Topic:	C2 Experimentation
Paper:	Elicitation of British Army Commanders' Personal Constructs

Author:	Paddy Turner
Organisation:	QinetiQ
Address:	Malvern Technology Centre, St. Andrew's Road, Malvern, Worcestershire. WR14 3PS
Telephone:	+44 (0) 1684 895478
Fax:	+44 (0) 1684 896767
E-mail:	**pjturner@QinetiQ.com**

Elicitation of British Army Commanders' Personal Constructs

Paddy Turner
Intelligence Systems Department, Knowledge and Information Systems Division, QinetiQ.
Alan Turing Building, Malvern Technology Centre, Malvern, Worcestershire, United Kingdom. WR14 3PS

Abstract

The design of information systems to aid the attainment of Situation Awareness would benefit from an understanding of the mental models that commanders reference in assessing situations. This paper presents novel experimental research conducted to elicit the personal constructs that constitute the fabric of British Army commanders' mental models. For a given scenario, it is shown that four out of the twenty-one identified constructs accounted for the majority of individual situation assessments made; three of these four were, in fact, common to all ten participants. There were, however, individual differences: four of the participants referenced a specific pair of constructs in the majority of their assessments whilst the remainder used a far wider set. The constructs were mapped onto Endsley's three levels of Situation Awareness and a significant positive correlation was found between experience level and the mean Situation Awareness level of the constructs referenced. This suggests that experience influences the fabric of commanders' mental models and indicates how this manifests itself in terms of constructs.

The design of information systems to aid the attainment of Situation Awareness (SA) requires an understanding of user information requirements. For British Army tactical commanders, doctrine and training provide a view on what these information requirements might be in terms of a set of factors that are to be taken into consideration in the decision-making process. However, doctrine and training only provide a template for SA; it is recognised that expert commanders, in particular, tend to deviate from recognisable doctrinal practices. Moreover, it is postulated here that individual commanders possess fundamentally different mental models that they reference in assessing situations. If all commanders were to be supported in the attainment of SA, it would be beneficial to understand these mental models by eliciting the factors that they actually use.

This paper details research funded under the UK MOD Corporate Research Programme (Human Sciences and CBD Domain). The aim of the research was to elicit the factors that British Army commanders *actually* reference in assessing situations, and the explore differences in their choice of factors[1]. To this end, an elicitation experiment was conducted in October 2002. The elicitation of land tactical commanders' mental models represents a novel piece of research; the application of the chosen elicitation method within this domain is also novel. The remainder of this

---

[1] In the remainder of this paper we shall refer to 'factors' only when discussing doctrine, and 'personal constructs' when discussing the factors that individuals actually use. Both, however, serve to characterise situations.

paper provides the experimental design, outlines the coding of the qualitative data, presents the analysis and interprets the findings.

Experimental Design

Personal Construct Elicitation

The experiment was based on an application of Kelly's Personal Construct Elicitation technique developed by George Kelly [1]. A characteristic of this approach is that it involves the presentation of *triads* of elements (stimulus material) to participants who are then asked how two of the elements are similar and thereby different to the third.

Experience scale

A team of current and recently retired military officers assisted the research team at all stages of experimental design, execution and analysis. In particular, the most senior military team member (ranked Lt Col) developed the experience scale. This scale measured participants' levels of experience relevant to the specific scenario and was based on scores against a number of component scales, including rank, level of academic achievement, military education, military training, number of exercises participated in, number of operations participated in, operational posts and non-operational posts. The military team prepared an experience questionnaire to collect qualitative data against each of the component scales, and conducted an analysis of the results after the experiment. The overall experience scale had four levels: no experience, some experience, experienced and considerable experience.

Preparation of snapshot elements

The experiment detailed here used a Battlegroup warfighting scenario as the source of elements. The forces on both sides comprised approximately half of a Brigade (two blue Battlegroups vs. four red Battalions). The context for the scenario was encapsulated in the wider Brigade missions for both forces: red had crossed the international border at H Hour with the objective of capturing a key bridge some 40km to the west before H+24; blue, who were defending the ground, were required to prevent them from doing so before H+24. Blue's Brigade plan was based on a mobile defence, slowing red forces down by attacking with limited force for limited periods of time, sufficient to force them to deploy from the advance to attack. This was to be achieved in four phases, the most aggressive of which was the third phase, to be launched against the red Brigade's second echelon forces. Red doctrine dictated a two-echelon advance; the second echelon forces would join up with the first echelon forces to meet the Brigade objective together. The two blue Battlegroups were given the mission of slowing the second echelon forces down sufficiently to prevent this from happening, at it would be achieved through a co-ordinated counterattack from hidden positions into open ground. The terrain was fairly constrained, although the Area of Operations did include a valley wide enough to offer a frontage for two Battalions with three single Battalion mobility corridors offering routes from the valley to the west.

