18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Consensus: a Comprehensive Solution to the Grand Challenges of Information Fusion

Dale A. Lambert, Adam Saulwick, and Kerry Trentelman

***Abstract*—The information fusion community faces at least five grand challenges. In this paper we describe how our revolutionary implementation of an automated high level information fusion system addresses these grand challenges. Our innovative system processes real-time heterogeneous information sources, including track data, as well as spoken and written English language, transforming all inputs into a rich canonical semantic form for deep automated reasoning. Users can engage in real-time Question Answering with Virtual Advisers and a Virtual Battlespace using spoken and written English and haptic devices. This solution offers the prospect of comprehensive situation awareness.**

***Index Terms*— Avatars, Computational Linguistics, Computer interfaces, Data Fusion, Information Fusion, Intelligent control, Knowledge representation, Multi-agent Systems, Natural Language Processing.**

I. Introduction

OVER a decade ago, Lambert [1] wrote a paper entitled "Grand Challenges for Information Fusion". The grand challenges identified were:

- **System Challenge**: How should we manage information fusion systems formed from combinations of people and machines?
- **Paradigm Challenge**: How should the interdependency between the sensor fusion and information fusion paradigms be managed?
- **Semantic Challenge**: What symbols should be used and how do those symbols acquire meaning?
- **Epistemic Challenge**: What information should we represent and how should it be represented and processed within the machine?
- **Interface Challenge**: How do we interface people to complex symbolic information stored within machines to provide decision support?

The following sections outline the implemented "Consensus" solution for these grand challenges of information fusion. The presentation accompanying the paper features a video of a live demonstration of human interaction with the Consensus system as it is processing both real-time track data and English text. A variant of the North Atlantis scenario [2] is used to demonstrate Consensus. The complex scenario involves Redland attempting a munitions resupply after seizing the Camrien Peninsula from Blueland, while Blueland is concurrently forming a military convoy to retake the Camrien Peninsula. This is supplemented by other military surface vessels, submarines, military aircraft, ground based radars, commercial shipping, commercial airliners and fishing boats, all going about their business.

Date submitted 26-Feb-15.
The authors are with the Defence Science and Technology Organisation, Edinburgh, Australia (e-mail: firstname.lastname@dsto.defence.gov.au).

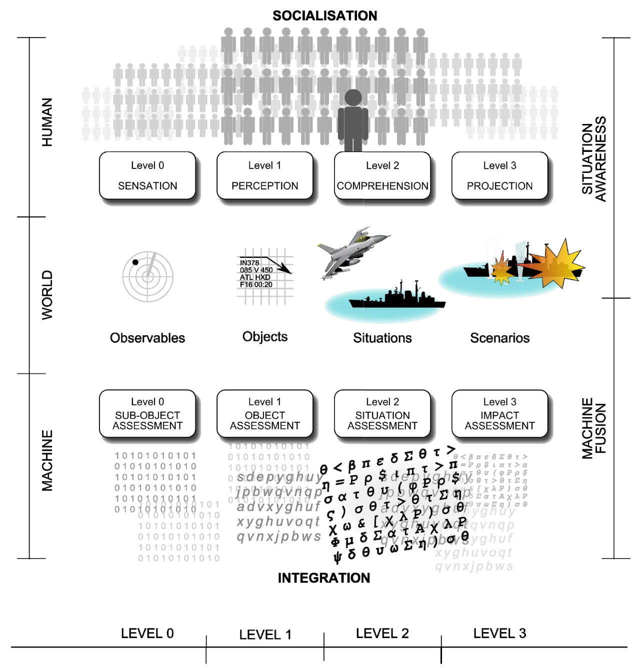

Fig. 1 Consensus fusion system schematic.

II. The System Challenge

In 2001 Lambert [3] noted the direct parallel between the levels in Endsley's prominent account of *human situation awareness* [4] and the levels in the dominant Joint Directors of Laboratories (JDL) *machine data fusion* model [5-9]. JDL level 0 *sub-object assessment* is about the *sensation* of observables; JDL level 1 *object assessment* is about the *perception* of objects; JDL level 2 *situation assessment* is about the *comprehension* of situations as relations between objects; and JDL level 3 *impact assessment* is about the scenario *projection* of relations between objects to expose consequences. Lambert concluded,

> *Thus situation awareness can be understood as the human counterpart to machine data fusion, while data fusion can be conceived as the machine counterpart to situation awareness* [10].

A complete fusion system needs to integrate both human situation awareness and machine data fusion, including a means of interfacing between the two. Fig. 1 illustrates the desired integration. The Consensus fusion system delivers consensus between humans and machines through the seamless horizontal and vertical integration of all of these components. This addresses the System Challenge by identifying how we should manage information fusion systems formed from combinations of people and machines. A general protocol for agent interaction has also been developed, be those agents human or machine [11], but it is not used in the current Consensus system. Fig. 2 shows the implemented Consensus fusion system in operation.

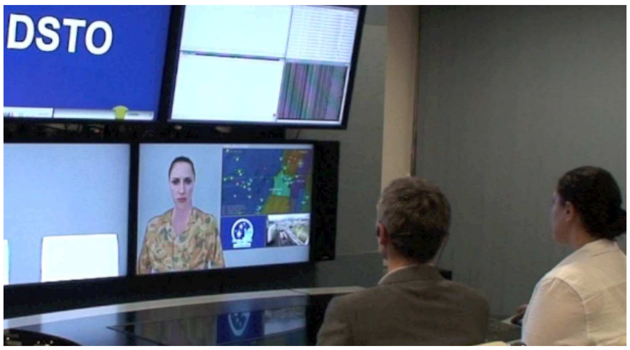

Fig. 2 Consensus fusion system in operation.

