2007

Towards a Framework for Supporting Professional Teamwork: Modelling Human Actions in Small Group Meetings

Terence Blackburn
University of South Australia, terence.blackburn@unisa.edu.au

Paul Swatman
University of South Australia, paul@swatman.net.au

Towards a Framework for Supporting Professional Teamwork: Modelling Human Actions in Small Group Meetings

Terence Blackburn, Paul Swatman
School of Computer and Information Science
City West Campus, University of South Australia
Adelaide, South Australia 5000
{terence.blackburn,paul.swatman}@unisa.edu.au

Abstract

Research into computer support for colocated small group meetings has not delivered any significant benefit in spite of significant investment and decades of research. We suggest that the focus has been on the wrong problem. Each time a process is enacted in small group work it will different (even if only slightly different) due to the situational context. Therefore the process for supporting the process should also be different each time and prescriptive, systematised support will be of little value in this situation. We propose that the processes must be exposed and identified dynamically so that useful support can be provided when it is required and one of the first steps is to understand what people are doing in these processes. The main contribution of this paper is a framework for helping researchers to understand the problem before considering a solution.

Keywords

Meeting, support, emergent, situated, actions.

Introduction

Research in Computer Supported Cooperative Work (CSCW) has moved away from workplace support to focus on alternative settings such as homes, games and museums (Crabtree, Rodden & Benford 2005). Technological advances in computing, such as increased bandwidth, more processor power and denser storage technologies, have allowed researchers to widen their scope to include mobile, ambient, ubiquitous, mixed reality, wearable and embedded computer technologies.

The move away from the workplace has happened prematurely. After decades of effort and significant investment, there are still no useful support mechanisms for ill structured, synchronous, colocated small group interactions; *meetings*, in other words - one of the key mechanisms for decision making and information dissemination in information systems. Where is the evidence for this claim? We only need to look at these meetings and (usually) the only visible technologies are slide shows on projectors (and mobile phones, which support individuals rather than groups). Where are all of the meeting support applications?

Our definition of the problem is different from previous research. Most of the *actions* that occur in meetings are emergent and situated, which means that they are influenced by their context and environment (Suchman 1987). The unpredictable nature of these actions – and, by definition, these actions are unstructured - makes it practically impossible to completely (or even partially) systematise technological support. We posit that work processes that are not (formally) structured are either semi structured or ill structured (where "structure" refers to predefined processes.) These ill structured processes are what we term "creative" processes because they cannot be fully systematised and *cognitive* effort – rather than technological effort - is required to execute them or find solutions for them. There is no premeditated focus on creativity per se, but instead a recognition that arbitrary work tasks will often be performed in a different way each time they are executed (Suchman 1987) and the path that people choose to navigate through their tasks is *emergent*. If we accept this, how can a systemised approach be flexible enough to cope with all of the variations in the processes? How can an application support an activity *now* when we don't know what is currently happening in a workspace and the current processes are likely to be different from those in our original problem domain?'

In our view, the CSCW community is trying to solve the wrong problem. (We also include complementary or associated domains such as GDSS and Groupware in this statement.) The problem that CSCW research tries to solve was defined with known pre and post conditions but without any knowledge of the contextual aspects of process enactment. The solution to the problem was designed in the mistaken belief that the *exact* same problem or task would reoccur when the support application was deployed. The enactment of the process in any task is context dependant and will be different from the same process of five hours ago or five weeks ago. Any two process enactments will be different in either subtle or obvious ways.

This paper proposes an alternative approach to providing support for synchronous, colocated group interactions but the scope of this work is limited to groups that have a "Western" cultural influence. It proposes that support for process enactment should be dynamically driven by the emergent actions of group participants. Existing tools and applications can be dynamically introduced as required to support the finer grained *actions* – as small parts of activities – rather than offering prescribed support for a sequenced set of actions where the application will be unable to keep in step with the emergent and unpredictable set of human actions.

The rest of the paper is structured as follows. Our background section explores some of the work and some of the criticism of the work that has tried to support ill structured tasks. We extend this in the next two sections to characterise the issues and nature of these tasks. Then we consider how to go about modelling human actions in small group settings such as meetings and this leads to our proposed framework. This takes up most of the rest of the paper although we include some discussion of our approach. We finish by summarising our work to date as well as providing a description of where we are taking this work.

Background

The claim that structured and prescriptive approaches have failed to support small group interaction is not a new observation. Writers have been articulating this problem either directly or indirectly for years. In 1990, for example, one view purported that modern technologies had failed to improve meeting efficiency or they were largely impractical (Robinson 1990). In 1997 it was suggested that although CSCW frameworks had contributed much to our understanding of the field, they could still be criticised for a lack of understanding of the complex social activities – perhaps too difficult rather than misunderstood - that make up group work and this often led to the failure of well designed CSCW applications (Ngwenyama & Lyytinen 1997). In 1998 developers were criticised for often focusing on the technology. They produced highly sophisticated software systems to address problems that may not have existed in real groups or they may not have effectively addressed their problems (Kamel & Davison 1998). In 2000 in a book detailing workplace studies, the editors argued that technologies explicitly aimed at supporting collaborative group activities missed their mark completely. They seemed to be ill suited to the contingent, emergent and collaborative aspects of the work they claimed to support (Luff, Hindmarsh & Heath 2000).

