Interdisciplinarity: an Emergent or Engineered Process?

Yvonne Rogers*, Mike Scaife* and Antonio Rizzo**

* School of Cognitive and Computing Sciences, University of Sussex, Brighton, BN1 9QH, UK

** Multimedia Communication Lab, University of Siena, 53100 Siena, Italy

Published as Cognitive Science Research paper 556
2003
School of Cognitive and Computing Sciences, University of Sussex,
Brighton, BN1 9QH, UK

"Cognitive science should be more than just people from different fields having lunch together to chat about the mind." (Thagard, 1996, p. 7)

Abstract

Interdisciplinarity – the integration of concepts and epistemologies from different disciplines – is often considered highly desirable as a way of gaining insight and furthering our understanding of a research problem. This is especially the case when an impasse is reached due the constraints of one's own discipline preventing any further progress. At the same time interdisciplinarity is very difficult to achieve: the positions adopted by the disparate disciplines are often incommensurable. We examine alternative ways of advancing understanding, from both interdisciplinary and multidisciplinary approaches. We point out how many successful breakthroughs, in both theoretical and applied research, come about through approaches that modify and reappropriate existing frameworks and concepts within disciplines, rather than those that try to create new ones by mixing and matching concepts, selected from different disciplines. In our critique we examine what it takes to develop new inter-disciplines, theoretical frameworks and methods.

Keywords: interdisciplinarity, multidisciplinarity, cognitive science, theoretical frameworks, external cognition

Introduction

There is a widespread view that interdisciplinary research is a good thing. By 'interdisciplinarity' is usually meant something like: the emergence of insight and understanding of a problem domain through the integration or derivation of different concepts, methods and epistemologies from different disciplines in a novel way. However, it is also widely believed that 'true' interdisciplinarity is very difficult to achieve and, more often than not, remains an elusive goal. In practice, many self-styled interdisciplinary enterprises actually work at the level of being multidisciplinary (or pluridisciplinary): where a group of researchers from different disciplines cooperate by working together on the same problem towards a common goal, but continue to do so using theories, tools, and methods from their own discipline, and occasionally using the output from each other's work. They remain, however, essentially within the boundaries of their own disciplines both in terms of their working practices and with respect to the outcomes of the work.

In this respect it is illuminating to look at the case of cognitive science, where interdisciplinarity was explicitly intended to be the defining feature of progress. Despite the aims of its founders (see later) to integrate disciplines in terms of differing levels of description and analyses, cognitive science has been predominantly a multidisciplinary activity where psychologists continue to work as psychologists, philosophers as philosophers and AI researchers as AI researchers, each continuing to use their own tools, methods and terminology.

The aim of this paper is to examine some of the main assumptions that lie behind the ideal of interdisciplinarity and to ask: when is it is really necessary? How might it be realized? What alternatives are there to help us progress research when we find ourselves frustrated and thwarted by the limits of our own discipline's conceptual and methodological armory? We shall pose these questions in the context of firstly, previous attempts to consider the issue and secondly, in terms of the challenge(s) for contemporary cognitive science. Finally we shall conclude by a brief examination of some of our own work, reflecting on how we dealt with bridging across disciplinary limitations.

Received views on the multi/inter- distinction

The terms multidisciplinarity and interdisciplinarity are often used interchangeably to refer to researchers from different disciplines or backgrounds coming together to collaborate on a common goal, be it basic or applied research (e.g. Brown, 1990). At a basic level of research, both terms have been used to describe the way a group of academics from different disciplines (e.g. psychologists, anthropologists, economists and computer scientists) may collaborate in order to develop a more extensive understanding of a situation or phenomenon, whilst at an applied level they have been used to describe the process of bringing a team of professionals together, from different backgrounds – such as designers, educational technologists, computer scientists and Human Factors specialists – in order to develop a product, such as building a software application (e.g. Kim, 1991). Using the two terms interchangeably is not problematic if they are being used simply to refer to some kind of cooperation or collaboration between different people. However the terms can have quite distinct meanings when used to denote different processes of collaborative activity. For example, bringing together a group of experts from different disciplines or professions to contribute to a single project, which would not be able to be accomplished by any one profession alone, is not necessarily the same process as when a group of researchers from distinct disciplines try to generate novel concepts and integrate different levels of explanation. The former may be considered a form of multidisciplinarity, where each person contributes their expertise to the project; the latter is a form of interdisciplinarity where novel research questions are addressed (cf. The Royal Society's position on the distinction between them, 1996). In this sense the main difference between multidisciplinarity and interdisciplinarity lies in the mechanism of the research process and, relatedly, its outcomes. Interdisciplinary approaches are assumed to derive *novel concepts, methods* and *theoretical frameworks* through the *melding* of concepts, methods and theoretical frameworks coming from different disciplines. An example is ecological economics where scientific aspects of ecological events have been integrated with their social consequences in order to make objective assessments of ecological aspects. By contrast, multidisciplinary approaches are assumed to evolve *new understanding* through *adapting and modifying* existing concepts, methods and theoretical frameworks within a discipline and occasionally borrowing ideas from others. Here we would classify the majority of research carried out in Cognitive Science (see also Schunn et al, 1998).

