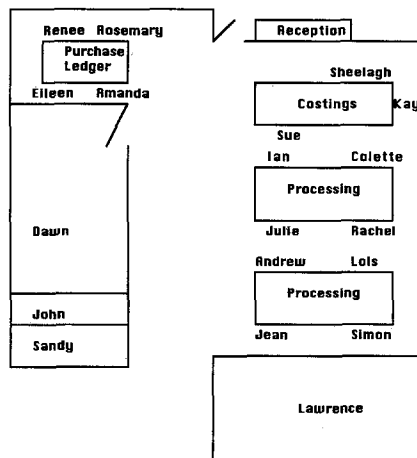


Fig. 1. Layout of LTC's main office.

Anyone looking around the office sees not Country Kitchens, Concessions, and Airports as locales of activity but a paper processing path or *trackway* (Gerson and Star 1986; Star and Greismer (1989) for any invoice no matter what it is about. The trackway has several well known phases or stages :

1. Invoices arrive in the post. Some are sent on from the units but this is not encouraged. When an invoice is sent direct to a unit, a letter is sent to the supplier concerned telling them to bill Telford. Sending the invoice to the unit makes two things possible. First, the unit manager could alter the figures. Second, much more likely, the manager will probably hold the invoice until the next lot of unit post is due – which means it may not go until the weekly time sheets or the stocksheets are sent. Given the length of time it takes anyway to process an invoice, this just adds to the delay. When it arrives, the invoice is stamped and passed on to Rachel.
2. Rachel separates the food and non-food invoices and pulls out any which she feels Sandy, the Financial Director, might want to see. The food invoices are sent to Sue for coding. Non-food invoices go to Rosemary who fills in the various codes using the boxes as set out in the invoice stamp (Fig.2). These codes are how the invoice is filed in the computer. SUPP No is supplier number as specified in the computer's list of authorised suppliers. Thus

L.T.C. Co No. 3		
SUPP. No.		
UNIT		
IND DATE	/ /	
IND TOTAL		
IND No.		
NET		
VAT		
COST	EXP	
AUTHORISED FOR PAYMENT		

Fig. 2. LTC invoice date stamp.

ASB 001 will mean the main account with Associated Biscuits. The Unit gives a designation for the division against which the invoice will be charged: 1000 = concessions; 1100 = airports; 2000 = Country Kitchen; 3000 = head office; 4000 = hotels. The logic of this system has since collapsed since "hotels" is used as a catchall and various items are coded to "Head Office" simply because they relate to Directorial expenses. Invoice Total is used to indicate whether an invoice (A) or a credit note (C) is appropriate. EXP is the expenditure code (e.g. food, liquor, cleaning, etc.) against which the item is to be charged. This is the difficult item to define and much of Rosemary's time is spent trying to categorise items which have not been processed before. If Rosemary cannot resolve the problem, she will ask Sandy to make a coding decision. The AUTHORISED FOR PAYMENT is where the Chairman (Lawrence) and the appropriate Director sign their initials. Without initials it won't be processed. Sandy occasionally pushes an invoice through before Lawrence sees it and holds it for payment (i.e. holds the cheque ready made out) until he has signed the invoice. Electricity and other utilities are the most common cases. Rosemary also checks the totals are correct.

3. Once it has been coded, the non-food invoice goes to the Divisional director for authorising. Here different strategies are adopted. However, all involve checking that the expenditure was authorised by the manager (in the case of repairs etc), that the work was done satisfactorily. For expense claims, these are processed by the Director himself.
4. The authorised invoices are returned to Sandy who passes them on to Lawrence. He signs them (or queries them) and passes them back. He generally goes through the invoices once or twice a week, when he has time or can't find anything better to occupy himself with.
5. The invoices are returned to Purchase Ledger where they are bundled with a cover sheet into 70's and sent to the computer room. The cover sheet gives the date and the codes of the invoices in the bundle. After being typed into the computer they are filed.

