An Exegesis of Data Fusion

Dale A. Lambert

Information Technology Division

Defence Science and Technology Organisation

PO Box 1500 Salisbury, South Australia 5108

dale.lambert@dsto.defence.gov.au

Summary:
Data fusion is the process of utilising one or more data sources over time to assemble a representation of aspects of interest in an environment. The Joint Directors of Laboratories (JDL) model is currently the most dominant data fusion model. It provides a coarse framework for the data fusion process. This paper attempts to further refine the JDL model by presenting a definitional, conceptual, and theoretical basis for that framework.

Keywords: *Data Fusion, Object Assessment, Situation Assessment, Impact Assessment*

1. Introduction

The Joint Directors of Laboratories (JDL) model is currently the most widely accepted model of the data fusion process (Steinberg et.al., 1998). In its current form, however, it provides only a weak framework in which its components enjoy varying degrees of conceptual and theoretical maturity, and in the author's opinion, this has impeded practical advancements.

This paper seeks to sketch a more mature, strategic foundation for understanding data fusion based around the JDL model. To that end it does not forward detailed algorithms for addressing specific fusion problems, but instead offers definitional, conceptual, and theoretical comments on components of the JDL model. In difference to much of the existing fusion literature, comparatively less attention is devoted to the theoretically mature components, in an effort to place greater emphasis on the less well understood aspects. So called "level 2" and "level 3" fusion emerge as the primary focus, with some discussion of "level 1". "level 0" is omitted. "Level 4" is partially addressed in a companion chapter within this text (Rutten and Lambert, 2000).

2. The JDL Data Fusion Model

2.1 JDL Data Fusion Definition

The 1987 JDL Data Fusion Subgroup defined data fusion as

> … a process dealing with the association, correlation, and combination of data and information from single and multiple sources to achieve refined position and identity estimates, and complete and timely assessments of situations and threats, and their significance. The process is characterized by continuous refinements of its estimates and assessments, and the evaluation of the need for additional sources, or modification of the process itself, to achieve improved results (White, 1987).

Steinberg et.al. (1998) revised a number of its elements and have instead proposed the following definition.

> Data Fusion is the process of combining data to refine state estimates and predictions (Steinberg et.al., 1998).

The revised definition removes the slight tracking literature bias of the earlier definition, though reference to "… state estimates and predictions …" remains a tracking literature description of proceedings. The omission of reference to data sources and time is curious, since both seem to be foundational to any real data fusion process. The author therefore prefers the following variation.

> Data fusion is the process of utilising one or more data sources over time to assemble a representation of aspects of interest in an environment.

2.2 JDL Data Fusion Model

Steinberg et.al. (1998) introduced a level 0 to the JDL model and defined level 0 to level 4 in the following manner.

> Level 0 – Sub-Object Data Assessment: estimation and prediction of signal/object observable states on the basis of pixel/signal level data association and characterization;
>
> Level 1 – Object Assessment: estimation and prediction of entity states on the basis of observation-to-track association, continuous state estimation (e.g. kinematics) and discrete state estimation (e.g. target type and ID);
>
> Level 2 – Situation Assessment: estimation and prediction of relations among entities, to include force structure and cross force relations, communications and perceptual influences, physical context, etc.;
>
> Level 3 – Impact Assessment: estimation and prediction of effects on situations of planned or estimated/predicted actions by the participants; to include interactions between action plans of multiple players (e.g. assessing susceptibilities and vulnerabilities to estimated/predicted threat actions given one's own planned

> actions);
>
> Level 4 – Process Refinement (an element of Resource Management): adaptive data acquisition and processing to support mission objectives.

The model offers only limited instruction beyond this.

2.3 Terminology of the JDL Data Fusion Model

The (revised) JDL model is a *process* model for data fusion in which the process of data fusion is decomposed into a number of component processes.

> The JDL distinction among data fusion "levels" … provides an often useful distinction among data fusion processes that relate to the refinement of "objects", "situations", "threats" and "processes". (Steinberg et.al., 1998).

The characterisation of the process components of the JDL model as "levels" engenders a hierarchical conception of data fusion, despite the bus architecture of Figure 1. This conception has some drawbacks, as noted in the conclusion.

The revised JDL nomenclature for these process components is also of some concern. To distinguish the *process* from its *product*, the author suggests that the term "fusion" be reserved for the process and "assessment" be used to refer to the product. Thus, for example, "situation fusion" is hereafter used to refer to the process of component 2, while "situation assessment" is retained to refer to the product of component 2.

3. Object Fusion

3.1 Object Fusion Definition

In phrasing their account of "level 1" fusion, Steinberg et.al.(1998) cite familiar tactical level military examples and utilise conventional tracking theory jargon in an attempt to anchor the reader's understanding. This is appropriate when exploring a theory for object fusion, but when defining object fusion, we are more interested in knowing what it is rather than how to do it. The author therefore advances the following two definitions.