This gloss has the benefits of polarizing the distinction between the two terms and we shall use it below to consider what is really important when trying to specify good working practices. However we are well aware that such descriptions as we have offered above do, of course, beg many questions. For example the very idea of a 'discipline' as something coherent in terms of its methodology and theory, and of the assumption of a single agreed internal language, is itself suspect. Within Psychology, for example, the chasm between the various 'branches', such as social and cognitive, makes the assumption of disciplinary inclusiveness an interesting one. We need also to be aware that the practices within a discipline may owe more to internal political agendas than we would like. In short, while we have referred to interdisciplinarity as an ideal, even the glossing of multidisplinarity as being based on positions firmly rooted 'within' disciplines is itself an idealization.

When Do You Need Interdisciplinarity?

Whilst it is relatively easy to be multidisciplinary, since everyone is good at championing their own areas of expertise, it is much harder to achieve interdisciplinarity. There are many epistemological obstacles and cultural differences that prevent cross-fertilization of ideas. These include incommensurability of concepts, different units of analysis, differences in world views, expectations, criteria and value judgements. For example, a traditional divide is between the social and cognitive sciences: social scientists (e.g. sociologists) do not accept cognitive concepts as having explanatory power whilst cognitive scientists have typically ignored the environment and social structures. The main reason behind this obduracy is that each views the other as self evident and hence simply takes it for granted (Cicourel, 1995). For example, in one of our European project[1] meetings a heated debate took place between the sociologists and the cognitive psychologists, whereby the latter was trying to convince the former of the need to understand how cognition works with respect to work practices. One of the sociologists simply retorted "why do you need to look inside the brain? It's all out there." There are still other problems, however, for those who are able to see beyond their disciplinary boundaries. Bannon (1993) warns of the danger of attempting to 'wed' different conceptual frameworks in order to develop a new unifed one – it inevitably results in the dilution of the contributing specific theoretical frameworks. The added value of such efforts can all too often fail to materialize and instead a mish-mash of ideas, methods and theory may be the ensuing result.

A critical and ultimately more valuable question might be, therefore, when do you really need interdisciplinarity? In many situations, it is possible to continue progressing research by adopting a multidisciplinary stance. For example, research into the use of language, which has made substantial advances, is an area where cognitive psychologists, computational linguists and sociolinguists, have come together in the same forum borrowing ideas from each other where they see fit, but primarily remaining as separate and distinct approaches. In other situations, however, an impasse may be reached where researchers find that working by themselves is too limiting and unable to address the problem at hand. Accordingly, they may reach a point where the constraints of their own discipline prevent them from making any further progress and as a consequence are forced to work at the fringes of their field and in so doing forge new ones (The Royal Society, 1996). It is at these junctures that striving for interdisciplinarity may be most promising. But how can we characterize what these junctures are and, perhaps,

[1] We are partners of an European research training network, called COTCOS, where cognitive ergonomists, psychologists, linguists, ethnomethodologists and computer scientists have been brought together to develop theoretical and methodological frameworks for analyzing complex work settings for cooperative technologies.

recognize the need for a change of tactic from the multidisplinary approach?

One way forward here is to identify what the impetus has been in cases where we can see an area that clearly demands input from more than one conventionally-defined discipline and where no one alone has a comprehensive set of theoretical frameworks or methodological tools to deal with it. For convenience we divide these into cases where an existing problem has simply seemed too large for a single discipline to cope with by itself and those where something external to the disciplines has forced itself on their attention. Here we shall consider examples of both of these: firstly a program to develop a more comprehensive account of cognitive science and secondly, the evolution of two related applied fields – Human Computer Interaction (HCI) and Computer Supported Cooperative Work (CSCW).

Interdisciplinarity: as an ideal

Cognitive Science is a classic example of the emergence of a new field that set itself up to be truly interdisciplinary. A main motivation behind its inception was to enable a number of different disciplines to come together in order to develop a better understanding of how the mind works. An overarching aim was to provide more extensive accounts than was possible from a single discipline, which would comprise interacting levels of description that covered social, behavioral, cognitive and biological aspects. For example, Norman (1980) argued that the study of cognition needed to be much more far-reaching, considering it in terms of interacting aspects of a phenomenon, including social, cultural and internal cognitive factors:

"I wish Cognitive Science to be recognized as a complex interaction among different issues of concern, an interaction that will not be properly understood until all the parts are understood, with no part independent of the others, the whole requiring the parts, and the parts the whole". (p 336, Norman, 1980).

