Using Saliency Detection and Anticipation to Predict Surprise Events

The Challenge

As evidenced by recent history, there is a need to predict "surprises." At the strategic level, surprises include events such as the Pearl Harbor attack, the 9-11 Terrorist Attacks, the global financial turbulence of July-August 2007 and the Arab Spring of 2011. Events such as D-Day or even everyday events such as the moment at which forest fires reach criticality are examples of tactical surprise. This paper will outline an array of technologies and methodologies that can be integrated to provide improved event prediction capabilities.

While many aspects of intelligence handling have been blamed for these unpredicted (or under predicted) surprises – such as delayed communications, failure to integrate multiple sources and misunderstood context – an underappreciated contributor is the lack of methods available for enabling analysts and policy makers to anticipate or predict surprises. While the phrase "anticipating surprise" could be considered an oxymoron, it also is a necessary capability that the intelligence community needs to develop for analysts at both strategic and tactical levels.

The 9-11 Terrorist Attacks and the Pearl Harbor Attack were not anticipated sufficiently because of several factors, including a lack of integration of intelligence data, a historical and cultural misunderstanding of the mindset of the adversary and a lack of anticipation of the potential for this type of surprise.

Similar lists of cause and effect can be produced for each of the events that were surprises when they happened, but that could arguably have been anticipated. Therefore, the Intelligence Community needs the capability to anticipate surprise that spans from strategic waypoints in our future to tactical events that could change our future, such as when/if China will attempt to take Taiwan, when/if a transnational Islamic Caliphate will arise at the strategic level or when the next major terrorist attack will occur at a tactical level.

Using the first one as an example, identification of which factors are important for predicting another nation's intent is the first task. For instance, the fact that China holds more foreign debt (approximately 1 trillion dollars of US debt), has a growing middle class, assumes a historical claim to Taiwan, and have a long established priority on the development of cyber weapons, provide several potential factors to analyze. These are all easy to quantify, but hard to put into perspective. On the other hand, the potential transnational Islamic Caliphate would appear to have more dependence on religious and tribal factors; which are hard to quantify or analyze. Thus, identification of the salient factors is a very complex task.

The Components of a Solution

The process of prediction will require the following major components:

1) Identification of salient factors and features
    a. Seamless integration of human and machine processing
    b. Multi-granularity inputs that span from signals to information

   c. Hierarchical outputs that span from small tactical to grand strategic

2) Tracking of salient factors and features

   a. On-line or dynamic learning processes

3) Prediction of when/if the combined factors will reach a tipping point

   a. Research into chaos, complexity, etc.

   b. Anticipatory systems research

The surprise event can be thought of as the avalanche or point at which several variables with stable trends come together in a way that causes a large non-linear change. The "knee of the curve," has been referred to as the "edge of chaos" or the point at which complexity is maximal. This is also known as the boundary between order and chaos. In order to predict the transition event, related phenomena must be sorted from all of the unrelated phenomena. Thus, by utilizing saliency detection, machine learning techniques and bio-inspired processing, very chaotic and highly complex information threads can be mined for trends or events that are salient but that may have been overlooked by a human reviewer.

The Solution

Salient Feature and Factor Detection

First, the solution will consist of a system that integrates human-derived and machine-learned inputs to find the salient factors and features within all available data that are relevant to a given prediction. As mentioned above in the Components of a Solution, this will require the integration of multiple granularities of inputs and hierarchical outputs. A prime example of the mixing of human-derived expert opinions and detailed measurements is given by the search for gold-laden ship that had sunken 130 years before it was found.[1] The research team "turned history into mathematics" by incorporating both objective and subjective information into a Data Correlation Matrix. This is a very good example showing that 33 subjective error-prone human accounts written 130 years earlier could be combined with modern measurements and sensors to determine a very small search space out of a literal ocean of possibilities. Noting that "the more minds from divergent areas of thought you can get to agree on a given subject, the sounder the idea," Bayesian Analysis techniques were used to cluster the estimates. This provides a good model of how to merge expert opinions with detailed measurements and how to span multiple granularities of information.

That effort was accomplished over 20 years ago. Combining that philosophical framework with recent advances in machine learning, clustering techniques and hierarchical output methodologies would provide an adequate framework for integrating outputs from multiple sources.

Some aspects will be easier than others to formalize. For instance, in the China/Taiwan analysis, it is relatively straightforward to find historic data to attempt to predict how a trend in debt holdings can lead to a change in a nation's aggressiveness, but it is not as straightforward to connect a change in electromagnetic signature profiles to a strategic political decision. This will require both new software tools and organizational changes such as specialized teams that include signals analysts, traditional analysts and policy developers.

[1] Kinder, G. (1998). Ship of gold in the deep blue sea. New York, NY: The Atlantic Monthly Press., pp. 160, 217-222

Robust training algorithms that distinguish wheat from chaff must be developed. As noted by one source, one reason that the Pearl Harbor attack was not anticipated was that "officials had to deal with too many 'signals' or warnings, and too much background 'noise'."[2] This same challenge is often paralleled in machine-learning problems that focus on mining measurements for individual sensor outcomes.

