

Modelling Expertise of Temporary Design Teams

John S. Gero and Udo Kannengiesser
University of Sydney
Sydney NSW 2006, Australia
{john, udo}@arch.usyd.edu.au

Abstract:

This paper proposes a model of expertise of temporary design teams. We start by outlining a situated view of expertise as being grounded in individual experience, which emphasises its applicability and teleology. We use the construct of a computational agent to model the individual expert. We then show the advantages of the function-behaviour-structure (FBS) schema to represent the knowledge of a situated agent, which includes its knowledge about other agents. The ability of an agent to construct knowledge about the behaviour and parts of the knowledge of other agents has been found to be needed to allow social interaction among agents. Based on individual agents and their FBS views of one another, we then develop a model of expertise of temporary design teams that emerges from the interaction of individual experts. This bottom-up model is sensitive to changes in the composition of the team: some experts might leave the team while others join. However, the team members can usually establish the expertise of the new team without much effort as they can build upon their generalised knowledge gained from previously working in similar teams. We claim that our model provides a better understanding of team expertise and a basis for computational simulations of design teams.

Keywords: "situated expertise", "design teams", "team expertise"

1. Introduction

Complex design tasks often necessitate the confluence of various special knowledge and skills. Typically, this heterogeneous set of expertise is embodied in a multitude of designers who come together in unique combinations to form temporary, project-specific teams. The mere collection of different experts, however, is not sufficient to establish a good team. In order to achieve good performance as a team, the individual team members must be able to interact using the results of individual expertise in a way suitable for the current task. This ability can be enhanced using various technologies and infrastructures, however its major part is generally formed through continuous processes of learning each other's roles, responsibilities, priorities, practices, etc. The composition of a design team often changes with every new project, i.e. new members join the team while others leave. As a consequence, the expertise of the team is constantly reformed. In most

cases, however, experts do not have to start building the team expertise from scratch. They can use their experience gained from being involved in previous teams to quickly integrate themselves and other experts into the new team. Expertise of teams can thus be viewed as emerging from the ability of individual experts to interact with one another using their generalised and specialised knowledge acquired over a series of team interactions.

A significant amount of research has been undertaken to address various aspects of expertise (Feltovich et al. 1997). This paper develops a model of expertise of temporary design teams making use of Gero's (1990) function-behaviour-structure (FBS) schema. It takes a situated stance linking the knowledge of an individual team member/expert to their purposeful interactions with the environment as well as a social stance accounting for interactions among them to emerge team expertise. This provides a better understanding of the abstract notion of team expertise, which goes beyond models aimed at simulating or configuring teams (Levitt et al. 1994). In particular, this paper addresses the idea of tacit team expertise that persists and evolves in an environment where teams constantly break apart and form anew. We will develop a model of such team expertise using the construct of a computational agent to model the individual expert. We claim that this model provides a basis for simulating the expertise of temporary design teams.

2. Situated knowledge

The (non-situated) view of expertise as a body of specialised knowledge has been used extensively in traditional Artificial Intelligence (AI) research, which has produced computational reasoning systems known as expert systems (Giarratano and Riley 1998). These systems embody knowledge from human experts and encode them into the computer in various forms. While expert systems are powerful tools in static environments, they often fail when used under varying external constraints. Some researchers have therefore abandoned encoding knowledge representations inside the computational system or the agent and have directly linked its behaviours to the environment (Brooks 1991). Brooks' (1986) subsumption architecture organises a set of simple, locally acting agents into a hierarchy to emerge robust global behaviour that is sensitive to changes in the environment. Similar bottom-up approaches have characterised the research area of Artificial Life (ALife) (Langton 1989). However, as the agents base their behaviour merely on the current state of the local environment, they cannot deal with global, longer-term goals and requirements, which are essential features of the environment of designers. In addition, their robustness vanishes in environments other than those for which their behaviour has been pre-defined beforehand via fixed stimulus-response couplings.

