Running Shared Mental Models as a Distributed Cognitive Process

Adrian P. Banks & Lynne J. Millward

Department of Psychology,

University of Surrey,

Guildford,

Surrey,

UK.

GU2 5XH

Running Head: Shared Mental Models and Distributed Cognition.

Requests for reprints should be addressed to Adrian Banks, Dept of Psychology, University of Surrey,
Guildford, Surrey, UK. GU2 5XH (email: a.banks@surrey.ac.uk).

Running Shared Mental Models as a Distributed Cognitive Process

Abstract

Shared mental models theory normally takes the individual as its unit of analysis. This paper proposes a theoretical framework for studying shared mental models in which the model is considered to be distributed amongst the team. From this framework a cognitive process is predicted which describes how shared mental models are run. A team reasoning task requiring planning was used to illustrate this framework and test predictions derived from it. Two aspects of sharing mental models were studied; the degree of overlap between team members' mental models and the organisation of the division of the model between team members. Experimental results showed that the cognitive processes used were distributed amongst the team and support was found for most, but not all, aspects of the proposed process of running a shared mental model. The organisation of sharing was found to influence this process.

Introduction

The concept of shared mental models (Cannon-Bowers, Salas & Converse, 1990; Cannon-Bowers, Salas & Converse, 1993) has been used to explain the performance of teams in a range of circumstances including piloting aircraft (e.g. Stout, Cannon-Bowers, Salas & Milanovich, 1999; Heffner, Mathieu & Cannon-Bowers, 1995), operating tanks (Minionis, 1995), in software engineering teams (Carley, 1997) and in new product development teams (Madhavan & Grover, 1998). The shared mental models approach proposes that the overlap of individuals' mental models leads to greater shared expectations and explanations within a team which leads to improved coordination, communication and other team behaviours which in turn leads to superior team performance (Rouse, Cannon-Bowers & Salas, 1992). In short the shared mental models approach explains team performance through the mental models of individuals.

There is some evidence supporting this model of team performance. Heffner, Mathieu and Cannon-Bowers (1995) and Heffner, Mathieu and Goodwin (1998) found that sharedness of mental models in a team operating a flight simulator significantly improved team process which in turn improved team performance. Minionis (1995) found that shared team mental models significantly enhanced coordinated performance in teams operating a tank simulator. Likewise Stout, Cannon-Bowers, Salas and Milanovich (1999) demonstrated that increased sharedness across individual mental models aided communication under high workload and resulted in improved coordinated team performance, again on a low fidelity flight simulator.

Shared mental models are most commonly operationalised as the correlation between link weighted networks using Pathfinder analysis (Schvaneveldt et al, 1985) which represent individuals' mental models (e.g. Heffner, Mathieu & Cannon-Bowers, 1995; Heffner, Mathieu & Goodwin, 1998; Stout, Cannon-Bowers, Salas and Milanovic, 1999). This all points to a relatively simple concept in which individuals' mental models are compared along a unidimensional continuum ranging from 'exactly the same' to 'the exact opposite', passing through 'completely different' when the correlation is zero. The hypothesised process by which shared mental models operate is that individuals run their mental

models independently of each other to form expectations and explanations of the task and team. Where these are the same, performance is optimum (Cannon-Bowers, Salas & Converse, 1993). Differences in the models will be reflected in differences in these expectations and explanations. Therefore complete overlap across individual mental models is the optimum arrangement to ensure that team members have similar expectations and explanations for the task and team (e.g. Cannon-Bowers, Salas & Converse, 1990; Cannon-Bowers, Salas & Converse, 1993).

However, this paper will demonstrate that by hypothesising a mechanism in which all cognitive processes begin and end within each individual, a large part of the cognition that is relevant to teamwork is omitted by shared mental models theory. As a result this perspective does not fully address the question of the role of mental models in teams.

Here an alternative approach is adopted from which to view the cognitive processes essential to the team's task. It is proposed that the cognitive process of running a mental model can be divided or distributed amongst the team and other artefacts. One benefit of a team is to distribute and not to duplicate labour. A fully shared mental model entirely duplicates labour so that each team member does all the work rather than dividing it between the team¹. This is clearly not always desirable as it is inefficient to replicate everything and is moreover not necessarily always possible, as in for example, teams comprised of experts representing different disciplinary or professional fields.

The perspective of dividing the cognitive process derives from notions of distributed cognition. The defining feature of the distributed cognition literature is that it attempts to redefine the boundaries of the unit of analysis for cognition. Traditionally within cognitive psychology the unit of analysis has been the individual; however, this is not always the most appropriate unit. For example, Hutchins (1995a) analysed the cockpit of a plane as one socio-technical system of pilot, co-pilot and aeroplane. He argued that the positioning of markers on a dial indicating the range that speed etc. must be kept within when landing a plane, constitutes cockpit 'memory'. Hutchins (1995b) presents a detailed case

¹ Acknowledgements for this point must be credited to an anonymous reviewer.

study describing the task of navigating a U.S. Navy ship as a distributed task. No single person completes this computation; one person takes the bearing, another configures the chart and so on. As the state of the computation (the location of the boat) is not bound by the individual it must move through the system of individuals and artefacts. The cognitive processes therefore involve the “propagation of representational state across a series of representational media” (Hutchins, 1995b:117). Thus the location can be a physical point when represented on the chart or a number when represented as a bearing. The representation varies depending on which medium the state is in. Similarly Tsoukas (1996) describes an organisation as a distributed system as again no single person is privy to all the knowledge used by the organisation. The action of the organisation is not overseen by one individual but is determined by the cumulation of a series of local actions based on local knowledge. This analysis is in line with traditional metaphors of organisations as brains (e.g. Morgan, 1986), but neither of these accounts provide adequate specificity.