In the beginning the disciplines that were brought in to develop the new field of cognitive science were cognitive psychology, artificial intelligence (computer science), linguistics, philosophy and neuroscience (Green, et al, 1996; Johnson-Laird, 1988; Thagard, 1996; Von Eckardt, 1993). Since then a whole host of others have been identified as important contributors, including anthropology, sociology, engineering, HCI and education (Schunn et al, 1998). With so many potential collaborators, the stage seemed set for a range of combined efforts to emerge. Indeed, a number of such collaborations have been reported throughout the potted history of cognitive science. For example, Schunn et al (1998) cite the early collaborative efforts of Simon, Newell and Shaw's when building their logic theorist program. It is claimed that their work involved combining ideas from economics, psychology, mathematics and computer science. Their output was a computer program that arguably had more explanatory power than what would have evolved from within a single discipline. Kosslyn's work on mental imagery is also viewed as a paradigmatic example of interdisciplinary research – whereby his early empirical research from the 80s, on how mental images work, spurred a number of other researchers from different disciplines to extend his research in relation to their own models, perspectives and empirical findings. For example, Farah (1984) extended Kosslyn's work into the area of imagery deficits in neurologically impaired people, developing a further explanation of the role of imagery in cognition by identifying its physiological localization (Von Eckardt, 1993). More recently, Green et al (1996) have pointed out how neuropsychological research (especially on brain damaged patients) has provided insight for models of cognitive functioning, and in so doing they claim enabling a better integration of biological and cognitive accounts. Schunn et al (1998) in a survey of the main publications arising out of cognitive science in the last 20 years note how there have been a significant increase in papers citing research carried out where the methods of computer science and psychology have been combined - the most notable being the presentation of computer simulations of previously published empirical data sets.

Whilst these examples have been set up as paradigmatic of interdisciplinary research in cognitive science, it is actually quite difficult to determine what form they have taken. In many instances the main form of collaboration is in terms of *borrowing* or *building* on another's findings or ideas to extend or support their own theoretical explanation of some aspect of cognition. Another form is the use of another discipline's methodology, the most common being psychologists using computer programs to simulate some aspect of cognitive behavior (see Schunn, 1998). By our criteria that is really a form of mutlidisciplinarity. Moreover, it is a far cry from the desideratum proposed by Norman (1980) for interdisciplinary research – that its focus be concerned with developing explanatory accounts that cover complex interactions between different levels of description of the phenomena.

So why has there been so little demonstrable interdisciplinary research? Ten years on from his original thesis of what was needed in cognitive science Norman (1990) acknowledges that part of the problem for there being no real interdisciplinary progress is that the key issues he thought were important to this kind of research (e.g. learning, memory, thought, language, emotion, consciousness) could actually continue to be studied within a single existing discipline, not requiring any new interdiscipline to frame the research questions and methods needed. Indeed this appears to have remained predominantly the status quo under the umbrella of cognitive science: psychologists have continued to study psychology, computer scientists computer science and linguists linguistics. Part of the problem for this resistance is that one can carry out research within a single discipline with much more certainty. Methodologies have been tried and tested and can thus be relied upon for carrying out the research in a rigorous and systematic fashion. Furthermore, there are still many unanswered questions within the separate disciplines to keep researchers busy for a long time. To move beyond the boundaries of one's own discipline into an unknown territory and to try integrating unfamiliar terms and concepts is much more of a high risk enterprise, where the process and outputs are always uncertain. Those who have tried even multidisciplinary collaboration will often report how difficult it is to reach any level of shared understanding between the different parties with regard to the referents and terms each is using. Such frustrations can leave researchers wondering whether the costs involved in such ventures outweigh the benefits of doing so. For example, in their survey of multidisciplinary research in cognitive science, Schunn et al (1998) found that multi disciplinary collaborations were not rated as being any more successful than mono-disciplinary collaborations. One of the biggest complaints was that multi collaborations generated too many different ideas. Similarly, Scaife et al (1994) note that one of the key problems arising from their collaborative research project, with partners from cognitive psychology, design and computer science, respectively, was the difficulty of communicating and knowing what to do with the different ideas generated between them.

As mentioned previously we are trying to identify occasions when there has been an effective impetus for interdisciplinary research to break out and make any headway. Cognitive science was 'founded' because of the perception that we had reached a point where no further progress could be made in understanding the mind because of the limitations and restrictions of the parent disciplines. But pursuing interdisciplinarity for its own sake in the hope that a better understanding would result has proven to be a route that very few researchers have been prepared to take. Instead most researchers in cognitive science have preferred to take an egocentric stance, remaining within the confines of their own discipline, importing ideas and methods from other disciplines where they see them as being useful for supporting their own research.

If true interdisciplinarity is ever to take off then what is needed is a paradigm shift whereby a whole set of new issues and research questions are framed that force new ways of conceptualizing and working. Norman (1990) recognizes this dilemma and with hindsight suggests that there are now such a set of interdisciplinary concerns emerging – that cannot be addressed within the confines of a single discipline as could the previous issues he had identified 10 years earlier. These included the need firstly, to develop an applied cognitive science and secondly, a way of overcoming the deficiencies of a disembodied theory of cognition.