> Object fusion is the process of utilising one or more data sources over time to assemble a representation of objects of interest in an environment.
>
> An object assessment is a stored representation of objects obtained through object fusion.

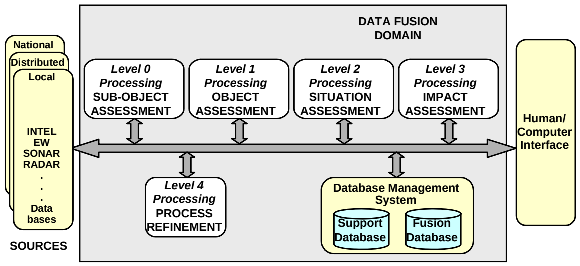

Figure 1: The JDL data fusion model

3.2 Object Fusion Conception

Object fusion exists because in our interaction with the world, we are inclined to associate bundles of near-coincident observable properties with objects, and to associate the persistence of those objects with the observed existence of those properties under periodic review. Object assessments allegedly document persistent objects having properties, and in an engineering context, these properties are usually measurable. This accounts for the numerical bias evident in the "level 1" fusion literature. In a radar environment, for example, signal and track processing is used to conclude the existence of objects associated with, *inter alia*, measured range, azimuth, elevation, Doppler, radar cross section, target type and target identity properties.

Object assessments are the product of a highly conceptualised view of the world. In the radar illustration, the affirmed persistence of an object with certain properties relies on the imposition of a number of essentially *a priori* conceptualisations, including: the choice of spatio-temporal geometry; the choice of kinematic theory and its association to target type; presumptions about measurability and equipment in relation to radar cross section; presumptions about how much an object's properties can change while remaining the same object; and the presumption that the world is a world of objects with properties.

The last of these, the conceptualisation of a world of objects with properties, is a fundamental characteristic of every object assessment, and its development can be tracked historically. From about 600 BC, the ancient Greeks had conceived of a world of objects (things) manipulated by the actions of gods. With the advent of the Presocratics, the assortment of gods gradually gave way to an understanding based in nature, but the newer outlooks remained rooted in a world of objects. Around 350 BC, Aristotle (Barnes, 1984) refined the view by proposing that objects were composed of both form (properties) and matter, *id est*, objects were formed matter. Matter is the "stuff" of which a thing is composed, the characteristic that makes a ship *this* ship, rather than that ship. Form is that which determines *what* a thing is, the characteristic that makes a ship, a *ship*. Objects therefore had properties or forms associated with them, and through the persistence of certain "essential" forms, an object could remain the same object while some of its non-essential forms changed. Modern object assessments are based upon this Aristotelian conceptualisation.

3.3 Object Fusion Theory

Object fusion is the maturest of the fusion components because of its reliance on established association and estimation processes. The association process forms new sensor observations from "level 0" signal and image processing before attempting to associate them with previously recognised objects. The Aristotelian conception of a world of objects possessing properties tends to be implemented fairly directly, with each sensor observation typically being a data structure or property vector, and with fields within that data structure for each property of interest. These measured properties may arise from different sensors.

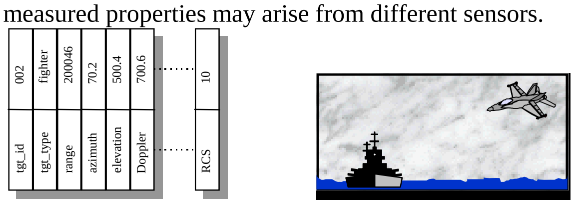

Figure 2: (a) Object Data Structure (b) A Scene

The outcome of the association process being applied to a collection of new sensor observations can include: the recognition of new instances of previously observed objects; the recognition of newly observed objects; or the recognition that some previously observed objects are no longer being observed. The estimation process assists with the association process by predicting the measured value of each observation property for the next observation, and by defining a bounded acceptance gate for association. By combining the association and estimation processes, the existence and properties of an object can be tracked over time via a linked set of sensor observations. Object assessments are sets of associated sensor observations (property vectors) so formed.