Overlapping with the literature on distributed cognition are concepts such as situated cognition (e.g. Clancey, 1997; Greeno & Moore, 1993) in which the role of external representations in cognition are explored. External representations are proposed to be more than inputs and outputs to cognitive processes which take place exclusively within the brain, they are part of the cognitive process. They can be used to offload some of the computation required, they can re-represent a problem making it easier or more difficult and they can constrain inferences about the representation by limiting the ways in which features can relate (Scaife & Rogers, 1996). Within the mental models literature O'Malley & Draper (1992) demonstrated that users' internalised mental model of a word processing package was not a complete model of the system. Rather, the internalised model consisted of what the computer display did not provide and also the information required to access what it does provide. The information is therefore distributed between “knowledge in the head” and “knowledge in the world” (Norman, 1988). In order to run the model information in the world is accessed and used, thus running the model cannot be understood independently of the interactions between the person and computer.

Thus the notions of distributed cognition, situated action and related concepts highlight what is lost by studying only cognition that occurs within the individual, as is the case in shared mental models theory. This paper will integrate these considerations into a new approach to shared mental models with a view to investigating the full range of cognitive processes harnessed in the completion of a team task.

In order to hypothesise about the implications of the distributed cognition concept for our understanding of shared mental models it is necessary to clarify what is meant by a mental model. The diversity of uses of this term prompted Wilson & Rutherford (1989) to propose that mental models are all things to all people and O'Malley & Draper (1992:73) to suggest "Talking about mental models can be a dangerous thing". However Moray (1997) attempts to draw all the differing approaches into one theoretical framework, highlighting that apparent differences in mental models are due to the range of tasks for which mental models can be used and reflect the flexibility of the construct rather than any fundamental differences. Moray's overarching definition of a mental model is:

"The canonical form of a mental model, as indeed of any model, is a homomorphic mapping from one domain to another, resulting in an "imperfect" representation of the thing modelled"
(Moray, 1997:238).

The basis of this definition is the assumption of a homomorphic or many-to-one mapping from the elements of the thing modelled to one element in the mental model. For example many elements in a pasteurisation plant such as the boiler, steam pump and so on could be mapped into one element in the mental model: 'heating system' (Moray, 1999). Commonly, as Holland, Holyoak, Nisbett & Thagard (1989) point out, such a model will not be valid in all situations. Therefore the model will contain exceptions for these situations and as such is technically a quasi-homomorphism. Both of these structures are plausible accounts of mental models. However this study can be applied in either case and as such the distinction between these two is not crucial here.

Rouse and Morris (1986) propose that mental models are used to describe, predict and explain system purpose, function, state and form. Of particular interest for this study is planning and so only the prediction of system states will be expanded upon. This function of mental models was first proposed by Craik (1943). Craik defined a model as:

“... any physical or chemical system which has a similar relation-structure to that of the process it imitates. By ‘relation-structure’ I do not mean some obscure non-physical entity which attends the model, but the fact that it is a physical working model which works in the same way as the process it parallels, in the aspects under consideration at any moment’.
(Craik, 1943:15).

Craik suggests the process of prediction involves simply using the model to parallel the real process. The outcome of the parallel run in the model when translated back into the terms of the real process is the prediction of what the outcome of the real process would be. This is sometimes referred to as running a mental model. Two key aspects of mental models identified by Rouse and Morris (1986) are crucial to this process: form and state. Form is the model itself, the homomorphic mapping consisting of elements and their relations which represents the thing that is being modelled. State is the dynamic configuration of the aspects of the model that can change as a result of running the model. For example, if the mental model is of a bath then it will have a certain size and therefore maximum volume, which does not change. This constitutes a ‘form’ feature. All of these form elements constitute the relation-structure described by Craik. But the amount of water actually in the bath at any point can vary over time according to the setting of taps, whether the plug is in and so on, so the amount of water in the bath is a ‘state’ feature. The running of the model is the computation of a future model state on the basis of a given state and the relation-structure of the model.

The idea of running a parallel simulation mentally remains the same in more recent papers on the topic. The process has been studied experimentally in a range of contexts, for example Gentner & Gentner (1983) propose that there are two main models of electricity, flowing waters and teeming

crowds. These make differing predictions about the behaviour of a simple circuit and Gentner & Gentner show that the differences in predictions were in alignment with the different mental models held, suggesting that the predictions were made using those models. Similarly Green (1997) elicited mental models of population change and found several different models existed. Subjects then made actual predictions of population change and the differences between these were found to be in alignment with the differences in the mental models held. It seems therefore that prediction can occur through running a parallel model of a process mentally and retranslating the outcome of this back into the terms of real process, as Craik suggested.

The Present Study

This paper will focus on only one function of shared mental models, planning future actions. This is a useful task for testing the broader issues of decision-making and problem solving by the team, which have been suggested to involve shared mental models (Cannon-Bowers, Salas & Converse, 1993). Decision-making and problem solving with mental models relies on running the mental model (e.g. Klein, 1993) and so a task involving the planning of future action tests this underlying theoretical process. Within the existing literature shared mental models are predicted to be greatly beneficial in planning future actions. It has been proposed that less overt planning time is required by teams with shared mental models because team members will be able to predict what others will expect of them and so it will not be necessary to verbalise this (Rouse, Cannon-Bowers & Salas, 1992). The less they can predict, the more they will need to verbalise and so planning will be less efficient. Therefore extant theory predicts that optimum planning will occur when mental models completely overlap and performance will degrade the less similar mental models are. This hypothesis is logical within the current view of shared mental models. However, through applying the distributed cognition perspective a different conceptualisation of the cognitive processes that are necessary for planning in a team emerges.

