

Probabilistic Argument Maps for Intelligence Analysis: Capabilities Underway

Robert Schrag*, Edward Wright*, Robert Kerr*, Robert Johnson*, Bryan Ware*, Joan McIntyre+,
 Melonie Richey§, Kathryn Laskey%, Robert Hoffman&

*Haystax Technology, McLean, VA and Las Vegas, NV (Schrag)

§Camber Corporation, Columbia, MD

+Innovative Analytics and Training, Fairfax, VA

%George Mason University, Fairfax, VA

&Florida Institute for Human-Machine Cognition, Pensacola, FL

Abstract

We describe enhancements underway to our probabilistic argument mapping framework called FUSION. Exploratory modeling in the domain of intelligence analysis has highlighted requirements for additional knowledge representation and reasoning capabilities, particularly regarding argument map nodes that are specified as propositional logic functions of other nodes. We also describe more flexible specifications for link strengths and node prior probabilities. We expect these enhancements to find general applicability across problem domains.

1 Introduction

Haystax has developed the FUSION framework (see our companion paper [7]) to facilitate creation of useful Bayesian networks (BNs) by non-technical subject matter experts (SMEs). Because we use exclusively binary random variables (BN nodes) over the domain {true, false}, it is natural to construe FUSION models as probabilistic argument

maps. Exploratory modeling in the domain of intelligence analysis—where argument mapping [1] is useful but (until FUSION [6]) has not been underpinned by mathematically sound probabilistic reasoning—has highlighted requirements for additional capabilities.

Forthcoming sections first present an early model motivating some of these capabilities, then describe our technical approach to each. Appendices describe key elements of our proposed intelligence domain-specific variant of FUSION called CRAFT.

2 Motivating FUSION model

Figure 1 is a screenshot of a FUSION model addressing the CIA’s Iraq retaliation scenario [4], where Iraq might respond to US forces’ bombing of its intelligence headquarters by conducting major, minor, or no terror attacks. The model emphasizes Saddam’s incentives to act. By setting a hard finding of false on the node SaddamWins, we can examine computed beliefs under Saddam’s worst-case scenario. See [7] for details regarding the scenario, our different models addressing it, and analyst SME model feedback, as well as review of related work.