If we look beyond the CSCW domain and instead consider the current status of Group Decision Support Systems (GDSS), Paul Gray has a similar message to impart. He says that GDSS "is an idea which has come and, sad to say, has gone". That is not to say that the research in this field was not useful as it has had spin offs in data mining, knowledge management and business intelligence (Gray 2007). Some of the benefits of GDSS include the ability to increase the functional size of a small group beyond what is normally considered productive through the use of technological infrastructures, and the ability to negate, to some extent, the effects of individual dominance in a group through the use of anonymous contributions during activities such as brainstorming. (The latter has beneficial social side effects.)

Social aspects of group interaction have not been accorded much more attention than noting them and agreeing that that are important. Few systems have been built to support social characteristics such as: group building; equity sharing (in terms of shared power); and managing truth or deception in human interaction. It is likely that we are far from addressing these issues but one aspect that has been recognised and studied is the contingent and emergent nature of the human actions in small group tasks (Suchman 1987). These ideas - where the execution of a task is determined to some extent by its context, the people involved and the environment - have been studied and argued about for some time. Our research accepts this characterisation of work and, even though Suchman's ideas may be correct and often cited, trying to develop technologies to support the concept of situated and emergent work may not be useful. The ideas don't lead anywhere and it has been difficult to do anything with them (from a technological point of view).

Perhaps it is appropriate to use Kuhn's model (Kuhn 1962) and suggest that it is too early to adopt these ideas because technologies have not evolved far enough to support them. (We recognise that Kuhn believed that the paradigm shifts he described only applied to scientific domains but we would not be the first people to extend the idea beyond and into sociological contexts.) Notwithstanding the possible premature nature of our research, we include *changing circumstances*, *ad hoc actions* and *ill structuredness* as being essential components in small group interaction. We seek to design applications that treat *change* and *ill structuredness* as mandatory components of a small group work model and not as *exceptions* that are dealt with as undesirable side effects. Many attempts have been made to go down this path and "emergent workflow" (Jorgensen 2004) is an illustrative example.

In this work, as in much other research, a graphical model is used to provide a flexible representation of the work activities of self directed groups alongside their established work processes. The workflow activities are coordinated through "mutual adjustments" and travel through decision points. The modelling language is accessible to non programmers and participants use contextual awareness to allow themselves to steer their activities and to control enactment of the process. In summary, an emergent workflow is a flexible, partially structured process model that emerges from the work. The process definition is interleaved with it's enactment and allows users to dynamically add tools and information types (Jorgensen & Carlsen 1999). However, one of the costs of using the application is speed. This cost might be too high where people in meetings want to work quickly and interacting with a workflow may seem unnatural and may represent a time constraint.

Social issues and supporting the contingent, emergent nature of work are problems that have been well documented but have not been solved. A recent platform had as one of it's specific goals support for the emergent and unpredictable nature of work but it did not meet this goal (Bogia, LTolone et al. 1993). Successes have consistently fallen far short of expectations and issues (in no particular order) include: the disparity between those who will *use* an application and those who must *support* it; the lack of management intuition which leads to decision-making failure and to ill-fated development efforts; and the extreme difficulty of evaluating these applications (Grudin 1988).

Other similar approaches have included layered policies in the definition of the business process (Bogia, D P & Kaplan 1995), "advisory" process models (Abbot & Sarin 1994) and "libraries of actions", which are invoked when participants in the activity consider them to be useful (Carlsen & Gjersvik 1997). From our knowledge of the literature, none of these systems is used in real small group meetings. One question that arises from this work is 'if a flexible workflow is made so flexible that it can't execute without human intervention, when does it stop being a workflow?' Our answer is 'when it has to dynamically respond to human caprice, which is a completely non prescriptive activity'.

Issues in Colocated Meetings

The CSCW field has been dominated by models of collaborative teams in academic or design environments. Such settings are characterised by relatively equal levels of personal power and a division of labour established by the collaborating parties (Eason 1996). Outside of these domains many groups have unequal levels of power and introducing technology that levels the playing field can only be good for the weaker parties. Why then would the stronger, more powerful parties be interested in utilising technologies? Politics enters the picture.

