Probabilistic Early Warning for National Security Crises

Elisabeth Paté-Cornell, Stanford University
David Blum, Stanford University
mep@stanford.edu
dmblum@stanford.edu

1. Executive Summary

The primary focus of the research for the CREATE project in the department of Management Science and Engineering at Stanford University has been to apply the tools of risk analysis to the 'Indications & Warning' function of intelligence analysis in order to improve the timeliness and accuracy of crisis warnings. We are building on past efforts to incorporate Bayesian assessment techniques into the analysis of intelligence, in particular efforts at the CIA in the 1970's, but also academic efforts, including those of Paté-Cornell (2001), McLaughlin and Paté-Cornell (2005), Stech and Elsaesser (2007), and Pinker (2007). We address three shortcomings identified explicitly by CIA methodologists that are central to the nature of warning intelligence, but were not adequately addressed in subsequent academic literature. These issues are:

I. Introducing dynamics into the Bayesian approach to both the evolution of the crisis and the probability assessment to account for the fact that signals are not received all at once, but rather, over time

II. Incorporating the probability estimates into a decision to issue a warning that optimizes the satisfaction of the policy maker with consideration of lead time

III. Capturing the dependences between intelligence signals without undue computational burden

Our approach utilizes a Partially Observable Markov Decision Process (POMDP) to model an underlying crisis scenario and integrate real-time Bayesian assessments of relevant intelligence data. It is a model that belongs to the larger class of Bayesian forecasting models explored by Harrison and Stevens (1976) in their seminal work. When solved using Bayesian assessment techniques assisted by computer software, the model yields an optimal decision whether to issue a warning (and at what level of urgency) or wait after each intelligence report is received and assessed. We also explore how to incorporate geolocational data regarding a crisis event into the POMDP in scenarios where uncertainty in the location of a potential crisis is a salient feature, as is the case when an analyst must assess local risk. Finally, we explore three other issues that are central to any formal model of crisis warning: effects of disaggregated decision making (multiple layers in the chain of command), means of effective scenario generation, and means of simplifying probability assessments (the latter being widely applicable to any model that employs Bayesian forecasting).

This research was supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under award number 2010-ST-061-RE0001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security, or the University of Southern California, or CREATE.

In the near term, we intend to apply this model to forecast both strategic-level and operational-level crises, possibly including nuclear weapons development and/or regime collapse (strategic level), and terrorist operations (operational level). The latter will exhibit uncertainty in the location of a crisis, not just the timing of a crisis. Eventually we expect the model to be incorporated into the set of tools in use at intelligence and law enforcement agencies to actively monitor the development of crisis scenarios in both a national and homeland security context, where analysts will use the tool to help them perform better their warning function.

Keyword 1: intelligence analysis
Keyword 2: indications and warning
Keyword 3: Bayesian forecasting

2. Research Accomplishments

2.1 Generalized model

We began by proposing a simple procedure to assess the dependent signals in the context of a static, traditional Bayesian network representing a crisis scenario. With the help of Dr. Thomas Fingar (Stanford CISAC), we refined the procedure through the use of a an illustration, a nuclear weapon entering a US port in a container, in which we performed Bayesian inference on a stream of notional intelligence containing a mixture of dependent and independent signals. The refined procedure consisted of the following steps:

- Create a list of intelligence signals received since the model was initiated, and when a new signal is received, add it to the list
- Review the list, and identify earlier signals on which the newest signal is plausibly dependent
- If no plausible dependences are identified, assess the probability of having received the newest signal conditional upon each node in the network being true or false --
  - Use the assessments to update each node according to Bayes' Law
- If one or more plausible dependences are identified, note these dependences next to the newest intelligence signal on the list --
  - Re-compute the posterior probability for each node in the network while *omitting* the conditional probability assessments corresponding to the dependent signals
  - Combine the dependent signals including the newest signal into a single "super-signal" that is a multi-dimensional combination of signals.
  - Assess the joint probability of having received the super-signal conditional upon each node being true or false
  - Use the conditional probabilities associated with the super-signal to update the probability distribution computed above

We found this procedure to be straightforward, though we note that it is not perfect. It relies on the analyst's intuition to rule out dependences without formal assessment. This step is necessary to prevent every possible combination of intelligence signals from joint assessment, which would lead to an undue computational burden. However, as analysts ultimately rely on intuition to perform probabilistic assessments, relying on their intuition to rule out possible dependences does not seem like an unwarranted step. The super-signal joint distributions themselves were fairly quick to assess for 2-dependent signal and 3-dependent signal combinations.