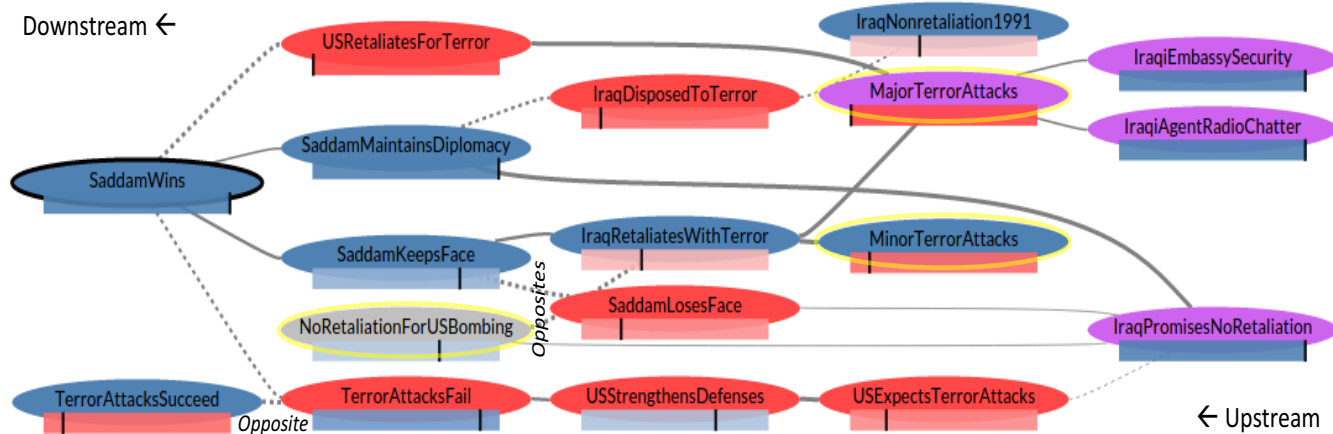


Figure 1. Statement nodes are connected by positive (solid grey line) and negative (dashed grey line) indication links of various strengths (per line thicknesses). Argument flow (from evidence to outcomes) is from right to left—e.g., SaddamWins is strongly indicated by SaddamKeepsFace. Outcome hypothesis nodes are circled in yellow. SaddamWins (hard finding false) captures Saddam’s incentives to act or not.

Node colors in Figure 1 capture whether Saddam’s disposition (or attitude) regarding a given statement is favorable (blue) or unfavorable (red). FUSION computes these dispositions from a single directive—to propagate the favorability of SaddamWins upstream (only), respecting link polarities. One node (gray, bottom left) is not touched by this propagation. Four nodes (purple, top right) have ambiguous status with respect to Saddam’s disposition. Belief bars augment spatial tick marks with colors chosen to reflect the degree of concern a given statement would pose given a node’s disposition. Lower belief (redder bar) poses less concern regarding a statement viewed unfavorably, more regarding one viewed favorably, higher belief (bluer bar) the reverse. Red/blue contrast thus draws attention to a statement that should be of concern to Saddam.

FUSION models can include argument map link types per Table 1.

Table 1. FUSION supports probabilistic argument map model link types (center column). Only the final two link types implementing propositional logic operators take an arbitrary number of input statements. All other link types are binary.

Downstream ¹ statement	IndicatedBy	Upstream statement(s)
	CounterIndicatedBy	
	MitigatedBy	
	RelevantIf	
	OppositeOf	
	ImpliedByConjunction	
	ImpliedByDisjunction	

In a FUSION model, every argument map statement is a Hypothesis. For the last two link types in Table 1, the downstream statement also is a Logic statement.

In developing the model in Figure 1, we identified the following representation and reasoning shortcomings for which we are now implementing responsive capabilities.

- Beliefs computed for the scenario’s three outcomes do not sum to 1.0. We can correct this by recasting IraqRetaliatesWithTerror as an exclusive-or (summary or constraint) Logic statement (see section 3) over the scenario’s outcomes, rather than as an absolutely supported Hypothesis, while also addressing the next issue.
- In current FUSION, a given node cannot be both a Logic statement and an indicator. Our FUSION spec-to-BN conversion software [6] treats these as distinct patterns of parents. An indicated node acts as a BN parent to an indicating one. A Logic statement’s input nodes act as parents to the Logic node itself. To implement indication by a Logic statement node L , we will create (under the hood) an auxiliary node I to serve as the indicator and assert a Logic constraint C (with parents L and I) establishing $L = I$.

¹ Per argument map convention, “downstream” is left, “upstream” right in the left-flowing argument map of Figure 1.

- TerrorAttacksFail (likewise TerrorAttacksSucceed) should be allowed to be true only when TerrorAttacks also is true. We are correcting this by adding the capability to assert Logic constraints, such as the one described for this model in section 3 to be used instead of the naïve OppositeOf link now connecting TerrorAttacksFail and TerrorAttacksSucceed.

3 Logic statements

We are working to make FUSION support any standard propositional logic expression using unary, binary, or higher arity operators². When a Logic statement has a hard true finding³, we refer to it as a Logic constraint, otherwise as a summarizing Logic statement.

Figure 1’s model would be better if TerrorAttacksFail were allowed to be true only if TerrorAttacks also were true. We know that an attempted action can succeed or fail only if it occurs. By explicitly modeling (as Hypotheses) both these potential action results and adding a Logic constraint⁴, we can force zero probability for every excluded truth value combination, improving the model (in tradecraft terms, correcting a “logic flaw”). See Figure 2.

Figure 2’s graphics are from the COTS Netica BN Application. Because all nodes in the generated BN underlying a FUSION model are binary, we also could render the full BN using more perspicuous single belief bars.