Returning to Craik's original arguments about prediction occurring through running a parallel model the key concept is that of the model. A model based in the 'nervous system', as Craik put it, of one

person is simply one medium in which modelling can operate successfully. He cites mechanical models such as Kelvin's tide-predictor - a mechanical device to forecast tide movement - to highlight the range of media in which models operate. It is therefore possible to conceive of an abstract model of a system, namely the elements of a system and their relations, which is not necessarily tied to any particular medium. In principle this model would be as effective as any other model in predicting future states as, in principle, it is equally able to model reality. If part of this model is placed in one medium and part in another, the model would be equally successful in predicting future states as an equivalent model in only one medium as long as there is free communication between the media or the propagation of the representational state across the two media, in Hutchins's (1995b) phraseology. If this can occur then the differences between the media are not of great relevance to the actual task of running the mental model. In this way a mental model can be distributed. The relevant unit of analysis for cognition is not only the individual, but simultaneously the model which is distributed across a system of individuals. In principle then, shared mental models are a distributed system which runs a model collectively through the propagation of representational states across representational media.

A key feature of distributed systems is that they are comprised of elements which can be decomposed into modules or relatively independent components of the whole system (Chandrasekaran, 1981). This allows different aspects of the cognitive process to occur simultaneously as the individuals in the system engage in different tasks. The important point is that the modules into which the cognition is decomposed must be able to operate autonomously for the parallel activities to take place. The outcome of these can then be combined at a later point. Depending on the nature of the model, it is possible for some mental models to be decomposed. For example, Moray, Looftseen and Pajak (1986) studied a model of hot water flowing through four interconnected containers. This model can be decomposed such that each tank is considered a module on its own and its behaviour estimated and then combined with the predicted state of the other tanks to discover the overall system state.

In the case of shared mental models parts of the overall model are divided or distributed amongst the team members. If individuals have some complete modules but not all of the modules of which the

model is composed it is not possible to run the model in its entirety. However the modules can be run independently so it is possible to place some constraints on the outcome of running the model. When these are communicated to other team members who have other modules, they can further constrain the range of possibilities. For example the water tank described above may be computed to have any amount of water in it up to half full. If this constraint is communicated to the other team members who have the remaining modules, they can compute further constraints on the state, for example the water tank cannot be less than one third full. This iterative process can be repeated until there is a single remaining state which is consistent with all of the modules of which the model is composed (i.e. assuming that all the mental model parts are consistent with each other). Thus the possible future state moves between the team members being continuously refined until the final solution is reached, the outcome of running the model.

This iterative process will eventually reach the final state if it is possible to place constraints. However, this depends to a certain extent on the division of the model. If the model is divided in a disorganised fashion so that individuals do not have entire modules, that is modules are split between people, then it is not possible to compute the outcome of that module and one aspect of the overall model is lost. To circumvent this, it would be necessary to transfer information about the module - its form rather than its state - to one person so that they can complete the computation with the iterative process described above.

Given the specific components of a mental model described above, such as the distinction of form and state, it is possible to code all of the communications used and describe all of the actions that can be undertaken with respect to a mental model. Firstly, communications can refer to the model form, either 'requests' which ask what the model form is and 'offers' which inform what the model form is. Secondly, communications can refer to the model state. When individuals have computed some constraints on the state with the modules they know they can 'propose' these constraints and so initiate the process of reaching a prediction or ask a 'question' about the mental model state. Once the initial constraints have been mooted the rest of the team can alter them in only three different ways. More constraints can be added by 'honing' which causes the range of possible states to be narrowed.

Constraints can be removed by suggesting other possible states which is 'widening'. Finally it is possible to not change the currently proposed state, but agree with it by 'supporting'. In the present study all communication relevant to running the mental model were coded into these categories. Other communication which referred to team or task functions not concerning mental models were coded as 'other' and finally irrelevant or nonsensical comments were coded as 'not relevant'.

In order to test this proposed process a laboratory task manipulating the mental model construct was designed. As discussed above mental models consist of both the internal and external representations used in the model, therefore the mental models were manipulated by altering the external representations used. This means that the contents of much of the model is explicit and so both the effect of the manipulation and the actual degree of sharing are clear. As a result this methodology enables a more rigorous test of theory which can be applied to more externally valid, naturalistic teams in future studies.

The experiment involved teams of three participants solving a computer based team reasoning task. An approximate representation of the computer display for a typical trial is presented in figure 1. The task required the team to navigate around six islands in their boat, collecting scrolls. On each island there was a scroll numbered one to six and these had to be collected in numerical order as quickly as possible, in the minimum number of moves whilst making as few errors as possible. To collect all the scrolls required a series of moves - some direct and some via one or two other islands. There were three rules constraining which island it was possible to move to from any particular island. Firstly, the six islands were coloured either red, green or blue, two of each colour. The only legal move was to an island of a different colour. Secondly, each island had a particular inhabitant; either headhunters, trolls, priests or vegetarians and two sacrifices; from heads, fish, devils or pumpkins. In order to move to an island it was necessary to have the appropriate sacrifice on the island at which participants were currently situated for the inhabitants of the island they wished to move to. To move to an island with headhunters, heads were required; trolls required fish, priests required devils and vegetarians required pumpkins. Finally, each island was protected by a specific god and it was necessary to invoke this god in order to gain access to the island. However, it was not possible to invoke specific

gods immediately following other gods. The god of love could not precede fire, fire couldn't precede war, war couldn't precede peace and peace couldn't precede love. Each of these were independent, for example it was possible to calculate the constraints imposed by island colour without the god or sacrifice information. Therefore each rule and associated information was a module and so the model consisted of three modules, islands, gods and sacrifices which could be computed independently but needed to be combined in order to check the legality of a move.