Each element, or snapshot, depicted a possible battle situation that might have occurred if the scenario had been fought for real. Snapshots were developed by the military team who war-gamed a range of feasible red COAs against the blue COA to determine likely outcomes at various times in the battle's development. These outcomes were developed into a set of 18 snapshots that, together, were representative of the envelope of expected battle situations as judged by the military team. Each comprised:

- a map overlay employing standard military symbology depicting red and blue forces at the level of command two below Battlegroup/ Battalion (this is standard practice);
- a written brief containing a time-stamp, status and activity information for units depicted on the overlay, relevant environmental information and a summary of recent events.

Map overlays and briefs were generated by the military team. Current ISTAR[2] capabilities were assumed and this dictated exactly what information was included on the overlays and briefs (i.e. the information was neither complete nor real-time). Further, each snapshot was viewed from the perspective of one of the two Battle Groups and this also shaped what information was deemed to be available; the Battle Group in question was indicated at the top of the written brief.

It should be noted that the scenario contained a number of simplifications that ensured that resultant snapshots were not overly complex whilst remaining within the bounds of military credibility. For example, only combat and combat support units (artillery and engineer) were depicted; combat service support units were excluded under the assumption that both the red and blue missions were sustainable for the duration of the battle. Details of flank and rear units, and air assets, were omitted under the assumption that these were involved in other parts of the wider operation and would not impact upon events in the Area of Operations for the duration of the battle.

Experimental context

It was necessary to provide participants with realistic, credible and equivalent context for the subsequent assessment of snapshots. This was achieved through:

- a military brief, including the Brigade plan and missions (COAs) for each of the two blue Battle Groups involved in the scenario;
- a short map-based mission planning exercise, during which participants were given the opportunity to assess both the Area of Operations and likely red COAs, and to understand the blue Battle Group COAs.

The 18 snapshots were set across a range of times, the earliest of which was 75 minutes after the beginning of the battle (for which a final situation update briefing had been provided). Participants were therefore required to assess situations without having been exposed to any precipitating events. In the field, commanders do spend periods away from the HQ (e.g. when sleeping or visiting subordinate commanders); upon returning to the HQ, commanders would expect to receive a situation update brief from their staffs and would subsequently receive updates as and when they arrive at the HQ. In the experiment, the written brief played the role of the staffs by providing the update of recent events; it also provided much of the information that would be conveyed verbally to the commander in the field and not encapsulated on the map or overlay. Hence both the style of snapshot presentation, and the process by which participants were prepared for elicitation, were realistic and credible.

---

[2] Intelligence, Surveillance, Target Acquisition and Reconnaissance

Experimental procedure

1. At the beginning of the experiment participants were assigned a random and unique participant ID. The 18 snapshots were randomly allocated to 6 triads.
2. Participants were verbally briefed on the context for elicitation and the experimental protocol. They were asked to assume the role of commander for one of the two blue Battle Groups (as indicated on the written briefs). They were asked to assess, concurrently, the situations presented to them and, in so doing, answer the question "How are two of the snapshots similar and thereby different from the other one?" They were not directed to follow any specific method for situation assessment.
3. Participants were presented with each triad for 15 minutes. At the beginning of this period, the snapshot overlays were laid out on top of identical maps, together with the corresponding brief, on three separate tables.
4. Each participant was provided with a data sheet for recording similarities and differences against specific snapshot IDs (pairs or singletons). They were advised that there was no limit to the number of responses they might provide.
5. At the end of each 15-minute period the data sheets were collected. The procedure was repeated for each triad and on the completion of elicitation against all 6 triads, the experiment ended.

Note that the experimental procedure described above deviated from the standard application of Kelly's Personal Construct Elicitation in its use of self-completion questionnaires, rather than structured interviews, to elicit constructs. This approach was chosen to increase the sample size, given the fixed amount of time available for experimentation.