FUSION implements a target belief spec either (depending on purpose) using a BN node like Figure 2’s ConstraintTBC or (equivalently) via a likelihood finding on the subject BN node. The FUSION GUI does not ordinarily expose an auxiliary node like ConstraintTBC to an analyst.

This example is for illustration. We can implement this particular BN pattern without target beliefs. We also could implement absolute-strength IndicatedBy links as simple implication Logic constraints. However, this would not naturally accommodate one of these links’ key properties—the ability to specify degree of belief in the link’s upstream node when the downstream node is true—relevant because we can infer nothing about P given $P \Rightarrow Q$ and knowing Q to be true. It also demands two target belief specs that tend to compete. We are working to identify more Logic constraint patterns that can be implemented without target beliefs and to generalize specification of belief degree for any underdetermined entries in a summarizing Logic statement’s CPT.

² See, e.g., https://en.wikipedia.org/wiki/Truth_table.

³ A likelihood finding could be used to implement a soft constraint.

⁴ This constraint can be rendered (abbreviating statement names) as (or (and occur (xor succeed fail)) (and (not Occurs) (nor Succeeds Fails))) or more compactly via an if-then-else logic function (notated ite) as (ite Occurs (xor Succeeds Fails) (nor Succeeds Fails))—if an attack occurs, it either succeeds or fails, else it neither succeeds nor fails.

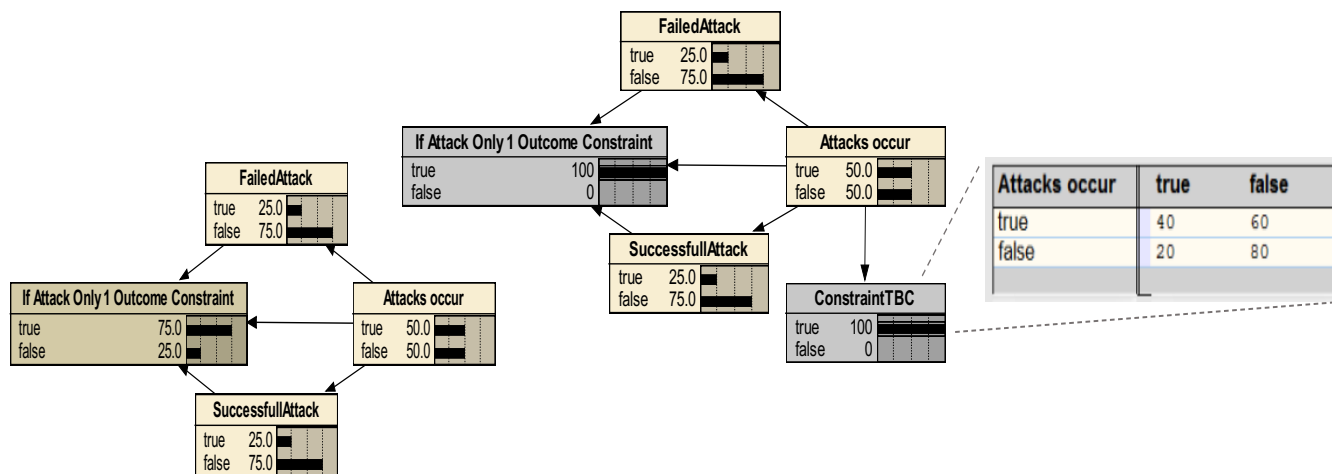


Figure 2. We are implementing Logic constraints to enforce sound logical reasoning. The constraint node (left, in right model fragment) ensures that the model will believe in attack success/failure only when an attack actually occurs. Setting the hard true finding on this node turns the summarizing Logic statement (left, in the left fragment) into the Logic constraint—but also distorts the model’s computed probabilities for the three Hypotheses. Presuming these probabilities have been deliberately engineered by the modeler, FUSION must restore them. It does so by specifying (bottom fragment) a target belief (implemented via the ConstraintTBC node) on one of the Hypotheses.