Politics provides another reason for the dearth of computerised technologies in meetings (Vogel 2006). Groups sometimes may not want to *remember* decisions and some participants may not want their positions to be explicitly articulated and disseminated. A facility for "active forgetting" is deemed to be a requirement in a meeting system (Maaranen & Lyytinen 1995). There is a documented account of a group of researchers who tried to implement a colocated meeting support system in a round of multilateral diplomatic meetings. The account detailed the endless problems associated with the implementation, which was eventually abandoned (Maaranen & Lyytinen 1995). Lessons learned during the process included: organisational and procedural constraints can affect user acceptance; the design process must be broadened to include social and psychological factors; the whole room should be used as a user interface instead of just tailoring the application; and personal factors, task related factors and social aspects all represent a part of the process and should form components in every meeting support framework. If the solution is even partially misaligned with the problem, then, people may reject the technology solution for a number of reasons. These may include: it doesn't solve the problem satisfactorily, the cost of altering the computer generated solution may be too high or it doesn't solve enough of the problem to warrant using the technology. (In the case of "politics" there may be no technological support solution.)

What are the main issues in meetings that need computer support? The following list identifies a few of the issues: too much time is spent arranging meetings; insufficient time is spent in analysing issues during meetings; too much time is spent travelling to meetings; and time is wasted collecting and tabulating anonymous votes during meetings. It is difficult to collect and analyse data from meeting members, there is a poor focus on the task resulting in poor use of time and too much time is spent filling in forms and searching for documents (Kamel & Davison 1998). (Even though this reference is twenty years old the problems haven't changed because they haven't been satisfactorily solved.)

This work does not seek to render unnecessary the research and applications that have been built over recent decades but instead it seeks to find more efficient and timely ways of deploying them. We have identified a gap in the research, which we describe as providing *contextually relevant and dynamic technological support for meetings*. We want to *identify and articulate* exactly what people are doing in meetings at any given point in time. If we can do this with a high degree of confidence, we may be able to build and provide useful supportive technologies that deploy *when* they are needed and in the format that is required. This approach is non prescriptive and takes its cue directly from the meeting participants.

Autonomic, Automatic and Semi-automatic Support

This project is a component in a larger field of research (Swatman, Blackburn et al. 2005) that seeks to provide autonomic, automatic and semi automatic, ICT mediated support to synchronous, colocated small groups of professional "Western" workers.

Autonomic support means self managing or instigating an action or series of actions in response to a set of states observed by *sensors*. The concept comes from human physiology where the body acts in a way that is self repairing, self configuring, self protecting and self optimizing. Automatic and semi automatic responses are a little different in that interventions in a workspace are motivated by people's activities and the responses are the result of artificial intelligence predictions of what is likely to happen next. It is worth illustrating the differences with examples.

In a meeting room designed for six people but populated by ten people, it would be likely that the temperature would rise over a period of time and it may become uncomfortable. An *autonomic* intervention would occur when the air conditioner activated and started cooling the room to make it more comfortable. (This same thing also happens in the human body on a hot day.) An *automatic* intervention might consist of the computer interrogating a database for information that is contextually relevant to the dominant conversation in the meeting and displaying the information on a public screen. It means that group members will not have to explicitly invoke a service but, instead, the computer infrastructure will make a decision to intervene or *offer* an intervention (*semi-automatic* support). This highlights an important characteristic of this research; the *human interface to the technological environment*.

Modelling Human Actions

This interface is composed of cameras, microphones and other sensor devices that passively observe group members as they undertake their tasks. Cognitive dust, the raw data that represents evidence of human communications, is manifested and mediated between humans and sometimes through the use of computer devices (Blackburn, Swatman & Vernik 2007). This raw data can be observed and persisted in syntactic models that become increasingly coarser grained as the data is abstracted to higher levels. We can reason about these models within our theoretical framework in order to extract semantic meaning and our abstraction process stops at the highest level; a collection of *concepts*.

As an example of this process, a person "pointing" is a semiotic, communicative action that conveys the "desire" of one person to focus the gaze of a second person in a certain direction. This is an example of cognition being mediated from one person through an arm and finger to another person and creating a linked set of dust elements. This pointing action is recorded through a camera as a set of frames and our model allows the infrastructure to reason that the first person, the pointer, wishes to transfer a piece of cognition, the "desire", to the second person, whose focus should change in accordance with the direction of the finger. (Different levels in the model have been omitted from this explanation such as joining a set of frames to represent an arm movement and another set of frames that show the direction of the second person's gaze.) Semantic meaning has been extracted from the pointing model and the infrastructure has "understood" the communicative intent. Automatically generated interventions may ensue at the focal point of the gaze (for example, a public screen) but this is dependant on verification by other evidence. For example, the second person's gaze model might support the new focal point or a speech model may contain the words, "Look there!" This sentence contains an imperative and a direction and confirms the desire to move the focus to an indicated point.