After devising and testing this procedure, we turned our attention to introducing dynamics into the Bayesian inference model. To accomplish this we devised a Hidden Markov Model whose embedded Markov process represents our understanding of some underlying crisis (e.g. the state of Pakistani nuclear missiles -- in storage, un-fueled on launch vehicles, fueled on launch vehicles, or erect on launch vehicles). The Hidden Markov Model allows us to explicitly capture the dynamic nature of a crisis scenario and use intelligence signals to infer the state of the crisis in the current time period. If we let one state of the embedded Markov process represent the occurrence of the crisis under consideration, the first hitting time of the crisis state represents the lead time given to a policy maker if a warning were issued in the current period. We use linear systems analysis to compute a distribution on first hitting times.

The obvious limitation of the Hidden Markov Model approach for representing the dynamics of a crisis scenario is the assumption of Markovian behavior. That is, it assumes that the future state of the crisis scenario does not depend on the past, given the present. This assumption limits the types of crisis scenarios for which the model is suitable. We suspect, however that because the model is being applied to a warning problem, for which the policy maker's horizon is finite in length, the Markov assumption need only be approximately true over periods of time of the same order of magnitude as the time horizon. Where the policy maker's horizon is 1 week, and the system is assessed to have non-Markovian behavior over a period of 6 months, the Markov assumption should be reasonable. Notwithstanding, it will be important to explain the Markov assumption and its implications in language that an analyst can readily understand to ensure the model's users are aware of its limitations.

The Hidden Markov Model may be a reasonable representation of crisis dynamics, but it does not incorporate a decision element to warn the policy maker. Thus, we reformulated the Hidden Markov Model as a Partially Observable Markov Decision Process where in every period, the analyst decides whether or not to issue a warning. A time slice of the generalized POMDP used to model an early warning decision is shown below.

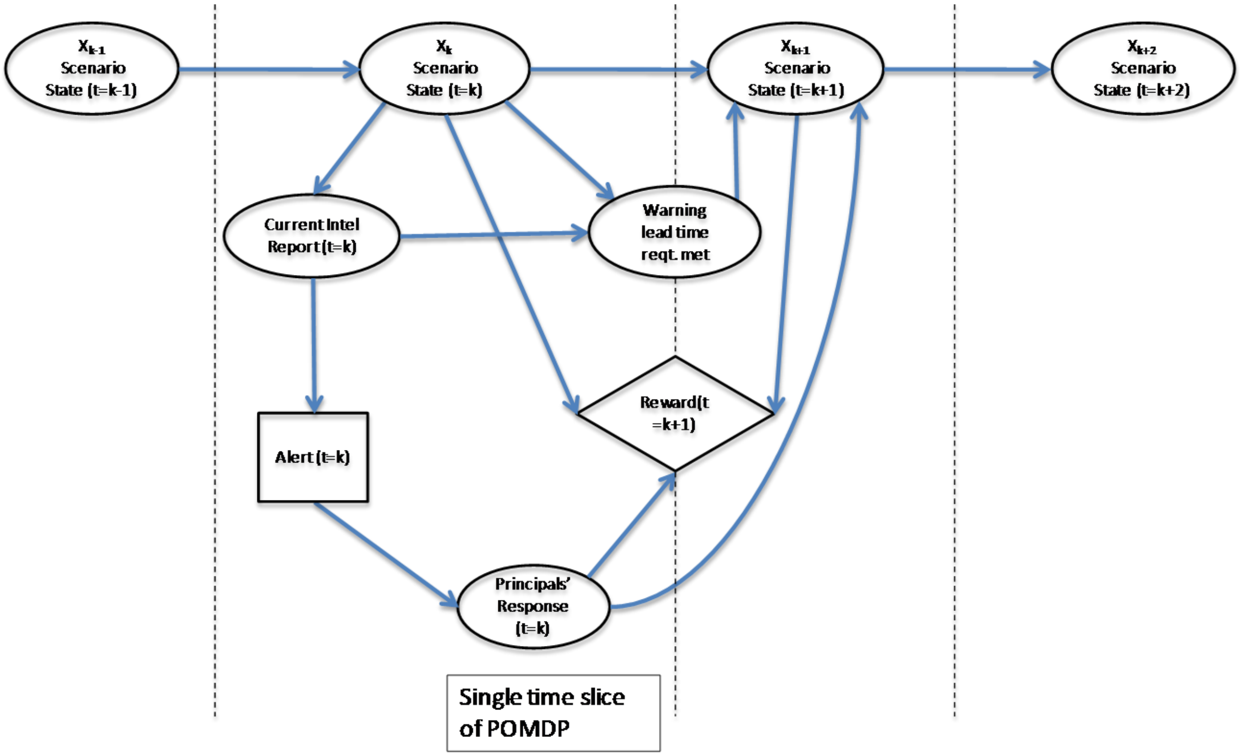

Figure 1: Influence diagram for a Partially Observable Markov Decision Process

2.2 Case studies

To date we have performed five case studies to explore, test, and demonstrate different elements of the generalized model. Each case study is listed below along with the component of the generalized model which it demonstrates:

