A BDI-Based Methodology for Eliciting Tactical Decision-Making Expertise

Rick Evertsz, John Thangarajah and Thanh Ly

Abstract There is an ongoing need to computationally model human tactical decision-making, for example in military simulation, where the tactics of human combatants are modelled for the purposes of training and wargaming. These efforts have been dominated by AI-based approaches, such as production systems and the BDI (Beliefs, Desires, Intentions) paradigm. Typically, the tactics are elicited from human domain experts, but due to the pre-conscious nature of much of human expertise, this is a non-trivial exercise. Knowledge elicitation methods developed for expert systems and ontologies have drawbacks when it comes to tactics modelling. Our objective has been to develop a new methodology that addresses the shortcomings, resulting in an approach that supports the efficient elicitation of tactical decision-making expertise and its mapping to a modelling representation that is intuitive to domain experts. Rather than treating knowledge elicitation, as a process of *extracting* knowledge from an expert, our approach views it as a collaborative modelling exercise with the expert involved in critiquing the models as they are constructed. To foster this collaborative process, we have employed an intuitive, diagrammatic representation for tactics. This paper describes TEM (Tactics Elicitation Methodology), a novel synthesis of knowledge elicitation with a BDI-based tactics modelling methodology, and outlines three case studies that provide initial support for our contention that it is an effective means of eliciting tactical decision-making knowledge in a form that can be readily understood by domain experts.

R. Evertsz (✉) · J. Thangarajah
RMIT University, Melbourne, VIC, Australia
e-mail: rick.evertsz@rmit.edu.au

J. Thangarajah
e-mail: john.thangarajah@rmit.edu.au

T. Ly
Defence Science and Technology Group, Perth, Australia
e-mail: thanh.ly@dsto.defence.gov.au

R. Sarker et al. (eds.), *Data and Decision Sciences in Action*,
Lecture Notes in Management and Industrial Engineering,
DOI 10.1007/978-3-319-55914-8_2

Keywords BDI · Cognitive modelling · Knowledge elicitation · Tactical decision making

1 Introduction

AI has been central to the computational modelling of tactical decision-making for over 20 years; early applications include air combat [10], and more recently, autonomous live-fire robot targets [6]. Previously, we have argued that these efforts have all been at the *implementation* level, and that tactics need to be represented at the *design* level in order to facilitate maintenance and sharing amongst developers, as well as comprehension by domain experts [7]. As part of that work, we proposed TDF (Tactics Development Framework) as a means of addressing this shortcoming and extended agent-oriented software engineering approaches such as Prometheus [14] to better handle the demands of tactical modelling.

Whereas TDF allows models of tactics to be captured, it does not support the elicitation of those tactics, and this is a challenging and important aspect that has not been addressed until now. In this work, we propose a methodology for the elicitation of tactical decision-making that maps directly to TDF concepts.

In theory, tactics can be automatically generated if the domain can be sufficiently well formalised. However, this is not practical in complex, real-world domains, and so the dominant approach to tactics modelling is to study domain experts. Unfortunately, eliciting knowledge from experts is a non-trivial exercise. The so-called *knowledge acquisition bottleneck* was recognised very early on as a key problem in the construction of intelligent systems [8]. Pioneering work on the knowledge acquisition problem was performed predominantly by those developing *expert systems*—computer models that were intended to be comparable in domain-specific competence to human experts.

Our experience, collaborating for over 20 years with military analysts modelling tactics in areas such as undersea warfare, air combat, tank battles, and special forces, suggests that knowledge elicitation continues to be a very significant problem. Across these groups of analysts, various techniques are employed, e.g. Cognitive Task Analysis [4], with varying degrees of success. It is difficult to critique these approaches collectively because they tend to be idiosyncratic, with a wide range of elicitation objectives and products. However, what they do have in common is the absence of a well-defined representation language for tactics that is readily understood by domain experts. The elicitation products tend to be static and result in an informal snapshot of the knowledge elicitation sessions that quickly becomes divorced from the tactics models. It is very difficult for domain experts to critique the resulting tactics implementation, because all they can do is observe the behaviour without a straightforward means of mapping that behaviour back to the elicited tactics.