The need for an applied cognitive science stems from a recognition that there are "massive gaps in our scientific knowledge... because there has not been sufficient study of real, naturally occurring behavior" (Norman, 1990, p.4). We shall come back to this later but first we shall briefly consider the external (to cognitive science itself) push for disciplinary collaboration generated by technological advances in computing and telecommunication technologies. These have provided us with much scope for new forms of collaboration, communication and computational support including the ability to manipulate and interact with information in a multitude of ways, together with interacting with each other in remote and virtual spaces. In turn there has been a growing expectation within the system design community that cognitive science should and could have practical application for understanding these developments. Existing tools, theories and methods from within the contributing disciplines, especially cognitive psychology, however, have proven to be largely disappointing, being inappropriate and largely unusable (e.g. see Barnard, 1991; Rogers, 2000). Here, therefore, was an opportunity for a breakaway group of researchers, frustrated by the limits of their existing disciplinary knowledge, to come together and create a new field that could evolve new knowledge and methods that could be applied to practical problems.

Interdisciplinarity: forming applied fields

The perceived need for a new form of interdisciplinarity was very much the driving force behind the emergence of two new applied fields – Human-Computer Interaction (HCI) and Computer Supported Cooperative Work (CSCW), in the late 70s and 80s, respectively. In both cases, evolving technological advances created new theoretical and applied challenges, which existing tools, methods and frameworks from the separate disciplines of psychology, computer science, engineering, human factors and design, were considered inadequate to deal with, by themselves. In HCI the goal was originally to bring together psychology and computer science to design more effective human-computer interfaces for single user applications. In CSCW, the goal shifted towards bridging the gap between the social sciences and computer science in order to develop more usable and useful collaborative computer systems for multi-user settings.

What we can see, however, from these kinds of more applied endeavors that have tried to attain interdisciplinarity, is that the process is very much an uphill struggle to break away from a multidisciplinary mindset. The jury is still out as to whether either HCI or CSCW have in fact been able to achieve any significant level of interdisciplinarity. In a critique of the interdisciplinary accomplishments of the two fields, Bannon (1992) argues that whilst there have been several laudable attempts to develop new frameworks that allow for a family of theories and different concepts to be incorporated (e.g. Kuutti and Bannon, 1993), there has yet to be any convincing research projects reported, where different disciplines have genuinely wedded together, and made mappings across concepts, that have resulted in the development of a common unified theory. However, rather than see this as a failing of such enterprises, he argues that the goal of true interdisciplinarity in these contexts is fundamentally flawed since the world views, backgrounds, research traditions, perspectives, etc. of each of the contributing disciplines are often so different that they are simply not commensurable with each other. Attempts to build such hybrid frameworks are likely to come up against this dilemma. Recent examples like Mantovani's (1996) model of social context – where he takes a wide range of concepts and research findings from the social and the cognitive sciences, combining top-down with bottom-up approaches for the purpose of analyzing social norms and mental models together – are witness to this. Different terms, ontologies and methods are mixed together, making it difficult to make sense or apply together the various strands and levels presented in the framework.

Interdisciplinarity: changing the unit of analysis

One of Norman's other proposals for progressing cognitive science was to study cognition as it occurs in the real world rather than the traditional

model of isolating and controlling it in a laboratory setting. The reason for this is based on a growing acknowledgement that the assumptions behind lab-based cognition do not necessarily hold true in the real world:
"In the tradition of disembodied intellect, the person simply cogitates. The assumption is that the person starts with full and complete knowledge of the world-state relevant to the issue at hand, selects a course of action, then plans and executes it. I argue that this is neither what people do nor is it possible." (Norman, 1990, p.6)

Thus, although a psychologist can try to study the behavior of subjects in an experimental lab – observing them interacting with environments that embody knowledge they can control – they cannot understand the behavior of, say, operators in a control room since they cannot extrapolate from the former setting to the latter. This is because they have no real understanding of the knowledge embodied in the external representations that the operators create and use in their work. The continuous interplay of internal and external representations is completely out of the psychologist's range of investigation unless they begin to study, together with engineers, physicists and others, the way in which artifacts are actually used in the control room work.