The estimation process tends to be dominated by the Kalman filter, which defines a general solution to minimised mean square estimation for linear estimators. It assumes that the observed property measurement of each object of interest can be modelled over time by a linear equation of the form $\underline{o}(k) = H\,\underline{u}(k) + \underline{v}(k)$, where: $\underline{o}(k)$ is the observed m-dimensional property vector of the object at discrete time k; $\underline{u}(k)$ is the unknown n-dimensional property vector or state of the object at discrete time k; $\underline{v}(k)$ is a m-dimensional zero mean, white, Gaussian noise vector with covariance R at discrete time k; and H(k) is an m×n measurement matrix at discrete time k. It also assumes that the state of each object of interest can be recursively modelled over time by $\underline{u}(k+1) = F\,\underline{u}(k) + \underline{q}(k) + \underline{d}(k+1|k)$, where:

$\underline{d}(k+1|k)$ is the known deterministic input change; $\underline{q}(k)$ is an n-dimensional zero mean, white, Gaussian noise vector with known covariance Q at discrete time k; and F is a given behavioural model of the object expressed as an $n \times n$ state transition matrix. Minimising the expected mean square error produces a predicted state estimate $\underline{\hat{u}}(k+1|k)$, defined by $\underline{\hat{u}}(k+1|k) = F\ \underline{\hat{u}}(k|k) + \underline{d}(k+1|k)$, where the predicted k^{th} state given the k^{th} observation is $\underline{\hat{u}}(k|k) = \underline{\hat{u}}(k|k-1) + K(k)\ \underline{\hat{o}}(k)$, for gain matrix $K(k) = P(k|k-1)\ H^t\ S(k)^{-1}$ and predicted observation (innovation) vector $\underline{\hat{o}}(k) = \underline{o}(k) - H\ \underline{\hat{u}}(k|k-1)$ with covariance $S(k) = H\ P(k|k-1)\ H^t + R$. The predicted state covariance $P(k+1|k)$ is defined by $P(k+1|k) = F\ P(k|k)\ F^t + Q$, where $P(k|k) = (I - K(k)\ H)\ P(k|k-1)$. Kalman filter usage can be sophisticated. The Interactive Multiple Model approach (Bar-Shalom and Xiao-Rong, 1993) involves a set of predicted state estimates and covariance matrices being used concurrently to model different possible object behaviours, that are combined through a Markov model describing transition probabilities between the behaviours.

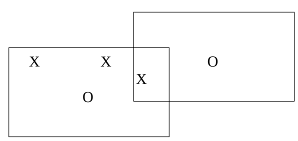

Figure 3: Data Association ambiguity

The association process begins by determining which sensor observations lie within which bounded acceptance gates of the predicted state estimates. The gates are often derived from the innovation covariance matrix S(k). Difficulties arise when the assignment of sensor observations to gates is ambiguous, as illustrated in Figure 3, with predicted state estimates O and sensor observations X. The conventional solution has been to assign a single observation to each existing object where possible, typically through some form of nearest neighbour association. More recently there has been a tendency toward multi-hypothesis data association, where alternative data association hypotheses are maintained in the presence of data association ambiguities. Blackman and Popoli (1999, Ch. 7) outline several multi-hypothesis approaches, including some based upon Lagrangian relaxation and dynamic programming.

4. Situation Fusion

4.1 Situation Fusion Definition

The significance of the Steinberg et.al.(1998) definition of “level 2” fusion lies in its recognition that “relations among entities” are a key aspect of interest. This encourages the following definitions.

> <u>Situation</u> <u>fusion</u> is the process of utilising one or more data sources over time to assemble a representation of relations between objects of interest in an environment.
>
> A <u>situation</u> <u>assessment</u> is a stored representation of relations between objects obtained through situation fusion.

4.2 Situation Fusion Conception

The significance of relations between objects is often not understood. Aristotle’s conception of a world of objects endured throughout medieval times, receiving heavy-handed endorsement from the Church during the scholastic period. Indeed, it was only little more than two centuries ago that challenges first began to surface. Olson (1987) documents the transition. The impetus for change was the emergence of *relations* as a conception over and above properties, and it arose because of limitations in the Aristotelian outlook. These same limitations arise in object fusion.

As we progress through the data fusion process we desire the fusion system to represent more than measurable properties of objects. For example, we might want the system to recognise that fighter2 is targeting ship3 in Figure 2(b). In the Aristotelian framework, this is represented by allowing fighter2 to have the property of targeted_at ship3, while ship3 has the property of being targeted_by fighter2. Figure 4 depicts the representation.

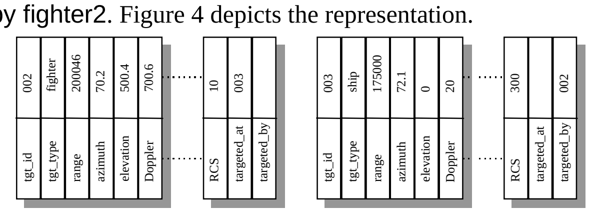

Figure 4: fighter2 targeting ship3

The difficulty with this object based conception is that the more general targeting relation existing *between* fighter2 and ship3 cannot be represented, as it *belongs* to neither fighter2 nor ship3 alone. It is a binary relation *between* two objects, and thus can reside in neither data structure. Properties are merely unary relations.