Our research objective is to develop a methodology that supports the efficient elicitation of tactical decision-making expertise and its mapping to a tactics representation that is intuitive to domain experts. To this end, we have developed TEM (Tactics Elicitation Methodology), which advances the state of the art in knowledge elicitation to better handle the demands of tactics modelling. Rather than treating knowledge elicitation as a process of *extracting* knowledge from an expert, our approach views it as a collaborative modelling exercise with the expert involved in critiquing the models as they are constructed. To foster this collaborative process, we have employed an intuitive, diagrammatic representation for tactics. We argue that knowledge elicitation is intimately bound to the target conceptual framework. For example, elicitation methods that work well when building an ontology may need to be quite different to those that suit a diagnostic expert system. The former is likely to focus on teasing out the fine detail of conceptual structures, whereas the latter will be more concerned with cues and problem-solving methods. Likewise, tactics elicitation focuses on particular factors such as the events that trigger the adoption of goals, the methods used to achieve those goals, and the conditions that indicate that a goal should be dropped in favour of a more pressing one.

This paper makes several contributions to the field of behaviour modelling in intelligent systems: (i) it offers a novel synthesis of knowledge elicitation with a BDI-based tactics modelling methodology; (ii) it focuses knowledge elicitation on the aspects that are important to tactics modelling; and (iii) the elicitation methodology's application to three domains is evaluated.

2 Background

Military science definitions of the term, *tactics*, tend to emphasise their adversarial nature. More generally, tactics can be viewed as the means of achieving an immediate objective in a manner that can be adapted to handle unexpected changes in the situation. Tactics are distinguished from *strategy* in that the latter is concerned with a general approach to the problem, rather than the means of achieving a more short-term goal. A submarine commander's use of stealth to approach a target is a strategy, whereas the particular method used, for example hiding in the adversary's baffles (a blind spot), is a tactic. Tactics are concerned with the current, unfolding situation—that is, how to deflect threats and exploit opportunities to achieve victory. This view of tactics, as the means of achieving a short-term goal in a manner that can respond to unexpected change, seems to be common to all definitions, whether in military science, game theory, or management science.

2.1 Knowledge Elicitation

Over the years, a significant proportion of the effort to construct intelligent systems has adopted a *knowledge-based* approach. In contrast to general problem-solving methods, the knowledge-based approach relies on encoding a large quantity of domain-specific knowledge. This was motivated by the observation that expert human performance achieves its efficacy by applying domain-specific strategies and knowledge to the problem at hand. The computational modelling of expert competence came to be known as *knowledge engineering*, with the term *knowledge elicitation* being used to denote the initial process of extracting the knowledge from the expert before encoding it in a suitable knowledge representation language.

Practitioners generally agree that the elicitation of expertise is a challenging endeavour. Musen identified a number of reasons [11]; perhaps the most significant is that expert knowledge is largely tacit. It is generally agreed that *knowledge compilation* [12] is central to the human cognitive architecture's ability to efficiently process large amounts of information as smaller, composite items that functionally expand working memory. The downside is that it is difficult, if not impossible, for experts to decompile such knowledge to understand the basis of their decision-making. Consequently, experts can be inaccurate when recounting their thought processes, and may unintentionally fabricate reasoning steps in order to bridge the gaps in their conscious understanding of their own decision-making [9].

2.2 BDI and the Theory of Mind

Our experience of eliciting tactical knowledge from experts, in domains including air combat, infantry tactics, air traffic flow management, and undersea warfare, suggests that the outcome is greatly influenced by whether the experts find the conceptual framework in which the knowledge is cast to be intuitive. Experts are better able to describe their reasoning if the emerging tactical model corresponds to how they *think they think*. In the field of cognitive science, it is generally accepted that human beings employ a *theory of mind* to explain and predict the actions of others, as well as to provide a narrative for their own reasoning processes. According to the *intentional stance* [5], we attribute intentions to ourselves and others, and believe them to arise from a combination of our beliefs and desires.