Designing is a process that inherently changes the world in which it operates (Gero 1998). Consequently, the environment of a designer is dynamic not only in terms of parameters but also of structures. This recognition has driven a third approach of modelling knowledge, which we term the situated approach or situatedness. It can be seen as a synthesis of both the symbolic and the behavioural paradigms, as it uses internal representations, but grounds them in the agent's interactions with the environment. A situated agent adapts its behaviour to changes in the environment based on its current goals, its knowledge and its interpretation of the environment (Clancey 1997). This claim has been supported by empirical studies of human designers (Schön and Wiggins 1992; Suwa et al. 1999), which have characterised designing as an interaction of the designer with their environment. After performing actions to change the environment (e.g. by producing sketches of the design object), the designer observes and interprets the results of these actions and then decides on new actions to be executed on the environment. This means that the designer's concepts may change according to what they are "seeing", which itself is a function of what they have done.

Gero and Fujii (2000) have developed a framework for situated cognition in an agent, which describes the agent's interpretation of its environment as interconnected sensation, perception and conception processes. Each of them consists of two parallel processes that interact with each other: A push process (or data-driven process), where the production of an internal representation is driven ("pushed") by the environment, and a pull process (or expectation-driven process), where the interpretation is driven ("pulled") by some of the agent's current concepts. This has the effect that the interpreted environment is biased to match the current expectations.

The environment that is interpreted can be external or internal to the agent. The situated interpretation of the internal environment accounts for constructive memory, which is a notion whose foundations can be traced back to the work of Dewey (1896). Quoting Dewey, constructive memory holds that "sequences of acts are composed such that subsequent experiences categorize and hence give meaning to what was experienced before". The implication of this is that memory is not laid down and fixed at the time of the original sensate experience but is a function of what comes later as well. Memories can therefore be viewed as being constructed in response to a specific demand, based on the original experience as well as the situation pertaining at the time of the demand for this memory. Each memory, after it has been constructed, is added to the agent's knowledge and is now available to be used later, when new demands require the construction of further memories. These new memories can be viewed as new interpretations of the agent's augmented knowledge.

3. Modelling situated expertise

Any model of expertise that covers the adaptiveness and flexibility of expert practitioners (Schön 1983) has to account for the groundedness of knowledge. This entails modelling its embodiment in a situated agent that interacts with its environment in an autonomous, rational way (Smith and Gero 2000). The following two subsections develop such a model using Gero's (1990) FBS schema.

3.1 The FBS schema to represent grounded knowledge

The FBS schema originally introduced by Gero (1990) provides a formalism to represent design knowledge, i.e. knowledge about existing (physical) or to be designed (imaginary) objects. Function (F), behaviour (B) and structure (S) are defined as follows:

- *Function* (F) describes the teleology of an object.
- *Behaviour* (B) describes the attributes that are derived or expected to be derived from the structure (S) of an object.
- *Structure* (S) describes the components of an object and their relationships.

A design agent that uses the FBS schema to represent a particular object constructs connections between function, behaviour and structure through experience with the object. Specifically, the agent ascribes function to behaviour and derives behaviour from structure. A direct connection between function and structure, however, is not established.

The FBS schema is sufficiently general to cover all interpretations of objects and thus to support processes of grounding as required by the situated paradigm. This is particularly important for situated agents, as their views of the same object are different when dissimilar goals and dissimilar prior knowledge are used as interpretation biases. This is most obvious for interpretations of the teleology (F) of an object, since it is closely connected to the agent's current goals that are likely to change for one agent as well as differ for different agents. However, the interpretation of behaviour (B) and structure (S) of an object is also situated. Take the example of a mobile phone; the structure (S) relevant for the user (or the marketing person that takes the stance of the user) may include its dimensions and the number of scroll keys. These structural parts usually establish important behaviours (B) such as the case volume and the degrees of freedom to scroll through the menu. In contrast, the structure (S) viewed by an electrical engineer may include the electronic circuits in the phone. This more technical knowledge about the structure can be used to account for an additional range of behaviours (B) of the phone, such as the ringing volume, the Specific Absorption Rate (SAR) or the ability to browse the Wireless Application Protocol (WAP). Many of these behaviours here are relevant also for the

