

A collaborative eye to the future

Kshanti A. Greene, Terrance Goan, Eugene R. Creswick¹

Stottler-Henke Associates, Inc.
1107 NE 45th St, Suite 310
Seattle, WA 98105

ABSTRACT

We present a novel approach to *predictive* situation awareness that leverages human insight to enhance the forecasting abilities of Fusion levels 2 and 3. Existent technologies fail to support predictive and impact modeling under realistic conditions, particularly when there exist few historic exemplars on which to base inferences or when full awareness of the situation includes unobservable elements. We report on our ongoing efforts to develop *FutureFusion*, a collaborative system that builds predictive awareness and enables futurists to visualize paths to possible futures and formulate predictions on the ultimate outcome of scenarios of interest. *FutureFusion's* human interpretable knowledge representation is unique in its ability to capture qualitative descriptions of possible futures and quantify them to build computational models. Further, *FutureFusion* captures both popular consensus as well as high-risk outliers, thereby reducing the potential for surprise. Finally, by efficiently diversifying the modeling process across a heterogeneous and distributed community of experts, this approach avoids the common pitfalls of more traditional modeling approaches. *FutureFusion* helps to cast light on blindspots, mitigate human biases, and maintain a holistic, up-to-date predictive and impact awareness.

1 INTRODUCTION

Advanced situation awareness (SA) methods can improve our ability to avoid crises and seize opportunities. Fusion levels 2 (situation assessment) and 3 (impact assessment) [20] incorporate information from diverse sources to assess the current state of the world and reach into the unknown when considering the impact of actions on an environment. However, there are many challenges that inhibit our ability to capture full awareness in uncertain environments, in particular situations that involve forecasting the effects of actions on an environment or predicting possible futures. The first bottleneck is to capture sufficiently rich information to develop models. Relying on autonomous systems such as sensor networks alone will restrict collected information to historic or observable signals and events. While this is useful for capturing a snapshot of the world or for considering the patterns that have occurred in the past, autonomous systems have difficulty developing insights about situations that have never been observed. Impact assessment requires the ability to extract and represent the causal relationships between events and predictive anticipation requires insight into possibilities outside the realm of collected artifacts.

We can diversify and expand our network of information sources to capture a greater perspective on a situation. Leveraging insight from human experts and collected human intelligence (HUMINT) will increase the relevance of predictive models immensely. However, diversity will also increase the amount of conflict, redundancy and noise. In addition, human nature introduces many other challenges that need to be overcome. First, combining human

¹ [kgreene, goan, rcreswick]@stottlerhenke.com

contributions is complicated by the vagaries of natural language and the qualitative nature of human judgments. Humans also base their intuition on experience and often have difficulty being objective in the face of uncertainty. If those contributing knowledge have insufficient coverage of all factors involved, then models will be biased and only capture a partial awareness of the situation and its implications. A solution that leverages both human insight and the computer's ability to capture and sift through large amounts of data will enable a richer perspective from which to build a sufficient awareness to enable decision-making.

FutureFusion exploits the conflicts that arise from a diversified knowledge base to develop an awareness of *alternative futures*— including those that are more likely and those that represent a higher risk to a system's wellbeing. Our approach builds predictive awareness models by combining the insights of many experts. The first step is to provide an intuitive interface that enables individuals to describe the issues at hand and relationships between them in a natural way. The second step is to form computational models from these insights that provide assessments on the likelihood of events and the potential impact of changes in an environment. Our approach transforms the qualitative judgments of many individuals into a quantitative assessment of a situation and its implications. We extend the concept of *scenario maps*, introduced by [2] as an approach to help decision-makers visualize the driving forces that lead a situation to multiple alternative futures. Scenario maps utilize natural language in a graphical representation to illustrate how major clusters of activity are linked to one another and what their important consequences are. In their original form scenario maps are typically theoretical but very useful as a visualization and problem-solving tool. Our *executable* scenario maps extend these capabilities to form mathematically sound models using Bayesian representation and reasoning.

This paper is organized as follows. We first discuss technologies that inspire *FutureFusion* in Section 2. Then we discuss our approach in Section 1, beginning with our unique approach to eliciting predictive information to build awareness. Section 3.1.1 describes our scenario maps for situation awareness and 3.1.2 highlights our method for eliciting probabilistic information from humans. In Section 3.2 we discuss our approach to turning qualitative input into quantifiable computational models. In Section 4, we highlight the discoveries we have made with limited user testing and conclude with a summary of our findings in our on-going research in Section 5.

2 RELATED WORK

FutureFusion takes inspiration from many different research fields, including data fusion, scenario development, collaborative knowledge elicitation, and Bayesian belief aggregation. We highlight these technologies below.