In the same manner, utterances, writing and other representations each have semantic meaning that can be extracted from their respective syntactic models. (One of our future objectives is to aggregate as many of these data streams as possible in order to build a richer model of workspace processes to increase our confidence that the infrastructure correctly "understands" the state of workspace activities. This will provide the best chance of correctly inferring future actions or processes from observing past and current processes. It may also provide an opportunity to intervene in the activity and provide a useful and appropriate tool to the group members.) One of our challenges in this work has been to model the cognitive dust in a suitable schema and store it in data repository.

ThIS (schema) and RES (database)

In this section we describe our Theory Informed Schema (ThIS), which is a template for informing the design of the data store, along with our Real-time Ethnographic Store (RES), which is the database used to persist cognitive dust collected from the workspace. This section describes the concepts and relationships depicted in Figure 1.

ThIS

We are trying to understand, articulate and model human behaviour in our workspace in order to build a comprehensive framework for a computer infrastructure. While some low level data can captured with cameras, microphones and other devices and abstracted to model some of the activities within group interaction, the models usually lack meaning within the larger picture. For example, a person opens an application on a public screen. There is no information that might explain why the application was opened, what it might be used for and what might happen next. We believe that the only approach that might fill in these comprehension gaps is to look at existing work that attempts to provide explanations. Existing theories are the only source of illumination for this problem and there is an enormous body of instructive research in domains such as Social Psychology and Sociology. In these domains we have looked for and found theories that explain small components of the complex behaviours that takes place in human interaction.

Each theory can do any or all of the following functions. It may describe its concepts. It may elaborate in the form of rhetoric. It may predict or reason about its concepts or it may have application within its reference domain (Halverson 2002). Each of the theories in this work articulates some of these functions (but rarely all of them) although description and rhetoric are usually produced at a level of abstraction that is too imprecise to be of value to us. In some cases the theories also take a positivist position. These ones are most useful in allowing us to understand what might be happening in an activity while at the same time providing a clue about what might happen in the future.

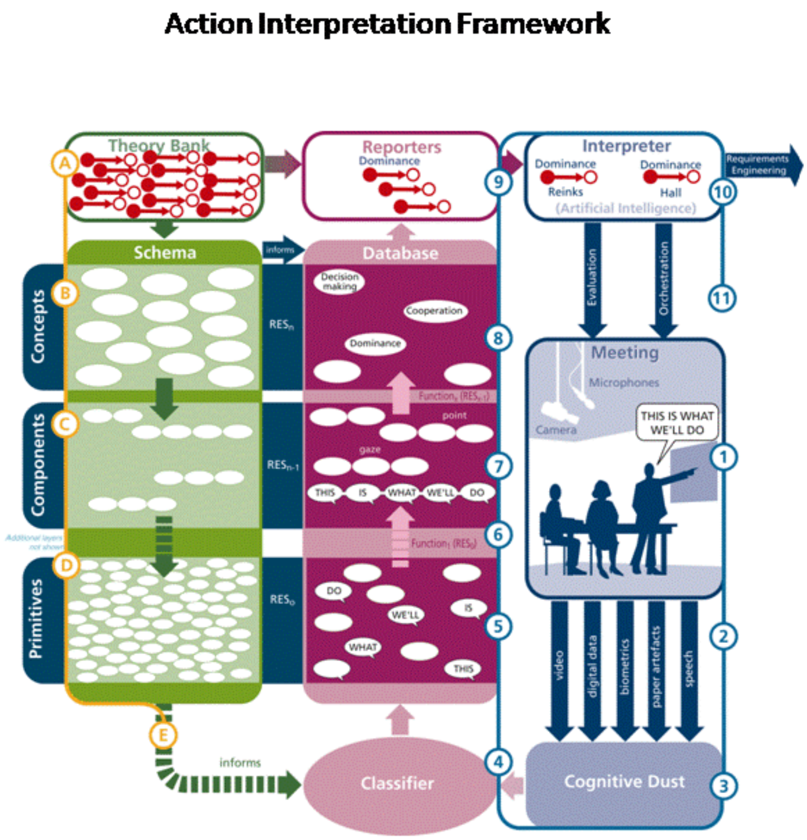

Figure 1: ThIS (Schema) and RES (Database)

Our collection of theories is stored in the theory bank (A in Figure 1) and it is the *concepts* in these theories that inform the design of our Theory Informed Schema (ThIS). In other words, ThIS is a repository of all of the concepts that can be identified in the theories of small group behaviour. It is constructed in a top down fashion where the concepts represent the highest layer (B in the figure). Examples of concepts are cooperation, competition, idea generation, dominance and (social) group support.

The concepts are decomposed into their component parts (C in the figure) which could abstract data structures. For example, concepts such as dominance, status or power can be manifested and described by a number of ways. Theories tell us that people who speak more (Rienks & Heylen 2006), smile less (Hall, Coats & LeBeau 2005) or try and control a larger volume of physical space (Tiedens & Fragale 2003) may be exhibiting signs of dominance in the group. Components in these concepts would be respectively: a count of each speaker's words; a count of each person's smiles and an estimate of the volume each person tries to command within their personal space (particularly those people that try and increase their space). The components in these simple examples are also our primitives in the schema (D in the figure) so no further decomposition is required. The set of primitives (E in the figure) is available for the classifier to refer to when encountering the cognitive dust (Blackburn, Swatman & Vernik 2007) collected in the workspace.