Several researchers within cognitive science have taken up the challenge of studying cognition as practiced in different cultural settings, providing alternative explanations that reconceptualise cognition as situated within its cultural, social and environmental context (e.g. see special issue of Cognitive Science, 1993, on situated cognition). Such attempts have tended to adapt and assimilate concepts from other fields to contextualise their existing theories about cognition. As such the process of evolving a new understanding arises through local adaptation. A more extensive form of adaptation is to seek ways of developing a new understanding by reconceptualising a domain area using a new unit of analysis. An example of this more global strategy is Hutchins (1995) distributed cognition approach, where he broadened the mainstream cognitive science unit of analysis – which focuses exclusively on the properties and processes inside the mind of a single person – to one which extends to a family of cognitive systems. As well as continuing to allow for a unit of analysis that accounts for the workings of the individual mind, Hutchins proposed that we really need to begin studying other more extensive kinds of cognitive systems, such as an individual interacting with a set of tools and even more comprehensively, sociocultural systems comprising a group of individuals interacting with each other and a set of artifacts over a historical period of time (e.g. the collaborative activity of flying a plane be described in terms of the interactions between the pilots and the air traffic controllers, and the pilots and their interactional use of the instruments in the cockpit).

The rationale for developing the distributed cognition approach was largely motivated by a deep dissatisfaction with previous efforts in the disciplines of cognitive science and anthropology to explain cognition and culture. Hutchin's notes how both disciplines have marginalized the role of culture in cognition and, conversely, the role of cognition in culture to such an extent that "history and context and culture will always be seen as add-ons to the system, rather than as integral parts of the cognitive process, because they are by *(their)* definition, outside the boundaries of the cognitive system" (Hutchins, 1995, p368). To put matters right he argues that researchers need to study cognitive systems in 'the wild', by carrying out a kind of 'cognitive ethnography'. This involves analyzing the processes and properties of a particular cognitive system in terms of the propagation of representational state across the different media in that system – which maybe inside or outside of the individuals. In other words, Hutchins advocates continuing to use the conceptual currency of classical cognitive science but in a manner that is modified to enable a much more flexible unit of analysis. He acknowledges, though, that broadening the scope of one's enterprise makes the work of carrying out the research much more difficult and the ensuing outcomes much more uncertain. However, he argues that one of the benefits of adopting this more precarious stance, is that it can result in a more accurate picture of the functional specification of cognition and culture – something that has been largely missing in existing accounts of cognition and culture. In particular, it is suggested that one outcome of analyzing different kinds of cognitive systems is that it can reveal cognitive properties that cannot be reduced to those of individual persons.

Here is an attempt, then, to obtain new understanding of a phenomenon by reconceptualising the domain of interest through using a modified unit of analysis. Such an approach is not new in itself. Indeed, a similar reconceptualisation of the unit of analysis was proposed by Vygotsky in his cultural-historical approach for analyzing psychological processes (Vygotsky, 1978). He argued that to understand cognition we need to go beyond analyzing the isolated mind in the lab to studying our everyday practices in relation to the nature and evolution of cultural artifacts in use. Further, he claimed that the more pervasive such artifacts are in our everyday life the more they mediate cognition.

What is significant about these kinds of theoretical developments is that they have made advances in our understanding through extending and adapting concepts rather than through creating novel ones from an interdisciplinary approach. In setting itself more modest objectives, the distributed cognition approach has been able to develop a more integrated view of human cognition that, arguably, overcomes the shortcomings of previous attempts (i.e. cognitive science and cognitive anthropology) to account for the relation between cognition and culture.

Interdisciplinarity as practiced: the challenge of working together

Whether we are discussing inter- or multidisciplinary collaboration there is a problem that is common to both: the evolution of a common ground that will allow the coordination of concepts and efforts. There is surprisingly little written about this process, perhaps because the lessons of such collaboration are usually of the 'we could have done better' variety. However where there are accounts they tend to come from applied research, where the situation can arise where difficult problems need to be solved and individual disciplines are not adequately equipped to do so by themselves. Here multidisciplinarity as basic working practice

– where novel outputs are yielded through different individuals coming together and working as a team – is generally viewed as desirable. To achieve this, however, requires the various individuals becoming more open to new ideas and ways of communicating with each other. It also means learning about and accepting the other discipline's way of working. It is becoming increasingly apparent, however, that to enable this kind of mutual understanding to occur also requires some kind of lingua franca (Green et al, 1996). In particular, what is needed is a way of representing and talking about new concepts, that can be readily exchanged between the participating disciplines.

A good example of where researchers from different disciplines can work together and develop a new method is a project carried out by a team of sociologists and software engineers at Lancaster University (Sommerville et al, 1993). They were interested in developing systems for multiple end-users, in particular, for the domain of air traffic control. Their starting point was to acknowledge that the conventional software engineering approaches to requirements capture and analysis were inappropriate in their current form for use in the design of these kinds of collaborative systems. This was because the software engineering methods – developed originally to support formal structures – were seen as being unable to cope with the dynamic and informal ways of working which groups of people invariably adopt in different work settings. An important step for the group was then to determine how groups actually work together and to then work out a way of using this information to inform the design of new systems. It was here where it was considered that the sociologists could help, by providing accounts of group working from their ethnographic studies. However, a dilemma emerged. Members of the two disciplines recognized a language problem: the sociologist's detailed descriptions did not fit in with the structured way of working required to do software engineering. They needed to determine, therefore, how the software engineers could usefully take on board the sociologist's findings and analyses. Their solution was to design a new tool that could enable the sociologists to represent their findings in a more structured form that the software engineers could more easily relate to and use when designing software.