The longstanding Aristotelian conception of a world of objects failed to accommodate the comparatively recent emergence of the idea of a relation, and so a new metaphysical conception surfaced. Ludwig Wittgenstein explicitly proposed a *world of facts* as the fundamental substrate, where facts are subsequently understood as the application of relations to objects. In his cryptic, unapologetic style, Wittgenstein (1922) launched his English publication of

Tractatus Logico-Philosophicus with,

> 1.The world is all that is the case.
> 1.1 The world is the totality of facts, not of things.

Wittgenstein supplanted a view that had persisted for well over 2000 years. The world could now be understood as a collection of facts developed from basic facts, the latter being expressed symbolically by associating a relation symbol with some object symbols. Thus (targeting fighter2 ship3) expresses the fact that fighter2 is targeting ship3. That something is targeting somethingelse can be expressed by ∃x ∃y ((targeting x y) & (≠ x y)). The shift to a metaphysics of symbolically presented facts delivers a more expressive capability, and accounts for the symbolic bias evident in the "level 2" fusion literature.

When engaging the world, we rarely attend to individual facts in isolation. Figure 2(b) depicts a typical mental snapshot picture of the world over a limited time frame and region, and in assessing this picture we will naturally be inclined to represent it as a *collection of facts*. Barwise and Perry (1983) proposed that situations were the fundamental building blocks of our assessment of the world.

> Reality consists of situations - individuals having properties and standing in relations at various spatiotemporal locations. (Barwise and Perry, 1983).

Situations are essentially collections of related spatiotemporal facts, where facts consist of relations between objects. This is a step up from Wittgenstein's world of facts. Here the world is a world of situations, and assessing the world involves individuating situations. Situation assessment involves assessing situations, not facts or objects *per se*.

In clarifying "level 2" fusion, Steinberg et.al. (1998) remarks, "Level 2 assignment involves associating tracks (i.e. hypothesised entities) into aggregations". The aggregates referred to are not simple clusters of assessed objects, such as a fleet formed from individual ships, or a political faction formed from individual politicians. These are artefacts of object fusion. It is the inclusion of *relations* between objects that leads to situation assessments, as Steinberg et.al. (1998) recognise when they continue with "The state of the aggregate is represented as a network of relations among its elements". They too attach the term "situation" to this kind of aggregate.

> As the class of relationships estimated and the numbers of interrelated entities broaden, we tend to use the term 'situation' for an aggregate object of estimation. (Steinberg et.al., 1998).

4.3 Situation Fusion Theory

Situations are *abstractions* of the real world. They can be represented formally using expressions. The set of expressions, **Expr**, is defined (here) by:

Expr ::= **BoolExpr** | **IntExpr** | **RealExpr** | **SymExpr** | **VarExpr** |

SitExpr | (**Expr** … **Expr**)

The situation expressions comprising **SitExpr** allow us to describe the real world abstractly in at least four important senses.

Under the Barwise and Perry (1983) formulation, the vast majority of propositions are not identified as true or false in a situation because they are not reported. *Situations are abstract in the sense of being incomplete*. For example, one might abstract the scene of Figure 2(b) as the scene of Figure 5(a) comprising just a ship. The symbol ship3 might be selected to represent this object, so that for interpretation function A, A(ship3) denotes the object of Figure 5(a). Figure 5(b) presents a less abstract scene in which both the ship and fighter objects feature. Historically, mathematics has dealt with such occurrences through the use of the set membership relation $\in$ and *admitted a set of objects as an object* within our conceptual universe. In naïve set theory, {fighter2, ship3} represents Figure 5(a) as the set of objects {A(fighter2), A(ship3)}.

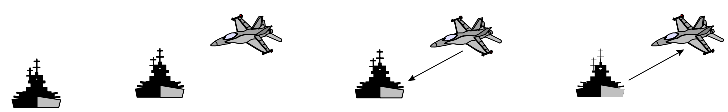

Figure 5: (a) (b) (c) (d)

Level 2 fusion is about identifying relations between objects. For example, we may wish to represent the fact that the fighter is above the ship in the scene of Figure 5(b). We can introduce an above relation symbol, but by the Axiom of Extensionality[1], {above, fighter2, ship3} = {above, fighter2, ship3}. The order of elements matters in assertions, but not in sets.