A more complex example provides counterpoint to the previous one. Idea generation is a concept that could also be a component in a more complex concept such as *planning*. On its own, it can be broken into components such as verbalising, writing, typing or arguing (amongst others). Verbalising can consist of sentences, questions, answers and suggestions. Each of these components (or concepts) can be further broken down to words; our primitives in this case. Typing on a computer can be deconstructed into keyboard events and mouse events, which is another example of the cognitive dust (3 in the figure) that can be observed in the workspace (1 in the figure).

The dust components at this lowest level have been defined indirectly or inferred by the theories through the different levels of the schema and these categories are used to classify the dust. The classifier's job is to identify, classify and persist the cognitive dust into RES in its various native formats (4 in the figure).

RES

RES is a multi-layered repository for cognitive dust in its unmodified (5 in the figure) and abstracted forms (7 in the figure). At the lowest level, RES_0, the data formats include: WAV or MPEG files for captured video, WMA or MP3 for audio, textual representations of computer events, transcripts of speech, paper documents, electronic files in various formats and biometric data in native formats (2 in the figure). Functions are employed to abstract this raw dust into representations with higher level semantic meanings (6 in the figure). Words, for example, can be aggregated and processed to form sentences. (RES is currently implemented in the relational format.) The sentences are then stored at a higher level in RES_{0+m} where m may or may not equal 1 as levels in RES can be bypassed. For example, a sentence can represent a concept directly and this might need immediate interpretation or it could be part of a conversation which could be part of a brainstorming session; another concept. These aggregations of lower level representations are built up through an iterative process and using different functions with results stored at successively higher layers in RES.

Ultimately, at layer RES_n, the cognitive dust has been processed and abstracted into the coarsest grains possible and this level represents concepts (8 in the figure), many of which can be found in the theory bank. Interestingly, RES can contain concepts that do not exist in the theories or have not been identified yet. For example, we are unable to a find useful theory that explains the meaning of hand actions that people use when they are talking. (We are hopeful that new theories can be abduced from our future empirical work).

It is at this level - a *collection of concepts* - that people consciously comprehend their world and reason about people's activities. They can tell who has been accepted as the leader in the group, who has superior work skills, who feels comfortable taking minutes and who sometimes likes to act like a clown. People observe instances of concepts such as competition, cooperation, conflict and harmony. At a semi conscious level they may infer succeeding actions from a series of current and previous states. This is based on learning and experience. For example, previous observations of Bob when he has been angry lead Mary to infer that Bob's next action is to pick up an object and throw it somewhere. However, when Mary observes Philip displaying anger, she may reason that Philip will calm down quickly and apologise for his behaviour. The paired concepts of *anger and violence* and *anger and contrition* are each causal and show that Mary has been theorising at some level in a positivist style. From Bob's anger she has inferred his violence and from Philip's anger she has inferred his contrition. Mary has, subconsciously or otherwise, reasoned about the concepts and built two production rules. She may not have created the productions on her own, of course, as she may have adopted them from elsewhere. These productions, if they could be identified in theories, could help to infer future actions from current and previous actions or states.

The concepts that Mary observes need to exist in her personal version of RES, which is, of course, her memory. In the same manner these concepts also need to exist in our version of RES. The concept of anger could be represented in a number of ways. It might be manifested as: violent physical actions which have no associated sounds; a series of words spoken with enough volume and *angry* content; or a lack of action but instead a fixed gaze along with a set of face contortions. Each of these sub concepts is represented at a lower level in RES by smaller components. To take the easiest example from above, the verbal violence would contain sentences or words that would normally be associated with conflict. These could include contrary points of view, fixed points of view or dominating phrases. The change in volume would be easy to detect.

If we can identify these concepts from the data in RES and if we can find theories that link the concepts causally, it may be possible understand what is happening in a small group. This understanding is represented in our model as the best candidate in a series of *stories*. A story is a description of an activity or process that has been derived from a theory. It is told from one perspective and many theories can each tell a slightly different story of the same data instance from its own perspective. This is because each theory has usually been generated by different data, in a different domain and with different subjects. It follows, then, that the theories and, consequently, the stories may be contradictory and they may be incomplete. The latter may come about due to a lack of evidence in the data to populate the story sufficiently.

Each story or even a set of stories can be generated from a suitable theory and it is the theories that provide the basis for the process of interpretation. Separate candidate theories that focus on any arbitrary concept might help to explain the meaning of that concept or a set of linked concepts; evidence of which exists in the data. These candidate theories are selected from the theory bank by a single reporter (9 in the figure) which is responsible for monitoring the development of instances of that concept in the data. The reporter either interrogates RES_n or data is pushed to it as soon as it appears in the database. The reporter is responsible for some pre-processing of candidate theories. It excludes those that are clearly irrelevant and sends those candidates with enough similarity to the interpreter for further processing (10 in the figure).