Through the work of Bratman [2] on practical reasoning, this philosophical perspective gave rise to the BDI paradigm, a modelling approach that has been popular for many years in the multi-agent systems community. It has been claimed that the BDI paradigm is a good *folk psychology* model of human reasoning, i.e. a model of how we think [13].

The BDI model is a particularly parsimonious conception of rational agency, characterising agents as reasoning in terms of their *beliefs* about the world, *desires* that they would like to achieve and the *intentions* that they are committed to. Apart from its intuitive appeal to domain experts, it is a powerful computational abstraction for building sophisticated, goal-directed, and reactive reasoning systems, and consequently is well suited to tactics modelling.

Put very succinctly, a BDI agent performs a continuous loop in which it updates its beliefs to reflect the current state of the world, deliberates about what to achieve next, finds a plan for doing so, and executes the next step in that plan. Each time around this cycle, it effectively reconsiders its options, yielding goal-oriented behaviour that is also responsive to environmental change.

Our research objective is to develop a tactics elicitation methodology that is easier to master than Cognitive Task Analysis, and offers an intuitive representation of decision-making—one that is compatible with experts' intuitions about how they reason tactically. With this in mind, we have developed a BDI-based elicitation methodology.

2.3 Tactics Development Framework (TDF)

TEM targets TDF, a BDI-based tactics design methodology founded on agent-oriented software engineering principles. In TDF, the tactical environment is represented in terms of the important *percepts* received, *actions* performed, and the *actors* the system interacts with. Methods of achieving *goals* using *plans* are encapsulated in terms of *tactics*. These in turn are encapsulated by the *missions* that the system can undertake. The other concepts that are relevant to TEM are *messages* passed between *agents*, *scenarios* outlining the different ways that a mission can play out, the *roles* embodied by the system, and the *data* maintained by the system.

TDF divides modelling into three major stages: (i) *System Specification* of system-level artefacts, (ii) *Architectural Design* of the internal structure of the tactical system, and (iii) *Detailed Design* of the internals of the agents in the system. In addition, *tactics design patterns* span the three stages, providing a reusable template from which new tactics can be built. TDF has been used to model undersea warfare, Apache attack helicopter tactics, UAS decision-making, and infantry tactics. Further detail on the methodology can be found in [7].

3 Tactics Elicitation Methodology (TEM)

TEM's artefacts all map directly to the corresponding TDF concepts (apart from TEM's *concept maps*, which TDF does not include). Some of the TDF artefacts have been renamed so that they are more meaningful to domain experts within the context of knowledge elicitation, namely *mission* becomes *vignette*, *agent* becomes *character*, *actor* becomes *environmental entity*, *scenario* becomes *storyline*, and *data* is now *belief*.

The TEM artefacts are as follows: (i) *Concept Map*—a diagrammatic graph-like structure that informally represents concepts and their interrelationships; (ii) *Vignette*—a succinct description of a tactical situation, as per [3]; (iii) *Tactic*—specifies how to achieve a goal in a situation that might change in unexpected ways, and comprises a goal/sub-goal structure, the methods used to achieve goals, and the conditions under which goals are adopted, dropped, suspended or resumed; (iv) *Storyline*—a specification of the key cues and decision points that occur as the vignette unfolds; (v) *Goal Structure*—a goal/sub-goal hierarchy; (vi) *Plan*—a method of achieving a goal, or responding to an event; (vii) *Character*—a behavioural entity embodying some combination of tactics, goals, and plans; (viii) *Team*—an aggregation of characters and/or teams; (ix) *Role*—functionality assigned to a character or team; (x) *Percept/Action*—a character receives information in the form of percepts and affects its environment by performing actions; (xi) *Environmental Entity*—source of percepts or the target of actions; (xii) *Belief*—each character or team has a set of beliefs that defines its declarative knowledge.

3.1 TEM Steps

TEM is an interview-based elicitation methodology, as this is the most efficient method in terms of time and resources [4].

The following account is necessarily abbreviated to fit within the confines of this paper, but all of the steps are included, if not in great detail. The examples are based on a study [1] in which an HVU (High Value Unit) is protected by four destroyers (Screen Ships) from attack by a submarine.