III. THE PARADIGM CHALLENGE

To provide a unifying paradigm for data fusion, Lambert proposed the State Transition Data Fusion (STDF) model [10, 12, 13]. Under STDF, each JDL level characterizes the world in terms of transitions between states, with increasing JDL levels being associated with increasingly more sophisticated concepts of "state". Level 0 is a world of observables represented by transitions between feature vector states; level 1 is a world of objects represented by transitions between state vectors; level 2 is a world of situations represented by transitions between states of affairs; and level 3 is a world of scenarios represented by transitions between situation states. Fig 1 depicts these states for each JDL level. The STDF model also asserts that the same basic fusion process applies at each JDL level. Fig 3 depicts the general form of that basic fusion process. It aims to *explain* the world through *prediction* and *observation*, and therefore also characterizes all science as applications of data fusion.

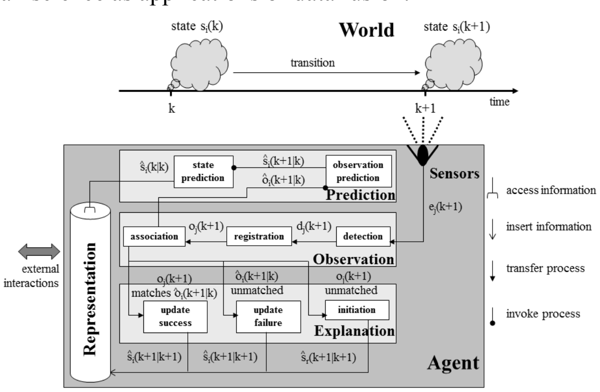

Fig. 3 General STDF Model applied at each JDL Level.

The Paradigm Challenge is met by having a common paradigm across all the JDL levels and then by having some of the levels function as capability elements within other levels (see section V below).

IV. THE SEMANTIC CHALLENGE

If the information sources are broadly categorized as signal, textual or imagery, then lower-level (level 0 and level 1) fusion often relies on separate signal, textual and image STDF implementations to identify observables at level 0 from which separate signal tracking, parsing and image tracking STDF implementations respectively identify objects at level 1. Multi-source lower-level fusion is now starting to become more prevalent. By the time we move to higher-level (level 2 and level 3) fusion, however, the requirement for multi-source fusion becomes overwhelming, since we are seeking an explanation of a situation in the world, not separate signal, textual and imagery explanations of observables or objects in the world. Moreover, since level 2 explanations involve relations between objects, some form of *symbolic* representation is required for level 2 explanations. The most economical solution is to adopt the *same canonical symbolic representation* for all higher-level fusion explanations. In general these symbolic representations will need to be quite expressive, reporting not only positional and kinematic information, but also information that is geographical, biographical, sociological, political, economic, *et cetera* in nature.

The Mephisto Semantic Framework [1, 14-17] was developed as a *canonical symbolic representation* for all higher-level fusion explanations. It is implemented in the Consensus solution. Mephisto involves five layers of conceptualisation. The *metaphysical layer* includes concepts like existence, space and time; the *environmental layer* has concepts like air, water surface and submerged land; the *functional layer* supports concepts like move, strike and sense; the *psychological layer* provides concepts like belief, perception and anger; and the *social layer* delivers concepts like authority, ownership and conflict. DSTO is expanding the social layer into formalising social influence [18]. It is believed that this will increase the capacity of Consensus to fuse a variety of soft data sources. The conceptual aspect of the Semantic Challenge is illustrated by the first column of Fig. 4.

The Mephisto framework is called a "semantic" framework because it is designed to allow the machine to operate in accordance with the meaning of these concepts. This is achieved through a three step process, the first being to *conceptualise*. The second step is to *formalise* each layer, represented by the second column in Fig. 4.

Each of the Mephisto concepts is represented in a symbolic form to produce a formal language with expressions like **exists**(x), **water_surface**(x), **strikes**(x, y), **believes**(X, α) and **authority**(X, Y, α). This addresses the first part of the Semantic Challenge – "what symbols should be used".