Expressions are introduced to textually represent conceptualisations. Symbols are primitive expressions while composite expressions respect order and are delimited by parentheses. Primitive expressions are each interpreted as a tuple <i,v,j> comprising: an index i; an interpreted object v; and the index j of the next element. The interpretation E of the symbol expression ship3 might be E(ship3) = <4,A(ship3),0>. Composite expressions are interpreted as a set of linked expression elements. The composite expression of Figure 5(b) gives rise to two possible expressions, (fighter2 ship3) and (ship3 fighter2), which are illustrated in Figure 5(c) and Figure 5(d) respectively. (fighter2 ship3) might be interpreted by E((fighter2 ship3)) = {<6,A(fighter2),4>,<4,A(ship3),0>}. Introducing a relation symbol into the linked scene can then represent relations. E((above fighter2 ship3)) is illustrated in Figure 6(a). It is also convenient to introduce integer, real and Boolean objects with their conventional symbolic representations. Figure 6(b) illustrates.

[1] Formally the axiom is expressed by $\forall x\ \forall y\ (x = y \Leftrightarrow \forall z\ (z \in x \Leftrightarrow z \in y))$.

An interpretation E of an expression is defined as a conceptualisation of the real world with indexing set $\Im$, such that: $\mathsf{E}(x) = \mathsf{X}(x)$ for symbol x and $\mathsf{E}((\alpha_1 \ldots \alpha_k)) = \{\mathsf{X}(\alpha_1), \ldots, \mathsf{X}(\alpha_k)\}$for expressions $\alpha_1, \ldots, \alpha_k$, where:

A. $\mathsf{X}(x) = <i,\mathsf{A}(x),j>$, for any symbol $x \in$ **BoolExpr** $\cup$ **IntExpr** $\cup$ **RealExpr** $\cup$ **SymExpr**, where $\mathsf{A}(x)$ is that which is referred to by symbol x and $i(\neq 0)$, $j \in \Im$ are some selected indices, such that A[**BoolExpr**] = {*true*, *false*}; A[**IntExpr**] = Z; A[**RealExpr**] = R; and A[**SymExpr**] is the set of admissible real world objects and relations;

B. $\mathsf{X}(x) = <i,\mathsf{A}(v(x)),j>$, for any symbol $x \in$ **VarExpr**, where v : **VarExpr** $\rightarrow$ **Expr** is a contextual variable assignment function and $i(\neq 0)$, $j \in \Im$ are some selected indices;

C. $\mathsf{X}((\alpha_1\ldots\alpha_k)) = <i,w,j>$, for some $i(\neq 0)$, w, $j \in \Im$, such that $\exists a \exists n(\mathsf{X}(\alpha_1)=<w,a,n>)$ and $\forall u$ ($u\in\{1,\ldots,k\text{-}1\} \Rightarrow \forall n \forall p \forall q$ (($\mathsf{X}(\alpha_u) = <m,a,n>$ & $\mathsf{X}(\alpha_{u+1}) = <p,b,q>) \Rightarrow n = p$));

D. $0 \in \Im$ and for any expression α, $\mathsf{E}(\alpha) = <i, a, 0>$ if and only if α is not linked to a next expression element.

The model theoretic construction (see for example Chang and Keisler, 1977) allows us to represent complex incomplete abstractions of the world.

Figure 6: (a) (b)

Probabilistic propositional expressions comprising **PropExpr** can be used to express uncertain claims about the world. *Situations are abstract in this second sense of accommodating abstract notions of uncertainty.*

PropExpr ::= **BaseExpr** | **CondExpr**;
BaseExpr ::= **PosLitExpr** : **ProbExpr**;
CondExpr ::= **ConseqExpr** if **AntecExpr** : **CondProbExpr** such that whenever $x \in$ **AntecExpr** has k terms then the associated $y \in$ **CondProbExpr** has 2^k terms;
ConseqExpr ::= **PosLitExpr**; **AntecExpr** ::= **LitExpr** and … and **LitExpr**
LitExpr ::= **PosLitExpr** | (not **PosLitExpr**);
PosLitExpr ::= (**SymExpr Expr** … **Expr**);
CondProbExpr ::= <**ProbExpr** … **ProbExpr**>; **ProbExpr** = [0,1].

Using this we might form a base proposition like (blue sky) : 0.8 to indicate that the probability of a blue sky is 0.8 or express a conditional proposition like

(heavy ?x) if (solid ?x) and (not (small ?x)) : <0.95, 0.2, 0.1, 0.001>

to identify relative conditional probabilities, such as the conditional probability of something being heavy given that it is not solid but is small, as being 0.1.

A probabilistic inference relation for **PropExpr** has been implemented in ATTITUDE. $\Sigma\ J_{A,p}\ \alpha[v]$ holds if for interpretation A and variable assignment v, the probability of base expression α is p, given that some set of propositional expressions Σ holds. Lambert and Relbe (1998) define a similar fuzzy inference relation. Fabian and Lambert (1998) explore the probabilistic issues. Probabilistic inference allows us to manage uncertain claims in our abstractions of the world.