Coding the results

Content analysis

The data was analysed for content by the author and codes were developed based on the recurrence of words and themes. This process was necessarily subjective, and was based on perceived references to factors being considered. Although codes were not predetermined, they were doctrinal in nature. The hierarchical coding scheme is illustrated in the coding map below.

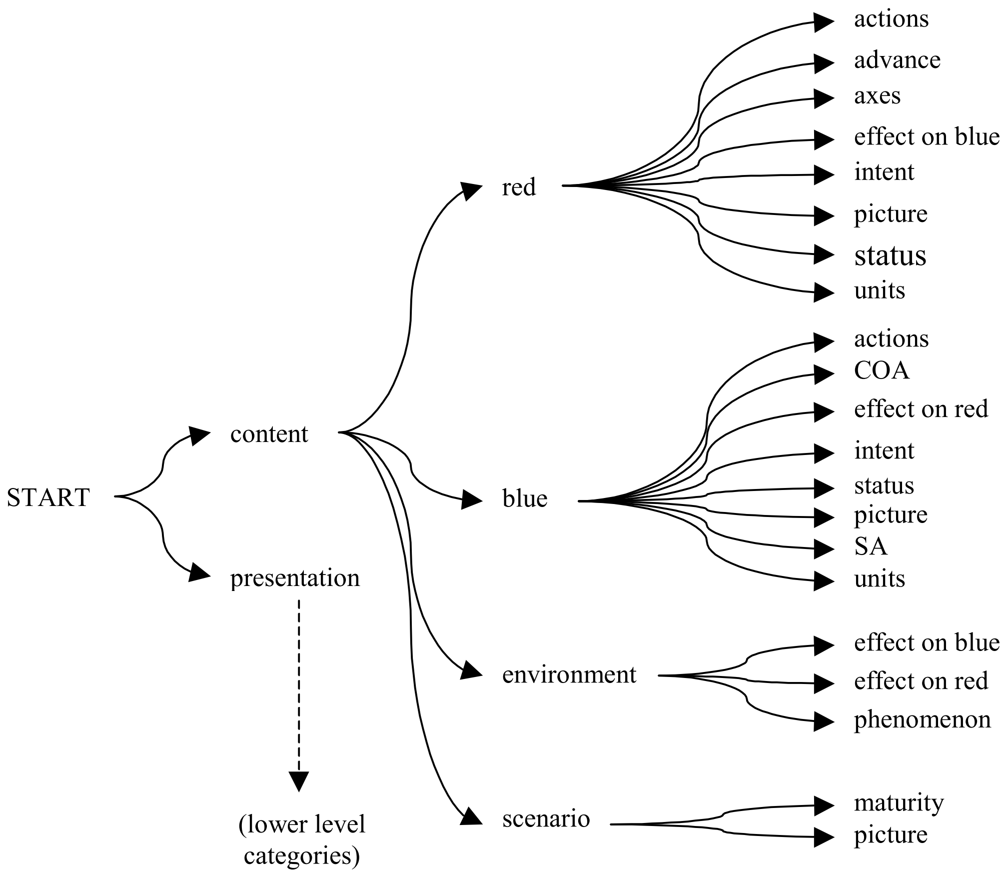

Figure 1; Coding map

The code definitions are provided in Appendix A.

The level 1 codes *content* and *presentation* distinguish between those constructs that describe some aspect of the meaning within the snapshots and those that describe some aspect of its presentation. The latter were of no interest to this study and were excluded from further analysis. Further, 5 participants provided an extremely high proportion of *presentation* responses (between 56% and 67%) and it is believed that they misconstrued the experimental procedure. Consequently, their data (including *content* responses) were also excluded from the analysis.

Classification of constructs by Situation Awareness level

Endsley [2] defines Situation Awareness (SA) as “the perception of elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future.” This definition is expanded to define three levels of SA, each representing a greater level of awareness than the last:

- Level 1 SA – Perception of the Elements in the Environment: the status, attributes and dynamics of relevant elements in the environment;
- Level 2 SA – Comprehension of the Current Situation: a synthesis of disjointed level 1 elements to form a holistic picture of the environment, including the significance of objects and events;

- Level 3 SA – Projection of Future Status: the ability to project the future actions of elements in the environment, based on both level 1 and level 2 SA.