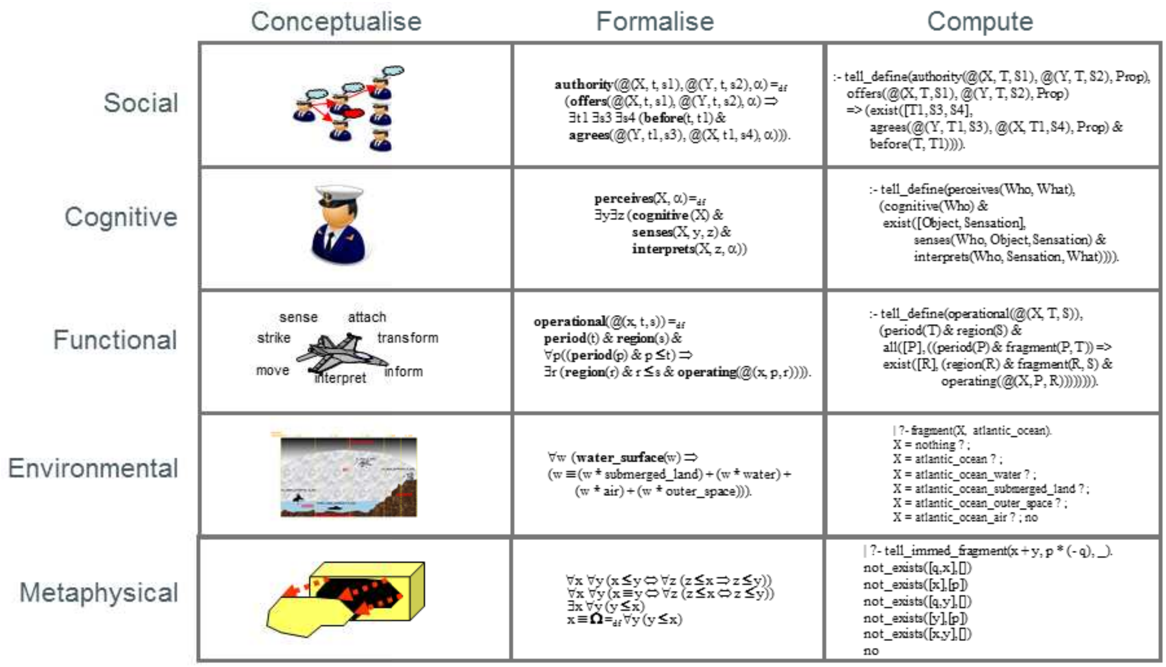

Fig. 4 Mephisto Semantic Framework.

The remaining part of the Semantic Challenge is "how do those symbols acquire meaning". This is achieved by extending the formal language of Mephisto symbols $\mathbb{L}$ into a formal logic <L, $\models$> through the introduction of a truth preserving inference relation, relation $\models \subseteq \mathrm{P}(\mathbb{L}) \times \mathrm{L}$, for powerset $\mathrm{P}(x) = \{u \mid u \subseteq x\}$, and by then constraining the allowable interpretations of those symbols by specifying axioms and definitions. For example, Fig. 3 contains the formal axiom $\exists x\ \forall y\ (y \leq x)$ to say that there is a fragment of the world that contains all fragments of the world, which is then defined as the universe Ω by $x \equiv \Omega =_{df} \forall y\ (y \leq x)$. By mathematically imposing inferential truth conditions on the symbols in this way, the symbols acquire meaning [15].

Mephisto is a perdurant ontology, thus everything in the world is treated as a process [19]. We use late-alphabet letters *x*, *y* and *z* to denote processes. Processes occupy spatiotemporal regions. We use the letters *t* and *s* to denote temporal and spatial projections of processes, respectively. Every process can be thought of as a four-dimensional space-time chunk. We use mereological product to consider segments of these chunks. The relational assertion @(*x*, *t*, *s*) represents the process *x* at time *t* at space *s*. We use mereological part-hood to denote parts of space-time chunks. The relational assertion $x \leq y$ states that process *x* is a part of process *y*. The spatial components of a process obey the axioms of the region connection calculus [20], whereas temporal components obey Allen's interval algebra [21].

The final step is to implement the set of Mephisto symbols $\mathbb{L}$ within a machine and to implement the inference relation $\models$ as an inference engine. This *compute* step is illustrated in the third column of Fig. 4. This transforms the mathematical imposition of truth conditions on Mephisto symbols into a computation imposition of truth conditions on Mephisto symbols. As a consequence the machine is constrained to compute in accordance with the truth conditions associated with the symbols, and thus computationally respects the meaning of those symbols. The result is a machine that will correctly reason with concepts like existence, water surface, strikes, belief and authority.

V. The Epistemic Challenge

Fig. 5 depicts the STDF model for JDL level 2 with states as states of affairs. In Consensus the JDL level 2 STDF model is implemented through the ATTITUDE Multi-Agent System [15, 22, 23]. The ATTITUDE Psychological Model, illustrated

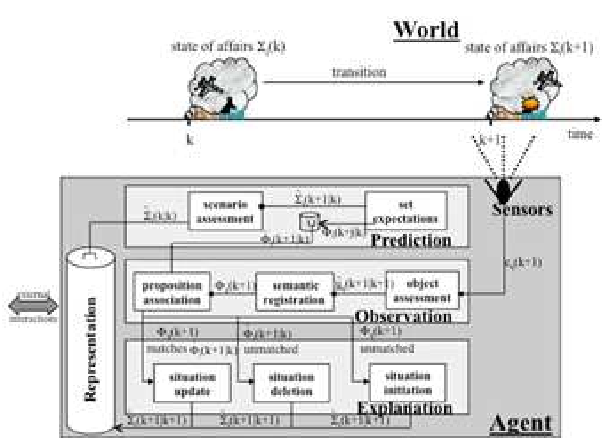

Fig. 5 JDL Level 2 STDF Model.

in Fig. 6, provides the basis for the implemented agents, with Mephisto as the agents' "language of thought".

Folk Psychology ascribes mental states to cognitive individuals as beliefs, expectations, hopes, *et cetera*. These mental states are termed *propositional attitudes* because in everyday language they are represented by *propositional attitude expressions* of the form <subject> <attitude> that <proposition-expression>. Fred believes that the sky is blue is a sample propositional attitude expression. In a propositional attitude expression: the subject, e.g. Fred, expresses which individual has the propositional attitude; the propositional expression, e.g. the sky is blue, expresses some assertion about the world; and the attitude, e.g. believes, expresses the kind of response the subject has toward the proposition. The ATTITUDE Psychological Model uses formal propositional attitudes as a basis for programming. The subject, e.g. Fred, identifies which agent is being referenced; the attitude, e.g. believe, identifies a kind of internal memory; and the propositional expression, e.g. blue(sky), identifies Mephisto content stored in that memory. The formal *propositional attitude instruction* believe(Fred, blue(sky)) instructs agent Fred to store the propositional content blue(sky) in its belief memory.