software engineer, however this kind of specialist usually works with the structure (S) of software elements. Another, less mechatronic example is a simple house as the design object: For an architect, the structure (S) of a house may be composed of a configuration of spaces, whereas a structural engineer may view the structure (S) of the same house as a configuration of walls and floors (spaces, as they are derived from this structure, would then be interpreted as behaviour (B)). However a particular agent interprets an object, the resulting internal representation can always be modelled using the FBS view.

Once a number of experiences with design objects has been gained and represented in the FBS form, the agent is able to generalise by clustering sets of like experiences. When an agent needs to access its knowledge about a particular object, it can derive a large part of this knowledge without much computational effort from its generalized experiences. This derived knowledge may even add (default) assumptions about an object where information gained from directly interacting with that object is missing.

The FBS schema provides a uniform set of constructs to model objects at all levels of generality and thus significantly supports this generalisation. Gero (1990) has introduced so-called design prototypes that represent generalised design knowledge based on the FBS schema, from which specific design objects can be instantiated. Design prototypes are useful for a design agent to start designing even if only incomplete information is available about the function, behaviour and structure of the object to be designed.

3.2 An FBS view of situated expertise

The FBS view is sufficiently general to comprehensively represent all kinds of objects. Even processes can be easily thought of as objects (this is what object-oriented software engineering does) and can similarly be represented in terms of F, B and S. One particular class of "objects" comprises agents. Figure 1 shows how an agent can be represented using the FBS view (Gero and Kannengiesser 2003a).




Figure 1 – An FBS view of an agent.

The function (F) of an agent is the purpose that an observer ascribes to its behaviour. It typically refers to the agent's role in some environment.

An obvious interpretation of an agent's behaviour (B) is how the agent acts given a set of conditions (which are shown in Figure 1 as input from the environment). This corresponds to the notion of a "black-box" or "input-output" view of the agent.

As illustrated, we distinguish two kinds of structure (S) of an agent. One refers to the "fixed" parts of the agent (S^f), i.e. those components or processes that are given at its inception and that are not subject to significant change. This type of structure is the same as for objects that have no agency, and typically includes "visible" components such as the sensors and effectors of the agent. The other kind of structure refers to the "situated" parts of the agent (S^s), i.e. those internal representations or processes that are constructed by the agent's interaction with its environment. The situated structure of an agent may be interpreted as its grounded concepts, beliefs, goals etc., which are often hidden from the observer (and thus depicted "inside" the fixed structure in Figure 1).

The specific instantiation of the FBS schema depends, just as for all other "objects", on the particular observer or the particular point of view. This is especially true for the structure (S), consisting of various fixed and situated parts, most of which are not readily visible for the observer. If the observer is not the creator of the agent, knowledge about the structure can be quite small.

The FBS schema can be used to model any kind of agent. Purely behaviour-based agents can be modelled from a "black-box" perspective, i.e. using the behaviour (B) and perhaps some parts of the fixed structure (S^f) while

ignoring the situated structure (S^s) part. Static symbol-processing agents, such as theorem provers and expert systems, as well as situated agents can be modelled using both behaviour (B) and structure (S). Static agents, however, can exhibit only a limited range of possible behaviours (B), as all of their structure (S) is encoded and fixed. Situated agents, in contrast, have the potential to construct new knowledge (S) and consequently perform new behaviours (B), which in turn can drive the construction of further knowledge (S).