Insert figure 1 about here.

Three different types of sharing were studied. The 'Complete' condition in which each participant had all three modules and thus the whole mental model. The 'Organised' Condition in which each participant had two of the three modules and none of the third, for example all of the god and colour information but none of the sacrifice information. These were arranged so that each participant had a different combination of modules ensuring that all the information necessary for the task was given to the team. Finally the 'Disorganised' condition in which the participants had two-thirds of each of the three modules, for example colour information for four out of the six islands but blank for the remaining two, four out of six gods and so on. Again this sharing was arranged so that all the information necessary for the task was given to the team. Thus the degree of overlap between individuals was higher in the Complete condition than the Organised and Disorganised conditions, in which the overlap between individuals was the same. But these conditions differed because the overlap in the Organised condition was more coherent as modules were presented in their entirety whereas they were not in the Disorganised condition.

According to the existing theory of shared mental models the following outcomes would be predicted. In the Complete condition participants have a total overlap of the mental models whereas the Organised and Disorganised conditions have a partial overlap. Therefore the Complete condition would be predicted to require less planning time than the other two conditions. The Organised and Disorganised condition have the same amount of overlap and these would therefore both be predicted to require the same amount of planning time as each other. There would be no predicted

differences in the number of errors or the number of moves as shared mental models theory does not directly address the accuracy of the models. However, the predictions made by the mechanism proposed above differ both in the greater specificity of what is predicted about the planning process and as a result of this the predictions concerning time taken directly oppose those of the existing shared mental models theory.

Firstly, it is predicted that the cognitive process is indeed distributed. This can be ascertained from the difference in the number of statements referring to model form in the different conditions. If the reasoning that occurs is not distributed, then participants would need to gather all of the model themselves in order to plan the next move. To do this requires 'request' and 'offer' communications. If this were the case then the Complete condition would have significantly less requesting and offering than the other two as all the information is given to all the participants and the Organised and Disorganised conditions would have equal requesting and offering as they both lack the same amount of the overall mental model. However, if cognition is distributed in the manner proposed then neither the Complete condition or the Organised condition will request or offer information as the model is distributed in such a way as to allow constraints on the predicted state by individuals independently. However in the Disorganised condition this is not possible and so there will be more request and offer communications. If the cognition is distributed, it is hypothesised that:

Hypothesis 1: The number of request and offer communications in the Complete and Organised conditions will be the same. The number of request and offer communications in the Disorganised condition will be significantly more than either of these two conditions.

Secondly it is predicted that the process used by the team to complete the planning will differ between the conditions. As participants in the Complete condition have all of the mental model, they will not need to use the refining process of honing and widening as much as in the Organised condition in which individuals cannot place all of the constraints on the predicted state independently. The Disorganised condition will also use honing and widening, but because of the lack of complete modules this process will be inefficient and so they will use honing and widening statements more than the Organised condition. If this mechanism is correct it is hypothesised that:

Hypothesis 2: The number of hone and widen communications in the Organised condition will be significantly more than in the Complete condition. The number of hone and widen communications in the Disorganised condition will be significantly more than in the Organised condition.

Due to the differences in the amount of communication predicted to be required and the patterns of communication in the different conditions, it is hypothesised that the trials will take different times to complete. This prediction is different to that derived from the existing shared mental models theory described above.

Hypothesis 3a: The time taken to complete the trials in the Complete condition will be significantly shorter than the Organised condition which will be significantly shorter than the Disorganised condition.

However, although the process hypothesised may vary in time, it does not predict that this will necessarily lead to a lower quality of running the mental model. Therefore, it is not predicted that the quality of prediction will alter.

Hypothesis 3b: The number of errors will be the same in all the conditions.

Hypothesis 3c: The number of moves will be the same in all conditions.

Method

Participants

Participants were 36 students at the University of Surrey who volunteered to take part in the study. There were 25 female and 11 male participants. Their mean age was 28.2 years, with a standard deviation of 8.7 years.

Apparatus

The experimental task was a team reasoning game based on three networked Pentium II PCs. Presented on each participant's screen was the information allocated to that participant for the trial. Also there were a series of buttons with which participants interacted with the computer. There was one set of buttons on their screen pertaining to one module; either gods, sacrifices or islands, so that they each controlled one of the modules, but each had a different one. By clicking on a button this choice was relayed to each of the other participant's screen to a text box which displayed their current choice of island, god and sacrifice. Participants had a 'go' button which they could click when they were satisfied with the choice of island, god and sacrifice on their screen. If any change was made to the choices on the screen, all the participants were required to click go again. In this way consensus was ensured. When all the participants had clicked go, the computer would check the legality of the move. If legal, the boat would move to that island. If illegal, a message box would appear telling them the move was illegal. The timing of all of the buttons clicked and movements of the boat was recorded in a database on the computer.

In total there were nine trials and so nine different games were needed in order that the solution was novel in each. It was essential that these were equivalent. Therefore nine different games were constructed which each had a total of twelve possible legal moves. The minimum number of moves required for each was also twelve moves and in each game this consisted of one direct move, four moves via one other island and one move via two other islands. These steps were in different orders in different games though. The total number of each type of god, sacrifice and island colour was the same in order that no particular element was practised more than another. Likewise the positioning of these elements and the scrolls were equally spread through all the trials in order to prevent familiarity with certain combinations. Therefore the games were different, but equivalent.

During the experiment all the verbal communication of the team was recorded using a Genexxa tape recorder.