An ATTITUDE agent's sensors can be added dynamically and can be distributed across platforms. Interaction processes and memories are used for sensor and effector processing. The sub-object and object fusion processes and memories in Fig. 6 provide the "object assessment" capability shown in Fig. 5. They produce detection state vectors at time step k+1. A sample track related state vector is $\underline{\hat{u}}_q(k+1 \mid k+1)$ =
<t_821, 7200665, 40193.1, -108826, -215.141, 209.048, 1.004414, -0.465357, 0, 30627.3, -4925.19, 1326.65, -154.695, -4925.19, 41766.1, -154.19, 1676.42, 1326.65, -154.19, 116.128, -6.58736, -154.695, 1676.42, -6.58736, 130.768>.

The "semantic registration" process in Fig. 6 then provides the "semantic registration" capability in Fig. 5 by converting each state vector $\underline{\hat{u}}_q(k+1 \mid k+1)$ into a set of Mephisto propositional perceptions $\Phi_q(k+1)$ reporting position, speed, course, environment and allegiance. The previous $\underline{\hat{u}}_q(k+1 \mid k+1)$ generates perception $\Phi_q(k+1)$ =
{at(t_821, timestamp(2001, 6, 16, 13.0 ,45.0 ,51.28),

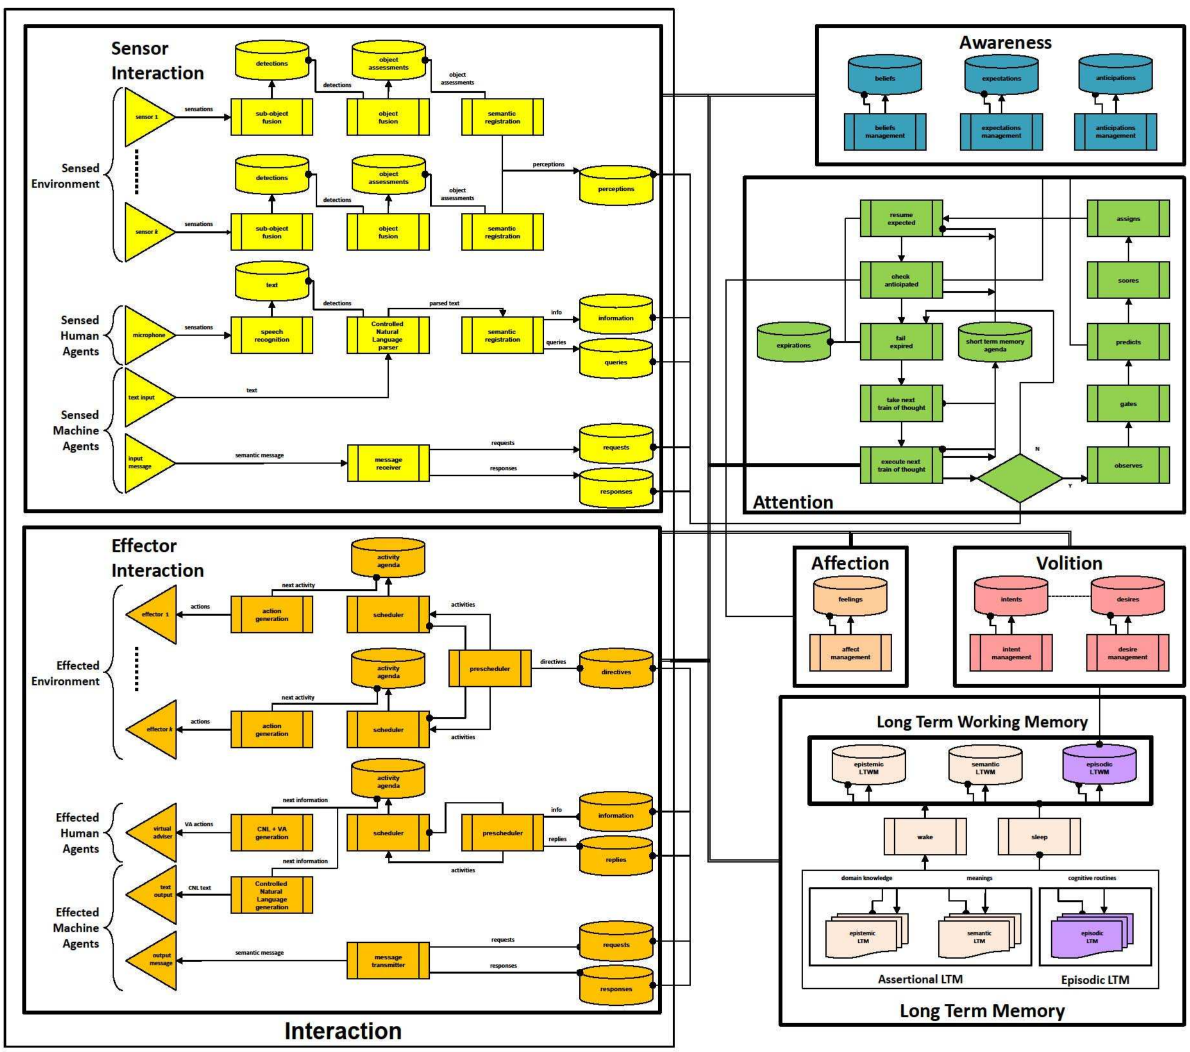

Fig. 6 ATTITUDE Psychological Model.