The attribution of function (F) to a behaviour (B) of an agent depends solely on the observer. A particular behaviour can have one or more functions or no function at all. However, for any behaviour to be useful an observer has to attribute a function, which connects that behaviour to some goal of the observer. Therefore, for a situated knowledge structure (S) to be useful, it has to be both applicable, via appropriate behaviour (B), and valuable, by ascribing a desired function (F) to that behaviour.

As different situated agents are likely to have grounded different experiences by their interactions with the environment, there will often be significant differences in their situated knowledge. The perception of these differences is a condition for the knowledge of a particular agent to be judged as distinctive or specialised within a certain population of agents. Our model of situated expertise (as specialised situated knowledge) based on the FBS schema can account for the ability to represent different agents in a uniform way to allow their comparison.

4. Modelling interactions of experts

The FBS schema can now be employed to develop a model of interacting experts. Figure 2 shows the models, represented as FBS, that an agent (0) constructs of itself as well as of every agent it has interacted with (1 – 6). The FBS models of the other agents (1 – 6) are represented in the figure as nested in the FBS model of the agent (0) itself, since they are part of an agent's knowledge and thus its situated structure S^s . The FBS models are not only the results of previous interactions but also the starting point for new interactions. This is most obvious for those interactions directed by goals that are directly mapped to the function (F) of an entity or agent. For example, if an architect's goal is to get the results of a cost analysis (which this architect is not capable or willing to produce), they will address a costing expert, whose function (F) is to provide these results (via the behaviour (B) of computing, say, the total cost of a building). After identifying the right expert to carry out a certain task, the following interaction with that expert – in the form of communication – becomes a more difficult task. The reason for this is the commonly differing terminology and domain knowledge of the communicating experts, which can lead to problems of mutual understanding. This is where the structure part (S) of the FBS model of an

expert becomes important, as it is used to adapt the communication to that expert's knowledge. This adaptation is usually done by both the communicating and the "listening" agent (Clark and Murphy 1982).

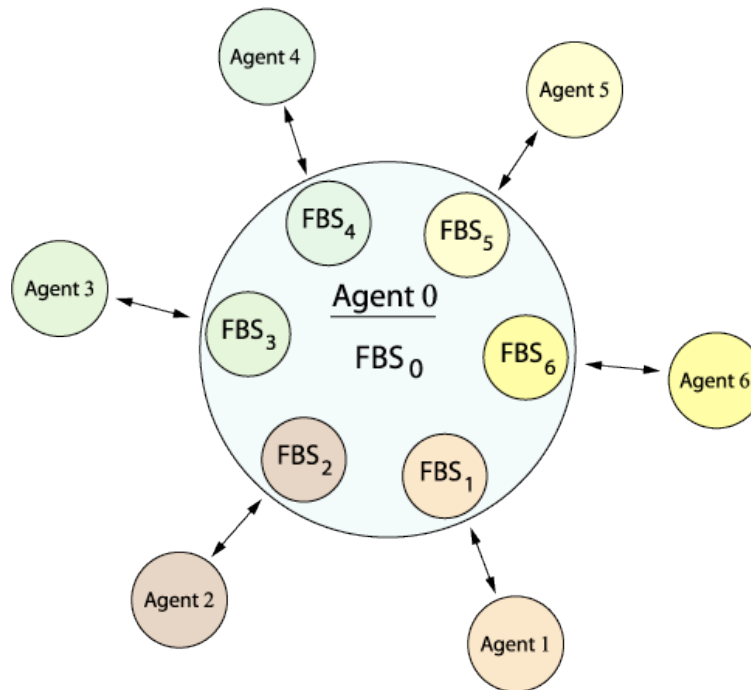


Figure 2 – FBS models of itself and other agents constructed by Agent 0 in an agent society.

For communication between interacting agents to succeed there needs to be a common ground (Clark 1992). This notion affirms that these agents have constructed appropriate models of each other to an extent sufficient for the purpose of the current interaction. Common ground is thus an emergent property of the current interaction. Figure 3 depicts the pairs of FBS models that have to be adequately consistent to establish the common ground between two agents.

