Computationally assessing conceptual coherence in teamwork

A. Dong[a] and A. Vande Moere[a]

[a] *Key Centre of Design Computing and Cognition, The University of Sydney, Sydney NSW 2006 Australia*

Abstract

Research and commercial development in collaboration and project management have established tools to support experiential cognition, giving teams the tools to deal with deadlines, deliverables, and gathering information to maintain ‘group awareness.’ Supporting the ability to plan, monitor progress toward, and reflect on the quality of teamwork and associated work processes has normally been the role of team managers. The viewpoint explored in this paper is that teams need to develop skills to render explicit the cognitive and social processes needed for collaboration and that software can play a central role in developing such skills. A computational linguistic and information visualization toolkit for rendering explicit social mechanisms in team collaboration is presented. Using the toolkit, analyses of a design teams’ conversations in relation to exposing these social mechanisms are described.

Keywords: teamwork, communication, information visualization

1. A role for software in evaluating teamwork

In this paper, we will study teamwork in the context of design. Because design is a social activity characterized by information exchange, compromise and negotiation, the development of a shared, organized understanding and mental representation of knowledge about key elements of the designed work is critical for a successful outcome. Social cognition in design concerns with the generation, transmission and evaluation of information and knowledge to create collective sense-making of the function, behaviour, structure and meaning of a product. Collective sense-making is regarded as a central aspect to facilitate the coordinated action that is required for successful team-based design.

For the purposes of this paper, teamwork is regarded as the activities and interactions of people actively sharing knowledge, perceptions, and ideas when working together toward a common purpose. Our view is that teams need to develop conscious, explicit understandings of the team’s progress towards collective sense-making and commonly held knowledge, perceptions, and ideas in order to modify and improve the teamwork processes that affect them. Improving teamwork requires both the skills to recognize and diagnose teamwork problems as well as behavioural changes necessary to rectify them.

The aim is to establish alternative methods for teams to assess the nature of their collaboration through characterizations of their communication. Developing teamwork assessment skills can enable teams to engage in reflections about the nature, purpose, and utility of their collaboration processes and to thereby come to understand them better, use them more effectively, and improve them. Software can play a role in developing the skills necessary for a team to engage in the reflections about their collaboration processes and behaviour change in order to improve the likelihood of successful teamwork.

One way is that the software could give strategic advice on the nature of collaboration. But, our view is that this may be obtrusive and unlikely to persuade team members to adapt their behaviour to improve teamwork. Rather, we believe that the principles of self-monitoring and surveillance from the field of computers as persuasive technology [1] offer other ways of conceiving how software could enable team members to learn about their teamwork behaviour and to learn about the teamwork behaviour of others. The principle of self-monitoring asserts that people can change their behaviour if a tool allows them to monitor their behaviour in order to modify their behaviour towards a predetermined goal, i.e., effective teamwork. There are several studies [2] which indicate that the analysis of one's own behaviour is a prerequisite for modifying inadequate social processes. The principle of surveillance is based on the finding that people are likely to change their behaviour to meet others' expectations when they know that they are being monitored [3]. The software system presented in this paper renders explicit one dimension of teamwork, conceptual coherence, by depicting the effect of team members on conceptual coherence.

2. Computationally assessing conceptual coherence

This paper presents conceptual coherence as a computationally derived measurement of shared knowledge, one of the social dimensions of teamwork. In lay terms, conceptual coherence is a measurement of the degree of inter-relatedness of the team's ideas, or what is called 'being on the same page.' Conceptual coherence affects teamwork cohesiveness, normally defined as "the resultant forces which are acting on the members to stay in a group" [4, p. 274]. Cohesiveness has been shown to be highly significant in the performance of a group [5]. Conceptual coherence is one way to quantify one of the 'forces' that may drive team towards cohesiveness because a team which exhibits conceptual coherence is likely to be cohesive, but a team which is not cohesive is not likely to exhibit conceptual coherence. Teams which exhibit similar patterns of constructive thinking are likely to be cohesive [6]. The effect of shared knowledge is a prerequisite for effective collaborative work [7]. Other social accounting metrics of teamwork include leadership and esteem. We start with the concept of conceptual coherence because similarities between ideas may indicate shared knowledge in the sense of pragmatic intersubjectivity, a crucial precursor for effective teamwork.

The software toolkit operates by text mining the communication of the team participants in order to assess conceptual coherence. There may exist a danger in assuming that group communication and group shared knowledge could be understood as public versions of private thought. Nonetheless, there is value in examining the way in which communication about designing unfolds in the documentation within the context of the activity of designing. In particular, aligning the dénouement of the design documentation with other measurable outcomes such as the quality of the design or with protocol studies offers a means to study the ways in with group shared knowledge unfolds in a socially mediated setting.