The only differences to this occur when the invoice is urgent. Should this be the case then a manual cheque (known as a "kalamazoo cheque") will be made out. This requires Renee to make a cash book entry, a journal amendment entry for the computer and to process the cheque and get Sandy to sign it. Generally, all electricity and gas bills are paid this way and it involves a lot of extra work.

Anyone who is familiar with the setting and the stages through which invoices pass, can look around the office and see the division of labour; what is going on where and who is carrying it out. So much is obvious. But that organisational knowledge is made available by the physical arrangement of the desks and people, knowledge of their allocation to roles within the division of labour and an understanding of the overall process which is being carried out. If people move desks, push their chairs together, re-arrange the furniture, these will all be under-

stood in terms of the standard frame of reference. Departures from it, such as when people 'cover' for one another, even though they might be signalled by such changes, will be accountable. People will feel obliged to explain what they are doing and why. In that sense, the normative distribution of desks and people affords the standardised division of labour. It is directly available to anyone who knows the firm and its financial procedures well enough to be able to see it. Since this is a known-in-common way of seeing what is going on, changes in the division of labour have to be marked in some clear way. This can be done by moving to another chair; picking up the files of the person one is covering for; making a public announcement; setting up the arrangements in advance; etc. etc. All of these provide public resources which allow the lay-out of the office to provide a means of seeing a (somewhat modified) division of labour. That they are done is testimony to the institutionalised character of this practice.

But more than this. Knowing the organisation means being able to see whereabouts any particular invoice might be in the temporal sequencing of activities in the room. If Rosemary knows that she checked the coding of an invoice five days ago, then she knows that it is likely to be "with Dawn" or "filed". The same is true for everyone. The layout is a time-table. The ability to see the time-table of activities from the spatial distribution enables the office to be used as a search protocol. Bottlenecks occur in well known places, on Lawrence's desk, in Sandy's office, at costings. And they occur for well known, standardised reasons. Lawrence and Sandy never have time. Lawrence is finicky about some details. They have particular interests in specific projects, etc. Not only, then, does looking round the office provide the normal scheduling for anyone who knows the office procedures. The office also tells them where to look when such scheduling goes wrong and how to account for what has happened.

Closely related to the above is the way in which the locality can be used to define horizons of relevance. Knowing where an invoice might be does not necessarily mean one ought to know (or care) why it is there, or that one can account for the logic of the specific sequence of steps which brought it there. Knowing that there will be reasons which could be discovered if one had the need to find them is enough.¹⁶ Furthermore, it is enough to know that someone else has to evaluate these reasons and that the work that entails is also institutionally taken care of. Thus, the well known bottleneck on Lawrence's desk is an organisational fact of life, but the rationale for passing all non-food invoices by Lawrence is not necessarily known, nor is what he does with them, how he makes the decisions he does, and so on. Not having this knowledge will make no difference to the effectiveness of the search which Sue or Amanda may undertake to find "a query". They know where it is likely to be and, should they need them, the reasons for its being there can be discovered. This asymmetry of relevances is how the reciprocity of perspectives within the office is managed. To put the point a great deal more crudely, Sue and Amanda have far more idea how to go about each others' tasks than they do Lawrence's.

An orientation to this shared understanding enables the mutual intelligibility of their actions. Much of what Lawrence does, and why, is treated as a mystery. They and the rest of the office account for it as simply 'the way he is'.¹⁷ But they know, if they were him, they would probably do things much the way he does.

In that one can see the hierarchy of decision making expressed in the distribution of knowledge displayed by the office layout, then it can be taken to represent the social distribution of rights, responsibilities, authority and freedom of action as well. Looking around the office, you can see LTC's social structure.