TEM comprises six sequential steps that are repeated until the requisite information has been acquired: (i) Scope Requirements; (ii) Expand Modelling Objective; (iii) Construct Storylines; (iv) Populate Concept Map from Storylines; (v) Develop TDF Design; and (vi) Plan Next Session.

3.1.1 Scope Requirements

- *Elicit Stakeholder Objective.* If required, interview the stakeholder to determine the question to be answered by the modelling exercise, for example, *'How many screen ships are required to protect an aircraft carrier from submarine attack?'*.
- *Elicit Modelling Objective.* The stakeholder objective is usually not sufficiently specific for the purposes of modelling and needs refining to derive the *modelling objective*—ideally a testable hypothesis relating independent to dependent variables. If the modelling objective can be expressed quantitatively, determine the MOEs.[1]
- *Study Background Material.* Ensures interviewer has sufficient grounding in the domain to efficiently interview the domain expert.
- *Scope Interview Sessions.* Purpose is to determine how much time is available, and to plan the sessions so that the stakeholder objective is adequately addressed.

[1] An MOE (Measure of Effectiveness) is a quantitative impact assessment, e.g. *'Dollar value of aircraft carrier damage saved per screen ship deployed'*. In simulation applications, the values for a dependent variable are typically aggregated across multiple simulation runs to yield an MOP (Measure of Performance), e.g. *'average dollar value of aircraft carrier damage'*. The MOE combines the MOPs into an impact statement.

3.1.2 Expand Modelling Objective

The modelling objective is central to the TEM process and is used to ensure that the elicitation does not drift off into interesting but irrelevant detail. This step results in a concept map representation of the modelling objective properties and related attributes, such as independent variables, dependent variables, and MOEs. This step comprises three activities:

- Create a concept map with the modelling objective at its root.
- Link MOEs to the modelling objective in the concept map.
- Add any modelling objective preferences to the concept map, e.g. variables that define the optimality of a given solution.

The concept map is a graph that acts as a shared diagrammatic representation of the domain. It does not constrain the interviewer to a particular set of artefacts, and this notational freedom is instrumental in fostering rapid elicitation. However, structure is added to the concept map after each elicitation session, by identifying relevant TDF artefacts (Fig. 1 shows part of a bushfire concept map after the TDF artefacts have been identified).

3.1.3 Construct Storylines

- *Elicit Relevant Vignettes.* Prompt the domain expert for vignettes that he/she has personal decision-making experience of, and that address the modelling objective. Identify those vignettes that are most likely to provide useful information about the modelling objective. Add the selected vignette(s) to the concept map. Where appropriate, employ diagrams to document the spatial layout of the key entities in the vignette, for the purposes of discussion with the domain expert.
- *Generate Relevant Storylines.* The domain expert selects the vignette that reveals the most about the modelling objective, and recounts a related storyline, for which he/she was the primary decision-maker. Failing this, the domain expert should generate a hypothetical storyline.
- *Verify Selected Storyline.* Verify that the storyline addresses the modelling objective *before* too much time is spent eliciting the details. Establish that the outcomes (as characterised by the MOEs) will depend upon the identified dependent variables. If the relationship between the storyline and modelling objective is not convincing, choose a new storyline. If the domain expert cannot come up with a more effective storyline, select a new vignette.
- *Fill Out Storyline.* Map out the key decision-making cues and decision points in the storyline. Link the storyline to the vignette in the concept map. Elicit missing information and significant storyline variations that will be needed to build the tactics, i.e. (i) key decision points, (ii) the goal being tackled in each episode of the storyline, (iii) conditions that need to be maintained during storyline episodes, (iv) goals adopted/dropped/suspended/resumed and the conditions under which this

happens, (v) the name of the tactic adopted to achieve the goal, (vi) other courses of action considered, (vii) actions performed as part of the tactical approach, (viii) significant roles/characters/teams involved in the storyline, (ix) communications between characters, teams, and environmental entities, (x) timing information (relative and/or absolute) that relates to the events in the storyline.

3.1.4 Populate Concept Map from Storylines

This step's purpose is to identify concepts that arise from the storylines and elicit further detail about those concepts. To do so, label concepts in the storyline that relate to goals, tactics, roles, teams, characters, key conditions, and domain-specific terminology. Figure 1 shows an annotated concept map from a firefighting case study (with actions, beliefs, goals, percepts, plans, and roles).