However there are a number of persistent difficulties preventing multidisciplinary teams from successfully communicating ideas between each other as can be seen in another research study, this time on producing a tool for fashion designers (Scaife et al, 1994). This was a study of collaboration between computer scientists and a field research team with backgrounds in psychology/cognitive science. They encountered a number of problems: how to relate observation of work practices to design decisions; how to manage responsibilities within the project; what priorities to adopt in development; how to involve users in prototyping. Many difficulties stemmed from the differing backgrounds and concomitant assumptions about good practice and what the correct way of working should be. Goldstein and Alger (1992) make a useful distinction between (i) software development, which involves a model of how the world operates and (ii) software development methodology, which is a model of how people behave in a software project. Just this kind of dual perspective was operating in the project. The field workers were attempting to map their model of the (fashion designer) world into a series of incomplete prototypes. They were not thinking in detail about software issues. The software developers, with their own model of how development should occur, were resistant to this process. Thus initial debates about the completeness/functionality of the software conflated the issue of development methodology with the issue of what constituted an adequate model of the fashion design process itself. This is important because, as Bond (1992) points out, one problem of collaboration in joint design is that each party may have a "private justification language" (p.463). In this project the field team produced summary models and lists, believing that this was a format that software developers would be happy with. But there were problems with these summary formulations since they were still very close to the original data: they reflected the designers' terminology and concepts. This is likely to be an issue in any study where the need to provide abstractions for system design competes with the requirement to phrase models in a way that users can understand and comment on. This is a separate problem from that of understanding terminologies across disciplines.

What the above examples demonstrate, therefore, is that even multidisciplinarity in an applied domain can be highly problematic. It suggests that for multidisciplinary teams to have the best chance of succeeding then both the nature of the problem space has to be clear to all and that all parties concerned are prepared to act upon it by being willing to change how they do their research. This may mean taking a radical departure from what is prescribed in their parent discipline, but in breaking away novel solutions can emerge.

Interdisciplinarity as emergent: Understanding how external representations work in relation to human cognition

We have seen something of the problems (and opportunities) that can occur in the process of collaboration and communication in multidisciplinary team work. But in the cases we've discussed the aim was to produce a new artifact. What occurs when the goal is the more nebulous one of 'promoting understanding' or 'evolving new ideas' within a domain? We can examine something of this process by looking at our own efforts with a particular research problem: the role of external representations in cognition. We consider ourselves as cognitive scientists and so we shall briefly present this work through the lens of Norman's (1980, 1990) desiderata for cognitive science which we have quoted previously. Amongst other things, it will be recalled, he argued for the necessity of understanding interactions between issues and for a more applied and situated orientation. One question here is whether, in so doing, we can therefore allow ourselves the label of interdisciplinarists?

The impetus for our research was the lack of any generalisable theories in this domain. By this we mean explanations that could enlighten us on how people interact with different kinds of external representations – be they diagrams, animations, multimedia or virtual reality – for a variety of cognitive activities (e.g. learning, problem-solving, reasoning). In an extensive review of the literature, including cognitive science, education, psychology, instructional science, HCI and art history, we discovered a fragmented and poorly understood account of how graphical representations work, thereby exposing a number of assumptions and fallacies (see Scaife and Rogers, 1996). The main reasons for this state of affairs seemed to be twofold. Firstly prior empirical work had a strongly parochial flavor, with investigations largely limited to the efficacy of this or that representation within a particular context, such as a pump animation for teaching high school physics. Secondly such mainstream cognitive theories as were available focused primarily on internal representations (e.g. Bauer and Johnson-Laird, 1993; Hegarty, 1992), missing out much of the cognitive processing that goes on when interacting with external representations. Consequently there was little or no integration of levels of description in such process models as were offered. The result of this situation was a chorus of disaffection from the community of designers and educational advisers that there was little of general relevance coming out of the research to date.

This led us to recognize that we needed to (i) explain more adequately the interplay between internal and external representations – and (ii) use this analysis to better understand, design and select graphical representations, which are appropriate for the learning environment, problem-solving task or entertainment activity in question. In particular, we believed that the value of adopting this alternative approach would be to focus our attention more on the *cognitive processing* involved when interacting with graphical representations, the *properties of the internal and external structures* and the *cognitive benefits* of different graphical representations. What we were not sure of, however, was how best to characterize and operationalise this relationship between the internal and external. In essence, we needed a new set of concepts that were generalisable in both a theoretical and applied context.