2.1 Multi-source Fusion

Existing data fusion systems typically merge output from sensor networks and other input streams (such as SIGINT) to develop a picture of an environment and help assess the situation [18], [14]. Several levels of abstraction iteratively build from low-level numerical input to high-level assessment about the behavior of entities in the environment and the impact of their actions. Human intelligence (HUMINT) fusion is a relatively nascent aspect of this technology as it is more difficult to integrate natural language due to semantic vagaries and the higher level of abstraction typically associated with human intelligence.

Wright and Laskey [18] use a Bayesian extension they developed for merging multiple sources that considers the validity of the source, incorporating *credibility models* containing information about evidence for a belief and the quality of the source into their Bayesian representation. While our approach also utilizes a Bayesian-like model, we focus on the challenges of eliciting and building models from human intelligence. We relax our reliance on Bayesian mathematics by making a number of independence assumptions, allowing users to build larger, more realistic models that incorporate varying beliefs. Similar to their representation of evidence, we also enable contributors to submit their sources as evidence, or to supply other reasoning behind their contributions, but we leave this out of the models for simplicity. Users instead can read this information and provide feedback on their agreement with another user's contribution. While Laskey's approach utilizes extensive modeling to determine the competence of an agent, we utilize the power of democratic collaboration to validate the accuracy of input. The credibility models rely more on building the evidence for and consensus of a particular belief versus enabling multiple alternative futures to emerge.

General Dynamics developed a system, called the *Human Processing Subsystem*, for the Air Force Research Laboratory [14]. This system builds awareness about targets by extracting information about the entities involved and the physical characteristics of an environment from multiple structured and unstructured text-based sources. Their approach uses pattern matching and natural language techniques to extract specific types of information from

a document. Their system is intended to be used in a forward/backward chaining reasoning system that makes high-level inference from information extracted based on information received and queries made by operators. This system and other information extraction approaches represent a complementary approach to our human-built models

2.2 Scenario Development

Most strategy and scenario development approaches are entirely human-driven and do not incorporate a computer-based system. *FutureFusion* is the first to attempt to utilize this story-driven approach to build computational models. Peter Schwarz, the cofounder of the Global Business Network, has developed an approach to scenario building that leverages human insight to extract indications of *inevitable surprises* [15], [16]. His approach involves enumerating key driving forces and determining what their ultimate effects are on an organization or environment. The result is a set of *alternative futures*. Eden and Ackermann [2] utilize a similar approach to build strategy based on visualizing scenarios as graphical maps, called concept or *scenario maps*.

2.3 Collaborative Knowledge Elicitation

A large body of research and development exists related to collaborative knowledge elicitation and sharing. In particular, Collaborative Innovation Networks (or CoINs) [3] have emerged as a highly effective method to connect like-minded individuals to cooperate on knowledge and system development. CoINs follow the philosophy that the power of the people is greater than a small, closed group of “experts”. Two excellent artifacts of CoINs are Linux and Wikipedia. Over time, the quality and richness of both of these tools has increased considerably because their popularity encourages many knowledgeable individuals to contribute, and consensus helps to remove spurious contributions and fix inaccuracies. However, enforcing consensus also can have some negative side effects. In some cases, particularly when dealing with future unknowns, the events that present the greatest risk often arise as outliers to the more popular beliefs.

Peak-a-boom, created by Carnegie Mellon University [17], is a web-based tool that integrates many contributions to build knowledge systems using free human labor that would otherwise be too time- and cost-intensive. Peak-a-boom builds a vision database by incorporating tags submitted by users on image content. A unique aspect of their approach is that they attempt to encourage more contribution by turning their tool into a game. While users may find the game fun to use, they do not realize that they are actually helping to build a rich body of knowledge. Encouraging contributions is a very difficult task in collaborative knowledge systems [19]. People need to feel like they are getting a return on investment when they commit their own knowledge to a system shared by others. This can be through the appearance of monetary gain, power and popularity, or simply “fun.”

Google has taken on a market-based approach to prediction. They have conducted an internal experiment using bets placed on the likelihood of some event [1]. They have made several discoveries relating to the social aspect of collaboration; for example, employees who work closely together often make similar predictions. Google’s approach focuses only on the prediction of single events, not on developing complete models integrating causal relationships between entities and events. Therefore, it is significantly simpler than *FutureFusion* but is interesting as a study in social networks and market-based prediction.