3.1.5 Develop TDF Design

This step occurs after the elicitation session is over.

- *Develop Storyline Diagrams.* Results in a diagrammatic representation of each storyline, to be used in subsequent elicitation sessions, and as input to the TDF design process. Figure 2 shows an example of an *episode* of a storyline (a storyline can be made up of a number of episodes).

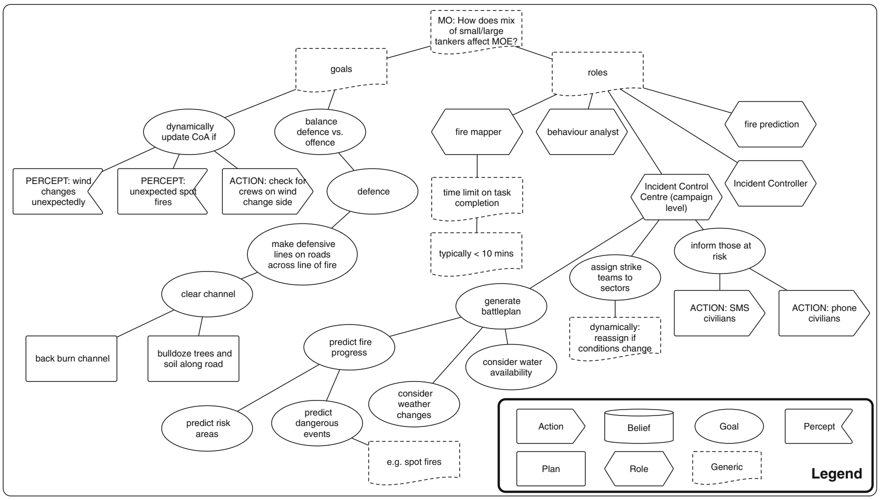

Fig. 1 Part of bushfire concept map (after artefact classification)

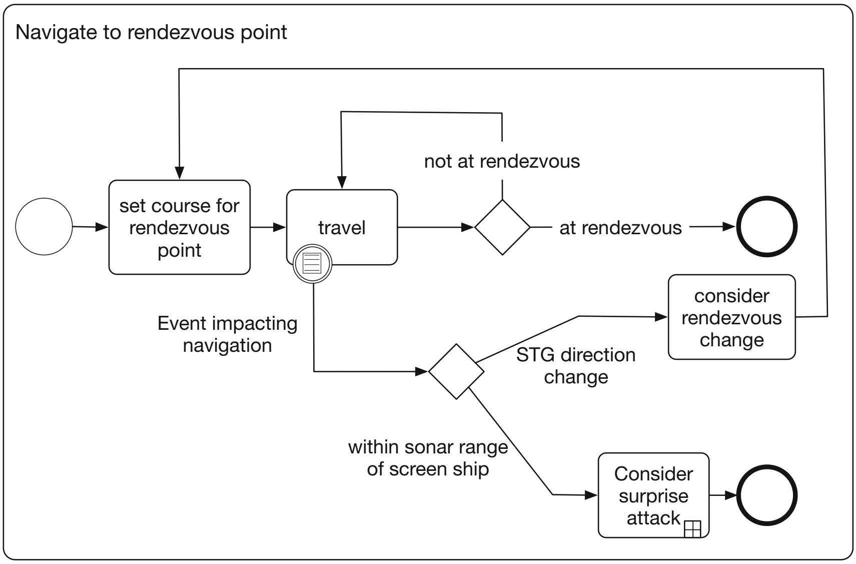

Fig. 2 Navigate to rendezvous (example storyline episode)

- *Develop Initial TDF Design.* This involves identifying and labelling TDF artefacts in the concept map and storylines (goals, plans, roles, percepts, etc.), using a palette of predefined TDF icons. Underspecified concepts are highlighted so that they can be addressed at the next elicitation session. An initial TDF design is constructed as outlined in the *System Specification*, *Architectural Design*, and *Detailed Design* stages, described fully in [7].