A sufficient amount of structure (S), notably the "hidden" situated structure (S^s), is critical in the construction of FBS models to reach common ground in communication. The agent can generally use two sources of information to access S^s . The first one includes those parts of S^s that the other agent makes directly available by communicating them. The second one includes generalisations over a set of experiences with other agents. Cues for constructing these generalisations are often provided by observations of the other agent's behaviour (B). Usually both sources of information are employed, with generalisations typically providing default assumptions when only incomplete information is available from direct communication.

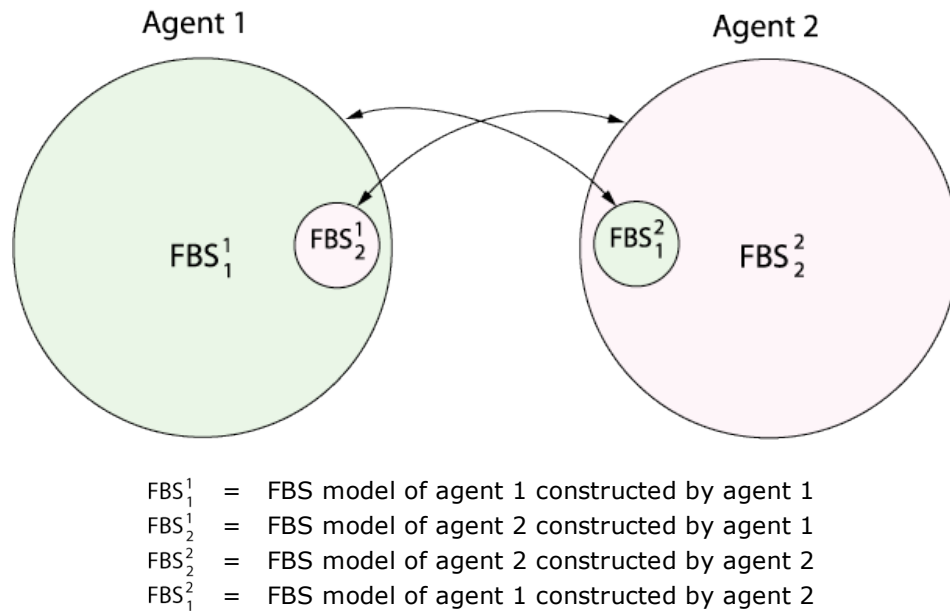


Figure 3 – Pairs of consistent FBS models that establish the common ground of two agents.

5. Interacting experts in temporary design teams

Temporary teams of experts are characterised by constant changes in their composition, as team members are selected according to the specific requirements of every design project. The criteria for this selection are usually the kinds of expertise of the individual designers that are expected to add up to form the desired kind of team expertise. Although this approach is an obvious one to follow in practice, it cannot provide a model of team expertise as it does not capture its emergent nature. If we want to develop such a model, we have to found it on the interactions of individual experts.

Expertise in temporary teams has to be reformed with every new project. However, as the individual experts have gained experience from previous work in similar teams, they have a potential to build up team expertise with less and less effort. The key for this potential is their ability to learn by forming generalised knowledge of different team environments and individual experts. Large parts of these models can be derived from generalised experiences. The FBS schema has been shown to allow generalisation in form of design prototypes (Gero 1990). We now illustrate how individual experts can use the FBS schema to flexibly emerge team expertise via generalisation.