Design

A repeated measures design was used to examine the effects of three conditions, '*Complete*', '*Organised*' and '*Disorganised*'. The order of the games, the module for which participants had buttons for and the order of the conditions were randomised. These were then rotated using an extension of a Greco-Latin square design involving three rotating factors in order to counterbalance the experiment. Games were rotated at a rate of one position per team, the buttons used were rotated at a rate of two positions per team and the order of the conditions were rotated at a rate of four positions per team. In this way the three factors moved relative to each other as well as to their order in the study and so all combinations of the three factors were used giving a more complete counterbalancing.

Procedure

Participants completed the team reasoning task in teams of three people. There were three trials in each condition and so a total of nine trials were completed by all the teams plus a practice trial prior to the experimental trials in which subjects learnt the rules of the task. In total this took approximately two and a half hours and so the trials were split into two sessions of five trials each on consecutive days lasting approximately one and a quarter hours which reduced the effect of fatigue on performance.

In the Complete condition all the information was present on all the participant's screens. In the Organised condition the information for gods and colours but no sacrifices were given to one participant, information for gods and sacrifices but no colours were given to another participant and information for sacrifices and colours but no gods were given to the third participant. In the Disorganised condition two thirds of the information on all of the three modules was given to all the participants. They each lacked a different third so that in total all the necessary information was present. In each of three trials within each condition the participants had the buttons pertaining to a different module; islands, sacrifices and gods. They therefore controlled each of the modules in each of the conditions. All of buttons pressed by the participants and the moves of the boat were timed and

recorded on the computer. All of the communication amongst the team was recorded using a tape recorder.

All of the verbal interaction of the teams was transcribed and coded according to the coding scheme outlined above. There were nine codes in this: 'request', 'offer', 'question', 'propose', 'hone', 'support', 'widen', 'other' and 'not relevant'. One of the authors coded all of the transcripts. A second coder, naïve to aims of the study, coded one quarter of the transcripts in order that inter-coder reliability could be assessed. Cohen's Kappa (1960) was calculated to be 0.80.

Results

Equivalence of the Trials

In total nine games were completed by each subject. Although these games were equivalent in the number of possible moves, minimum number of moves to solution and a counterbalanced design was employed, a one way ANOVA was used to investigate differences in performance between the nine trials, across the different conditions. This analysis revealed no significant differences in the time taken ($F(8,99) = 1.18, p > 0.05$) or in the number of errors ($F(8,99) = 1.06, p > 0.05$). But there was significant difference in the number of moves made ($F(8,99) = 2.47, p = 0.02$). Tukey's HSD test was used to identify which trials differed and significant differences were found. Game 1 required significantly less moves than game 3 ($p = 0.05$) and game 2 required significantly less moves than game 3 ($p = 0.05$). These findings suggest that the games were similar with respect to time and number of errors, but not with respect to number of moves.

Insert table 1 about here.

Effect of Counterbalancing

As several trials were administered to each team using a repeated measures design it is necessary to check for any order effects on the time taken to complete the trials, the number of moves made and number of errors made. A two way ANOVA was used to test for an interaction of condition with trial number. This analysis revealed no significant interactions for the time taken ($F(15, 82) = 0.86, p > 0.05$) or for the number of moves made ($F(15, 82) = 0.68, p > 0.05$) or for the number of errors made ($F(15, 82) = 0.64, p > 0.05$). These findings suggest that the main effects of condition were not moderated by order.

Of related interest is whether the number of statements changed over the duration of the experiment, i.e. whether the efficiency with which teams communicated changed over trials. As these data were

separated by trial number and condition, the total number of data in each cell of the analysis was quite small. Therefore a Kolmogorov-Smirnov test was used to test for normality. The samples were not normally distributed ($p > 0.05$) and so a non-parametric test, the Kruskal-Wallis, was used. This analysis revealed a significant difference in the number of widen statements in the disorganised condition $\chi^2 = 17.2$, d.f. = 8, $p = 0.03$. No significant differences were found in any other the other conditions for any type of statement ($p > 0.05$). These findings suggest that order affected only one type of communication in one condition. As the mean number of widen statements in this condition was very low (see table 3) the impact of this order effect was probably limited.

Time

Table 2 displays the mean time taken to complete the trials collapsed across the trials and participants. A one way repeated measures ANOVA revealed an effect of condition ($F(2,34) = 18.4$, $p < 0.0001$). Tukey's HSD revealed that the Complete condition was not significantly different to the Organised condition ($p > 0.05$) but the Disorganised condition was significantly slower than both of these ($p < 0.05$).

N° of Moves

Table 1 displays the mean number of moves taken to complete the trials in each of the three conditions. A one way repeated measures ANOVA revealed no effect of condition ($F(2,34) = 1.0$, $p > 0.05$).

N° of Errors

Table 1 displays the mean number of errors made in the trials in each of the three conditions. A one way repeated measures ANOVA revealed no effect of condition ($F(2,34) = 1.8$, $p > 0.05$).

Insert table 2 about here.

Codes

Table 3 displays the mean number of communications according to code in each of the three conditions. One way repeated measures ANOVAs were used to analyse the differences between conditions. There was a significant effect of condition on the number of 'request' statements ($F(2,22) = 100.0, p < 0.0001$); the number of 'offer' statements ($F(2,22) = 73.5, p < 0.0001$); the number of 'hone' statements ($F(2,22) = 10.8, p = 0.001$); the number of 'question' statements ($F(2,22) = 20.8, p < 0.0001$); the number of 'propose' statements ($F(2,22) = 10.2, p < 0.001$); the number of 'support' statements ($F(2,22) = 17.1, p < 0.0001$); the number of 'other' statements ($F(2,22) = 38.6, p < 0.0001$) and the number of 'not relevant' statements ($F(2,22) = 8.9, p < 0.002$).