3.1.6 Plan Next Session

This step identifies missing information so that gaps can be filled during the next elicitation session. It comprises the following activities:

- Work through any artefacts that were marked as underspecified while developing the initial TDF design.
- Verify that there is sufficient information to build a goal structure for each of the decision points in the storylines.
- Verify that there is enough information to build the plans required to achieve the goals.
- Identify missing artefacts, for example, percepts that are needed to discriminate between alternative plans for achieving a goal.
- Identify areas of focus for the next elicitation session.

4 Preliminary Case Studies

The evaluation of knowledge elicitation methodologies continues to be a problematic and challenging endeavour. An extensive discussion of the issues can be found in [15]. Impediments include the difficulty of accessing a large enough sample of domain experts for statistical significance, the lack of homogeneity in domain expert opinions, how to quantify coverage and depth of knowledge elicited, and how to compare the content of disparate knowledge representations. Although it would be feasible to conduct a quantitative evaluation using a restricted and formal board game, such as chess, this would be unenlightening with respect to the domains TEM is designed for. TEM targets dynamic, real-world domains and so focuses on eliciting the goal-directed and reactive aspects of decision-making that rely on complex and varied contextual cues and broad repertoires of tactical actions. With these caveats in mind, a preliminary evaluation of TEM was performed, using three distinct domains to broaden the coverage: (i) firefighting of bushfires, (ii) undersea warfare, and (iii) the League of Legends real-time, multi-player tactical game.

4.1 Bushfire Tactics (Pilot Study)

The objective of the initial evaluation of TEM was to investigate its efficacy. The domain expert had many years of experience fighting bushfires, both on the ground and from the control centre. After a number of knowledge elicitation sessions, totalling 5 h, a TDF design was produced from the elicited storyline and concept map. The domain expert judged the design to be a clear and accurate representation of the elicited firefighting tactics. He volunteered that he found the whole process to be very rewarding and was highly motivated to provide accurate and detailed information about tactical alternatives because he felt that he fully understood how his expertise was being represented and so felt ownership of the models as they were being developed. This feedback provides encouraging support for TEM's collaborative approach to elicitation.

4.2 Undersea Warfare Tactics

The undersea warfare environment is characterised by a high degree of uncertainty about the true tactical situation. Situation awareness is time consuming to acquire and often unreliable, and this makes it difficult to predict the outcome of a particular course of action. Many factors have to be considered when trying to eliminate hypotheses, and this, taken together with the high level of uncertainty, makes it difficult to elicit the submarine commander's tactical knowledge. An analyst, attempting to elicit and build a model of submarine commander decision-making, is faced with a very wide and potentially confusing array of options, and this was the original motivation for developing TEM.

A number of 90-min undersea warfare tactics elicitation sessions were conducted over the course of three consecutive days, amounting to about 18 h. Six domain experts took part in the sessions, which focused on submarine evasion of helicopter(s) equipped with dipping sonar. Some of the experts had extensive knowledge of the performance of active sonar, which is highly relevant to the modelling objective (dipping sonar is active, not passive), and one of the domain experts, a submarine commander, had direct experience of operating against helicopters with dipping sonar. The two interviewers had experience of informally interviewing domain experts, and one had been an observer in the bushfire elicitation sessions. Neither had any experience leading a TEM session, and relied solely on reading the TEM documentation beforehand.

The interviewers provided general feedback on TEM, but due to security considerations, were not able to provide us with the concrete products of the elicitation sessions. They reported that:

- the methodology was straightforward and very successful in guiding their elicitation of the requisite information;
- the methodology should include a step recommending diagramming of the spatial layout of the problem;
- the methodology's focus on working immediately from the vignettes/storylines was vital because the domain experts resisted discussing abstractions, preferring instead to discuss specific instances;
- the elicitation sessions exposed novel information that was not in any written material about the domain;
- the storyline mechanism was very effective in eliciting key perceptual cues, decisions, and sequences of actions; and
- the identification of tactical goals was instrumental in the elicitation of alternative courses of action.