1. State-run nuclear programs (Markovian scenario evolution)

2. Midair collisions under National Airspace automation (simulation design and event warning)

3. Smuggling of WMD device inside of a container (Bayesian inference of dependent signals)

4. Multi-stage decisions for operating fractionated satellites (dynamic programming to solve for optimal strategy)

5. Pearl Harbor retrospective analysis (combining all elements into a warning decision model)

The first of these case studies has been accepted for publication paper in scholarly journal. The second and fourth were presented at AIAA and IEEE conferences respectively, and will be published as part of the proceedings of those conferences. The fifth case study was recently completed and will be submitted to a scholarly journal shortly.

In the Pearl Harbor retrospective analysis, we successfully demonstrated a Matlab-based model with an ability to perform ongoing inference related to a dynamic crisis scenario using a stream of real-world intelligence signals, some of which were dependent. Furthermore, after tuning the transition matrix depicting the underlying scenario evolution, we demonstrated that this model would issue warnings upon receipt of sufficiently diagnostic intelligence signals, but would otherwise wait for additional collection. The difficulty of properly tuning of a POMDP such that it neither immediately warns nor perpetually waits is a concern that was expressed to us by CIA-associated researchers who were involved in the 1970's era experimentation. We are therefore encouraged and believe that our contribution can be made operational as we continue to move forward.

3. Applied Relevance

3.1 Improving Intelligence 'Indications & Warning' Methodology

The US intelligence community has suffered several major intelligence warning failures in recent times, the most noteworthy being the September 11 terrorist attack and the failure to warn of political instability that would topple regimes in Tunisia and Egypt. Warning is one of the fundamental functions of US intelligence, and it has become even more critical in an era where attacks can be carried out with short planning timelines, and where the internet spreads news virally so that unrest in one region can inflame unrest in another practically overnight. Without warning, US policymakers are forced to respond to crises reactively. If they have courses of action under consideration that require some lead time for successful implementation (as most do), those courses of action are not viable without warning.

As a result, many different parts of the US intelligence community have reached out to academia in recent years to sponsor research aimed at improving existing warning methodology, known as 'indications & warning'. O'Brien (2010) summarizes several of those efforts. Surprisingly, however, dynamic Bayesian forecasting is not among them. Our research fills this void. It is commonly accepted by those in the intelligence community responsible for teaching analytic methodology that there is no single correct tool or method to be applied to all problems. Dynamic Bayesian forecasting may not be ideal in all situations, but we do believe that it will provide more relevant warnings to US policy makers, with fewer false alerts and with improved lead time for several types of potential national security crises.

Two other approaches to warning that have attracted significant attention are statistical regression to improve indications and warning, and prediction markets. Statistical methods can be quite powerful where historical data exist regarding crises that analysts believe accurately represent how future crises of the same type will unfold. The CIA-sponsored Political Instability Task Force has taken this approach to the study of political instability. Prediction markets promise to improve how individuals with different expertise aggregate their data, but they require broad participation in order to make for a functioning market. Our POMDP is likely to prove most effective in scenarios were little historical data exist or are available – brainstormed or hypothetical scenarios (a surprise air attack on Hawaii in 1941, for instance) – and where one analyst or a small team are responsible for monitoring the scenario.

We envision that the POMDP that we are developing will form the core of a computer software tool that US intelligence analysts will use to monitor potential crises identified as of interest to policy makers and their staffs. This will not be the only tool used by warning analysts, but one of many. Assuming that analytic departments are sufficiently staffed, analysts will run multiple tools in parallel to monitor and warn of potential crises, and where the models conflict, they will turn to methodology tutorials along with academic literature to understand the advantages and disadvantages of each model. We intend to author at least one such tutorial to facilitate the use of our POMDP for this purpose. Where analytic departments are not staffed sufficiently to enable parallel use of models, the tutorials and literature should guide model selection.

3.2 Other Uses

The warning model illustrated in Figure 1 of Section 2.1 is not specific to national security. It is an influence diagram representing a general warning decision problem where an analyst must decide whether or not to escalate a problem up his chain of command where authority to respond resides. Scenarios that are also represented by this influence diagram include evacuation decisions for natural disasters, intervention decisions for financial markets, intervention decisions for medical conditions, and recall decisions for consumer products. In all four of these scenarios, the individuals with the expertise to monitor a potential crisis (climate scientists, staff analysts at regulatory agencies, doctors, and production managers, respectively) are not the same individuals as those charged with deciding on a response (senior government officials, patients, and corporate executives). We therefore expect that this model will find several applications outside of the national security realm.