There was no significant differences between the conditions on the number of 'widen' statements ($F(2,22) = 2.7, p = 0.49$).

Tukey's HSD was used to investigate the differences between the conditions. For the 'offer', 'request' and 'other' codes the Disorganised condition was greater than both the Complete ($p < 0.05$) and the Organised condition ($p < 0.05$) but the Complete and the Organised condition did not differ ($p > 0.05$). For the 'question' and 'support' codes the Complete condition was less than the Disorganised condition ($p < 0.05$), but Organised condition did not differ from either the Complete condition ($p > 0.05$) or the Disorganised condition. No significant differences were found between the conditions for 'propose', 'hone', or 'not relevant' statements ($p > 0.05$).

Insert table 3 about here.

Excerpts of Verbal Interaction

All verbal interaction was recorded, transcribed and coded. Typical examples of the communication in each condition, which was similar across all the groups, are presented below. Following each statement is the code ascribed to it in square brackets.

Complete condition

G (Initial letter of participant's name): We can't go straight away cos it's the same colour. [Hone].

D: We cannot go back to east island because the gods are not right. [Hone].

G: Can we go to fourth island? [Question].

D: Yeah. [Support].

G: Yeah, so devils. [Support].

H: Yeah. [Support].

Organised Condition

I: Can we go straight to east island? [Question]. Colours are OK. [Support].

S: No, we need a pumpkin. [Hone].

I: Oh, we need a pumpkin. [Other].

D: We can go to north can't we? [Question].

S: We can't go back to north because of the gods. [Hone].

Disorganised Condition

A: Can we go to 6 or 2? [Question].

L: They're the only two we can go to, aren't they. [Support].

S: OK, you'll have to fill me in then because I... [Other].

L: We can't go to fourth island because we haven't got a head and it's headhunters. [Hone].

A: Must be south island then. [Propose].

L: So we'll have to go to south island. [Support].

S: South island: Is that a troll or is it a priest? [Request].

A: Trolls. [Offer].

Discussion

The results of this study supported the idea that the cognitive processes used for running a shared mental model were distributed between the team members. They also provide some but not unequivocal support for the proposed mechanism of running a shared mental model. Finally the results supported the idea that different patterns of sharing according to degree of overlap and organisation affected the running of the model in some ways but not in others. Thus hypotheses 1 and 3a,b and c were fully supported and hypothesis 2 was partially supported.

Hypothesis 1 stated that the number of request and offer statements would be the same in the Complete and Organised conditions, but that there would be more in the Disorganised condition. This was the case. Thus the Organised team was able to reach a solution with less information than the Complete condition, as they began with less and requested and offered the same amount, but the Disorganised condition could not. This suggests that at least in the Organised condition the model was run by individuals with only a part of the full model each, and so the cognitive process used in planning was distributed amongst the team. It is not clear whether the Disorganised condition involved distributed cognition as it is not known whether the amount of information gathered was used to create a complete mental model in one individual or whether it was used to create whole modules which could then be used to run the model distributively. Likewise, it is unknown whether the Complete condition used distributed cognitive processes to run the mental model. Simply because it was possible to run the model individually does not mean that this was approach adopted. But there is convincing evidence that shared mental models were run using a distributed cognitive process in at least the Organised condition.

Hypothesis 2 also addressed the proposed process of running a shared mental model using a distributed cognitive process by comparing the differences in the types of communication between the conditions. A significant difference was found in the number of honing statements used in the conditions, but the post hoc tests did not reveal where these differences lay. Although it cannot be determined which conditions differed, this finding suggests that the iterative refining process central

to the distributed cognition was affected by either organisation or overlap or both. The second part of hypothesis 2, that there would be more widen statements in the Disorganised condition than the Organised condition and also more hone statements in the Organised condition than the Complete condition, was not supported. This does not support the proposed process, but does not rule it out either. It is possible that this particular task does not lend itself easily to widen propositions for some reason, possibly due to it's narrowly circumscribed nature. Evidence from a range of tasks is required to investigate this further.

Hypothesis 3 stated that conditions would vary in time, but not in number of errors or moves. As hypothesised the Complete condition and the Organised condition were not significantly different, but the Disorganised condition was significantly slower than both of these. This finding is the opposite to that predicted by existing shared mental models theory which suggests the greater overlap in the Complete condition would cause the team to be quicker than in the Organised and Disorganised conditions which had the same overlap and so would be expected to take the same amount of time. However there were no significant differences between the conditions in the number of errors or the number of moves made. Therefore the quality of running the model did not suffer as a result of the different patterns of distribution. This suggests that the extra time taken in the Disorganised condition was not due to lengthy routes or time consuming mistakes and so was therefore due to longer times between moves. This was most likely to be caused by the greater amount of communication needed to run the model indicated by the significantly greater amount of request, offer and hone statements in the Disorganised condition. Therefore the degree of organisation, which differed in the Disorganised condition compared to the other two, affected the time taken to run the model. This finding supports the mechanism proposed here over the existing theory of shared mental models which does not predict a role for organisation.