The feedback from the interviewers and domain experts was overwhelmingly positive, and the interviewers reported that the methodology was easy to follow, and the outcome was successful. There were two recommendations for change: (i) add a methodological step for creating one or more diagrams of the spatial layout of the tactical situation, and (ii) emphasise the need to begin with concrete vignettes and storylines, rather than abstract questions about the domain.

4.3 League of Legends Tactics

The final case study applied TEM to the elicitation of League of Legends tactics, and apart from investigating the efficacy of the methodology, it focused on the comprehensibility of the resulting TDF designs.

League of Legends[2] is a fast-paced, multi-player, online strategic/tactical game, with tens of millions playing it daily. Each player takes on the role of a *summoner* who controls a champion. Each game comprises two teams of five players, with each team battling the other to destroy their *nexus*. The League of Legends domain is particularly rich, with a wide range of characters possessing very different capabilities, many classes of weaponry, only partial situation awareness, team tactics, and frequent unexpected changes in the tactical situation. Although the tactics are goal-directed, there is a large reactive component and a good player must be able to respond quickly to significant changes, based only on partial information about the situation.

Three domain experts took part in the study; one was interviewed over the course of seven 90-min sessions, and the other two assessed the comprehensibility of the resulting TDF tactics diagrams. All three experts have a world ranking in the top few per cent of the League of Legends Elo rating system, but none of them were familiar with TEM or TDF. During the storyline phase, TEM helped elicit cues for adopting and dropping goals, with some complex interactions between *achievement goals* and *maintenance goals* (the latter are goals that are eligible for adoption when their maintenance condition is violated).

Comprehension was assessed by giving the two experts, who were not involved in the elicitation sessions, a two-page introduction to TDF's goal diagram notation, 13 goal diagrams representing some of the elicited League of Legends tactics, and a questionnaire comprising ten questions. The questions probed understanding of the meaning of the diagrams and how to change the diagrams to produce a different tactical outcome. One of the diagrams included a deliberate omission, which both experts were able to identify; they were also able to provide suggestions on how to extend the tactics to handle a wider range of cases. The expert used for the elicitation sessions assessed their level of comprehension as high, with a mean score of 85% correct. Their high scores were very encouraging, given that they had no previous exposure to TDF.

5 Discussion

This paper introduced TEM, a new approach to tactics elicitation that targets the BDI-based framework, TDF. The BDI paradigm was chosen because it is intuitive to domain experts, and the literature suggests that it maps well to their *theory of mind*. A major contribution of this paper is TEM's specific methodological steps and how they focus the interview process on the information required to build BDI-based tactics, namely roles, characters, teams, goals adopted/dropped/suspended/resumed, tactics, key decision points, and so on. The central role played by the modelling objective is another important contribution of this work to tactics modelling because it ensures that the elicitation process remains focused on the desired outcome,

[2] http://leagueoflegends.com/.

thereby ameliorating a common problem in knowledge elicitation—that of drifting off point and wasting time on interesting detail that ultimately does not contribute to the requirements of the study. This was a significant factor in the efficiency of the elicitation sessions in the three case studies.

The evaluation of TEM, though preliminary, was quite broad in scope and provides some confidence in the efficacy of the methodology. The bushfire case study revealed a few important deficiencies in the original methodology, and these were addressed before the undersea warfare case study. The undersea warfare study provides initial support for our claim that TEM is easy to learn and apply effectively; one of the two interviewers had only been an observer in the bushfire case study, and the other had no experience of TEM but had many years of experience in eliciting and modelling tactics for undersea warfare. His feedback on the effectiveness of TEM is particularly encouraging (see the five points at the end of Sect. 4.2). The League of Legends case study provides support for the claim that the TEM/TDF combination produces models that are easily understood by domain experts that were not part of the elicitation sessions. TEM's emphasis on collaboration between the modeller and domain expert is an important contribution of this work. The domain experts in the case studies all expressed that they found this aspect highly motivating and that they felt confident that the resulting models were a valid reflection of their tactical expertise.

A key contribution of this work is the fact that the elicitation methodology is closely tied to an effective and intuitive tactics modelling framework (TDF). This helps keep the sessions focused on the desired modelling products and provide continuity from elicitation, through to modelling, and ultimately implementation. We believe that, after two years of development and evaluation, this work is now sufficiently mature to warrant wider dissemination to those building intelligent decision-making systems that need to be based on human expertise. However, further work remains.