Figure 4 shows an expert (0) having constructed FBS models of other experts (1, 2, 3 and 4). As the differently sized FBS models in the figure suggest, some experts (1 and 2) are better known (grounded) than others (3 and 4), and the best-known expert for expert 0 is certainly itself. When the expert wants to interact with one of the other experts but has too little knowledge about that expert (here 4) to establish sufficient common ground for this interaction, it complements the existing FBS model with assumptions

reflecting its generalised knowledge about similar experts. This generalised knowledge is derived mainly from those instances the expert (0) is most

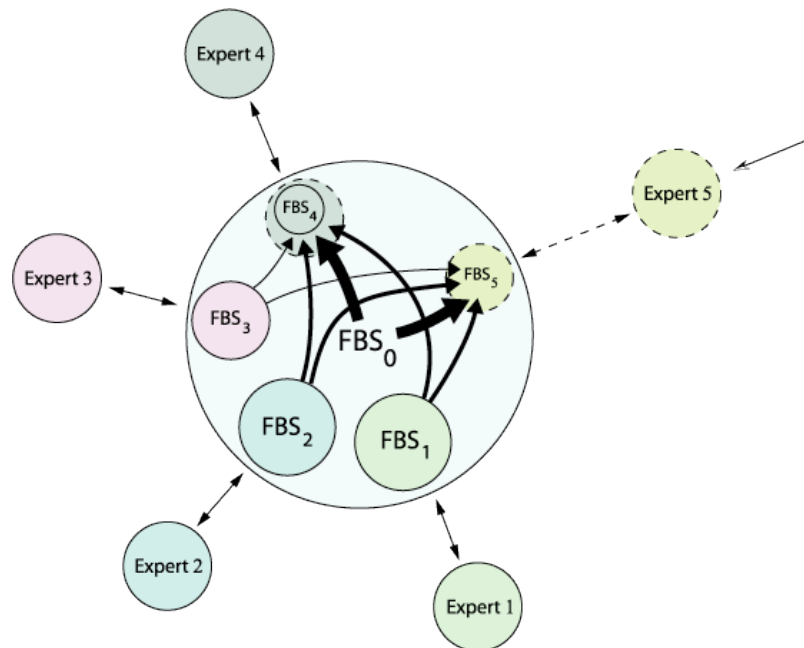


Figure 4 – New FBS models are constructed using generalisations of previously constructed FBS models. The size of the circle for each FBS is an indication of the amount of grounding of this FBS model of the other agent. The width of the arrows is an indication of the confidence of the potential applicability of the originating FBS model in constructing or supporting the FBS model of a new agent.

familiar with, as indicated by the different weights of the arrows in Figure 4, which principally includes the expert (0) itself. When a new, previously unknown expert (5) enters expert 0's environment, the generalised knowledge may still suffice to construct an adequate FBS model of that expert using the generalised knowledge about F, B and S individually and their relationships. If there is a conflict between the generalised knowledge and the interactions with a specific expert then a specialised FBS view of that expert needs to be constructed and added to the agent's knowledge of other experts.

6. Discussion

The model of expertise of temporary design teams outlined in this paper is based on situated agency that allows individual experts to adapt to changes in their environment. When placed into the environment of temporary teams the experts have to adapt to new combinations of team members. Consequently, they reform team expertise for every new team via emerging new common ground from their interactions. The FBS schema provides a means for them to maintain the potential for establishing the necessary common ground for team expertise. Figure 5 shows how an expert (0) uses its previous experiences of being a member of Team A to form the common

ground within the newly formed Team B. As Team B consists of former members (1 and 4) of Team A as well as of new members (5 and 6), both specialised and generalised knowledge is used.