coordinate(radians(0.9873294202374645), radians(0.4539469822836761), (metres(0.0), metres(100000.0)))), speed(@(t_821, timestamp(2001, 6, 16, 13.0, 45.0, 51.28), _1865), metres_per_second(299.9778594913298)), course(@(t_821, timestamp(2001,6,16,13.0,45.0,51.28), _1880), radians(2.3418315976755615)), in_air(@(t_821, timestamp(2001,6,16,13.0,45.0,51.28), coordinate(radians(0.9873294202374645), radians(-0.4539469822836761), metres(5.0E+04))), celtic_sea_ext*redland_region), unknown_allegiance(@(t_821, timestamp(2001,6,16,13.0,45.0,51.28), _3318)), [[40193.1, -108826.0, -215.141,209.048], [40216.767919900805, -109061.98584052833, 213.95539528911104, 204.80762301955747], …]}.

Sensors for speech and written text are similarly semantically registered into a Mephisto propositional form. For example, the written sentence $\underline{\hat{u}}_q(k+1 \mid k+1)$ = The missile hit the frigate is semantically registered as $\Phi_q(k+1)$ = {frigate(@(skc0002, t_0003, s_0002)), missile(@(skc0001, t_0003, s_0001)), before(t_0003, invl(timestamp(2001,6,16,13.0,45.0,51.38), timestamp(2001,6,16,13.0,45.0,51.28))), hits(@(skc0001, t_0003, s_0001), @(skc0002, t_0003, s_0002))}.

Sub-object words are detected by a lexical analyser and the detected words are tracked by a parser to form sentence objects [10].

The Epistemic Challenge states "What information should we represent and how should it be represented and processed within the machine?" The Consensus response is determined by the *background information* of the long-term memory depicted in Fig. 6 and the *contextual information* that arises in various working memories as an agent interacts with its environment. Each ATTITUDE agent's long-term memory is composed of: *semantic memory* that stores the Mephisto axioms and definitions in a computational form; *epistemic memory* that records facts about the world in a computational form; and *episodic memory* housing cognitive routines, each of which specifies a recipe for behaviour as a transition network of propositional attitude instructions.

Execution of a cognitive routine involves the agent attempting to successfully transition from the starting propositional attitude instruction of the network through to one of the final propositional attitude instructions of the network, by performing each propositional attitude instruction encountered. Performing a propositional attitude instruction either succeeds or fails, with success or failure determining which propositional attitude instruction in the network is attempted next, and with performance of a propositional attitude instruction resulting in side effects in the working memories of the agent. The cognitive routines specify the psychological processes that monitor patterns of behaviour in the world and so when operating, deliver the "scenario assessment" capability in Fig. 5. At time step k, $\hat{\Sigma}_i(k+1|k)$ is the predicted state of affairs for the state of affairs $\Sigma_i(k+1)$ of situation Σ_i in the world at time step k+1. A cognitive routine will often periodically halt and load an expectation into awareness memory. The expectation is about what is expected to be perceived in the world according to the pattern of behaviour being monitored. At time step k, expectation $\hat{\Phi}_i(k+1|k)$ is the predicted perception for situation $\Sigma_i(k+1)$ at time step k+1. This provides the "set expectations" capability in Fig. 5.

The execution of each cognitive routine is managed by the Attention process in Fig. 6. It provides the agent's focus of attention and short-term memory as it juggles execution of multiple competing cognitive routines. The "proposition association" capability in Fig. 5 is also handled by the Attention process in Fig. 6. It matches incoming perceptions to current expectations. There are three possible outcomes.

(a) When an *expected perception* occurs, the cognitive routine associated with that expectation $\hat{\Phi}_i(k+1|k)$ resumes execution, but with the new contextual information of the matching perception $\Phi_q(k+1)$. The cognitive routine will then typically update its explanation of the unfolding situation Σ_i as a new set of beliefs $\hat{\Sigma}_i(k+1|k+1)$ in Awareness memory in Fig. 6 about the state of affairs $\Sigma_i(k+1)$, thus providing the "situation update" capability of Fig. 5.

(b) When an *unexpected perception* $\Phi_q(k+1)$ occurs, it may match an anticipation $\hat{\Phi}_i(k+1)$ in the Awareness working memory of Fig. 6. Anticipated perceptions launch execution of a new cognitive routine to monitor a previously unmonitored situation. The launched cognitive routine execution typically records the initial beliefs $\hat{\Sigma}_i(k+1|k+1)$ about the new situation. This delivers the "situation initiation" capability of Fig. 5. Unanticipated, unexpected perceptions are ignored.

(c) When an *unmatched expectation* $\hat{\Phi}_i(k+1|k)$ occurs, it will typically be allowed to persist for a fixed number perception cycles, after which that expectation propositional attitude instruction fails. Depending on the structure of the cognitive routine, this may signal the discontinuation of situation Σ_i in the world, thus capturing the "situation deletion" capability in Fig. 5.