5.1 Further Work

This research has ongoing support from our undersea warfare sponsor, and now has funding from a different sponsor to investigate its application to team-based air combat tactics. Wider adoption of TEM would be greatly facilitated by the provision of appropriate tool support, and tool development is now underway. The case studies were conducted by adapting publicly available BPMN and concept mapping tools, but this is less than ideal because they are not sufficiently tailored to TEM, and there is no integration with the TDF tool. Our TEM tool development will support the construction of concept maps and storylines. Linking the TEM tool with the TDF tool will provide an opportunity for verification and validation of the elicitation and design artefacts, as well as TDF-generated code. Artefacts identified in the concept maps and storylines could be linked to their counterparts in the TDF design; storylines could also be mapped to TDF plan diagrams. Integration of this sort would help with traceability from elicitation right through to the behaviour that results from running code.

References

1. Akbori, F.: Autonomous-agent based simulation of anti-submarine warfare operations with the goal of protecting a high value unit. Master's Thesis (2004)
2. Bratman, M.E.: Intention, Plans, and Practical Reasoning. Harvard University Press, Cambridge, MA (USA) (1987)
3. CLAMO.: Rules of Engagement (ROE) Handbook for Judge Advocates. Center for Law and Military Operations, Charlottesville, Virginia (2000)
4. Crandall, B., Klein, G.A., Hoffman, R.R.: Working Minds: a Practitioner's Guide to Cognitive Task Analysis. MIT Press (2006)
5. Dennett, D.C.: The Intentional Stance. MIT press (1987)
6. Evertsz, R., Lucas, A., Smith, C., Pedrotti, M., Ritter, F.E., Baker, R., Burns, P.: Enhanced behavioral realism for live fire targets. In: St. Amant, R.S., Reitter, D., Stacy, E.W. (eds.) Proceedings of the 23rd Annual Conference on Behavior Representation in Modeling and Simulation. BRIMS Society, Washington DC (2014)
7. Evertsz, R., Thangarajah, J., Yadav, N., Ly, T.: A framework for modelling tactical decision-making in autonomous systems. J. Syst. Softw. **110**(C), 222–238 (2015). 10.1016/j.jss.2015.08.046
8. Feigenbaum, E.A.: Themes and case studies of knowledge engineering. In: Expert Systems in the Micro-Electronic Age, pp. 3–25 (1979)
9. Johnson, P.E.: What kind of expert should a system be? J. Med. Philos. **8**(1), 77–97 (1983)
10. Laird, J.E., Jones, R.M., Jones, O.M., Nielsen, P.E.: Coordinated behavior of computer generated forces in Tacair-Soar. In: In Proceedings of the Fourth Conference on Computer Generated Forces And Behavioral Representation (1994)
11. Musen, M.A.: An Overview of Knowledge Acquisition. Springer (1993)
12. Neves, D.M., Anderson, J.R.: Knowledge compilation: mechanisms for the automatization of cognitive skills. In: Cognitive Skills and Their Acquisition, pp. 57–84 (1981)
13. Norling, E.: Folk psychology for human modelling: Extending the bdi paradigm. In: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1, pp. 202–209. IEEE Computer Society (2004)
14. Padgham, L., Winikoff, M.: DEveloping Intelligent Agent Systems: a Practical Guide, vol. 1. Wiley (2004)
15. Shadbolt, N., O'Hara, K., Crow, L.: The experimental evaluation of knowledge acquisition techniques and methods: history, problems and new directions. Int. J. Hum. Comput. Stud. **51**(4), 729–755 (1999)

http://www.springer.com/978-3-319-55913-1

Data and Decision Sciences in Action
Proceedings of the Australian Society for Operations Research Conference 2016
Sarker, R.A.; Abbass, H.A.; Dunstall, S.; Kilby, P.; Davis, R.; Young, L. (Eds.)
2018, IX, 414 p. 93 illus., 63 illus. in color., Hardcover
ISBN: 978-3-319-55913-1