We were fortunate enough to discover that we were not alone in our endeavor, that others in the field of cognitive science had begun to recognize the need to broaden and situate the base from which to explain cognitive behavior (e.g. see special edition of the journal of Cognitive Science, 1993). A few researchers had also specifically been giving external representations a more central functional role in relation to internal cognitive mechanisms (e.g. Cox and Brna, 1994; Kirsch and Maglio, 1994; Larkin 1992; Norman, 1993; Zhang and Norman, 1994). Others, too, had begun putting forward alternative concepts like re-representation and expressiveness – originating from philosophy and logic – to explain why certain graphical representations were more effective than others (e.g. Stenning and Tobin, 1997). Finally, we were much inspired by Green's (1989, 1990) work on cognitive dimensions, where he has sought to develop a set of high level concepts that are easy to use by academics and designers, alike, for evaluating the design and assessment of informational artifacts, such as software applications.

On the basis of these ideas and some of our own empirical research (Rogers and Scaife, 1997) we have been able to identify novel concepts that we consider as an useful analytic framework from which to explicate aspects of external cognition. These are: *computational offloading, re-representation*, *graphical constraining* and *temporal* and *spatial constraining*. Each of these refer to different aspects of the relationship between internal and external representations, and which can be operationalised further in terms of empirical predictions and specific design decisions. For example, these include design concepts of visibility, explicitness and annotatability (see Figure 1 for more details of this operationalisation). Our approach, therefore, has been to develop a better understanding of a domain by rejecting 'old' computational theories that focus exclusively on the internal mind and constructing an alternative framework that conceptualizes the problem space using a different kind of framework. This, we believe, conforms at least partially to what Norman and others would consider a broader approach. It tries to map between levels of description/analysis; has a 'purely' cognitive level (computational offloading) and an applied (design) level; allows the analysis of situated use of external representations, for example by identifying trade-offs between design decisions for clarity of information presentation.

At the highest conceptual level, cognitive interactivity refers to the interaction between internal and external representations when performing cognitive tasks (e.g. learning). At the next level this relationship is characterized in terms of the following dimensions:

• **computational offloading** - the extent to which different external representations reduce the amount of cognitive effort required to solve informationally equivalent problems
• **re-representation** - how different external representations, that have the same abstract structure, make problem-solving easier or more difficult
• **graphical constraining** - this refers to the way graphical elements in a graphical representation are able to constrain the kinds of inferences that can be made about the underlying represented concept
• **temporal and spatial constraining** - the way different representations can make relevant aspects of processes and events more salient when distributed over time and space.

For each of these dimensions we can make certain predictions as to how effectively different representations and their combinations work. These dimensions are then further characterized in terms of design concepts with the purpose of framing questions, issues and trade-offs. Examples include the following:

• **explicitness and visibility** – how to make more salient certain aspects of a display such that they can be perceived and comprehended appropriately
• **cognitive tracing** – what are the best means to allow users to externally manipulate and make marks on different representations
• **ease of production** – how easy it is for the user to create different kinds of external representations, e.g. diagrams and animations
• **combinability and modifiability** – how to enable the system and the users to combine hybrid representations, e.g. enabling animations and commentary to be constructed by the user which could be appended to static representations

Figure 1: A theoretical framework of cognitive interactivity (adapted from Rogers and Scaife, 1997)

Conclusion

So, are we inter-disciplinary? Well we're not really sure how to answer this and whether it is really important to do so. As researchers, we are all currently working within several environments (cognitive psychology, cognitive ergonomics, multimedia), moving back and forth between fields. De facto we are 'doing' interdisciplinary work since we are constantly having to translate between disciplinary-based concepts and, perhaps most importantly, producing a framework that is explicitly aimed at being intelligible to people from several disciplinary areas. In practice, we believe, it may be of far greater practical significance to recognize the moment for a new 'look' (our search for 'impetus') than to worry too much about precisely how we engineer it. Much of the writing on taxonomies of types of disciplinary collaboration (multi, trans, pluri, inter ....) reflects an abstraction that, while useful in an analytic sense, may have little use value in the lab. The emergence of radically new forms of knowledge in individuals is a persistent problem for epistemology but what is generally accepted is that collective cognition can proceed very effectively through the mutual awareness of other points of view. It is in this sense that we see the commonality of forms of disciplinary collaboration.

Acknowledgements

The authors gratefully acknowledge support from the European Commission Training and Mobility of Researchers programme, project COTCOS (Cooperative Technologies for Complex Work Settings), number ERBFMRXCT960014.

References

Bannon, L. (1992). Interdisciplinarity or interdisciplinary theory in CSCW? *Workshop Proceedings of CSCW'92 Workshop on Interdisciplinary Theory for CSCW Design*, Toronto, Canada.

Barnard, P. (1991). Bridging between basic theories and the artifacts of Human-Computer Interaction. In J. Carroll (Ed.), *Designing interaction:*

psychology at the human-computer interface (pp. 103-127), New York: Cambridge University Press.

Bauer, M.I., & Johnson-Laird, P.N. (1993). How diagrams can improve reasoning. *Psychological Science,* 4, 372-378.