2.4 Bayesian belief aggregation

Bayesian models are often used to represent situation awareness in fusion systems [11]. Our particular approach involves integrating beliefs by many experts into a graphic model based on Bayesian reasoning. Bayesian belief aggregation [9] involves combining probabilistic beliefs of a group of individuals to generate a consensus belief. It is used to calculate a probability distribution from the joint probability distributions of n agents. The function that combines beliefs is called an *opinion pool* function. The most common opinion pool functions are LinOP and LogOP. LinOP (Figure 1) is a weighted mean of the contributor’s probabilities. LogOP (Figure 2) is a weighted geometric mean. In the figures, $Pr_i(\omega)$ represents each contributor’s opinion of probability and α_i represents the weight of the opinion, such that all weights sum to 1.0. LogOP has been found to more accurately represent conditional relationships as LinOP destroys some of the independence assumptions inherent in Bayesian networks [10].

$$Pr_0(\omega) = \sum_{i=1}^N \alpha_i Pr_i(\omega)$$

Figure 1: LinOP opinion pool function.

$$Pr_0(\omega) \propto \prod_{i=1}^N [Pr_i(\omega)]^{\alpha_i}$$

Figure 2: LogOP opinion pool function.

3 APPROACH

3.1 Leveraging human insight for holistic predictive SA

3.1.1 Scenario maps for situation awareness

FutureFusion was developed for an Air Force Phase II SBIR program entitled *Effective Behavioral Modeling and Prediction Even When Few Exemplars are Available*. Our approach is motivated by the inability of automated systems to anticipate events that pose a high risk to an organization. The nature of surprises implies that few observable indicators exist for automated systems to capture. In contrast, the same insight that enables humans to adapt to new situations allows us to imagine extensive what-if scenarios from a single event. To overcome the limitations of purely automated approaches, we leverage human insight to brainstorm and refine predictive awareness models. We utilize the wisdom of the crowd to expand awareness and overcome any biases as well as to temper overactive imaginations and focus on issues of concern. This iterative process of divergence and convergence causes the likely alternative futures to come to the forefront but also allows less plausible, but high-risk, outliers to remain in our awareness.

FutureFusion incorporates concepts from cutting-edge collaborative knowledge systems such as *Wikipedia* and *Yahoo! Answers* with strong foundations in both human-centered and mathematical predictive modeling approaches. We take inspiration from the field of scenario development [15], [2] in which futurists generate written and sometimes graphical scenarios that enumerate key driving forces and their effects on an outcome of interest. The result is a set of *alternative futures* often graphically represented by a *scenario map*. Figure 3 shows a scenario map representing a hypothetical situation following the 2005 Indian Ocean tsunami. This approach is more natural for human decision makers to undertake than mathematical predictive approaches such as Bayesian modeling [4] because mathematical models require modelers to understand the more complicated implications of relationships and also require quantitative estimates of likelihood for each combination of influential variables. We use information retrieval techniques to enable reuse and avoid information redundancy.

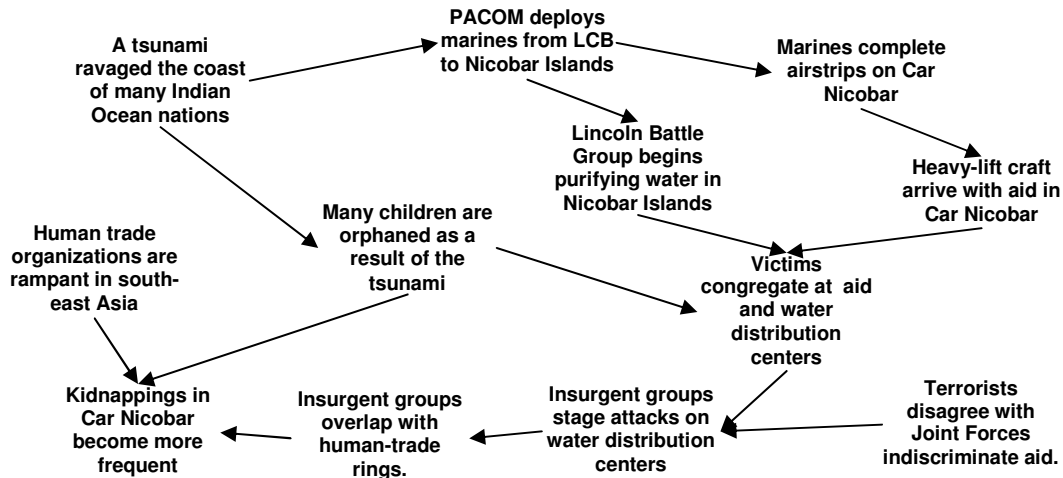


Figure 3: Scenario map showing how the Asian tsunami could have led to human trafficking.