6. The invoice as a stratified record of work

As it moves on its journey around the paperwork socio-technical system, the invoice accumulates a record of the work that has been done to it. From the moment it first arrives and is date stamped, everything which is done to it leaves its mark, either in the spaces provided within the date stamp or as appended comments, memos, queries, questions stapled to it or stuck on it. The recognisable regularity of the record is taken to represent the regularity in the work tasks being performed. But, for those whose work invoices are, the record is treated as self-explaining. That is, it is held to be self-explanatory, perspicuous, unproblematic. The record of work appears as the ticks beside items whose prices have been checked; question marks against those which are unknown; the initials in the various boxes; and so on. The record is composed in standardised ways. A limited range of symbols, initials and abbreviations are used. Different colour inks denote different types of amendments. In addition, permanent or temporary explanatory notes are attached in different ways (Post-its or stapled). Since these notes may be single words ("holds") or extremely gnomic ("orig chk GER") recognising the handwriting is often crucial to be able to say what has happened and why.

To anyone coming to an invoice at any moment of its path through the accounting system, the record of what has and has not been done, is seen as a stratified sedimentation of the sequence it has passed through and the actions taken. Since this sequence and these activities are standardised, a glance at the stratified record is enough to be able to tell what has happened, where things are up to and what the possible problems might be even though this may be the first (and only) time the person in question may have come upon this invoice. Marking up the invoice is not just a way of recording actions and sharing knowledge, although it is obviously these. Reading the invoice stamp is not just a way of accessing such stored knowledge, though that too is true. Both the writing and the reading are ways of re-producing within the day to day courses of normal activities the social institution, the cultural practice, of using the invoice as a socially available, stratified record of work. And it is the existence of this practice which Amanda, Sue,

Rosemary and the others rely upon treating invoices and their records in the ways that they do.

For those engaged in the production of “accountant’s objects” such as invoices, the orderliness of the record is the orderlines of the tasks. The organisation of one expresses the organisation of the other. As the boxes are filled in, as the correct codes are written in, the amounts checked and the authorisations given, the normal unproblematic routine working of the system reproduces itself. Work upon the invoice is, therefore, a distinct sphere of operations for those whose tasks it is to ensure the paperwork is completed “properly”. Its horizons, its internal organisation, and its structures are given to them as local and contextual knowledge about how things are done at LTC and what from the invoice one can say about what has been done and what yet needs to be done. Learning to read an invoice as a record of its production work is learning the paperwork division of labour.

7. Some preliminary reflections.

In what ways can we say that the office layout or an invoice *affords* an organisational division of labour? Here we have to take Gibson’s notions of affordance and information pick-up and expand them in order to use them as makeshift bridges between perception and meaning. For Gibson, clearly, the problem is simply how do we see things the way we do. That is, what is it about either our perceptual apparatus, or the perceptible world (or both) which enables us to see things as we do? The theories of affordance and information pick up seek to explain this capacity without resorting to psychological categories (i.e. “inner processes”). In a sense, then, Gibson is deliberately breaking with the ‘subjectivist’ tradition in psychology—a tradition that looks to explain perception of the outer world in terms of (at least some) inner processes. But, of course, with Gibson, the break is not a wholehearted one, for there is still room to ask what it is that we, the perceiving subjects, bring to perception in that we see this such and such (our earlier box example) in the context of that environment (the study floor). A couple of remarks of Wittgenstein’s point in the direction in which our thinking is going at this point.

Here is the first:

I see that an animal in a picture is transfixed by an arrow. It has struck it in the throat and sticks out the back of the neck. Let the picture be a silhouette. Do you *see* the arrow – or do you merely *know* that these two bits are supposed to represent part of an arrow?... “The phenomenon is at first surprising, but a physiological explanation of it will certainly be found.” Our problem is not a causal but a conceptual one. (Wittgenstein 1958: 203)

The second is:

“Just now I looked at the shape rather than the colour”. Do not let such phrases confuse you. Above all, don’t wonder “What can be going on in the eyes and the brain?” (ibid. p. 211)

Both remarks are reminders of how we can allow ourselves to think that the problems of perception and meaning ought (in the sense of “ought first”) to be addressed by psychological or physiological explanations. We allow ourselves to think this way (what Wittgenstein called our tendency to be “fascinated” by certain kinds of pictures) because we have adopted the subjectivist point of view. And, having done so, we then ask about the inner processes which enable us to “represent” the world as an environment or ecology of perceptible or meaningful objects, and the correspondence between these *inner representations* and the outer *perceived objects*.