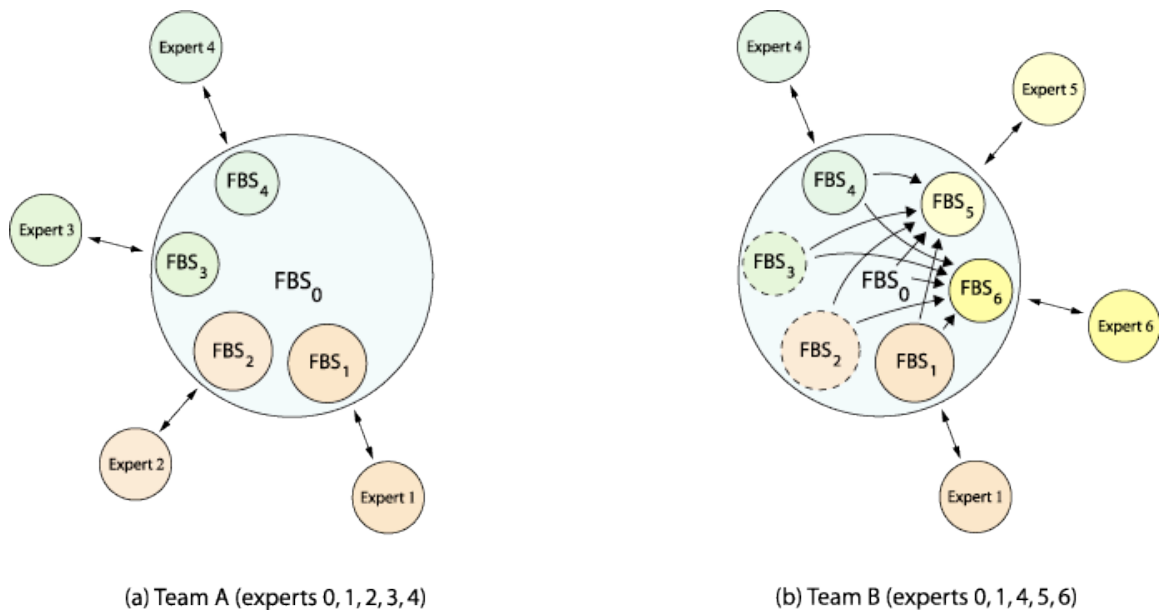


Figure 5 – Expert 0 integrates in Team B (b) based on its previous experience with being involved in Team A (a).

During the interactions within Team B expert 0's FBS views are modified again, which will affect the construction of further FBS views in future interactions. For example, after its participation in project Team B, expert 0 will possibly change its (FBS) views – usually by acquiring additional specialised and general knowledge. If Team A is then formed for a second time, expert 0 possibly constructs modified FBS models of the members of that team leading to differences in its (communicative) behaviour. The other experts of Team A, in turn, will then update their FBS model of expert 0 in response to these behavioural changes.

Modifications of FBS knowledge frequently originate from the interactions within a team itself, which then drives new interactions that would not have occurred without them. This accounts for the view of team expertise as emerging from team interaction rather than as the mere sum of the individuals' expertise. If the team (and its expertise) is considered as an entity itself, the observer can adopt an FBS view of the team rather than of the individual experts. This is commonly done when one is interested only in the collective output of the team and not how individual experts have contributed to that output.

Our model of expertise of temporary teams not only accounts for the ability of situated experts to interactively construct team expertise, it can also deal with organisational features of teams. For example, a team leader can be

modelled as an agent whose function (F) is to coordinate team interaction. More complex hierarchical structures can similarly be modelled using functions (F) referring to the specific privileges and competences of the particular agent in the respective hierarchical position. These functions can either be assigned beforehand or emerge during team interaction.

Finally, we contend that our model based on the FBS schema is useful for clarifying the very notion of expertise. It captures and distinguishes the different connotations of expertise as a social role, a skilled behaviour and a body of knowledge. The conflation of these different connotations has been a common cause for confusion when using the term expertise (Gaines 1994). The FBS schema also supports a social view of expertise that accommodates interactions of individual experts resulting in the emergence of team expertise. We see this as a basis for simulating temporary teams of situated design agents.

Acknowledgements

This research is supported by a Sesqui Research and Development grant from the University of Sydney and by an International Postgraduate Research Scholarship. Parts of this paper were presented at the *Design Thinking Research Symposium 6, Expertise in Design*, in Sydney, Australia (Gero and Kannengiesser 2003b).