The intuitive *FutureFusion* interface allows users to build scenario maps containing events and the influential relationships between them. Our scenario maps are composed of two main components: *predictions* and *influences*. *Predictions* have an associated likelihood that *FutureFusion* computes from user input and influences on the prediction. While “prediction” may imply only a future unknown, we also include facts and historical events as predictions since they are open to interpretation by different observers and experts. *Influences* represent the relationships between predictions. The relationships will typically be causal, temporal, or otherwise define a dependency between two entities. An influence has a direction and strength associated with it that quantifies the impact one entity has on another. The impact can be positive (increases likelihood) or negative (decreases likelihood).

Figure 4 shows the web-based *FutureFusion* interface displaying details of the prediction, *reduced vulnerability of coalition forces*. The Influence Panel (bottom half of figure) lists the prediction's incoming and outgoing influences separately. Incoming relationships (in the white area of the influences tab) show other predictions that change the likelihood of the current prediction. Outgoing relationships (in the gray area) list any predictions that the current prediction influences. The strength of each relationship is represented as a verbal description (e.g., "Slightly increases chances") to be consistent with the qualitative nature of human judgment. Figure 5 shows the graphic representation of the scenario map. We have found that users prefer to edit scenario maps utilizing the localized, textual interface shown in Figure 4, while the graph shown in Figure 5 is more conducive to visualization. The textual interface enables users to focus on local relationships without becoming overwhelmed by the size of the model, while the graphic interface shows holistic system behavior.

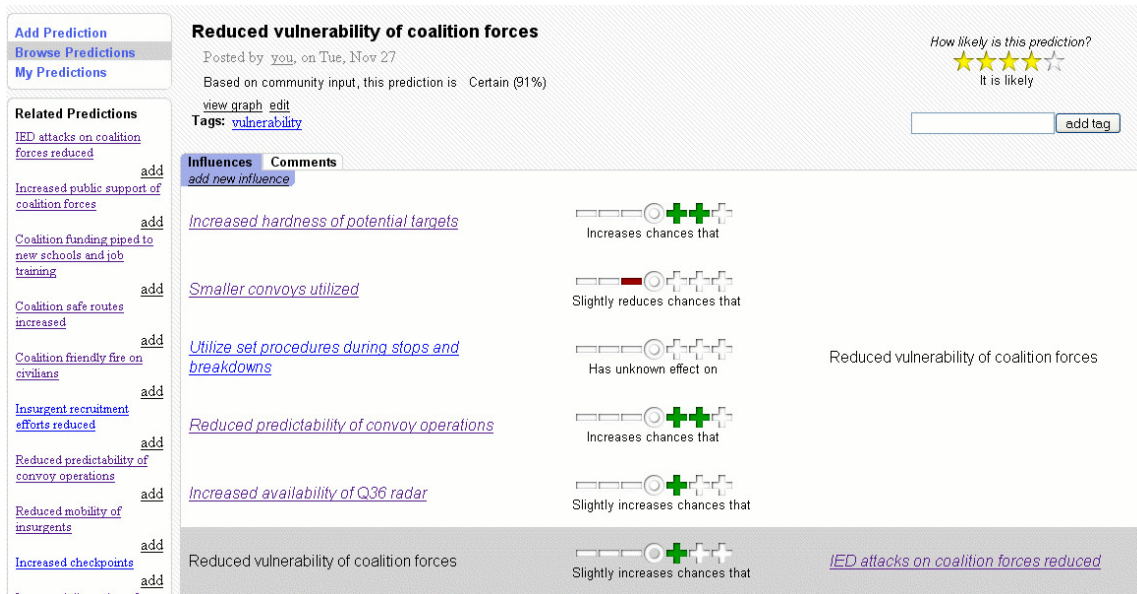


Figure 4: The FutureFusion interface showing a prediction and its influences.

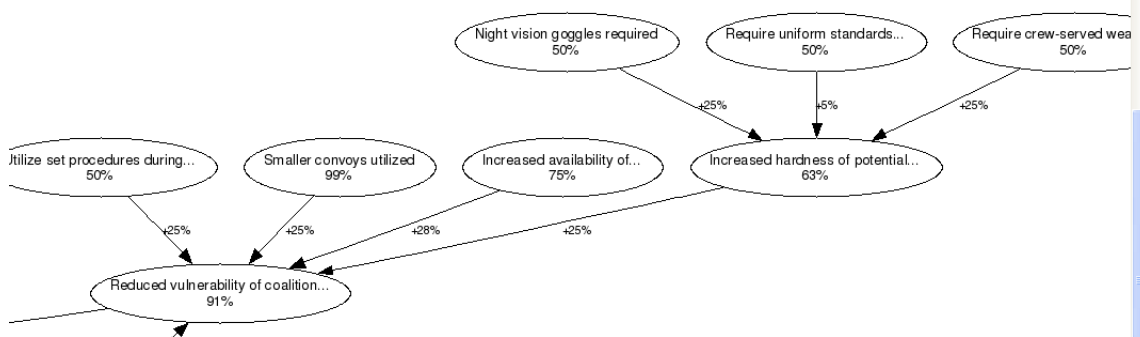


Figure 5: Segment of graph generated from the predictions and influences in Figure 4.