Interestingly there were not the differences between the Complete condition and the Organised condition that were predicted by hypotheses 2 and 3a. This could be for two reasons. Either the conditions were too similar and so differences were not significantly different as the effect size was small. This is possible, although the difference between complete overlap and an overlap of two thirds

in the mental models seems objectively quite large. Another possible reason is that even in the Complete condition participants used a distributed cognitive process even though it was possible to complete the task individually. The data supports this to some extent, for example there were 394 hone statements and 29 widen statements. In principle none of these are required as the individuals could make all of these deductions independently. However they do not. A distributed strategy requires only the running of one module at a time rather than three and so for an individual this requires less cognitive work. It is conceivable then that in the Complete condition teams used a distributed strategy so that the cognitive work was spread amongst the team rather than any individual completing it all. If this were the case then the number of hone and widen statements would not be expected to differ, which they did not. As a result, the amount of time communicating would be the same leading to the amount of time to complete the task being the same, which it was. It is possible that the lack of differences between the Complete condition and the Organised condition is caused by the teams treating the Complete condition in the same way as the Organised condition, rather than the manipulation of the Organised being too similar to the Complete condition.

Overall, all the hypotheses received total or partial support. Generally it was found that the Complete and Organised condition were similar. These conditions differed on overlap but were both organised. However the Disorganised condition was generally dissimilar to both the Complete and Organised condition. This conditions had the same degree of overlap as the Organised condition but was not organised into modules. Therefore, the organisation into whole modules seems to have more of an effect on the running of shared mental models than the degree of overlap.

Limitations and Future Work

The major limitations of this study are related to the way in which the process and degree of distribution are assessed. The degree and nature of distribution in all conditions should be measured more accurately, for example by asking participants to what extent they used this strategy. Some evidence for this can be taken from the transcripts, for example this quote from the disorganised condition:

S: OK, so shall we do it so that like I'll inquire to you guys as to what the thingies are and L, you can inquire as to what the...

A: I think we've all got different on each one.

S: Yes, we do, but if we need to know, if it's our responsibility kind of thing.

Quotations of this kind, which is not unique, indicate a distributed strategy is being used in the disorganised condition. But this analysis is post hoc and not sufficiently rigorous to do more than suggest how the mental model was distributed. It would be useful to know more about how and when this occurs.

Likewise the hypothesised process used merits further development, in particular the role of widening statements for which the predicted results were not found. Also the degree to which team members check each other's work was not controlled or measured and so it is not known to what extent this occurs or how it is influenced. For example it may be influenced by social factors or the expertise of the team members - team members cannot fully check the work of experts in a different field but must deal with it in some fashion. Thus this study indicates that the nature and process of the distribution of shared mental models is a useful area of research, but it requires more work to further specify the theory.

The Distinct Contribution of a Distributed Cognition Perspective and Its Implications For Shared Mental Models.

These findings justify the adoption of a distributed cognition approach to studying shared mental models. The distributed cognition approach allowed the possibility that using the mental model required a process undertaken by the whole team rather than the cognitive processes being solely undertaken by the individual and then the team performance being solely determined by the similarity or dissimilarity of the individual processes. Once the possibility of a team process had been considered, the nature of such a process could be hypothesised from the underlying form of a model. Finally, from this hypothesised process a series of predictions could be formed which were more

specific about the role of planning in shared mental models. These predictions were largely supported suggesting that the distributed cognition approach contributes a distinct new understanding of the role of shared mental models in planning. As the process of planning is a useful test bed for the broader cognitive processes of decision-making and problem solving, this approach may have significance for areas of team work which are currently poorly explained by shared mental models theory.

References

- Cannon-Bowers, J., Salas, E. & Converse, S. (1990). Cognitive psychology and team training: Training shared mental models of complex systems. *Human Factors Bulletin*, 33, 1-4.
- Cannon-Bowers, J., Salas, E. & Converse, S. (1993). Shared mental models in expert team decision making. In N. J. Castellan (Ed). *Individual and Group Decision Making: Current Issues*, pp. 221-246. Hillsdale, NJ: Lawrence Erlbaum Assoc.
- Carley, K. M. (1997). Extracting team mental models through textual analysis. *Journal of Organizational Behavior*, 18, 533-558.
- Chandrasekaran, B. (1981). Natural and social system metaphors for distributed problem solving: Introduction to the issue. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-11, 1-5.
- Clancey, W. J. (1997). *Situated Cognition. On Human Knowledge and Computer Representations*. Cambridge: Cambridge University Press.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20, 37-46.
- Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge: Cambridge University Press.
- Gentner, D. & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.). *Mental Models*. pp. 99-130. Hillsdale, NJ: Lawrence Erlbaum Assoc.
- Green, D.W. (1997). Explaining and envisaging an ecological phenomenon. *British Journal of Psychology*, 88, 199-217.

- Greeno, J. G. (1993). Situativity and symbols: Response to Vera and Simon. *Cognitive Science*, 17, 49-59.
- Heffner, T. S., Mathieu, J. E. & Cannon-Bowers, J. A. (1995). *The impact of mental models on team performance: Sharedness, quality or both?* Symposium presented at the annual meeting of the Society for Industrial and Organizational Psychology.
- Heffner, T. S., Mathieu, J. E. & Goodwin, G. F. (1998). *Team training: The impact on shared mental models and performance.* Symposium presented at the annual meeting of the Society for Industrial and Organizational Psychology. Dallas, TX.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E. & Thagard, P.R. (1989). *Induction: Processes of Inference, Learning and Discovery.* Cambridge, MA: MIT Press.
- Hutchins, E. (1995a). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.
- Hutchins, E. (1995b). *Cognition in the Wild.* Cambridge, Massachusetts: MIT Press.
- Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision-making. In G. A. Klein, J. Orasnu, R. Calderwood & C. E. Zsombok (Eds.) *Decision Making In Action: Models and Methods.* pp 138-147. NJ: Ablex Publishing Corporation.
- Madhavan, R. & Grover, R. (1998). From embedded knowledge to embodied knowledge: New product development as knowledge management. *Journal of Marketing*, 62, 1-12.
- Minionis, D. P. (1995). Enhancing team performance in adverse conditions: The role of shared team mental models and team training on an interdependent task. *Unpublished PhD, George Mason University, USA.*