VI. THE INTERFACE CHALLENGE

Natural interfaces are essential components for improving and integrating human and machine situation awareness. They address the Interface Challenge of higher-level fusion [24]. Consensus utilizes three relatively mature high-level interfaces: Lexpresso—a Controlled Natural Language for natural human-machine interaction via spoken and written English [25]; Virtual Advisers—dynamic programmable animated avatars who can report on situations and respond to questions [26, 27]; and a Virtual Battlespace—an interactive 3-D geospatial display [28, 29]. Consensus coordinates all three interfaces in a seamless environment to provide a natural dynamic real-time Question Answering capability over the domain of interest with particular regard to space and time. Crucially, these interfaces are tightly coupled to Consensus's Mephisto canonical semantic representation and ATTITUDE automated reasoning. As a consequence, each interface appears aware of the situation to which it attends. For instance, the Virtual Advisers 'comprehend' situations on which they report; the Virtual Battlespace 'perceives' the movement of vessels through the environment; and Lexpresso 'comprehends' the meaning of natural English questions and answers. Fig. 7 shows Consensus' major processing steps.

Lexpresso transforms ingested English text into Mephisto constructs. Explicit and implicit spatiotemporal information is converted and stored along with referential identifiers to entities sensed and perceived at JDL levels 0 and 1. At this

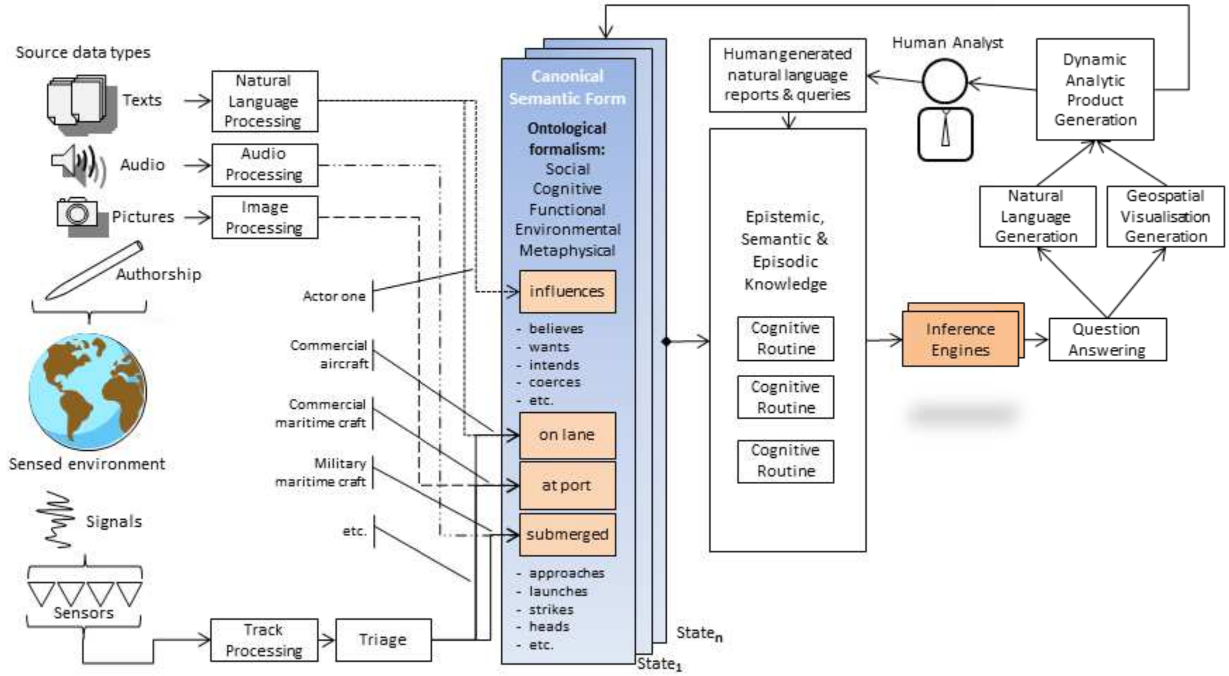

Fig. 7 Consensus major processing steps.

stage indexicals are resolved according to local anaphora resolution algorithms. Ambiguities in natural language are automatically identified and removed via user interaction choice. These processes are designed to ensure that the knowledge base is consistent.

Users can query ingested and stored knowledge for fused and inferred information. On request Consensus can generate automated situation reports on the live activities of various platforms within the North Atlantis scenario, i.e. merchant ships, commercial aircraft, radars, and military aircraft. Situation Reports are generated by sorting information constructed by various cognitive routines from knowledge sourced from sensor inputs. The situation reports are presented in natural English texts, spoken by Virtual Advisers via synthesized text-to-speech, and replayed with events visualized within the 3-D Virtual Battlespace.

VII. An Example from the Scenario

In the Consensus demonstration video Blueland Intelligence has intercepted an email from a person named Paula Lands sent to Redland Intelligence, the enemy. In the email Paula says she is on board a yacht at some coordinates, and that she can see a passenger ferry travelling in a convoy. Blueland Intelligence wants to know whether the convoy Paula sees is their covert Operation Liberty convoy, or the munitions resupply convoy. The following email is ingested into Consensus:

Subject: Convoy sighted
From: Paula Lands
To: Redland command
Sent: Saturday, 26 January 2015 12:41:50
I am on board a yacht at 56.8075N 26.8753W.
I see a passenger ferry in a convoy.
The convoy has a military vessel.
I recognise the ferry as the Sea Queen.
It has 500 soldiers on board.
It is approaching the Camrien Peninsula.