The calculation of conceptual coherence is based on the semantic coherence of the team's language-based communication. Communication is often defined by the creation of shared understanding through interaction among people. The purpose of communication in a team is to establish a set of coherent ideas from which team members can develop shared knowledge. From there, the team is able to articulate clearly to one another within the team, as well as to those outside the team, its goals, purpose, design process and product. Borrowing the definition of mental models as "the mechanisms whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions (or expectations) of future system states." [8, p. 351] conceptual coherence has been proposed as a measurement and indicator of shared knowledge [9].

Computational methods for identifying conceptual coherence (cohesiveness) in groups are not "new." Most methods for identifying cohesiveness in groups use statistical models of sociometric data to identify blocks of stochastically similar actors, calculate cohesion based on statistical measures of communication frequencies, or measure an individual's psychological feelings of belonging and morale as a result of belonging to the group through questionnaires. What we bring is an approach that imputes team cohesiveness from the *content* of language-based communication among team members using computational techniques which do not rely on explicit keyword matching or counting interactive behaviour between team members. In order to diagnose teamwork issues *in situ*, the method does not rely on *post hoc* surveys.

This article will describe two techniques for calculating semantic coherence, one based on direct calculation using latent semantic analysis and one based on indirect interpretation using the visual processing of humans. In these following sections, we present latent semantic analysis (LSA) and information visualization using the swarm intelligence metaphor of flocking to calculate and visualize the formation of conceptual coherence in design teams. A description of how collective sense-making was calculated for a team is presented as a demonstration of the toolkit AgoraProbe (from the Greek word *agora* for an assembly of people).

3. Analysis methods

3.1 Latent semantic analysis

Latent semantic analysis is a method which maps out the coherence of meanings of a series of words and how language use models knowledge acquisition and representation [10]. Latent semantic analysis (LSA) is a text analysis method that characterizes the semantic similarity between texts using a high-dimensional semantic space. The mathematical foundation for LSA lies in singular value decomposition (SVD), a matrix approximation method for reducing the dimensions of a matrix to the most significant vectors. By looking at the entire range of words chosen in a wide variety of texts, patterns emerge in terms of word choice as well as word and document meaning. The principal advantage of LSA over other standard document analysis techniques, such as keyword analysis, is that LSA examines context and removes obfuscation created by "noise" in the documentation.

To apply LSA to the analysis of a team's conceptual coherence, it is assumed that the psychological similarity between each designer's own mental representation and the socially held representation of the designed artefact is reflected in the semantic coherence between words in the way they co-occur in dialog and other language-based communication. Semantic coherence between two communicative acts $\mathbf{d}_p$ and $\mathbf{d}_q$ is calculated using the standard cosine similarity measurement for textual coherence. An abstract representation of the team's socially held group knowledge is calculated as the centroid of the 'latent' vectors representing the team's communication in the k-reduced LSA space. The complete details of the calculation of conceptual coherence using LSA are provided in [9].

Because conceptual coherence is dynamic and not reducible to a single value that could represent the state of group knowledge over a long period of time, the conceptual coherence is calculated over time. The semantic coherence between each team participant's communication and the group's communication is calculated. Each participant's semantic coherence should ideally follow the same trajectory as the group's and be located proximally close to the group's communication. The group's semantic coherence progresses from 0 to 1 at the end of the design project or at the end of a design group meeting. The slope or rate of progression towards 1 depicts how quickly the team's ideas (as reflected in their communication) are cohering. The limits of the semantic coherence are 0 and 1. At the extremes, a constant conceptual coherence of 1 would indicate that no new information is being added – clearly undesirable repetitions of the same ideas. At the other limit, a conceptual coherence of 0 would indicate dialectic semantics. In the aggregate, however, it is argued that the team should exhibit broad semantic agreement. Prior studies have theorized and empirically shown a direct correlation between such semantic agreement and quality design outcomes [11].

3.2 Information visualisation

In addition to latent semantic analysis, methods from information visualization, the representation of data graphically rather than textually using the high bandwidth human perceptual and cognitive capabilities, are applied to detect patterns and draw inferences from visual form. Specifically, we use the metaphor of flocking [12]. We create a mathematical simulation of flocking birds or swarming fish – that is, animals which tend to move together. Because design ideas and concepts are introduced and changed dynamically during the life cycle of the design process, we believe that self-organisation by flocking is a more apt metaphor than classical concept clustering algorithms which mostly operate on a priori non-time-varying datasets. We use humans' in-built cognitive capabilities of the visual system to directly relate simple motion typologies to complex behavioural reasoning – whether teams are coming to conceptual coherence or not. For example, when we see two people moving toward one another, we might assume that the two are coming to an agreement or coming together to discuss and work; when the two separate, we might assume that the two are moving in different directions or have completed the shared task and are moving onto new tasks which may require collaboration with others. We exploit this inclination toward these assessments to graphically render the formation of conceptual coherence in the design teams.