Over time, and with different reporters populating the interpreter with candidate theories to suggest explanations of the group activities, *pictures* start to emerge. The pictures do not need to be high fidelity. People are sometimes able to function with very low fidelity impressions of group actions. This is particularly true when people have sufficient domain knowledge and / or experience to make accurate inferences.

Deciding on the meaning of the pictures is the job of the interpreter. It will form its own view of the world within the group workspace. It may be an incorrect view or an incomplete view but over time and with a history of observations to consult, it may be more observant and more accurate that any individual in the group. It will have more eyes, more ears and more sensors to collect data that will, ultimately, generate stories.

The interpreter represents the last line of processing in the data chain. This is the point where if it has a high enough level of confidence, the story is compelling, and it has something useful to do, it may make decisions. The decision might be to: make an entry into a record, for requirements engineering purposes; initiate an intervention in the meeting, in the case of orchestration; or to declare some type of confirmatory action, in the case of evaluation (11 in the figure).

Even at this point, the multi-perspective story is likely to be ambiguous, imprecise and contradictory. This is an acceptable position, though, given that these characteristics sometimes also describe human perspectives on group actions.

Motivating Example

Much of our research focuses on teams of people but in this example we examine the actions of a dyad purely to illustrate a theoretical perspective. It is not difficult to extend the example to a group setting but this would add complexity without additional insight. The example focuses on some of the dominant behaviours or concepts that can be observed in group interaction.

The scenario opens with two people sitting at a table in a small meeting room. Betty and Tom work in the same department of a large organisation and Betty is Tom's boss. Tom has been working hard over the last six months. He has been putting in a lot of overtime and feels that he is overdue for a salary review. He wants to ask Betty to give him a raise but he is a submissive character and does not feel comfortable asking for the extra money. Tom smiles and opens the conversation but doesn't say much that is meaningful. He tries to open the conversation with a comment about the football last week. Before he can change the subject and establish what he wants to say, Betty interrupts with a laugh, nods her head and makes a comment of her own. She doesn't agree with Tom about the football and says so. She is wondering to herself why she is meeting with Tom when she has so much outstanding work of her own. Tom is immediately uncomfortable and tries to think of a suitable response to Betty's comment. Tom was hoping the conversation would start on a note of agreement so that he could move straight into his request for a higher salary. Betty, who wants to finish the meeting as soon as possible, can see that Tom is a bit stuck for words. She pats him on the arm, gets straight to the point and asks why he wanted the meeting. Tom is on the spot now. He knows that Betty doesn't have a lot of patience and has a bit of a reputation for being quick tempered. He has to do it now. He has to outline his case for the pay rise. Tom looks down and starts to talk in a soft voice about how hard he has been working. He is starting to sweat a little. Betty is silent and waits for Tom to make his point. She stares at him with a hint of annoyance on her face. Tom is really making a mess of this. The situation is getting a bit out of control. He is not saying what he had practiced. In fact, his words are not making much sense to Betty. Before Tom can finish his sentence, Betty gets the idea that Tom wants to ask for a raise. She starts speaking over Tom who immediately shuts up. She speaks in a calm voice and finishes what Tom wanted to say for him. Tom is glad the issue has been raised. He smiles but says nothing and lets Betty do all of the talking. The problem is that Betty is under pressure from her boss to keep costs down and so has her own ideas about giving her people pay rises. She is under a bit of pressure now and starts to get a little more animated in her speech. She starts to use more hand actions and gestures as she tells Tom that now is not a good time to discuss a pay rise.

There are many ways of analysing the scenario but we will stay with the theme of dominance. All of the examples are taken from the work that was introduced earlier (Hall, Coats & LeBeau 2005; Rienks & Heylen 2006). It's quite clear who is the boss and Betty's behaviour reinforces her position. At the same time, Tom doesn't help his own case by being fairly submissive. The first observation is that Tom is smiling. This has a meaning along the lines of "I won't attack you" and when two people smile at each other they are establishing some common understanding that they want harmony rather than conflict. (Smiling is not related to laughing which has its roots in aspects of play.)

The next observations are all behaviours that can be dominating. Nodding can mean that the nodder already knows what is being said (or pretends to know) while laughing implies that the *laugher* is relaxed. Fast talking means that more words are produced relative to a slower talker. Tom's gaze is wandering which means that he is not looking directly at Betty. This may or may not indicate that Betty can out stare Tom, who may be the first to look away in such an activity. It may mean that Tom does not want to face an issue or it may mean that he is unable to respond to Betty.