The experimental data may be considered as evidence that level 1, 2 or 3 SA had been attained by the participants. In particular, the constructs identified under the level 1 code *content* in the coding scheme above may be classified by SA level. All constructs except one can be mapped uniquely onto SA level. The remaining construct, *content-blue-SA*, was effectively a meta-construct that participants used to describe their *own* assessments of their SA. It is not, therefore, mapped onto any SA level as it does not constitute direct evidence of SA. Table 1, below, indicates the mapping of constructs onto SA levels.

SA Level	**Constructs**
Level 1 (Perception of elements)	*Content-red-units*
	Content-red-status
	Content-red-actions
	Content-blue-units
	Content-blue-status
	Content-blue-actions
	Content-environment-phenomenon
Level 2 (Comprehension)	*Content-red-effect on blue*
	Content-red-advance
	Content-red-axes
	Content-blue-effect on red
	Content-blue-picture
	Content-blue-COA
	Content-environment-effect on blue
	Content-environment-effect on red
	Content-scenario-maturity
	Content-scenario-picture
Level 3 (Projection)	*Content-red-intent*
	Content-blue-intent

Table 1; Mapping of constructs onto Endsley SA level

Analysis of results

Constructs

The content analysis generated a set of 21 constructs for the Battlegroup scenario across the 10 participants; the constructs themselves constitute the initial set of results and are of interest to the study. They are comparable to the factors included in the Evaluation of Factors stage of the Combat Estimate (conducted at Battlegroup level). British Army doctrine [3][3] provides guidance on the factors to be considered and the elicited constructs are broadly representative of these factors. All of the elicited constructs appear in some form in the list of factors, although some are not stated *explicitly* in [3]. They also expectedly omit those factors that were not present, or purposefully accounted for, within the scenario - Air and Flanks, Combat Service Support (CSS) and Operational Security (OPSEC). Four factors that were *not*, however, elicited as constructs are Relative Strengths, Surprise, Protection, and Deception. The construct perhaps closest to any of these was *red-picture*, which describes the overall pattern of forces in the Area of Operations, yet the associated responses did not describe the picture in terms of 'relative strengths'.

Frequency of constructs by participant

568 responses were provided by the 10 participants entered into the analysis. Participant 1 was forced to withdraw from the experiment early (for reasons unconnected with the experiment) and only provided responses against 4 of the 6 triads. To facilitate comparisons, it was decided to weight the frequency of participant 1's responses by a factor of 1.5, under the assumption that the responses primed by the 4 triads he encountered were a good predictor of those he would have provided across all 6.

The distribution of the 590.5 weighted responses across constructs are shown in table 2 below.

---

[3] This publication was updated in January 2003 to reflect the new '7 Questions' approach to Battlegroup planning. This new approach advocates that an *explicit* Evaluation of Factors stage is not conducted, in favour a more *implicit* consideration of factors in answering 7 general questions; yet the factors that require consideration are the *same* as those provided here. Importantly, the publication referenced here was current at the time of the experiment, and is certainly the doctrine that the participants had been taught to follow.

Construct	Frequency	%	%
Content-red-units	101	17.1	69.01
Content-red-status	13	2.202	
Content-red-actions	54	9.145	
Content-red-effect on blue	21.5	3.641	
Content-red-advance	24	4.064	
Content-red-axes	122	20.66	
Content-red-picture	43	7.282	
Content-red-intent	29	4.911	
Content-blue-units	11	1.863	21.17
Content-blue-status	4	0.677	
Content-blue-actions	49.5	8.383	
Content-blue-effect on red	10	1.693	
Content-blue-picture	8	1.355	
Content-blue-COA	28	4.742	
Content-blue-intent	3	0.508	
Content-blue-SA	11.5	1.948	
Content-environment-phenomenon	13.5	2.286	7.028
Content-environment-effect on blue	12.5	2.117	
Content-environment-effect on red	15.5	2.625	
Content-scenario-maturity	14.5	2.456	2.794
Content-scenario-picture	2	0.339	
Content total	590.5	100	100

Table 2; Frequency of responses by construct

Table 2 shows that *red* constructs were the most dominant, accounting for 69% of all responses. *Red-axes* (20.66%), *red-units* (17.1%) and *red-actions* (9.1%) were the most commonly referenced constructs, with *blue-actions* (8.4%) the next most popular construct. These four constructs accounted for 326.5 of the 590.5 responses (55.3%), suggesting that a relatively small set of constructs account for most of the salient characteristics of the scenario.