The first sentence after the email header ‘I am on board a yacht at 56.8075N 26.8753W.’ is converted into the following representation:

```
perceive(cnl_sensor,tells([teller(@(skc000002,invl(t
imestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11
,50)),s_000006),Paula_Lands),info_classification(unc
lassified)],[animate(@(skc000002,invl(timestamp(2015
,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_00000
2)),female(@(skc000002,invl(timestamp(2015,1,26,2,11
,50),timestamp(2015,1,26,2,11,50)),s_000002)),inanim
ate(@(skc000003,invl(timestamp(2015,1,26,2,11,50),ti
mestamp(2015,1,26,2,11,50)),coordinate(degrees(56.80
75),degrees(-
26.8753),metres(0.0)))),nouns(@(skc000002,invl(times
tamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)
),s_000002),Paula_Lands,[animate,definite,female,fir
st_name,gendered,last_name,proper_name,singular,inv(
male)]),nouns(@(skc000003,invl(timestamp(2015,1,26,2
,11,50),timestamp(2015,1,26,2,11,50)),coordinate(deg
rees(56.8075),degrees(-26.8753),metres(0.0))),
yacht,[inanimate,indefinite,singular,prep(on),SPACE(
at,coordinate(degrees(56.8075),degrees(-
26.8753),metres(0.0)))]),
verbs(directly_attached,[@(skc000002,invl(timestamp(
2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_0
00002),@(skc000003,invl(timestamp(2015,1,26,2,11,50)
,timestamp(2015,1,26,2,11,50)),coordinate(degrees(56
.8075),degrees(-
26.8753),metres(0.0)))],[pos(2),present,surface(be),
head_verb])]))
```

This representation contains meta-information which registers the time and place of the teller, Paula Lands, and the information classification. Drawing on lexical knowledge Lexpresso identifies the two referents and associates them with particular linguistic properties; Paula Lands with proper_name, animate and female; yacht with inanimate, indefinite, singular.

Lexpresso performs anaphor resolution and resolves the pronoun 'I' from the phrase 'I am on board' with Paula Lands, the teller. The solemnised constants `skc000002` and `skc000003` are introduced to name new processes.

The process `skc000003` at `timestamp (2015, 1, 26, 2, 11, 50)`, and at `coordinate (degrees (56.8075), degrees (-26.8753), metres (0.0))` is asserted to be a yacht.

The representation encodes that Paula Lands is `directly_attached` to the yacht at the specified time. A Mephisto axiom asserts that any process *x* directly attached to any other process *y*, is at the same coordinates as *y*. Hence Consensus infers that Paula Lands is at the same coordinates as the yacht.

The second sentence 'I see a passenger ferry in a convoy' is asserted as two propositions 'a passenger ferry is a member of a convoy' and 'Paula Lands sees a passenger ferry'.

```
[…,member(@(skc000009,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000009),@(skc000010,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000010)),
nouns(@(skc000009,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000009),passenger_ferry,[...]),
nouns(@(skc000010,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000010),convoy,[…]),
verbs(sees,[@(skc000002,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000002),@(skc000009,invl(timestamp(2015,1,26,2,11,50),timestamp(2015,1,26,2,11,50)),s_000009)],[…])]
```

A Mephisto axiom states that if any process *x* sees another process, then *x* is a cognitive process.[1] Another epistemic axiom states that for any cognitive process *x* and any process *y*, if *x* and *y* are less that 22km apart, then *x* can see *y*. From these axioms and the contents of the knowledge base, the system infers that Paula Lands is cognitive, and that if any process is within a 22km radius of 56.8075 degrees north and 26.8753 degrees west, then Paula Lands can see that process. This calculation may require consideration of environmental factors such as the curvature of the earth, altitude of entities and possible visual obstructions. A user can now query Consensus. A query such as ''Can she see it?" is converted into a query representation whereby the pronouns 'she' and 'it' are both resolved to Paula Lands and the Operation Liberty convoy respectively. Consensus can calculate the distance between two coordinates. Blueland Intelligence knows where its Operation Liberty convoy is at any given time. Hence Consensus infers that Paula can see the convoy.

The parser identifies the query as phrased in *dynamic modality* and transforms this into:

[1] We allow for a non-cognitive sensor to 'see' a target, but this another sense of 'see' and is not discussed here.

```
[...([Operation_Liberty_convoy(@(_359739,_359741,_359743)),nouns(@(skc000002,_359741,_359759),Paula_Lands,[animate,definite,female,first_name,gendered,last_name,proper_name,singular,inv(male)]),verbs(can_sees,[@(skc000002,_359741,_359759),@(_359739,_359741,_359743)],[pos(1),q(yes_no),able,surface(see),head_verb])])]
```

or informally 'is Paula Lands able to see the operation liberty convoy'. To answer this question, Consensus first needs to interpret numerous linguistic queues and convert them into symbolic forms. Once confirmed, the variable slots are filled and response generation commences.

The generation process transforms the canonical semantic forms into a natural English language response. The response formulation takes into account appropriate linguistic forms according to expected conventions of the English language, including tense, aspect and mode of verbs; number and definiteness of nouns; grammatical agreement between nouns and verbs, sentence polarity and the scope of any quantifiers.

VIII. Conclusion

We have outlined five grand challenges facing high level information fusion and presented DSTO's implemented innovative solution called Consensus. Consensus is under active development and currently processes real-time heterogeneous information sources, including track data, as well as spoken and written English language. It transforms all inputs into a rich canonical semantic form and performs deep automated reasoning for situation awareness. Users can engage in Question Answering with Virtual Advisers and a Virtual Battlespace using spoken and written English and haptic devices. Consensus provides a revolutionary solution to the grand challenges of information fusion.

Acknowledgment

The authors would like to thank Zhuoyun Ao, Mike Broughton, Takeshi Matsumoto, Chris Nowak, Marcin Nowina-Krowicki, Greg O'Keefe, Martin Oxenham, Steve Wark, Andrew Zschorn, for their significant contributions to developing Consensus.