Betty is indicating her dominant status by touching Tom on the arm. She is initiating the action rather than responding to a touch from Tom, who is still floundering in the rest of the observations and is unable to maintain a position of equality with Betty. Evidence of some level of equality, for example, might include speaking an equal number of words, gazing directly at Betty, not allowing interruptions to his speech or creating a physically large presence, such as sitting up straight or using expansive hand gestures when talking (Tiedens & Fragale 2003).

The theories of dominance used in this example are part of a larger set that covers many different aspects of human behaviour in small group interaction. Some of the theories may contradict each other but does that make one right and another wrong? Not necessarily. Any arbitrary theoretical perspective can seem to be usefully descriptive or accurately predict a change in dependant variables in one situation but be dysfunctional in another.

Discussion

A number of components in this model currently remain undefined. A substantial amount of work will be required to identify and define all of the different functions used to transform *concrete* data from one level of RES to a higher level. For example, the function used to aggregate words into a sentence (low level primitives to a higher level data structure) is quite different from that which is used to define a *head nodding*. This example takes a series of different specific head positions and aggregates them into the nodding sequence. For other examples, a series of hand and arm gestures may resolve into someone pointing their finger at an object and a series of mouse events representing a "drag and drop" action on a computer file uses another function again although, since we have found no theory to represent the latter, we have had to make one as we go.

Really useful theories make well defined claims, as in the positivist manner, about some aspect of group interaction. For a simple example, if person 1 is consistently speaking more than person 2, then person 1 may be exhibiting dominant behaviour relative to person 2 (Hall, Coats & LeBeau 2005; Rienks & Heylen 2006). (Dominance is used as a theme in this paper to promote some continuity and to make it more readable.) In the example, theories in this format allow us to expose production rules that can be used to generate stories in much the same manner as humans generate stories to explain conditional situations. In this case the reporter responsible for dominant speaking is looking for word counts in the data to populate a production rule such as:

If (person 1, word count) > (person 2, word count) then person 1 ==> dominant

This is a simple example of a production that exists among many in either of the two dominance theories mentioned earlier. At any point in time, many productions may be being populated and stories generated. Some of these stories are gong to be wrong and some are going to be contradictory. This creates a problem because we need to provide an interpretation mechanism that is as accurate as we can make it but we are required to work with data that may be incomplete and data that may be complete but inconclusive. This provides a whole new set of problems, which are beyond the scope of this paper.

It would be reasonable to ask why should we go to the effort of developing theory when there is already a research domain for observing meetings and analysing (with limited success) their audio and video records

(Nijholt, Akker & Heylen 2004). The reason is that these projects are accumulating huge repositories of data from which little *meaning* can be extracted. One of our goals is to distil human comprehension from our observations and identify the important workspace activities that will inform our support efforts. For example, if two people are both talking and pointing in an activity, we will have stories which suggest who has control of the floor and where we should direct support.

There are several problems that need to be surmounted before we can consider using the stories in a productive and supportive role. One of these issues is the inherent problems or contradictions within theories. To use Rienks et al and their simple dominance theory as an example, let us consider an alternative group of people to Betty and Tom. Let's refer to Table 1 and examine how Rufus, Bert, Lex and Maya are developing their dominant behaviours in a hypothetical meeting.

Table 1: A simple set of counts for dominance features

	Rufus	Bert	Lex	Maya
Word count	3300	5500	1200	3500
Speaking turns	18	8	35	16
Floor grabs	13	6	30	19
Less times interrupted	5	2	12	5
Results	???	???	???	???

How do we extract meaning from the numbers in Table 1? How do we attribute relative total values to each of the dominance characteristics? Rienks et al. used a set of eight meetings to generate their dominance features and they also tried to optimise their approach by minimising the number of features without losing their descriptive ability. Their most informative three features were: number of speaking turns, number of floor grabs and the number of successful interruptions. However, meetings have many different styles and they may range from very formal to very informal. In Table 1, Bert spoke a lot but only had eight speaking turns. Was Bert the secretary reading minutes from the previous meeting? Lex had the most floor grabs and speaking turns but didn't say much. Was Lex a heckler and was he really the most dominant person? At best, the theories are a guide to understanding human behaviour but they are the only guide that we have.

Conclusions and Future Work

The main contribution of this paper is a framework that enables researchers to identify, understand and articulate in a fine grained manner the actions that people perform in small group meetings. This framework is modelled on known theories of human behaviour and is constructed from cognitive dust, which is the observable evidence of human communication within work tasks. These tasks are usually carried out in a different manner each time they are executed and the framework allows a computer infrastructure to remain synchronised with the processes as they emerge from the task. In this manner, and this is one of the ultimate goals of this research, useful support or tools may be deployed when they are required and in the manner that people want to use them.

The scope of this particular work extends to the activities of the reporters, which match concepts emerging from within the workspace activities with existing theories from the theory bank although some commentary has also been made on the function of the interpreter, which draws conclusions about the workspace activities and decides whether or not to provide support. Future work will explore inferencing approaches within the interpretation mechanism as well as rule mining (Park & Lee 2007) within the theory bank.