Thus, the goal here is *not to pre-calculate* conceptual coherence but rather for the viewer to interpret and infer team cohesive by visualizing the motions of boid particles representing each member in a team. Each boid represents one unique team member and is placed in a three-dimensional virtual space. Each boid is capable of perceiving other boids in its close vicinity. In fact, the direction and speed of a boid A with position $\vec{p}_A$, is dependent on all the boids X with position $\vec{p}_X$ in its neighbourhood. A boid will attempt to move towards the centre of the flock as the boid perceives it. For all boids X in its neighbourhood, if the distance between $\vec{p}_X$ and $\vec{p}_A$ is smaller than the flock centring limit, boid A should try to direct itself towards the perceived centre of gravity of all boids X combined.

Once the boids have flocked, they will tend to stay together if the data values remain similar. A boid will attempt to stay close to boids that are in its neighbourhood and are interconnected in the context of the data values that it represents. Data similarity, in this case the connectedness between design collaborators, is determined by calculating the difference between the data values that the boids represent. Consequently, the strength of the attracting force is proportional to the distance between the boids and the connectedness between design participants.

The position $\vec{p}_A$ of a boid in the three-dimensional space is calculated based on the information content of each speaker's utterance using Shannon's formulation of information entropy [13]. The information entropy is intended to encode the similarity of the information sources (the team members) that produced the information (utterances). Similar information sources (team members) should produce similar information. Suppose we have a set of possible design concepts, ideas and issues, what we propose to call the *design space*. Further, suppose that the design space is expressed semantically as lexicalized concepts in each communicative act. We can calculate the probability of occurrence of a lexicalized concept p_i as proportional to the number of occurrences of a lexicalized concept in a communicative act i and inversely proportional to the number of times a lexicalized concept appears over all recorded communicative acts. We can then quantify the amount of information contained in a message using Shannon's measure of information entropy.

4. AgoraProbe example analyses

Analyses of one design team is presented here to show the capabilities of AgoraProbe. The transcripts come from the Bamberg Study [14] in which the teams designed a planetary gear train set. The full details on data collection and the rationale behind the use of students in the study are described by Stempfle and Badke-Schaub. Native German speakers with mechanical engineering backgrounds translated the Bamberg Study transcripts into English for analysis by the computational methods. These transcripts were chosen to enable qualitative comparisons between the results of our computational analysis methods with previously published studies of these design teams [14]. Such comparisons between computational approaches and methods such as protocol analysis can support the validity of the rationale behind the computational approaches describe above.

The team described in this analysis is Team 1102 which consisted of 6 participants (A-F) who contributed 15%, 21%, 9%, 20%, 18% and 16% of the content-bearing utterances. It is known that this team exhibited a lack of strong conceptual coherence (J. Stempfle, personal communication, January 22, 2004).

First, we present the analysis using the LSA component of the AgoraProbe toolkit. In the analysis shown in Fig. 1, all participants in a team were analysed. It is evident from the figure that Participant E, indicated by the (green) boxes, is not contributing to the construction of shared knowledge by the team.

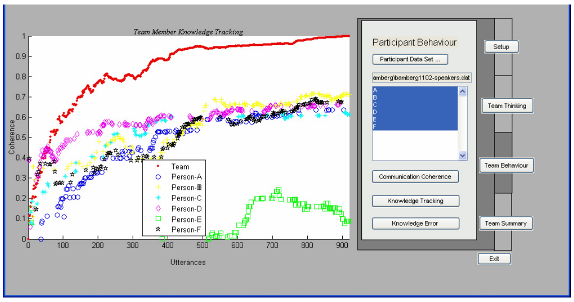

Fig. 1: Bamberg team with all participants analysed.

In other analyses [9], we have shown that Participant D is the most influential person in this group. Fig. 2 shows the effect of removing this person from the group. First, we can note that each participants' knowledge coherence in relation to the group's is more distant than before. This would imply that Participant D may have played a role in connecting the group's ideas. The participant may be what Sonnenwald [15] described as an "interdisciplinary star," a knowledge integration role.