As we say, Gibson partially breaks with this. He asks what it is about objects in the perceptible world which enables beings like us to see them. We would like to extend this line of thinking to ask what it is about the social worlds in which we all live which enable us to see organisational objects in the ways which we do? Now this is most definitely *not* a psychological question in the sense usually ascribed to that notion. We are not asking about individual meanings and understandings – individual subjectivity. We are, instead, asking about *intersubjectivity*, the ways that we organise our social and organisational worlds so that we can find them understandable, meaningful, significant in the standard, patterned, institutionalised and hence shared ways that we do. This is where the parallel with Gibsonian theory begins to break down. We begin not by asking how – as an individual cognitive process – we see the things around us, but how as a collective endeavour we organise the world of meaningful objects. How is the world made recognisable for us as the world which we unproblematically inhabit? What we want to say is that knowing how to look is like knowing how to speak – knowing the language games, the culture, the practices embedded in any environment.

This is where the connection is made back to the broader issues. That this understanding, this social competence, is intersubjectively shared is a premise of the lifeworld of any organisational (or other) setting. In our view this premise could form the departure point for the kind of sociology of cognition we seek. But such competence does not consist in lists of tacit, taken for granted or expert knowledge. Actors do not presume that others have precisely the same bodies of knowledge as themselves, only that a reciprocity of perspectives is possible. This “could not” is not an empirical prohibition but a conceptual one. Social life could not begin if, first, we had to align bodies of knowledge, frames of reference, meaning structures. Rather what is presumed is procedural knowledge – how to find things out in this setting, for these purposes etc. What social competence consists in is

the knowledge of how to ask questions, pick up, locate, see at a glance, and all the “normally thoughtless” looking and seeing that we all do day in and day out.

An important consideration is, of course, whether an orientation to the affordance relation will act as a conservative force in the design of, say, computational artifacts which might be introduced into settings. We can see that this might be so, if the focus was to be on re-instituting existing practices. However, designing for the possible affordance relations which might be made available does not mean innovation is impossible or more difficult. One could very well imagine systems of electronic filing or form-filling which used the affordance of the existing (paper) forms as a way of easing the transfer from one technology to the other by picking up on very “formal” similarities. The logic of those new technologies will undoubtedly change the patterns of activity which ensue. Another way of viewing this whole problem is to see that what affordance in this context forces us to think about is the panoply of organisational relationships into which technological artifacts are introduced and knowledge of which is embedded in them once they are in use. We may, then, want to take the opportunity of the introduction of new technologies to re-think the form which these organisational relationships might take and seek to express these changes in and through the designed artifacts.

In sum, we do not have to ask whether organisations can afford knowledge. Simply by being organisations they do. The questions are (a) how is any local organisational environment organised so that the participants to that setting can find their way around in it and do what they do in it, in the routine, unproblematic ways that they do and (b) in what ways can this process of sharing, training or ‘enculturation’ be made easier, more flexible, and more broadly based? Our aim in this paper has been to try to find ways to pose these questions as clearly as possible. Once we are able to do that, finding answers to them may become somewhat less daunting a task than it seems just now.

8. Acknowledgements

This paper began life as a response to a conversation with John Seely Brown and Tom Moran and, while neither would necessarily agree with the final outcome, we would like to thank them both for provoking us enough to want to straighten out our thoughts on these matters. A version was presented at the Annual Winter Workshops on HCI held at the University of Michigan in January 1991. We would also like to thank the participants for their comments, and Gary Olson, John Bowers and Bill Gaver for many incisive observations.