Grouping responses by both construct and participant, we find that only 3 of the 21 constructs were common to all participants (*red-units*, *red-axes* and *blue-actions*). Further, four of the participants dominated on a specific *pair* of constructs that accounted for approximately half of their *content* responses. They were:

- Participant c: *red-axes* (27.8%) and *red-picture* (19.4%);
- Participant f: *red-axes* (28.4%) and *red-actions* (21%);
- Participant j: *red-units* (53.1%) and *red-axes* (18.8%);
- Participant n: *red-axes* (35.4%) and *red-units* (17.7%).

Participant l was the only one who provided more instances of *blue* constructs than *red* and dominated on *blue-COA*. The remaining five participants provided a fairly even spread of responses across the 21 constructs, although there were many non-*red* constructs that were not referenced by each participant. All participants referenced *red*, *blue* and *environment* constructs, except participant j who provided no *environment* or *scenario* constructs.

A full table of results by participant can be found in Appendix B.

SA constructs by experience level

When the constructs were recoded by SA levels and compared with participants' experience levels, an interesting pattern emerged. Table 3 details the frequencies of constructs by participant, this time using the Endsley SA level coding. Note that the total number of (weighted) constructs is less than in table 2 because *blue-SA* was omitted from this coding scheme. The experience levels have been recoded from the original categorisation to 1=no experience, some experience; 2=experienced, considerable experience. The mean SA level has also been calculated for each participant, weighted by the frequency of responses at each level. This mean figure provides a measure of the depth of processing, with respect to SA, that each participant was conducting.

Participant	**Experience Level**	**Frequency**			**Total**	**Mean SA level (weighted)**
		Level 1 (Perception of elements)	**Level 2 (Comprehension)**	**Level 3 (Projection)**		
c	2	15	49	7	71	1.8873
e	2	16	20	6	42	1.7619
f	1	36	41	4	81	1.6049
g	2	20	28	7	55	1.7636
i	1	27	24	3	54	1.5556
j	1	42	22	0	64	1.3438
k	2	15	16	0	31	1.5161
l	2	15	42	3	60	1.8
n	2	42	34	2	78	1.4872
p	1	18	25	0	43	1.5814
Total		246	301	32	579	

Table 3; Experience levels and frequency of responses by SA level

The correlation between experience level and mean SA level (weighted) is r=0.557, p=0.094, which is significant at the 10% level.

Interpretation of results and validity

The use of self-completion questionnaires rather than structured interviews was chosen on the basis participant availability; given the time available, it ensured a larger sample size. A disadvantage of this approach is that participants were not challenged to develop the constructs they initially provided and so might have tended either to repeat them or revert to doctrinal factors. (Use of doctrinal factors is a key aim of training of commanders.) Further, it is believed that the deviation of the experimental procedure from any recognised Battle Group situation assessment process affected the type of response that some participants gave. Fundamentally, assessing three different situations concurrently was not something that any of the participants had previously been asked to undertake. This may explain why some participants provided a majority of *presentation* responses. Both these issues impact upon the validity of the elicitation process.

The coding process unavoidably introduced a degree of subjectivity into the analysis. Although constructs were not coded with explicit reference to doctrinal factors, it is recognised that the author focused on the detection of factor-like constructs within the responses, as this was the aim of the study. Hence a doctrinal structure was imposed upon the results. In the context of this study, the relative frequencies of constructs have been taken as a measure of their relative importance in assessing the snapshots. This measure should be treated with caution, however, not least for reasons given in the previous argument about construct repetition.

The elicited constructs reflect those aspects of participants' mental models activated by the Battlegroup scenario only. The fact that the elicited constructs did not include any that would not be expected given the scenario characteristics argues, in part, for the validity of the method. Four doctrinal factors were expected but not present in the elicited constructs. These were Relative Strengths, Surprise, Protection, and Deception. This may be partly due to participants' levels of immersion in the scenario; it is extremely difficult to create realistic operating conditions in an experimental setting. It is not known exactly how a lack of immersion may manifest itself in terms of elicited constructs, although it should be noted that the factors unaccounted for require level 2 and level 3 SA (by Endsley's definitions). Assuming that a lack of immersion disrupts the attainment of increasing levels of SA, these types of factors would be exactly those that would be less likely to be taken into consideration and therefore be elicited as constructs.