References

Abbot, KR & Sarin, SK 1994, 'Experiences with Workflow Management: Issues for The Next Generation', paper presented to Computer-Supported Cooperative Work.

Blackburn, T, Swatman, P & Vernik, R 2007, 'Cognitive Dust: A Framework that Builds from CSCW Concepts to Provide Situated Support for Small Group Work', *Lecture Notes in Computer Science*, vol. 4402.

Bogia, DP & Kaplan, SM 1995, 'Flexibility and Control for Dynamic Workflows in the wOrlds Environment', paper presented to Conference on Organizational Computing Systems, Milpitas, California.

Bogia, L, Tolone, W, Bignoli, C & Kaplan, S 1993, 'Issues in the Design of Collaborative Systems: Lessons from ConversationBuilder', paper presented to Schaerding International; Workshop on Task Analysis and CSCW, Schaerding, Austria.

Carlsen, S & Gjersvik, R 1997, 'Organizational Metaphors as Lenses for Analyzing Workflow Technology', paper presented to GROUP.

Crabtree, A, Rodden, T & Benford, S 2005, 'Moving with the Times: IT Research and the Boundaries of CSCW', *Computer Supported Cooperative Work (CSCW)*, vol. 14, no. 3, pp. 217-51.

Eason, K 1996, 'Division of labour and the design of systems for computer support for cooperative work', *Journal of Information Technology*, vol. 11, no. 1, pp. 39-50.

Gray, P 2007, *Reflections of Decision Support*, DSSResources.COM, viewed May 2007 <http://dssresources.com/reflections/gray/gray04072007.html>.

Grudin, J 1988, 'Why CSCW applications fail: problems in the design and evaluation of organization of organizational interfaces', paper presented to Computer Supported Cooperative Work, Portland, Oregon.

Hall, J, Coats, E & LeBeau, L 2005, 'Nonverbal Behavior and the Vertical Dimension of Social Relations: A Meta-Analysis', *Psychological Bulletin*, vol. 131, no. 6, pp. 898-924.

Halverson, C 2002, 'Activity theory and distributed cognition: Or what does CSCW need to DO with theories?' *Computer Supported Cooperative Work*, vol. 11, no. 1-2, pp. 243-67.

Jorgensen, H 2004, 'Interactive Process Models', Norwegian University of Science and Technology.

Jorgensen, H & Carlsen, S 1999, 'Emergent Workflow: Integrated Planning and Performance of Process Instances', paper presented to Workflow Management, Munster, Germany, 1999.

Kamel, N & Davison, R 1998, 'Applying CSCW technology to overcome traditional barriers in group interactions', *Information & Management*, vol. 34, no. 4, pp. 209-19.

Kuhn, TS 1962, *The Structure of Scientific Revolutions*, University of Chicago Press.

Luff, P, Hindmarsh, J & Heath, C (eds) 2000, *Workplace Studies: Recovering Work Practice and Informing System Design*, Cambridge University Press.

Maaranen, P & Lyytinen, K 1995, 'Designing Meeting Support Systems in a User-Centered Manner- The Case of the Helsinki Prototype System', paper presented to COOCS, Milpitas, California.

Ngwenyama, O & Lyytinen, K 1997, 'Groupware Environments as Action Constitutive Resources: A Social Action Framework for Analyzing Groupware Technologies', *Computer Supported Cooperative Work*, vol. 6, no. 1, pp. 71-93.

Nijholt, A, Akker, R & Heylen, D 2004, 'Meetings and Meeting Modeling in Smart Surroundings', paper presented to Social Intelligence Design: Third CTIT Workshop, Enshede, Netherlands.

Park, S & Lee, J 2007, 'Rule identification using ontology while acquiring rules from Web pages', *International Journal of Human-Computer Studies*, vol. 65, no. 7, pp. 659-73.

Rienks, R & Heylen, D 2006, 'Dominance Detection in Meetings Using Easily Obtainable Features', *Lecture Notes in Computer Science*, vol. 3869, pp. 76-86.

Robinson, M 1990, 'Computer Supported Cooperative Work & Informatics for Development', paper presented to Informatica, Havana, Cuba.

Suchman, L 1987, *Plans and Situated Actions*, Cambridge University Press.

Swatman, P, Blackburn, T, Nguyen, V, Raisey, D & Vernik, R 2005, 'Supporting Creative Teams Through ICT', paper presented to Australian Workshop on Requirements Engineering, Melbourne, Aus.

Tiedens, L & Fragale, A 2003, 'Power Moves: Complementarity in Dominant and Submissive Nonverbal Behavior', *Journal of Personality and Social Psychology*, vol. 84, no. 3, pp. 558-68.

Vogel, D 2006, Personal Communication.

Copyright