References

- Anderson, Robert, John Hughes and Wesley Sharrock. 1984. Wittgenstein and Comparative Sociology, *Inquiry*, 27(2): 268–276

- Anderson, Robert, John Hughes and Wesley Sharrock, 1989. *Working for profit*. Farnborough: Avebury.
- Austin, John 1962. *Sense and Sensibilia*. Oxford, UK: Oxford University Press.
- Barnes, Barry and David Edge. 1982. *Science in Context*. London: MIT Press.
- Brown, John and Susan Newman. 1985. Issues in Cognitive and Social Ergonomics: From our House to Bauhaus. *Human-Computer Interaction*. 1 (4): 359–41
- Cicourel, Alan 1990. The integration of distributed knowledge in collaborative medical diagnosis. In *Intellectual Teamwork*, eds. Joel Gallagher, Robert Kraut and Carmen Egido, 221–242. Hillsdale, NJ: Lawrence Erlbaum.
- Costall, Alan 1991. Graceful Degradation: Cognitivism and Metaphors of the Computer. In *Against Cognitivism*, Arthur Still and Alan Costall, 151–170. Falmer: Harvester.
- Coulter, Jeff and Ed Parsons. 1990. The Praxeology of Perception. *Inquiry* 33: 251–72.
- Gaver, William 1991. Technology Affordances. In *Proceedings of CHI '91*, 79–84. New Orleans: ACM Press.
- Geertz, Clifford. 1983. *Local Knowledge*. New York: Basic Books.
- Gerson, Elihu and Susan Star. 1986. Analysing due process in the workplace. *ACM Transactions on Office Information Systems* 4(2): 257–70.
- Gibson, John. 1979. *The Ecological Approach to Visual Perception*. New York: Houghton Mifflin.
- Giddens, Antony. 1989. *The Constitution of Society*. Oxford: Polity Press.
- Hutchins, Edwin. 1990. The technology of team navigation. In *Intellectual Teamwork*. eds. John Galegher, Robert Kraut and Carmen Egido, 191–220. Hillsdale, NJ: Lawrence Erlbaum.
- Lave, Jean. 1988. *Cognition in Practice*. Cambridge, UK: Cambridge University Press.
- Lave, Jean and Etienne Wenger. 1992. *Situated Learning*. Cambridge, UK: Cambridge University Press.
- Lombardo, Thomas. 1987. The Reciprocity of Perceiver and Environment: *The Evolution of James J. Gibson's Ecological Psychology*. Hillsdale, NJ: Lawrence Erlbaum.
- Von Mises, Ludwig. 1959. *Human Action*. Chicago: Contemporary Books.
- Nardi, Bonnie and James Miller. An ethnographic study of distributed problem solving in spreadsheet development. In *Proceedings of CSCW'90*, 197–208. Los Angeles, CA.: ACM Press.
- Norman, Donald and Edwin Hutchins. 1988. *Computation v. Direct Manipulation*. Final report to the Office of Naval Research. Contract No N00014-85-C-0133. University of California, San Diego, CA.
- Olson, Gary and Donal Atkins. 1990. Supporting collaboration in advanced multi-media electronic mail: the NSF EXPRES project. In *Intellectual Teamwork*, eds.; John Galegher, Robert Kraut & Carmen Edigo, 429–452. Hillsdale, NJ: Lawrence Erlbaum.
- Olson, Gary and Judith Olson. 1991. User Centred Design of Collaboration Technology. *Journal of Organizational Computing* 1: 61–83
- Sharrock, Wesley and Robert Anderson. 1985. Understanding Peter Winch. *Inquiry* 28(1): 119–121
- Schutz, Alfred. 1959. The Stranger. In *Collected Works, Volume III*. Dordrecht: Martinus Nijhoff.
- Schutz, Alfred. 1967. *The Phenomenology of the Social World*. Evanston: North West University Press.
- Schutz, Alfred and Thomas Luckman 1974. *The Structures of the Lifeworld*. London: Heinemann.
- Star, Susan and Judith Geisner. 1989. Institutional ecology, translations and boundary objects. *Social Studies in Science* 19: 487–420.
- Tjosvold, Djarard. 1990. Making technological innovation work. Collaboration to solve problems. *Human Relations*. 43 (11): 1117–1131.
- Van Maanen, John and Stephen Barley. 1985. Cultural Organisation: Fragments of a theory. In *Organisational Culture*. Peter Frost, Larry Moore, Meryl Louis, Craig Lundberg, Joanne Martin. 31–54. Beverley Hills: Sage.
- Van Maanen, John. 1986. Power in the bottle: informal interaction and formal authority. In *Executive Power*. Srinvas. ed. Srivastva, 204–238. San Francisco: Josey Bass.
- Wenger, Eienne. 1990. *Toward a Theory of Cultural Transparency*. Ph.D. diss., University of California, Irvine, CA.
- Wittgenstein, Ludwig. 1958. *Philosophical Investigations*. Oxford: Basil Blackwell.
- Whyte, William. 1991. *Participatory Action Research*. London: Sage.
- Young, Ed. 1991. On the Naming of the Rose. In *Reframing Organisational Culture*. Peter Frost, Larry Moore, Meryl Louis, Crsig Lundberg and Joanne Martin. 90–103. Beverley Hills: Sage.