A scenario map is a connected network of predictions held together by their influences. Users can add new predictions to a scenario map by adding them as an influence or they can create a new scenario map by making a new prediction with no initial influences. When predictions in unconnected scenario maps are joined with an influence, they form larger scenario maps that represent the interactions between smaller systems, also known as *systems of systems*. To help encourage merging scenario map fragments to build larger collaborative representations, we use information retrieval (search) tools to find existing predictions related to those currently being viewed. *FutureFusion* displays these related predictions on the left side of the screen in Figure 4, which can easily be added as incoming or outgoing influences. Similar retrieval techniques are used to help avoid information redundancy, which could easily become a problem in a multi-user system. *FutureFusion* will flag newly submitted predictions that seem to match existing ones. Figure 6 shows the dialog that notifies the user if a similar prediction already exists.

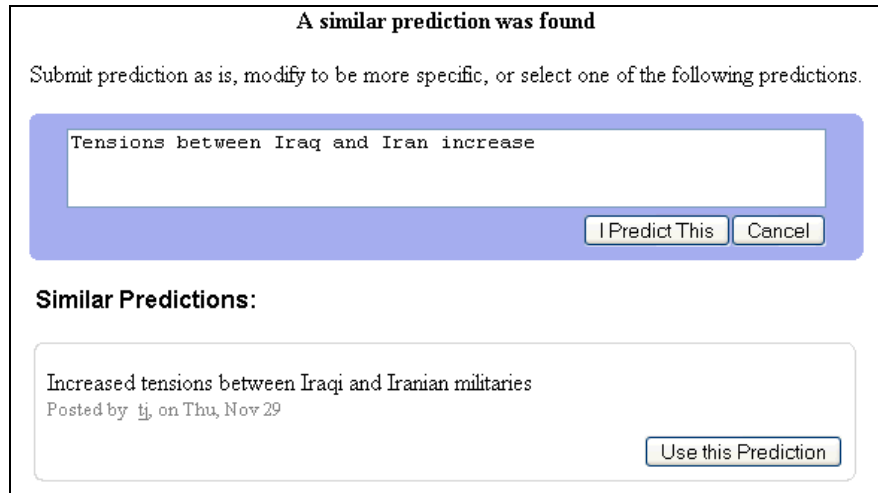


Figure 6: Dialog indicating that a similar prediction already exists.

3.1.2 Eliciting qualitative estimates of likelihood

Our goal is to reduce the cognitive load on human contributors and enable them to supply as few probability estimates as possible, while still maintaining rich quality of information. The elicitation approach used in *FutureFusion* is informed by research that shows that non-mathematically oriented people tend to become intimidated by requests for probabilities and are more comfortable using qualitative terminology that allows them to provide analysis in words using relative terms [12], [5]. We can also reduce the number of value estimations needed from users to infer likelihood by using the independence assumption [6] that considers each individual influence’s effect on a prediction independent of all other influences. This is in contrast to full Bayesian models that require a probability estimate for each possible combination of influences. In the example in Figure 4 and Figure 5, the five incoming influences would require $2^5 = 32$ estimates (exponential on the number of incoming influences) for a Bayesian conditional probability table. In reality, many influences are independent, for example, *smaller convoys utilized* is unrelated to the *increased availability of Q36 radar*. Using independence assumptions, we reduce the number of estimates in our example to $5+1=6$ (linear on the number of incoming influences).

FutureFusion users supply feedback on the *a priori* prediction likelihood and the *strength* of each influence. A priori likelihood represents the probability estimate of a prediction before a new situation is considered. In this case, the new situation includes the impact of any incoming influences. Thus, a priori likelihood represents the probability that the prediction occurs without knowledge of the known influences, or alternatively, the probability of *unknown* influences on a prediction. This distinction points to another limitation of Bayesian models. Typical Bayesian networks only represent *known* variables in a limited world. However, in reality it is very difficult to capture all variables that bring about an effect. For example, in Figure 4, we could easily be missing other influences that increase or decrease the likelihood of *reduced vulnerability of coalition forces*; however, users are only able to enumerate influences that they are aware of. Eliciting a priori probability allows users to qualify their uncertainty about the causes of a prediction, enabling a richer, more realistic model.