Brown, G.D.A. (1990) Cognitive Science and its relation to psychology. *The Psychologist,* 8, 339-343.

Cognitive Science, (1993). 17(1) Special Issue: Situated action.

Cicourel, A. V. (1995) Cognition and cultural belief. In P. Baumgartner and S. Payr, (eds.), *Speaking minds: Interviews with twenty eminent cognitive scientists.* Princeton Academic Press, NJ. pp. 47-58.

Cox, R. & Brna, P. (1994). Supporting the use of external representations in problem-solving: the need for flexible learning environments. *(Research Paper No. 686A)* Edinburgh: Dept. of AI, University of Edinburgh.

Farah, M.H.J. (1984) The neurological basis of mental imagery: A componential analysis. *Cognition,* 18: 245-272.

Green, D.W. and others. (1996) *Cognitive Science: An Introduction.* Blackwell Publishers, Oxford, UK.

Green, T.R.G. (1989). Cognitive dimensions of notations. In A. Sutcliffe & L. Macaulay (Eds.). *People and computers V,* (pp. 443-459), Cambridge: Cambridge University Press.

Green, T.R.G. (1990). The cognitive dimension of viscosity: a sticky problem for HCI. In: D. Diaper, D. Gilmore, G. Cockton & B. Shakel (Eds.) *Human-Computer Interaction - INTERACT'90.* (pp.79-86) North Holland: Elsevier Publishers, B.V.

Hegarty, M. (1992). Mental animation: inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Language, Memory and Cognition,* 18(5), 1084-1102.

Hutchins, E. (1995). *Cognition in the wild.* Mass: MIT Press.

Johnson-Laird, P.N. (1988) *The Computer and the Mind.* CUP: Cambridge.

Kim, S. (1990) Interdisciplinary collaboration. In B. Laurel, B. (ed) *The Art of Human Computer Interface Design,* p31-45.

Kirsch, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science,* 18, 513-549.

Kuutti. K. & Bannon, L. (1993). Searching for unity among diversity: exploring the interface concept. In *Proceedings of INTERCHI'93,* (pp. 263-268). New York: ACM.

Larkin, J. H. (1989). Display-based problem solving. In D. Klahr & K. Kotovsky (Eds.) *Complex Information Processing: The impact of Herbert A. Simon.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Larkin, J. H. & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science,* 11, (1), 65-100.

Mantovani, G. (1996). Social context in HCI: a new framework for mental models, cooperation and communication. *Cognitive Science,* 20, 237-269.

Norman, D. (1980) Twelve issues for Cognitive Science. *Cognitive Science,* 4, 1-32

Norman, D. (1990) Four (more) issues for Cognitive Science. *Cognitive Science Technical Report, No. 9001,* Dept of Cognitive Science, UCSD, USA.

Norman, D. (1994). Cognition in the head and in the world, *Cognitive Science,* 17(1), 1-6.

Rogers (2000). Recent theoretical developments in HCI: their value for informing system design. (in prep).

Rogers, Y. & Scaife, M. (1997). How can interactive multimedia facilitate learning? In J. Lee (Ed) *Intelligence and Multimodality in Multimedia Interfaces: Research and Applications.* Menlo Park, CA: AAAI. Press.

Scaife, M. Curtis, E. and Hill, C. (1994) Interdisciplinary collaboration: a case study of software development for fashion designers. *Interacting With Computers,* 6, 395-410.

Scaife, M. & Rogers, Y. (1996). External Cognition: how do graphical representations work? *International Journal of Human-Computer Studies,* 45, 185-213.

Schunn, C.D., Crowley, K. & Okada, T. (1998) The growth of multidisciplinarity in the cognitive science society. *Cognitive Science,* 22(1), 107-130.

Stenning, K. & Tobin, R. (1997). Assigning information to modalities: comparing graphical treatments of the syllogism In E. Ejerhed & S. Lindstrom (Eds.) *Logic, Action and Cognition.* Dordrecht: Kluwer Academic Publishers.

Sommerville. I., Rodden, T. Sawyer, P. Bentley, R. & Twidale, M. (1993). Integrating ethnography into the requirements engineering process'. In A. Finklestein, & S. Fickas (Eds.) *IEEE Symposium on Requirements Engineering* (pp. 165-173). New York: ACM.

Thagard, P. (1996) *Mind.* MIT Press: Cambridge, MA.

The Royal Society (1996). Interdisciplinarity - Transport and the Environment.
http://www.royalsoc.ac.uk/st_pol08.htm

Von Eckardt, B. (1993). *What is Cognitive Science?* MIT Press: Mass.

Vygotsky, L.S. (1978). *Mind in society: the development of higher psychological processes.* (M. Cole, V. John-Steiner, S. Scribner, E. Souberman, Trans.) Cambridge, MA: Harvard University Press.

Zhang, J. and Norman, D.A. (1994). Representations in distributed cognitive tasks. *Cognitive Science,* 18, 87-122.