Notes

1. It should be possible from this formulation to see the inklings of where we will part company with Gibson. In essence, since Gibsons concerns are with "the way things are" he is necessarily (albeit sotto voce) positing an ontology of sorts together with an associated epistemology – how we know how things are. We are not concerned with this order of question. Rather we want to focus on the consequences of the fact that the ways we know how things are and, hence, how things are themselves, are socially institutionalised. It is this fact which underpins, we would suggest, the possibility of CSCW systems since it is nothing less than the social distribution of knowledge.
2. Gibson himself has done this at length e.g. Gibson (1979), as have many others, e.g. Lombardo (1987).
3. This is what is meant by calling perception "Brentano-intentional", that is intentional in the phenomenological not the motivational sense.
4. We have borrowed this unfortunate neologism from Giddens (1989) but without any of the theoretical underpinnings which he gives it.
5. Which is not to say everything is quite so purposeful as we may have it seem. Yet even 'daydreaming' or 'enjoying the experience' of a surface, for instance, is a learnable and shareable pattern of action.
6. On the analytic possibilities of 'the stranger' as a social type see Alfred Schutz (1959).
7. These are not the only two "paradigms" available to us, nor are they necessarily mutually exclusive. Cf. Schutz and Luckman (1974).
8. The proper epistemology for an effective sociology of cognition is well beyond the scope of this paper. Some very preliminary thoughts on what it might contain are to be found in Sharrock and Anderson (1984).
9. Which is why, earlier, we said '... we see the object is a box' not *as* a box.
10. Representation thinking is one and only one of the ways in which we grasp the world around us. As we hinted earlier, the concept of representation is associated with disambiguation. If, therefore, we presume, as some theories of perception do, that what has to be explained is how we fix the objects around us as those objects and not some (possible) others – that is, how we disambiguate them – then it is not surprising that we reach for a representational account. On this account, fixing the object is a process of comparison between the 'external' and 'internal' representations.
11. The origin of this use of "praxeology" is to be found in Von Mises' (1959) attempt to build a general theory for the social sciences.
12. Schuts (1967) mentions fiction, theatre, day-dreaming, science, routine daily life as just some which are available to us.
13. Within a different vocabulary, these are sometimes referred to as "genres" c.f. C. Geertz (1983).
14. Note that we are not adopting a disciplinary position here. To take a different example, human reproduction is a biological process which is socially organised in as much as there are norms and values associated with how reproduction should take place. But no discipline, be it biology, sociology or psychology, has the first or the last word on it.
15. See Anderson, Hughes and Sharrock (1989) for detail of the work of the office.
16. Schuts (1967) makes the distinctions between what is "at hand", "within reach", "within restorable reach". In that they apply to knowledge instantiated in artifacts such as documents and files, these distinctions are critical for understanding the distribution of knowledge in social settings like offices.
17. Indeed, the efficient functioning of the office could well be threatened if it were not.