Users also supply their opinion on the *direction* and *strength* of each influence. This approach is based on Causal Strength (CAST) logic [13]. These influences combine with the a priori likelihood to form the prediction likelihood given the current situation. An influence can either be implicative, meaning that it implies or causes the prediction, or it can prohibitive, meaning that it reduces the chances of a prediction occurring. The strength of the influence indicates how much it increases or decreases the likelihood of a prediction. In Bayesian logic, one would supply a conditional probability indicating the likelihood of one event given another, or in mathematical notation; $p(y | x)$ meaning the probability that y is true given that x is true. Causal strength is roughly proportional to: $p(y | x) / p(y)$, where $p(y)$ is the a priori likelihood of a prediction. Again, we aim to reduce the cognitive load on non-mathematically inclined users, therefore we enable them to think in more natural qualitative terms such as “increases the probability” instead of requiring conditional probabilities.

FutureFusion is a collaborative application and allows all users with appropriate permissions to provide feedback on any prediction. Based on the combined feedback on a priori likelihood and influences, *FutureFusion* computes the ultimate likelihood of a prediction (indicated below the prediction text in Figure 4 as “Based on community feedback, this prediction is Certain (91%)”).

3.2 Inferring quantitative executable scenario maps

We use the qualitative feedback elicited from users to make our scenario maps *executable* by enabling them to output quantified probability estimates for predictions. Instead of eliciting continuous, mathematical values, we allow users to select from a set of discrete values that represent increments between the minimum and maximum values (in the case of probability; between 0.0 and 1.0). Rarely will even highly educated experts be able to quantify likelihood within a few percentage points. Based on experimentation and analysis, we map the discrete values to appropriate continuous values. These values can be configured for special situations that require more refinement, for example, when the likelihood of most events of interest is very low.

As stated previously, we currently use a form of causal strength logic [13] to compute the likelihood of a prediction based on its a priori likelihood and the likelihood and strengths of its influences. The first step is to find the a priori likelihood and strength for each influence based on the feedback from all users. Merging probability estimates is a key aspect of *belief aggregation* research [9] described in Section 2.4. We initially are using the simpler LinOp method, which is essentially a weighted average. However as we refine our computation we may utilize the more structurally sound LogOp formula. While currently we weight each contributor equally, it would be a trivial extension to weight contributors by experience or correctness of contributions. The second step is to compute the ultimate likelihood of a prediction from the aggregated a priori likelihood and influence strengths. Figure 7 shows the algorithm used to compute $P'(X)$, or the posterior probability of a prediction X . Each influence on X , $I(Y_k, X)$, where Y_k is the causal prediction in X 's k^{th} incoming influence, moves the probability of X closer to 1.0 (or *certain*) if it is positive, or closer to 0.0 (or *impossible*) if it is negative.

$$P(X) = \frac{\sum_{i=1}^n P(x_i)}{n} \quad \text{where } x_i \text{ is user } i\text{'s a priori belief in } P(X), \text{ } n \text{ is the number of user feedbacks}$$

$$P'(X) = P(X)$$

for each influence $I(Y_k, X)$ on $P(X)$ where Y_k is the causal prediction in the k^{th} incoming influence on X . $I(Y_k, X)$ is the causal strength of Y_k on X

$$I(Y_k, X) = \left(\frac{\sum_{j=1}^m I(Y_k, X)_j}{m} \right) * P'(Y_k) \quad \text{where } I(Y_k, X)_j \text{ is user } j\text{'s belief in the strength of } I(Y_k, X), m \text{ is the number of user feedbacks. } P'(Y_k) \text{ is the posterior probability of } Y_k.$$

if $I(Y_k, X)_j < 0$ then
 $P'(X) = P'(X) + I(Y_k, X) * P'(X)$
else
 $P'(X) = P'(X) + I(Y_k, X) * (1 - P'(X))$

Figure 7: Algorithm used to calculate the probability of a prediction given its a priori probability and probability and strength of its influences.