4 Target beliefs

We can set a target belief for any FUSION node whose computed belief under a baseline situation (say, before the application of a set of evidence items or assumptions) deviates from a modeler’s expectations or requirements. We have formerly done this to address unacceptably small probabilities for far-upstream nodes—an issue we hope to see less of with new flexible link strengths (section 5). A target belief also can serve an exogenous variable that should be informed by real-world data statistics.

As with BN belief inference, there is no closed-form solution for target belief satisfaction, which requires a gradient descent optimization over individual BN inference invocations, measuring in each step nodes’ differences between observed beliefs and targets. Our implementation [8] computes differences on a log odds (vs. linear) scale (reflecting actual belief impacts and reducing gradient descent oscillation), initially adjusts all involved nodes in a single optimization step (effectively, in parallel—saving steps compared to strictly serial optimization), and saves the work from previous satisfaction processes over a given model (e.g., under edit) for fast incremental operation.

In Haystax’ primary risk assessment FUSION application (called CARBON) including hundreds of BN nodes and dozens of target beliefs, processing takes just a few seconds. Our results compare favorably with those of related

published algorithms. Per the recent survey by Mrad et al [6], who refer to this capability as “fixed probability distribution” specification, the only published results are for much smaller problems.

5 Flexible belief and strength specifications

A given Bayesian network requires the specification of many individual, point probability values, but intelligence analysts generally work with probability ranges, as institutionalized in Intelligence Community Directive (ICD) 203 [2] per Figure 3 below. Beyond perfunctorily associating a uniform distribution with the range [45, 55]% (zero elsewhere) with the designation “roughly even chance,” a user may select a standard distribution (such as our grey one) or specify his/her own mode, inflection points, and non-zero probability range—perhaps all by dragging a few handles on a control. Only the curves’ shapes matter here. FUSION normalizes height to agree with unit area (= 1.0).

While several of FUSION’s parameters are probability-valued, the most prominent way probabilities enter a FUSION-generated BN is via under-the-hood encoding of on-the-dashboard-specified indication strengths. FUSION now encodes strengths using fixed odds ratios per Figure 4. More flexible specifications may be more appropriate for many problem types.




Figure 3. FUSION will bridge the gap between analysts’ familiar probability ranges (shown here with preferred natural language labels per ICD 203) and Bayesian networks’ exact probability requirements by supporting explicit probability distributions (shown here in grey). Beyond perfunctorily associating a uniform distribution with the range [45, 55]% (zero elsewhere) with the designation “roughly even chance,” a CRAFT user may select a standard distribution (like our grey one) or a customized one, perhaps (in a future GUI) by dragging a few handles (small circles) on the distribution curve. Thus, CRAFT supports analysts’ best intuitions about argument probabilities without forcing commitments. By sampling over input distributions, develop (under the hood) and display (on the dashboard) output belief distributions for argument map statements/nodes.




Figure 4. We are generalizing FUSION’S current fixed indication strengths (labeled vertical lines) as user-specifiable probability distributions similar to those planned for probabilities (Figure 3), supporting any odds ratio or distribution thereof. Absolutely is intended as logical implication. Fusion does not otherwise commit analysts to absolute certainty. An odds or log odds scale may be more salient than standard probability. This approach develops finest warranted beliefs.

Robust statistical sampling of belief distributions could, in large models, exceed GUI near-real-time thresholds for to-user feedback. Another attractive option, given reasonably tight belief distributions, is to develop only bounds for statement beliefs, exploiting FUSION’S per-statement actor disposition framework (described at Figure 1). Develop a statement’s favorable bound (i.e., the lower bound for an unfavorable statement, upper bound for a favorable one) by considering just other favorable bounds in an unambiguous-statement-only subgraph, and address limited combinatorics over disconnected subgraphs. Belief bars resulting would thus include just two tick marks, rather than a richer distribution shape.