The pattern of frequencies in table 2 may partly be explained by the scenario itself. The timing and focus of blue's counterattack was dictated largely by red's chosen COA and the state of their advance, whilst blue units remained fairly static until the conditions for attacking were met. Playing one of the two blue Battlegroup commanders, the participants also had a relatively good situation awareness of blue forces involved in the battle. This would explain the predominance of red constructs among the set of responses. Given the nature of the terrain and the forces available, red's COAs were dictated largely by their chosen axis/ axes of approach. This would explain the high frequencies of both *red-units* and *red-axes* constructs among the responses, with the former encapsulating assessments of the locations of individual red units/ sub-units and the latter providing assessments of the red COA.

It could also be argued that *blue* and *environment* constructs were less common since there was only a single blue COA and the participants had gained an understanding of both the terrain and likely blue forces through the brief and the map-based exercise. Further, blue forces remained relatively static until the conditions for attacking red were met, and were less dynamic than red across the scenario as a whole. This would explain why *blue-actions* were more prevalent than *blue-units*.

One of the more surprising aspects of this table is the low occurrence of *status* constructs (essentially those referring to combat effectiveness). As with relative strengths, combat effectiveness is widely regarded as a key determinant of the situation, especially at the tactical level. The scenario *may* offer an explanation in both cases, since blue are not concerned with rendering red non-combat effective and therefore may be far more concerned with red dispositions. However, blue combat effectiveness is key to blue's ability to succeed in their mission, since they need to remain an effective fighting force for a set period of time in order to slow red down.

The construct frequencies for the whole group indicated that 4 of the constructs accounted for over half of the responses. 3 of these 4 constructs were common to all participants, accounting for 272.5 (46%) of all responses. This is an interesting result, since it suggests that British Army commanders' mental models are dominated by a small set of common constructs, but can only be applied to the scenario in question.

It is the individual differences in construct systems that are of most interest to the study. The results suggest that there is a great deal of variability across the 10 participants, with 4 dominating on two constructs each (accounting for approximately 50% of their *content* responses), and 5 with a fairly even spread across the 21 constructs. Further, the minimum number of constructs referenced by any one participant was 8 and the maximum was 16. These results suggest that the participants possessed different *types* of construct systems, although it must be conceded that the dominance of one or two constructs may be due to repetition in the elicitation process (see above).

The significant correlation between experience and mean SA level suggests that participants with greater experience possess mental models dominated by higher level constructs (in terms of Endsley SA levels) than those with less experience. It also suggests that, for a given situation, more experienced commanders can attain better Situation Awareness than less experienced commanders.

The results cannot be readily generalised either across British Army commanders or situations, due to the small sample size and the limits of the scenario, respectively.

Summary and recommendations

The potential worth in understanding British Army commanders' mental models in terms of the factors they actually use, and how these differ across people and situations, is in the design of information systems that provide increased benefit to the situation assessment process. An assumption of the research is that expert commanders make good decisions because they are able to achieve a high level of SA, and that they have learnt to do this by implicitly choosing a set of constructs that optimally (or sub-optimally) characterise the situations they encounter. The main aim of the research was to explore the construct systems that British Army commanders possess. The results indicate that although there is a common set of constructs for the experimental scenario across the participating commanders, there are notable individual differences in the number and distribution of constructs used, and their relative importance, as measured by the frequency with which they were accessed. The most interesting result is that commanders with differing levels of experience use different sets of constructs in situation assessment.

The relative importance of specific constructs would be better measured, however, in other ways. Kelly's Personal Construct Elicitation is a precursor to Kelly's Repertory Grid technique that, for a set of elicited constructs, asks individuals to rate the degree to which constructs are 'activated' by specific elements. Having elicited the constructs within this study, a sensible next step would be to elicit repertory grids from a set of commanders against these constructs for same scenario. Further, SA itself is only of benefit in its capacity to prime effective decisions; another suggested method for measuring the relative importance of constructs is through their relationship with selected COAs, the measure itself dependent on the perceived effectiveness of those COAs.

The research must also be judged on the validity of the methods used, particularly in the command decision-making domain. Kelly's Personal Construct Elicitation is recognised as a valid method outside this domain and questions remain about its applicability here given the necessary complexity of elements. Further, the use of self-completion questionnaires impact on the validity of the results; in future experimentation the method of data collection may be dropped in favour of one-to-one interviews. Finally, it is recognised that the constructs elicited here are a product of both the participants' mental models and the scenario. The use of different scenarios is required so that the results may be interpreted in a wider context.