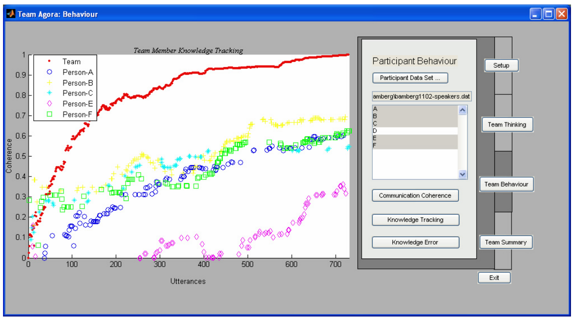

Fig. 2: Bamberg team with influential member D removed.

For the information visualization module, the AgoraProbe engine calculates the information entropy for each participant's linguistic input. Each participant is then represented as a boid in a 3D space, as shown in Fig. 3. Along with a text indicating the identity of the participant, the visualization shows the current value of the information entropy.

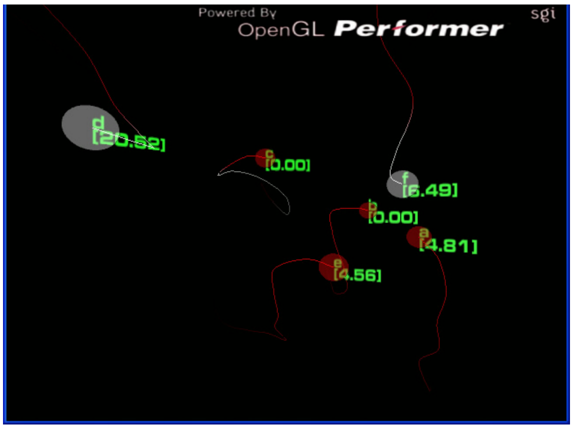

Fig. 3: Each team member is represented as a boid.

The overlap between the team members ideas (calculated as information entropy values) is visible when the sphere of influence is calculated and displayed as 'meshing' as shown in Fig. 4.

When the participants are apparently similar in the sense of the production of information (ideas), the boids seem to flock. In Fig. 5, the boids are in the process of flocking. Note that Participant D was originally 'far' away from the team but is now moving towards the flock centre since the average value of the flock is becoming closer to the information entropy values of Participant D.

When a few participants have differing ideas, then the boids will repel from the flock. In Fig. 6, Participants D and F are moving away from the flock to form a new flock whereas the trajectory of A, B and E are attempting to centre. Note that because Participant C has values which are not close to the values for the centre of any flock, it moves in a random direction away from the flocks.

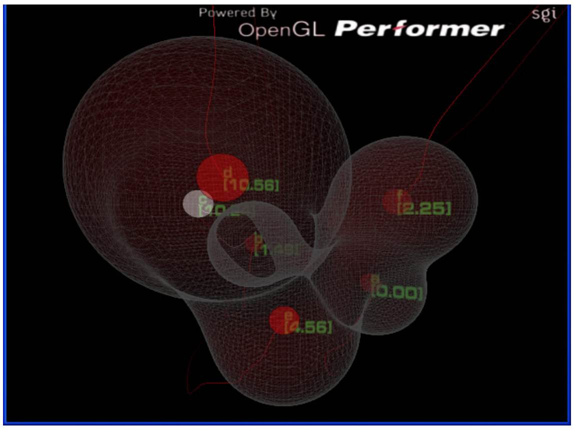

Fig. 4: Boids with mesh shown to give further visual data about spatial proximities.

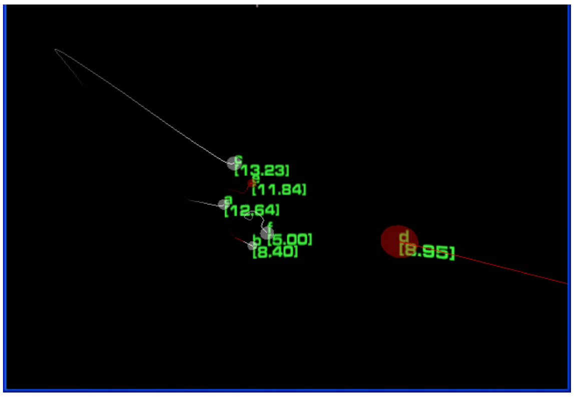

Fig. 5: The boids move towards each other and flock.

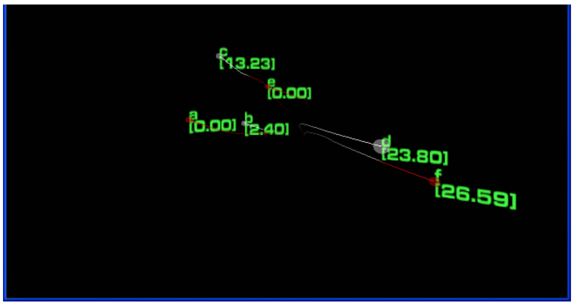

Fig. 6: The boids are moving away from the flock.