6 Conclusion

We are working to make FUSION, already a powerful framework for SMEs to author Bayesian network-based models, more complete and versatile for probabilistic argument map applications—in intelligence analysis and in other domains, generally. Allowing summary Logic statements to serve as probabilistic indicators of hypotheses will afford uniform expressive power across a model’s downstream and upstream levels. Supporting arbitrary Logic formulas as constraints as well as summary statements will increase expressiveness and conciseness. Supporting more flexible belief and strength specifications to capture finer

probabilistic distinctions will enable more precise modeling by individual SMEs and SME teams or crowds—values for numeric model parameters might be aggregated using crowdsourcing methods.

Our proposed CRAFT variant designed for intelligence analysis will express all of FUSION’S capabilities in analyst-friendly vocabulary (appendix A) and develop mechanisms and methods explicitly supporting the modeling of source credibility per Intelligence Community standards (appendix B).

References

- [1] CIA Directorate of Intelligence, “A Tradecraft Primer: The Use of Argument Mapping,” *Tradecraft Review* 3(1), Kent Center for Analytic Tradecraft, Sherman Kent School.
- [2] Intelligence Community Directive (ICD) 203, “Analytic Standards,” January 2, 2015. <http://fas.org/irp/dni/icd/icd-203.pdf>
- [3] Intelligence Community Directive (ICD) 206, “Sourcing Requirements for Disseminated Analytic Products,” October 17, 2007. <http://fas.org/irp/dni/icd/icd-206.pdf>
- [4] Richards J. Heuer, Jr., *Psychology of Intelligence Analysis*, Central Intelligence Agency Historical Document. <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/psychology-of-intelligence-analysis> (Posted: Mar 16, 2007 01:52 PM. Last Updated: Jun 26, 2013 08:05 AM.)

[5] Frank Hughes, “A Tutorial on Credibility Assessment for Various Forms of Intelligence Evidence,” Joint Military Intelligence College, Defense Intelligence Agency, Revised 5 April, 2003.

[6] Ali Ben Mrad, Veronique Delcroix, Sylvain Piechowiak, Philip Leicester, and Mohamed Abid, “An explication of uncertain evidence in Bayesian networks: likelihood evidence and probabilistic evidence,” *Applied Intelligence*, published online 20 June 2015.

[7] Robert Schrag, Joan McIntyre, Melonie Richey, Kathryn Laskey, Edward Wright, Robert Kerr, Robert Johnson, Bryan Ware, Robert Hoffman, “Probabilistic Argument Maps for Intelligence Analysis: Completed Capabilities,” *16th Workshop on Computational Models of Natural Argument*, 2016.

[8] Robert Schrag, Edward Wright, Robert Kerr, and Robert Johnson, “Target Beliefs for SME-oriented, Bayesian Network-based Modeling,” *13th Annual Bayesian Modeling Applications Workshop*, 2016.

[9] David A. Schum, *Evidential Foundations of Probabilistic Reasoning*, Wiley and Sons, 1994.

[10] Edward Wright, Robert Schrag, Robert Kerr, and Bryan Ware, “Automating the Construction of Indicator-Hypothesis Bayesian Networks from Qualitative Specifications,” Haystax Technology technical report, 2015, <https://labs.haystax.com/wp-content/uploads/2016/06/BMAW15-160303-update.pdf>.

A CRAFT node and link types

Table 2 and Table 3 describe the roles of different node and link types in our proposed intelligence domain-specific variant of FUSION called CRAFT. We anticipate distinct GUI icons for these types.

Table 2. CRAFT will distinguish among model statement types in ways that are salient to intelligence analysts.