References

1. A Manual for Repertory Grid Technique. Fransella, F and Bannister, D, 1977, Academic Press.
2. Toward a Theory of Situation Awareness in Dynamic Systems, in Human Factors Vol. 37, No. 1. Endsley, M, 1995.
3. Army Field Manual Volume 1 - Combined Arms Operations Part 2: Battlegroup tactics, Army Code No 71648. 1998, British Army.

Appendix A: Code definitions

Code	Definition
Presentation	Responses that do not relate to the meaning within the snapshots themselves, but to aspects of the presentation of the situations in snapshot format. These include references to overlays and symbology. *Presentation* responses were not subjected to any further categorisation since they are not of interest to this study.
Content	Reponses that encapsulate an assessment of the situations contained within the snapshots.
Content-red	All responses under the *content-red* code are direct assessments of the enemy (red forces). These are characterised by direct references to the enemy/ red force, or to specific red units or functions in combat. Most of these refer only to red forces although some discuss cause-and-effect relationships involving red forces and either blue forces or the environment. Where such effects are caused by red, they are coded under *content-red*; where they affect red, they are coded under either *content-blue* or *content-environment* as appropriate (see below).
Content-red-units	Responses that directly identify/ locate specific red units/ sub-units, but neither assess their status or activity, nor describe them in the context of the entire red force. Also includes identification/ location of man-made (red) obstacles.
Content-red-status	Responses that directly reference the combat effectiveness of the red force (at any level of aggregation from unit upwards); also responses that indicate that part of the red force has been destroyed (i.e. non-combat effective).
Content-red-actions	Responses that make reference to red activity, either by linking a type of activity (e.g. defending) to a specific unit or location, or by describing an perceived event instigated by red forces (e.g. helicopter landing of ground forces). Perceived future events are not included under this code, but under *content-red-intent*. No response coded in this way makes reference to effects, either recent or intended (these are coded under *content-red-effect on blue*).
Content-red-effect on blue	Responses that describe a perceived red-instigated action that both involves blue forces and has an effect on blue. Whilst not all responses *directly* describe the effect on blue (e.g. 'red have broken through') they all have

Code	Definition
	a clear implication for the blue force.
Content-red-advance	Reponses that describe the state of red's geographical progress, either comparatively or with reference to a specific location or blue Area of Responsibility (AOR).
Content-red-axes	This code covers all responses with a direct reference to red's axes of advance at unit or force level. Such references include existence of axes, number of axes, direction/ general area and named route. All responses providing perceived future axes are excluded from this code, included instead in *content-red-intent*.
Content-red-picture	Responses that described the layout of red forces on the terrain. Some responses described the level of understanding of the red picture, whereas most described some aspect of the geometry of the red picture in specific areas or at force level. Responses included here rather than in *content-red-axes* described the width of the front that red were advancing on without referencing axes.
Content-red-intent	All responses that either convey a level of understanding of red's likely future actions, or describe them. *Intent*, in this sense, includes both perceived immediate objectives and wider intent. Note that perceived future red actions, desired effects on blue and axes are included here rather than under other codes defined above.
Content-blue	All responses under the *content-blue* code are direct assessments of own forces (blue). These are characterised by direct references to own forces/ blue, or to specific blue units or functions in combat. Most of these refer only to blue forces although some discuss cause-and-effect relationships involving blue forces and either red forces or the environment. Where such effects are caused by blue, they are coded under *content-blue*; where they affect blue, they are coded under either *content-red* (see above) or *content-environment* (see below) as appropriate. Symmetric with *content-red*.
Content-blue-units	Responses that directly identify/ locate specific blue units/ sub-units, but neither assess their status or activity, nor describe them in the context of the entire blue force. Also includes identification/ location of man-made (blue) obstacles. Symmetric with *content-red-units*.