5. Discussion and future work

In demonstrating the AgoraProbe toolkit, this paper offers two main claims about the analysis of the social dimension of teamwork. The first claim is related to the computational analysis of the social dimensions of teamwork. Computational systems could provide depictions of teamwork with sufficient social lucidity to allow teams and team managers to troubleshoot problems. An added dimension of this work is that the analysis methods scale to deal with very large corpora and very large design teams. Neither is limited to keyword matching, and both operate on the content of communication rather than the flow of communication. Second, the toolkit claims, based on the principles of self-monitoring and surveillance from persuasive computing, that making it easy for team members to know how well they are performing to the target behaviour and knowing that others are observing how well they are performing to the target behaviour may induce behavioural change toward more effective teamwork. This claim is not yet verified but evidence from similar systems show promise.

However, one of the main assumptions of the work, that semantic coherence is an indicator of team cohesiveness and positive teamwork, is also its strongest inadequacy. This analysis ignored any influence that emotions or disagreements may have on the performance of the design team. In other words, a team may exhibit strong semantic coherence even though the team is actually in disagreement. That is, a conceptually coherent team may still not be cohesive. To latent semantic analysis and Shannon's information entropy model, the statements, "I want to design an ergonomic dog backpack." and "I think we should not design a dog backpack." are nearly equivalent in terms of semantic coherence and information entropy. Thus, the layer of understanding that is currently missing is an understanding of the sentiment expressed in the text alongside the semantics of the text.

Our future work is applying the computational linguistic technique of sentiment analysis to design team documentation to investigate how design teams appraise in design documentation and what effect those appraisals have on teamwork and the quality of the design outcome. Understanding appraisals and how these appraisals may communicate feelings of belonging or group conflict is an important direction in computationally assessing the social dynamics of teamwork using language-based communication.

Acknowledgments

This research was supported by an Australian Research Council grant DP0557346.

References

[1] Fogg, B.J., Persuasive technology: using computers to change what we think and do. 2003, Boston: Morgan Kaufmann Publishers. 283.
[2] Cannon-Bowers, J.A. and E. Salas, Reflections on shared cognition. Journal of Organisational Behavior, 2001. 22: p. 195-202.
[3] Turner, J.C., Social influence. 1991, Milton Keynes: Open University Press.
[4] Festinger, L., Informal social communication. Psychological Review, 1950. 57(5): p. 271–282.
[5] Mullen, B. and C. Copper, The relation between group cohesiveness and performance: An integration. 1994. 115(2): p. 210-227.
[6] Neck, C.P. and C.C. Manz, From groupthink to teamthink: Toward the creation of constructive thought patterns in self-managing work teams. Human Relations, 1994. 47(8): p. 929-952.
[7] Rouse, W.B. and J.A.S. Cannon-Bowers, Eduardo, The Role of Mental Models in Team Performance in Complex Systems. IEEE Transactions on Systems, Man and Cybernetics, 1992. 22(6): p. 1296-1308.
[8] Rouse, W.B. and N.M. Morris, On Looking Into the Black Box: Prospects and Limits in the Search for Mental Models. Psychological Bulletin, 1986. 100(3): p. 349-363.
[9]Dong, A., The latent semantic approach to studying design team communication. Design Studies, 2005. 26(5): p. 445-461.
[10]Landauer, T.K., Latent semantic analysis: A theory of the psychology of language and mind. Discourse Processes, 1999. 27(3): p. 303-310.
[11]Dong, A., A. Hill, and A.M. Agogino, A document analysis technique for characterizing design team performance. Journal of Mechanical Design, 2004. 126(3): p. 378-385.
[12]Proctor, G. and C. Winter, Information Flocking: Data Visualisation in Virtual Worlds Using Emergent Behaviours, in Virtual Worlds, J.-C. Heudin, Editor. 1998, Springer-Verlag GmbH: Berlin. p. 168-176.
[13]Shannon, C.E., A Mathematical Theory of Communication. The Bell System Technical Journal, 1948. 27: p. 379-423 and 623-656.
[14]Stempfle, J. and P. Badke-Schaub, Thinking in design teams - an analysis of team communication. Design Studies, 2003. 22: p. 473-496.
[15]Sonnenwald, D.H., Communication roles that support collaboration during the design process. Design Studies, 1996. 17(3): p. 277-301.