References

[1] D. A. Lambert, "Grand challenges of information fusion," in *Proceedings of the 6th International Conference on Information Fusion*, 2003, pp. 213-219.

[2] M. Blanchette, "Military Strikes in Atlantis - A Baseline Scenario for Coalition Situation Analysis," The Technical Cooperation Panel TR-C3I-TP-1-2005, 2005.

[3] D. A. Lambert, "Situations for situation awareness," in *Proceedings of the 4th International Conference on Information Fusion*, 2001.

[4] M. R. Endsley, "Design and evaluation for situation awareness enhancement," in *Proceedings of the Human Factors Society 32nd Annual Meeting*, 1998, pp. 97-101.

[5] J. Llinas, C. L. Bowman, G. Rogova, A. N. Steinberg, E. Waltz, and F. E. White, "Revisiting the JDL data fusion model II," in *Proceedings of the 7th International Conference on Information Fusion*, 2004.

[6] A. N. Steinberg and C. L. Bowman, "Rethinking the JDL data fusion levels," in *Proceedings of the National Symposium on Sensor and Data Fusion*, 2004.

[7] A. N. Steinberg, C. L. Bowman, and F. E. White, "Revisions to the JDL data fusion model," in *Proceedings of the joint NATO/IRIS Conference*, 1998.

[8] A. N. Steinberg, C. L. Bowman, and F. E. White, "Revisions to the JDL data fusion model," *Sensor Fusion: Architectures, Algorithms, and Applications,* vol. 3719, pp. 330-441, 1999.

[9] F. E. White, "A model for data fusion," in *Proceedings of the 1st National Symposium on Sensor Fusion*, 1998.

[10] D. A. Lambert, "The State Transition Data Fusion Model," in *High-Level Information Fusion Management and Systems Design*, E. Blasch, É. Bossé, and D. A. Lambert, Eds., Boston: Artech House, 2012, pp. 33-79.

[11] D. A. Lambert and A. G. Lambert, "The Legal Agreement Protocol," in *High-Level Information Fusion Management and Systems Design*, Boston: Artech House, 2012, pp. 173-190.

[12] D. A. Lambert, "A Unification of sensor and higher-level fusion," in *Proceedings of the 9th International Conference on Information Fusion*, 2006.

[13] D. A. Lambert, "STDF model based maritime situation assessments," in *Proceedings of the 10th International Conference on Information Fusion*, 2007.

[14] D. A. Lambert, "Formal theories for semantic fusion," in *Proceedings of the 9th International Conference on Information Fusion*, 2006.

[15] D. A. Lambert, "A blueprint for higher-level fusion systems," *Information Fusion,* vol. 10, pp. 6-24, 2009.

[16] D. A. Lambert and C. Nowak, "The Mephisto Conceptual Framework," Defence Science and Technology Organisation DSTO-TR-2162, 2008.

[17] C. Nowak and D. A. Lambert, "The semantic challenge for situation assessments," in *Proceedings of the 8th International Conference on Information Fusion*, 2005.

[18] A. Saulwick and K. Trentelman, "Towards a formal semantics of social influence," *Knowledge-Based Systems,* vol. 71, pp. 52-60, 2014.

[19] K. Hawley. (2010). *Temporal Parts (Stanford Encyclopedia of Philosophy)*.

[20] A. G. Cohn, B. Bennett, J. Gooday, and N. M. Gotts, "Qualitative spatial representation and reasoning with the region connection calculus," *GeoInformatica,* vol. 1, pp. 275-316, 1997.

[21] J. F. Allen, "An interval-based representation of temporal knowledge," in *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 1981, pp. 221-226.

[22] D. A. Lambert, "Advisers with Attitude for Situation Awareness," presented at the Proceedings of the 1999 Workshop on Defense Applications of Signal Processing, LaSalle Illinois, 1999.

[23] D. A. Lambert and K. Trentelman, "Cognitive Aspects of Higher-Level Fusion," in *Fusion Methodologies in Crisis Management: Higher Level Fusion and Decision Making for Crisis Management*, G. Rogova, Ed. Springer, in press.

[24] D. A. Lambert, E. P. Blasch, and É. Bossé, "Introduction," in *High-Level Information Fusion Management and Systems Design*, E. Blasch and É. Bossé, Eds. Boston: Artech, 2012, pp. 1–10.

[25] A. Saulwick, "Lexpresso: a controlled natural language," in *Controlled Natural Language, 4th International Workshop, CNL 2014*, Galway, Ireland, 2014, pp. 123-134.

[26] M. Broughton, O. Carr, P. Taplin, D. Estival, S. Wark, and D. A. Lambert, "Conversing with Franco, FOCAL's virtual adviser," in *Proceedings of the Conversational Characters Workshop, Human Factors Conference*, 2002.

[27] S. Wark and D. A. Lambert, "Presenting the story behind the data: Enhancing situational awareness using multimedia narrative," in *Proceedings of the Military Communications Conference*, 2007, pp. 1-7.

[28] S. Wark, M. Broughton, M. Nowina-Krowicki, A. Zschorn, M. Coleman, P. Taplin*, et al.*, "The FOCAL point -- Multimodal dialogue with virtual geospatial displays," in *Proceedings of the Simulation Technology and Training Conference*, 2005.

[29] S. Wark, D. A. Lambert, M. Nowina-Krowicki, A. Zschorn, and D. Pang, "Situational awareness: Beyond dots on maps to virtually anywhere," in *Proceedings of the Simulation Technology and Training Conference*, 2009.