Trend Analysis

This aspect of the solution can be more automated than the previous stage of the solution, given that the key factors and features of the data have been identified and, to the extent possible, quantified. However, the fact that it is more automated and can utilize machine learning techniques does not make it less complex. Dynamic learning techniques and technologies that improve machine learning over time are less mature than the more static clustering and feature selection technologies discussed above. This is yet another area that can draw from the way human's process and learn – bio-inspired signal processing. Human's have an amazing ability to "boot-strap" knowledge in a manner that appears to mix supervised and unsupervised learning techniques. These technologies will have to be incorporated with the feature selection capabilities in order to improve the fidelity, update and adapt to changing information.

[2] Crocker, C. (2008). The impenetrable fog of war: reflections on strategic surprise. Westpoint, CT: Greenwood Publishing Inc., p. 11 of Chapter 13.

Surprise Prediction

The most mathematically complex aspect of the Solution will be detecting the tipping point, avalanche or “knee in the curve,” such as the time at which a nation decides to take over another nation or dissolve a federation. The most controversial output of the Surprise Prediction methodologies will be determining “if” a given decision or commitment will occur. This will be coupled with the slightly less controversial output of predicting “when” it will occur. Any review of articles and books on this subject shows that this is a highly complex research area heavy on philosophical treatments and light on actual mechanisms for accomplishing them.

The actual prediction of the event is the most complex and hardest to quantify of all. Experts ranging from meteorologists to economists to political strategists often track trends in their respective variables with high fidelity and then make significant errors in predicting when/if an event will occur. It is the hope of this author that improving the fidelity of the previous steps (feature/factor selection, feature/factor training, historical comparisons) will make this final step more tractable. Within these, the most important stage is the development of a system that adequately combines human and machine-derived intelligence at a wide range of granularities. Previous systems that failed appeared to rely too heavily on expert opinions that led up to the region of a true tipping point without being supported by available measurable factors, or conversely systems that relied too heavily on detailed machine-derived intelligence without the context and feedback into the process that humans can provide.

The type of surprise detection required for this type of solution is called a “memory-rich” detection system, because it includes memory of the stimuli that caused the surprise. This approach borrows heavily from knowledge of biological processes in the human brain that learn from surprises. By backtracking from the surprise event to the stimuli that caused it, much can be learned about how the stimuli and surprises are related. This is supported mathematically by the field of inductive reasoning that attempts to find generalized models from sparse and “noisy” data. Translating these areas of research into applications that support prediction of surprise events will be the high-risk, high-reward part of the proposed solution.

Philosophical and Organizational Changes

In general, a philosophical commitment to pursuing the integration of human- and machine-derived intelligence must be undertaken in order to achieve a solution of this type. Of course, a philosophical commitment to the task of attempting to predict surprises is fundamental to any efforts described above. This has been described as a need for a wide array of inputs from a wide array of experts, who have means to process the information in a way that lead to decisive action.[3]

While this author’s knowledge of analytical organizations is limited, it appears that a more agile assignment of a wider array of experts is required. In particular, once policy makers identify an outcome of interest (e.g., when/if China will try to absorb Taiwan), teams that span from signals and data analysts to religious, economic, military and political experts that would add valuable context and could assemble for a limited time

[3] Ibid. p. 14 of Chapter 13.

(e.g., one year), direct the analysts to the best sources of historical data and some insights into factors, features and downselection criteria that should be of priority.

Conclusion

An unmet IC requirement to predict surprises has been identified. A solution has been outlined that integrates many proven technologies along with some emerging research areas. The three major aspects of the solution include 1) Feature/Factor identification and relevance assessment, 2) Feature/Factor tracking using dynamic learning and updated inputs and 3) Analysis that results in a prediction of when and if a given event is likely to occur. While the overall process has been utilized many times before, the wide array of granularities and the intense marriage of human- and machine-derived intelligence does not appear to have been achieved.

In addition to technological efforts, some philosophical and organizational efforts will be required. Namely, a philosophical approach that focuses on more tightly coupling human and machine-derived intelligence and organizational changes that allow more efficient assembly and re-use of experts as participants in teams that focus at the right times on the right portions of a prediction.

In this author’s opinion, the above described solution could be achieved in approximately two – three years from a technological point of view; since many of the enabling capabilities are already functional. The timeline to effect philosophical, organizational and funding changes is harder to assess, but should be at least partially achievable within the same timeframe.

Michael R. Moore has 24 years experience at Oak Ridge National Laboratory supporting National Security missions; spanning signals analysis and machine learning. He is the Deputy Director of both the Cognitive Radio Program and the Signals Solution Center. He has written several technology forecasts for DHS and DOD agencies.