References

- Brooks, R.A. (1986) A Robust Layered Control System for a Mobile Robot, *IEEE Journal on Robotics and Automation* 2(1), pp 14-23.
- Brooks, R.A. (1991) Intelligence without Representation, *Artificial Intelligence*, 47, pp 139-159.
- Clancey, W.J. (1997) *Situated Cognition*, Cambridge University Press, Cambridge.
- Clark, H.H. (1992) *Arenas of Language Use*, University of Chicago Press, Chicago, IL.
- Clark, H.H. and Murphy, G.L. (1982) Audience Design in Meaning and Reference, in Le Ny, J.F. and Kintsch, W. (eds.) *Language and Comprehension*, North-Holland Publishing Company, Amsterdam, pp 287-299.
- Dewey, J. (1896 reprinted in 1981) The Reflex Arc Concept in Psychology, *Psychological Review*, 3, pp 357-370.
- Feltovich, P.J., Ford, K.M. and Hoffman, R.R. (1997) (eds.) *Expertise in Context*, MIT Press, Cambridge, MA.
- Gaines, B.R. (1994) The Collective Stance in Modeling Expertise in Individuals and Organizations, *International Journal of Expert Systems* 7(1), pp 21-51.
- Gero, J.S. (1990) Design Prototypes: A Knowledge Representation Schema for Design, *AI Magazine* 11(4), pp 26-36.

Gero, J.S. (1998) Conceptual Designing as a Sequence of Situated Acts, in Smith, I. (ed.) *Artificial Intelligence in Structural Engineering*, Springer-Verlag, Berlin, pp 165-177.

Gero, J.S. and Fujii, H. (2000) A Computational Framework for Concept Formation for a Situated Design Agent, *Knowledge-Based Systems* 13(6), pp 361-368.

Gero, J.S. and Kannengiesser, U. (2003a) Function-Behaviour-Structure: A Model for Social Situated Agents, in Sun, R. (ed.) *Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*, International Joint Conference on Artificial Intelligence 2003, Acapulco, Mexico, pp. 101-107.

Gero, J.S. and Kannengiesser, U. (2003b) Modelling expertise of temporary design teams, in Cross, N. and Edmonds, E. (eds) *Expertise in Design*, University of Technology, Sydney, Australia, pp. 311-322.

Giarratano, J. and Riley, G. (1998) *Expert Systems: Principles and Programming*, PWS, Boston.

Langton, C.G. (1989) Artificial Life, in Langton, C.G. (ed.) *Artificial Life*, Addison-Wesley, Reading, MA, pp 1-47.

Levitt, R.E., Cohen, G.P., Kunz, J.C., Nass, C.I., Christiansen, T.R. and Jin, Y. (1994) The Virtual Design Team: Simulating how Organization Structure and Information Processing Tools affect Team Performance, in Carley, K.M. and Prietula, M.J. (eds.) *Computational Organization Theory*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp 1-18.

Schön, D.A. (1983) *The Reflective Practitioner*, Harper Collins, New York.

Schön, D.A. and Wiggins, G. (1992) Kinds of Seeing and their Functions in Designing, *Design Studies* 13(2), pp 135-156.

Smith, G. and Gero, J.S. (2000) The Autonomous, Rational Design Agent, in Fujii, H. (ed.) *Workshop on Situatedness in Design*, Artificial Intelligence in Design'00, Worcester, MA, pp 19-23.

Suwa, M., Gero, J.S. and Purcell, T. (1999) Unexpected Discoveries and S-Inventions of Design Requirements: A Key to Creative Designs, in Gero, J.S. and Maher, M.L. (eds.) *Computational Models of Creative Design IV*, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, Australia, pp 297-320.

This is a copy of the paper: Gero, JS and Kannengiesser, U (2004) Modelling expertise of temporary design teams, *Journal of Design Research* 4(3): <http://jdr.tudelft.nl/>