In a third sense of abstraction, situations can be counterfactual. Counterfactual abstractions can arise deliberately to facilitate hypothetical "what if" reasoning. We may deliberately form counterfactual abstract situations in order to understand what would happen if some event subsequently occurred. To capture this intent, situational expressions are defined as events or scenarios, where events are sets of propositional expressions and scenarios are sets of events.

SitExpr ::= **EventExpr** | **ScenarioExpr**;
EventExpr ::= **P**(**PropExpr**); **ScenarioExpr** ::= **P**(**EventExpr**), where semantically for any $x \in$ **SitExpr**, $E(x) = \langle i,\{X(u) \mid u \in x\},j\rangle$, with $\mathbf{P}(w) = \{v \mid v \subseteq w\}$.

The inference relation $\Sigma\ J_{A,S,p}\ \alpha[v]$ can be applied relative to any situation S by aligning Σ with S. If event A has occurred and either event B or event C might occur, then we can reason about both possible scenarios {A,B} and {A,C} through $\cup\{A,B\}\ J_{A,p}\ \alpha[v]$ and $\cup\{A,C\}\ J_{A,p}\ \alpha[v]$ respectively.

In a fourth sense of abstraction, situations are abstract sets and so can serve as objects within propositional expressions involving situations. If v(?sit) identifies a preferred situation with variable assignment function v, then we might choose to assert (preferred ?sit) to express this. This propositional expression could be included within a situation, and this situation could be the one identified by v(?sit). In that case, the situation v(?sit) would be self referential. It is equally possible to form self referencing expressions such as $({}_7\text{info (secret}^7))$ for which $E((\text{info (secret)})) = \{\langle 7,A(\text{info}),34\rangle, \langle 34,12,0\rangle, \langle 12,A(\text{secret}),7\rangle\}$. The admission of situations as objects accords with the revised JDL model ideology. But as abstract objects, situations and expressions engender a non well founded set theory (Azcel, 1988).

Formal situations can be used to represent our fusion agent's beliefs about abstract situations in the real world. The abstractions that we form in an automated system usually need to resemble those that are undertaken by humans. Humans tend to explain their decisions through mental attitudes, such as beliefs, desires, expectations, fears, hopes, *et cetera*., and when expressing these and other mental attitudes, engage a syntax of the form <subject> <attitude> that <propositional statement>. The following examples illustrate: Fred believes that the sky is blue; Tom expects that Mary will win lotto; and Mary hopes that Tom is insightful. Statements having this syntactic form are called propositional attitude statements and the beliefs *et cetera* that they denote are technically termed propositional attitudes. In a propositional attitude statement:

the subject, e.g. Fred, expresses which individual has the propositional attitude; the propositional statement, 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. With subtle modification, propositional attitude observations about Fred such as Fred believes that the sky is blue, can be transformed into propositional attitude instructions to Fred like Fred believe (blue sky). The latter is an instruction, commanding software agent Fred to believe that the sky is blue.

The multi-agent ATTITUDE system has been developed to facilitate attitude programming. Individual's verbal propositional attitude explanations of their own recorded behaviour can be implemented directly through ATTITUDE routines. Each routine consists of: (a) a basic expression α stating the goal that the routine is designed to satisfy; and (b) a network of propositional attitude instructions based upon beliefs, expectations, anticipations, routines, desires, intentions, sensors and effectors. Routines are executed whenever there is a desire or intention to satisfy the basic expression α specified by the routine. Routine execution involves navigating control through the network of instructions, with each instruction succeeding or failing. The routine is designed so that the goal α will be satisfied if a successful path of execution can be found from the start node to a terminal node of the network. At any given time, multiple desires, and hence routines, are competing for execution time.

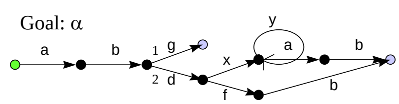

Figure 7: Schematic routine with goal α and propositional attitude instructions a, b, ..

Routines capture routine behaviour as a computational routine. They track an individual's interaction with both their environment and other individuals, as they form situation assessments. Routines can construct situation assessments in a number of ways: (a) by using the basic expression α of a successful routine to abbreviate the awareness arrived at through the routine's successful execution; (b) by the routine constructing one or more events to record awareness explicitly as it executes; or (c) by forming a situation mosaic accepting reports from the various routines that have been clipped together like lego blocks as the circumstances warrant. The probabilistic inference engine can be invoked across a set of events through a propositional attitude instruction.

Three kinds of belief must be considered when engineering automated situation fusion. *Environmental beliefs* arise from direct observation of the world at a given time, such as when I believe that there is a can on the desk in front of me. In ATTITUDE, a sensor can be used to sense object assessments and create the corresponding beliefs in an event associated with that sensor. *Definitional beliefs* derive from the meaning of terms, such as the belief that a son is a male child.