Most scenario maps will be composed of many influential relationships chained together. The ultimate likelihood of a prediction is the result of propagating the likelihood up from its root causes, a process known as *inference* in graphical models. *FutureFusion* utilizes Judea Pearl’s well-known *belief propagation* inference algorithm [8]. This algorithm uses the concept of message passing in which predictions send their updated probability to their effects, as opposed to recursively requesting probability from their causes. Pearl’s belief propagation algorithm is found to be very efficient as well as elegant. Our dynamic adaptation of the algorithm propagates changes in the model only to those predictions effected by the change, reducing computational load over re-computing the whole model. The final benefit of the belief propagation algorithm is that it works even on networks that contain loops (instances where a prediction is a cause of itself) [7]. This has significant positive implications for *FutureFusion*. While

typical causal models are required to be acyclic, feedback loops (cycles) occur frequently in real-world situations. This puts an unrealistic limitation on simulations created with most causal representations. However, because we are using loopy belief propagation in our *FutureFusion* system, we have successfully represented and inferred models containing feedback loops.

Since we are making several simplifying assumptions that approximate the approach of more knowledge- and computation-intensive Bayesian reasoning, we have not made any claims as to the accuracy of our approach. Our initial goal was to assess the validity of our process and elicitation techniques. Since we are collecting human supplied estimates of likelihood on typically abstract concepts or events that may or may not occur in the long-term future, it is difficult to judge accuracy. However, we aim to integrate observations from sensor networks and other external sources that will provide more concrete likelihoods and help validate results. Therefore, our continued goal is to reduce the computational error between our approach and formal Bayesian methods. We are working on methods to retain the benefits of causal strength logic (i.e., easier comprehension by users) while improving its posterior probability estimate. We are currently investigating a technique that derives conditional probability tables from causal strength and utilizes the conditional independence assumption from the Naïve Bayes classifier [6] to compute probability. In addition, our research in collecting human intelligence will enable us to determine where refinement will promote a more accurate computational model.

4 Discoveries

Our primary goal in developing *FutureFusion* has been to eliminate the bottleneck in collecting and integrating insights to build an awareness of potential futures. Autonomous extraction techniques are useful for mining existing information and sensor networks improve our ability to capture real-time observations; however only human insight can provide a vision of futures that have minimal observable evidence. On the other hand, existing modeling techniques for capturing human knowledge are tedious and too opaque for typical users. *FutureFusion* has opened predictive modeling to enable users from diverse backgrounds to collaborate on building a vision of the future. We made several key discoveries in our research and during our focused user testing.

In our initial research and development cycle, we found that providing simple and general model components is more effective than attempting to represent specific types of predictive information. We initially planned to provide two types of “nodes” in our scenario maps: *events* and *evidence*. However, we often found that it was difficult to distinguish between the two, since sometimes evidence for an event is itself an event. In addition, evidence is often so tightly correlated with an event that it is difficult to separate them for modeling purposes. Therefore, we provide only one type of node, called a *prediction*, that can be used to represent many types of information with temporal, physical, or conceptual characteristics. Instead of making evidence explicit in the models, we provide an interface associated with each prediction for users to submit comments and external evidence to validate their predictions.

We ran two focused tests with groups of 4-5 users working concurrently on a common topic. Our goal was to validate our elicitation approach and its ability to encourage user contributions. We found that in this focused, concurrent environment, users became engrossed in the process of submitting predictions and providing feedback. Our dynamic updates that immediately display new contributions to other users helped to increase community awareness and enabled users to *play off* each other’s submissions. We also found that *FutureFusion*’s computational representation helped discover discrepancies between user expectation of likelihood and the developed models. For example, a prediction’s a priori expectation sometimes varied drastically from the combined effect of its influences. This occurred with the prediction “Microsoft Zune market share outpaces iPod by 2010.” While the a priori likelihood was quite low, users tended to create positive influences on the prediction, increasing the overall likelihood of an event beyond expectations. This may have been due in part to biases that caused users to submit influences that would increase Microsoft’s market share because either they were playing the devil’s advocate or they wanted the iPod to have some healthy competition. In any case, this example demonstrates the difficulty of capturing all influences on a prediction. As a result, we have adjusted the a priori likelihood to explicitly represent belief in the *unknown* influences on an event. We are also improving our probabilistic inference algorithm to more accurately represent the combined effect of influences on a priori likelihood. These enhancements will increase the realism of our approach over traditional Bayesian models that typically only represent a limited view of the world, yet make assumptions of completeness. In a real world situation there will be many variables that cannot be made explicit.

We also released *FutureFusion* to a small group of web users that used the system on their own time, without a community of concurrent users. In these cases, users tended to submit their own scenario map fragments that did not connect to other user's contributions. We concluded that a greater awareness of the community is needed to encourage non-concurrent users to share and provide feedback on each other's contributions. We are currently integrating new components that display community feedback to help increase awareness. This will help pull users into a more interactive and engrossing discussion of possibilities. While we would like to see improvements with more casual individual users, we were pleased with the results of our focused group tests as these environments more closely match our intended user groups in such applications as strategic planning, counter-terrorism and crisis management.