Outcome	A Hypothesis distinguished as a primary intelligence analysis alternative—e.g., a possible future (or present, or past) situation.
Hypothesis	A statement whose probability of truth depends on upstream-neighbor links.
Evidence-Hypothesis	A statement that is in principle knowable from evidence available in the problem domain.
Evidence-Report	A statement reported by some source. Absolutely influences a like-content EvidenceHypothesis that is source-independent and accommodates any number of SupportedBy or RefutedBy links from EvidenceReports corresponding to different reports.
Credibility (of EvidenceReport)	A statement characterizing the credibility of an EvidenceReport. Per modeler discretion, either a single node or a conjunctive Logic statement summarizing relevant credibility factors.
Assumption	A statement posited by an analyst to fill a gap in available information. Plays a modeling role similar to EvidenceReport, when evidence is unavailable.

Validity (of Assumption)	A statement characterizing an analyst’s self-assessed legitimacy of a posited Assumption.
Logic	A propositional logic expression over other statement nodes, either summarizing or constraining them.

Table 3. CRAFT will support probabilistic argument map model link types (center column). All link types are binary.

Downstream statement	SupportedBy	Upstream statement
	RefutedBy	
	MitigatedBy	
	RelevantIf	
	LogicInput	

B CRAFT credibility reasoning

CRAFT will apply to a HUMINT EvidenceReport statement Schum’s framework [6] that assesses Credibility with respect to four attributes of the reporting agent: Veracity, Objectivity, Competence, and Opportunity to observe what’s been reported. A reported statement is believed credible only if all four of its factor statements are believed true, so we link the EvidenceReport as RelevantIf this conjunctive Logic statement holds. Note that the factor statements are themselves Hypotheses, subject to supporting and refuting statements bearing potentially rich argument structure. Note that direct or indirect corroboration will increase an EvidenceReport’s Credibility.

Below we outline how CRAFT will map (in *italics*) specific analytical tradecraft standards (in **bold**) regarding items of evidence per ICD 203 [2] to Schum’s factors or, as appropriate, other model elements.

Factors from ICD 203 D.6.e.1:

- Accuracy (use Credibility) and completeness (explicitly model omitted possibilities as Assumptions or other appropriate statements)**
- Possible denial (use Opportunity) and deception (elaborate a deceptive course of action as an Outcome or other Hypothesis—see SaddamMaintainsDiplomacy in Figure 1, e.g.)**
- Age and continued currency of information (use temporal relevance)**
- Technical elements of collection (apply true/false positive/negative sensor models, confusion matrices—ancillary to Schum’s framework—see also [5], discussed below)**
- Source access (use Opportunity)**
- Validation (model as corroboration by other sources)**
- Motivation (model as in Figure 1, e.g.)**
- Possible bias (use Objectivity)**
- Expertise (use Competence)**

More factors from ICD 206 [3] appendix A (glossary), “source descriptor:”

- Precision or technical quality (see 1, 4)**

11. **Context** (*connotes scope or socio-cultural setting—capture via other model statements, as appropriate*)
12. For human sources:
 - a. **Level of access** (*use Opportunity*)
 - b. **Past reporting record** (*associate with base rate/prior for any credibility factor—including/especially Veracity*)
 - c. **Potential biases**—e.g., political, personal, professional, religious affiliations (*counter-indicators for Objectivity*)

From ICD 206 appendix D.5:

13. **Potential strengths and limitations of available information** (*see 1*)
14. **Notable inconsistencies in reporting** (*impugn Credibility, per factor as warranted*)
15. **Important information gaps** (*see 1*)
16. **Other factors deemed relevant** (*model as appropriate*)

HUMINT is testimonial evidence. Hughes [5] addresses credibility also for tangible and sensor evidence and presents considerations that may inform argument structure affecting each of Schum’s testimonial evidence credibility factors. He also tenders the factor “observational sensitivity”—which may be as relevant to sensing by humans as it is to sensing by devices—and addresses authenticity, chain of custody, and primary vs. secondary (or tertiary, ...) sources. CRAFT might be engineered to support some of these refinements explicitly. CRAFT’s credibility reasoning will remain explicitly probabilistic, however, exploiting the FUSION foundation, rather than assessed ad hoc as in the later sections of Hughes’ report.