When using the above relation, we will usually implicitly take it to be a transitive relation[2]. Such beliefs need to be explicitly included within the machine before we turn it on. *Domain beliefs*, express how we presume the world to be independently of a direct observation, such as my belief that dark clouds usually produce rain. The domain beliefs can be expressed procedurally in terms of how something can be achieved (routines), or declaratively in terms of what is the case (inference). Situation assessments are situations comprising beliefs that are generated from the interaction of environmental, definitional and domain beliefs[3].

5. Impact Fusion

5.1 Impact Fusion Definition

Steinberg et. al. (1998) present "level 3" fusion in terms of the "… estimation and prediction of effects on situations …". In the author's opinion, "level 3" fusion is indeed about "effects on situations", but a definition cast in only these terms fails to capture the essence of "level 3" fusion. "Level 2" fusion is about belief. Via object assessments, situation assessments express beliefs (situations) about how the world might be. The consequences or effects of those beliefs are important to us, but only in as much as they impact upon what we want to be the case. This is the essence of "level 3" fusion - it is about how our beliefs impact upon our will. "Level 3" fusion is about the effect of situations on our intentions, and thus interprets the world in terms of opportunities and threats, with a view toward maintaining the satisfaction of our intent.

> Impact fusion is the process of utilising one or more data sources over time to assemble a representation of effects of situations in an environment, relative to our intentions.
> An impact assessment is a stored representation of effects of situations obtained through impact fusion.

5.2 Impact Fusion Conception

Impact fusion involves: (a) assessing the situations, that could or will arise; and (b) determining how those situations will impact upon our intentions.

---

[2] $\forall x\ \forall y\ \forall z\ (((\text{above } z\ y)\ \&\ (\text{above } y\ x)) \Rightarrow (\text{above } z\ x))$.

[3] In Kantian terminology (Kemp-Smith, 1929), *when we first turn the machine on*: environmental beliefs are synthetic and *a posteriori*, definitional beliefs are analytic and *a priori*, and domain beliefs are synthetic and *a priori*.

Dennett (1971) identifies three means by which people predict outcomes. The first is the Physical Stance, where one engages principles of Physics to predict outcomes. People employ this when playing snooker or assessing the trajectories of projectile weapons. The second is the Design Stance, where one engages principles of design to predict outcomes. People employ this when troubleshooting an automobile fault or coding and maintaining computer programs. The third is the Intentional Stance, where one engages principles of rationality to predict outcomes. People employ this when forecasting the actions of a fighter pilot or when competing with an advanced computer game. The Design Stance is used whenever the physics of the situation is too difficult or laborious. The Intentional Stance is used whenever the design underpinning the situation is too difficult or laborious. *A priori* definitional and domain beliefs for all three stances will be required for sophisticated impact assessments.

Assessing how predicted situations will impact upon our intentions requires self awareness of our own intentions and an understanding of how our beliefs and other mental states influence our intentions. This is rather like applying the Intentional Stance to oneself. The Intentional Stance proceeds by ascribing routines, beliefs, intentions and other mental states to individuals and then predicting their behaviour under a presumption that they will behave in accordance with our ascriptions. In ATTITUDE, these are explicit.

5.3 Impact Fusion Theory

A given situation S_t at time t can be effected by: the fusion system's behaviour; other agents; designed products; and physical effects. Rather than encode objective models of the physics, design and intentions of environmental elements, the ATTITUDE approach is to develop subjective behavioural routines that enable the fusion agent to monitor and respond to those aspects of the environment, as appropriate. The trajectory of a projectile, for example, will not be encoded by physical equations, but the consequences of those equations may well be incorporated into a routine that describes how the projectile will be experienced.

Routines perform a second role by being designed to satisfy particular goals. Intentions are the independent goals, and routines designed to satisfy them will typically contextually spawn dependent goals (desires) which other routines are designed to satisfy. The routines that are designed to satisfy intentions are therefore expected to do so while initiating routines to subjectively monitor and respond to the physics, design and intentions of environmental elements. This is of course what people do. In seeking to satisfy some intention, people consider possible events and monitor and respond to actual events of relevance. The dual capability of satisfying intentions and assessing situation effects, allows routines to provide the basis for impact assessments.

An intention can be formally represented by an ordered pair $<\alpha,t>$ comprising a basic expression α and a deadline t, meaning that we want expression α to be true at time t[4]. If W_t is the set of intentions having deadline t and our situation assessment at that time t turns our to be S_t, then a comparison between W_t and S_t will indicate how well our intentions were met. By defining a measure m : $(\mathbf{P}(\mathbf{SitExpr}) \times \mathbf{P}(\mathbf{SitExpr} \times \text{Time})) \rightarrow \mathsf{R}$, the extent to which our intentions were met can be measured by $m(S_t,W_t)$. The measure m is domain specific and I assume that W_t is satisfied if and only if $m(S_t,W_t) = 0$.