- Moray, N., Lootseen, P. & Pajak, J. (1986). Acquisition of process control skills. *IEEE Transactions on Systems, Man and Cybernetics, SMC-16*, 497-504.
- Moray, N. (1997). Models of models of ... mental models. In T. B. Sheridan & T. Van Luntern (Eds.) *Perspectives on the human controller: Essays in honor of Henk G. Stassen*. pp. 271-285. Mahwah, NJ: Lawrence Erlbaum Assoc.
- Moray, N. (1999). Mental models in theory and practice. In D. Gopher & A. Koriati (Eds.) *Attention & Performance XVII: Cognitive regulation of performance: Interaction of theory and application*. pp 223-258. Cambridge, MA: MIT Press.
- Morgan, G. (1986). *Images of Organization*. Beverly Hills, CA: Sage.
- Norman, D. A. (1988). *The Psychology of Everyday Things*. NY: Basic Books.
- O'Malley, C. & Draper, S. (1992). Representation and interaction: Are mental models all in the mind? In Y. Rogers, A. Rutherford & P. A. Bibby (Eds). *Models in the Mind: Theory, Perspective and Application*. pp 73-92. London: Academic Press Limited.
- Orasnu, J. & Salas, E. (1993). Team decision making in complex environments. In G. Klein, J. Orasnu, R. Calderwood & C. Zsombok (Eds). *Decision Making in Action: Models & Methods*. Norwood, NJ: Ablex.
- Rouse, W. B., Cannon-Bowers, J. A. & Salas, E. (1992). The role of mental models in team performance in complex systems. *IEEE Transactions on Systems, Man & Cybernetics, 22*, 1296-1308.
- Rouse, W. B. & Morris, N. M. (1986). On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin, 100*, 349-363.

- Salas, E., Cannon-Bowers, J. A. & Johnston, J. H. (1997). How can you turn a team of experts into an expert team?: Emerging Training Strategies. In C. E. Zsombok & G. Klein (Eds). *Naturalistic Decision Making*, 359-370.
- Scaife, M. & Rogers, Y. (1996). External cognition: How do graphical representations work? *International Journal of Human-Computer Studies*, 45, 185-213.
- Schvaneveldt, R.W., Durso, F.T., Goldsmith, T.E., Breen, T.J., Cooke, N.M., Tucker, R.G. & De Maio, J.C. (1985). Measuring the structure of expertise. *International Journal of Man-Machine Studies*, 23, 699-728.
- Stout, R. J., Cannon-Bowers, J. A., Salas, E. & Milanovich, D. M. (1999). Planning, shared mental models, and coordinated performance: An empirical link is established. *Human Factors*, 41, 61-71.
- Tsoukas, H. (1996). The firm as a distributed knowledge system: A constructionist approach. *Strategic Management Journal*, 17, 11-25.
- Wilson, J. R. & Rutherford, A. (1989). Mental models: Theory and applications in human factors. *Human Factors*, 31, 617-634.

Table 1

Mean time in seconds, mean number of moves and mean number of errors (SDs in parentheses) for each game

Game	1	2	3	4	5	6	7	8	9
	M(SD)	M(SD)	M(SD)	M(SD)	M(SD)	M(SD)	M(SD)	M(SD)	M(SD)
Time	516.3 (150.1)	467.5 (121.4)	732.5 (350.0)	540.2 (181.6)	649.5 (487.0)	523.7 (306.5)	542.1 (231.2)	488.8 (110.3)	576.4 (214.4)
Moves	12.1 (0.3)	12.1 (0.6)	14 (2.5)	12.4 (1.0)	13.6 (2.7)	13.1 (0.7)	12.8 (1.5)	12.4 (0.9)	13.2 (1.1)
Errors	1.1 (1.6)	0.6 (0.8)	0.9 (0.9)	1.4 (2.0)	1.3 (1.4)	1.3 (1.6)	0.7 (0.9)	0.7 (1.2)	0.9 (1.1)

Table 2

Mean time in seconds, mean number of moves and mean number of errors (SDs in parentheses) for each condition (Complete, Organised, Disorganised)

	Complete	Organised	Disorganised
	M (SD)	M (SD)	M (SD)
Time (s)	422.1 (102.6)	541.8 (181.3)	715.1 (361.6)
Nº of Moves	12.6 (1.5)	12.8 (1.2)	13.2 (1.9)
Nº of Errors	0.7 (1.0)	1.0 (1.5)	1.2 (1.4)

Table 3

Mean number of communications by code (SDs in parentheses) for each condition (Complete, Organised, Disorganised)

	Complete	Organised	Disorganised
	M (SD)	M (SD)	M (SD)
Request	1.6 (1.3)	11.7 (9.6)	108.3 (36.7)
Offer	7.7 (4.1)	27.9 (18.3)	175.4 (70.8)
Question	29.2 (13.2)	48.7 (17.1)	65.1 (31.7)
Propose	54.9 (12.6)	66.3 (15.8)	85.7 (32.2)
Hone	33.2 (10.6)	52.6 (12.0)	49.6 (16.8)
Support	50.1 (20.2)	70.6 (25.6)	77.8 (28.2)
Widen	2.4 (1.4)	2.2 (3.0)	1.5 (1.8)
Other	70.8 (38.9)	103.3 (52.4)	164.8 (64.2)
Not relevant	8.3 (5.3)	13.4 (9.2)	28.8 (25.0)

Figure 1: Representation of the screen presented to participants during the experiment.

N.B. The boat, pumpkin, head, fish and devil were represented with a picture in the experiment. Island colour was represented through colouring the island rather than using a label.