Code	Definition
Content-blue-status	Responses that directly reference the combat effectiveness of the blue force (at any level of aggregation from unit upwards); also responses that indicate that part of the blue force had been destroyed (i.e. non-combat effective). Symmetric with *content-red-status*.
Content-blue-actions	Responses that make reference to blue activity by linking a type of activity (e.g. defending) to a specific unit or location. Perceived future events are not included under this category, but under *content-blue-intent*. No response coded in this way makes any reference to effects, either recent or intended (these are coded under *content-blue-effect on red*). Symmetric with *content-red-actions*.
Content-blue-effect on red	Responses that describe a perceived blue-instigated action that both involve red forces and have an effect on red. Symmetric with *content-red-effect on blue*.
Content-blue-picture	Responses that describe the layout of red forces on the terrain. Some responses describe the level of understanding of the blue picture, whereas most describe some aspect of the geometry of the blue picture in specific areas or at force level. Symmetric with *content-red-picture*.
Content-blue-intent	All responses that either convey a level of understanding of blue's likely future actions, or describe them. *Intent*, in this sense, includes both perceived immediate objectives and wider intent. Note that perceived future blue actions and desired effects on red and axes are included here rather than under other codes defined above. Symmetric with *content-red-intent*.
Content-blue-SA	Responses that provide context for situation assessments encapsulated in other responses. Includes references to pieces of information gleaned from the snapshots that are specified as pivotal (confirmatory or confusing) in understanding the situations presented.
Content-blue-COA	These responses comment either on the state of the current blue Course of Action (COA); describe the urgency for a change in COA; or specify intelligence requirements for, or effects to be realised within, future Courses of Action.

Code	Definition
Content-environment	All responses under the *content-environment* code are direct assessments of aspects of the environment (predominately terrain and weather). These are characterised by direct references to environmental conditions. Some discuss cause-and-effect relationships, i.e. *environmental* causes of effects to either red or blue forces.
Content-environment-phenomenon	Responses providing a direct assessment of environmental phenomenon (terrain and weather) without describing any effect on either blue or red forces.
Content-environment-effect on blue	Responses describing the effect of some environmental phenomenon on blue forces.
Content-environment-effect on red	Responses describing the effect of some environmental phenomenon on red forces.
Scenario	All responses under the *content-scenario* code provide macroscopic assessments of the situations that cannot be uniquely attributed to either red forces or blue forces.
Scenario-maturity	Responses that made direct reference to the state of development of the battle, or to the scenario time.
Scenario-picture	Responses that describe the layout of all forces on the terrain. Some responses describe the level of understanding of the combined picture, whereas most describe some aspect of the geometry of the combined picture.

Table A.1; Code definitions

Appendix B: Frequency of constructs by participant

Construct	Participant										Total
	c	e	f	g	i	j	k	l	n	p	
Content-red-units	5	11	6	7	4	34	8	9	14	3	101
Content-red-status	0	1	0	1	5	1	0	0	1	4	13
Content-red-actions	3	0	17	5	8	1	1	3	9	7	54
Content-red-effect on blue	3	0	5	2	4	4	1	1.5	0	1	21.5
Content-red-advance	7	0	2	3	1	2	0	3	1	5	24
Content-red-axes	20	6	23	9	4	12	4	9	28	7	122
Content-red-picture	14	5	1	5	4	2	7	0	2	3	43
Content-red-intent	7	4	3	7	3	0	0	3	2	0	29
Content-blue-units	1	2	0	0	0	3	0	0	5	0	11
Content-blue-status	0	0	1	0	3	0	0	0	0	0	4
Content-blue-actions	5	1	7	7	7	3	6	1.5	9	3	49.5
Content-blue-effect on red	1	0	2	1	0	1	0	0	0	5	10
Content-blue-picture	0	1	0	0	1	0	0	6	0	0	8
Content-blue-COA	1	2	1	2	2	1	0	18	1	0	28
Content-blue-intent	0	2	1	0	0	0	0	0	0	0	3
Content-blue-SA	1	0	0	1	1	0	0	7.5	1	0	11.5
Content-environment-phenomenon	1	1	5	0	0	0	0	1.5	4	1	13.5
Content-environment-effect on blue	2	3	0	0	3	0	2	1.5	0	1	12.5
Content-environment-effect on red	0	3	2	5	3	0	0	1.5	0	1	15.5
Content-scenario-maturity	1	0	5	1	2	0	0	1.5	2	2	14.5
Content-scenario-picture	0	0	0	0	0	0	2	0	0	0	2
Content (total)	72	42	81	56	55	64	31	67.5	79	43	590.5

Table B.1; Frequency of constructs by participant