At any given time k, we will have a set of intentions with varying deadlines. Let $W_{k,t}$ be the set of intentions at time k with deadlines for time $t \geq k$. Then the set of all intentions at time k is given by $W(k) = \cup\{W_{k,t} \mid t \in \text{Time}\}$, while the set of all deadlines at time k is given by $D(k) = \{t \mid \langle a,t\rangle \in W(k)\}$. Each routine seeking to satisfy an intention is responsible for meeting that intention in the presence of a relevantly changing environment, but the side effects of multiple routines addressing different intentions could inadvertently impact on one another. The impact of a situation S_t is therefore not simply $m(S_t,W_{k,t})$, because the routines seeking to satisfy later intentions may be significantly advantaged or disadvantaged by S_t. The impact of a given situation must be assessed against current and future intentions, given the routines at our disposal. For a particular choice of routines generating a set of situations $\{S_t \mid t \in D(k)\}$, we can measure our overall dissatisfaction by $\text{dissatisfy}(W(k)) = \sum_{t \in D(k)} m(S_t, W_{k,t})$.

This is essentially a dynamic programming problem for which we must decide which routines to execute in order to minimise dissatisfy(W(k)). A new dissatisfaction minimisation will be required each time a new intention is added or if a routine seeking to satisfy an intention either fails or exceeds its deadline. To facilitate the computation, each routine designed to satisfy an intention will need to record an estimated duration and an expected situation outcome, with the latter being used together with generalised weakest precondition trees (Rutten and Lambert, 2000) to estimate when a routine is executable. A trace of higher level routines representing the best course of action will emerge as a by-product of the dynamic programming impact fusion process.

6. Conclusions

A hierarchical thread is evident as objects are parts of facts; that are parts of situations; that are the effects of situations. A non-hierarchical thread is equally prominent. QF005 should retain its object association if its airlane deviation can be explained by a storm situation. Two unfriendly object tracklets could be

[4] More sophisticated scheduling options can also be considered.

combined because the effect of that possibility is a significant threat. Data fusion requires an integrated and multi-disciplinary strategic framework.

References

Aczel, P. (1988). Non-Well-Founded Sets, CSLI Lecture Notes No. 14. Center for the Study of Language and Information, Stanford.

Barnes, J. (1984). The Complete Works Of Aristotle Vol. 1and 2. Princeton University Press, Princeton.

Bar-Shalom, Y. and L. Xiao-Rong. (1993). Estimation and Tracking: Principles, Techniques, and Software. Artech House, Boston.

Blackman, S. and R. Popoli (1999). Design and Analysis of Modern Tracking Systems. Artech House, Boston.

Chang, C. C. and H. J. Keisler (1977). Model Theory (2nd. Ed.) North-Holland Publishing Company, Amsterdam.

Fabian, I. and Lambert D. A. (1998). "First-Order Bayesian Reasoning", In Advanced Topics in Artificial Intelligence, 11th Australian Joint Artificial Intelligence, AI'98, 131 – 142. Springer-Verlag.

Kemp-Smith, N. (1929). Immanual Kant's Critique of Pure Reason. Macmillan and Company Limited, London (Translation of Kant's "Kritik der Reinen Vernunft" 1787).

Lambert, D. A. (1999). "Assessing Situations", In Proceedings of 1999 Information, Decision and Control, pp. 503 – 508. IEEE.

Lambert, D. A. and M. G. Relbe. (1998). "Reasoning With Tolerance", Proceedings of the IEEE Second International Conference on Knowledge-based Intelligent Electronic Systems. Vol. 3 pp. 418 – 427, Edited L.C. Jain and R. K. Jain. IEEE Inc., NJ.

Rutten, M. and D. A. Lambert (2000). "Automated Situation Assessment Refinement", *appearing in this text*.

Steinberg, A. N., C. L. Bowman and F. E. White (1998). "Revisions to the JDL Data Fusion Model", The Joint NATO/IRIS Conference, Quebec.

White, F. E. (1987). "Data Fusion Lexicon", Joint Directors of Laboratories, Technical Panel for C^3, Data Fusion Sub-Panel, Naval Ocean Systems Center, San Diego.

Wittgenstein, L. (1922). Tractatus Logico-Philosophicus. Routledge And Kegan Paul, London.

http://www.springer.com/978-3-540-00246-8

Soft Computing in Measurement and Information
Acquisition
Reznik, L.; Kreinovich, V. (Eds.)
2003, XIV, 284 p., Hardcover
ISBN: 978-3-540-00246-8