5 Conclusions

In this paper, we have described the progress we have made in developing the *FutureFusion* system – a collaborative environment for building predictive awareness. The evaluations conducted to date have demonstrated the value of this unique knowledge elicitation and predictive modeling approach. *FutureFusion* provides an intuitive and rapid story-based approach to collecting knowledge from potentially large and diverse sets of contributors. By diversifying the modeling effort across a broader community of experts, *FutureFusion* is able to mitigate the threats posed by human bias and blindspots. Further, by supporting synchronous and asynchronous collaboration (without the need for a facilitator), this approach to predictive knowledge modeling encourages the continuous revision of predictive models as events occur and new insights arise. Such 'living' models offer obvious advantages over the static models that typically result from the more costly war-room style modeling approaches.

6 Acknowledgements

This project was funded through the Air Force Research Laboratory (AFRL) Small Business Innovative Research (SBIR) contract FA8750-05-C-0042. We would very much like to thank AFRL's Michael Hinman and John Salerno for their guidance and assistance in this project.

7 References

- [1] Cowgill, B., Wolfers, J. and Zitzewitz, E., "Using prediction markets to track information flows: evidence from Google", white paper, <http://bocowgill.com/GooglePredictionMarketPaper.pdf> (2008).
- [2] Eden, C. and Ackermann, F., [Making Strategy, The Journey of Strategic Management], SAGE Publications, London (1998).
- [3] Gloor, P., A., [Swarm Creativity: Competitive Advantage through Collaborative Innovation Networks], Oxford University Press, USA (2005).
- [4] Jensen, F. V., [Bayesian Networks and Decision Graphs], Springer-Verlag, New York (2001).
- [5] Jenkinson, D., "The elicitation of probabilities-a review of the statistical literature", Bayesian Elicitation of Experts' Probabilities (BEEP) working paper (2005).
- [6] Lewis, D., D., "Naive Bayes at forty: The independence assumption in information retrieval", Proc. of European Conference on Machine Learning, 4-15 (1998).
- [7] Murphy, K., A brief introduction to graphical models and Bayesian networks, online tutorial, <http://www.cs.ubc.ca/~murphyk/Bayes/bnintro.html>, (1998).
- [8] Pearl, J., [Probabilistic Reasoning in Intelligent Systems], Morgan Kaufmann, San Francisco (1988).
- [9] Pennock, D., M. and Wellman, M., P., "Graphical representations of consensus belief", Fifteenth Conference on Uncertainty in Artificial Intelligence, 531-540 (1999).
- [10] Pennock, D., M. and Wellman, M., P., "Graphical models for groups: belief aggregation and risk sharing", Decision Analysis 2(3), 148-164 (2005).
- [11] Punska, O., "Bayesian approaches to multi-sensor data fusion", dissertation submitted to Cambridge University for the degree of Master of Philosophy (1999).
- [12] Renooij, S. and Witteman, C., L., M., "Talking probabilities: communicating probabilistic information with words and numbers", International Journal of Approximate Reasoning, 22, 169-194 (1999).

- [13] Rosen, J., A. and Smith, W., L., "Influence net modeling with causal strengths: an evolutionary approach", Proc. of the Command and Control Research Symposium, Naval Postgraduate School, 699-708 (1996).
- [14] Schiller, M., Gucwa, J., Neal, J., Crystal, M. and Weischedel, W., "HUMINT extraction and fusion system", Final technical report for General Dynamics contract #AFRL-IF-RS-TR-2005-12 (2005).
- [15] Schwartz, P., [The Art of the Long View: Planning for the Future in an Uncertain World], Doubleday, New York (1991).
- [16] Schwartz, P., [Inevitable Surprises], Gotham Books, New York (2003).
- [17] von Ahn, L., Liu, R. and Blum, M. "Peekaboom: a game for locating objects in images", Proc. Computer Human Interfaces (CHI) (2006).
- [18] Wright, E., J. and Laskey, K., B. "Credibility models for multi-source fusion", in Proceedings of the 9th International Conference on Information Fusion (2006).
- [19] Wu, S., Lin, C. and Lin, T-C., "Exploring knowledge sharing in virtual teams: a social change theory perspective", in Proceedings of the 39th Annual Hawaii International Conference on System Sciences (2006).
- [20] Yannone, R., M., "Exploring architectures and algorithms for the 5 JDL/DFS levels of fusion required for advanced fighter aircraft for the 21st Century", whitepaper found at <http://stinet.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA391672> (1999).