While our research focus remains on the national security application, because of the sensitive nature of intelligence and the inherent secrecy that will be a part of any effort to test our model's performance in a national security setting, we may leverage the above-mentioned fields for testing and peer review. If we undertake such an effort, we will identify the commonalities and differences between the field in which the model is tested and national security crisis warning. We think that the most similar context will be in the area of natural disaster warning. Indeed, some scholars have discussed natural disasters in national security terms, and the line distinguishing the two contexts is not an obvious one. However, the other contexts –financial, medical, etc.- also have enough in common with the national security context that we believe testing would generate some insights regarding how the model might perform in a national security setting.

4. Collaborative Projects

Over the last year we have engaged in direct collaboration with the Department of Defense and agencies in the US intelligence community to assist with implementation of methodology to carry out real world crisis early warning. To that end, we met with officials from ODNI, CIA, IARPA, DIA, and NSA. Due to sensitivities regarding these collaborative efforts, we cannot provide details of the collaboration in this report. Please contact David Blum for further information.

We also collaborated with NASA Ames Research Center on a related effort to analyze the risk associated with a proposed automation of air traffic control in the en-route portion of the National Airspace. The system is designed to detect and warn of future midair collisions, and to provide sufficient lead time to generate alternative aircraft trajectories, while system failure presumably would lead to midair collisions. Three Stanford Ph.D students participated in this project and worked side-by-side with two NASA Ames flight engineers for ten weeks. After presenting results to NASA Ames, the students and flight engineers met several times over the next few months to make improvements to the simulation model and to prepare a submission to the AIAA Guidance Navigation and Control conference, which was held in August 2010. NASA Ames has published the AIAA paper on its website, the url for which is below.

http://www.aviationsystemsdivision.arc.nasa.gov/publications/2010/ThipphavongGNC2010Final.pdf

5. Research Products

Research Products (Please detail below)		#
5a	# of peer-reviewed journal reports published	1
5a	# of peer-reviewed journal reports accepted for publication	2
5a	# of non-peer reviewed publications and reports	0
5a	# of scholarly journal citations of published reports	0
5b	# of scholarly presentations (conferences, workshops, seminars)	6
5b	# of outreach presentations (non-technical groups, general public)	4
5c	# of products delivered to DHS, other Federal agencies, or State/Local	0
5c	# of Paténts filed	0
5c	# of Paténts issued	0
5c	# of products in commercialization pipeline (products not yet to market)	0
5c	# of products introduced to market	

5.1. Publications and Reports

Note that in several cases, CREATE was only one of the sources of support for these publications.

Research Areas: Risk Assessment (RA); Economic Assessment (EA); Risk Management (RM)

CREATE PUBLICATIONS	Research Area	Referred	Not Referred	PDF Available for DHS
Paté-Cornell, M. Elisabeth – Stanford University				
1. Caswell, D.J. and Paté-Cornell, M.E., "Bayesian Networks, Semi-Markov Nodes and Nuclear Proliferation, accepted for publication in *Military Operations Research*, 2011	RA		x	
2. Blum, D.M, Thipphavong, D., Rentas, T.L., He, Y., Wang, X. and Paté-Cornell, M.E., "Safety Analysis of the Advanced Airspace Concept using Monte Carlo Simulation," 2010 American Institute of Aeronautics and Astronautics Meeting Papers on Disc, Vol. 15, No. 9, 2010	RA		x	x
3. Daniels M., Tracey, B., Irvine, J., Schram, W., and Paté-Cornell M. E. "Probabilistic Simulation of Multi-Stage Decisions for Operation of a Fractionated Satellite Mission," IEEE Aerospace Conference Proceedings, Paper #1323, 2011.	RM		x	x

Presentations

PRESENTATIONS - CONFERENCES

Paté-Cornell, M. Elisabeth – Stanford University

1. Paté-Cornell, M.E., "Risks and Frames: Three Models and Illustrations," *PSAM 10 Conference*, Seattle, WA, 2010

2. Paté-Cornell, M.E., "Risk Analysis when Intelligence Actors are Involved," Invited Lecture Series, National University of Singapore, Singapore, 2010

Blum, David – Stanford University

3. Blum, D., "Safety Analysis of the Advanced Airspace Concept using Monte Carlo Simulation," AIAA Guidance Navigation and Control Conference, Toronto, ON, 2010

4. Blum, D., "Probabilistic Warning Systems for National Security Crises," Information Sciences Doctoral Seminar, U.S. Naval Postgraduate School, Monterey, CA, 2010

5. Blum, D., "Mechanisms of Network Disruption Due to Stress," 78th Military Operations Research Society Symposium, WG 7 (Intelligence, Surveillance, and Reconnaissance), Quantico, VA, 2010

Daniels, Matthew – Stanford University

6. Daniels, M. "Probabilistic Simulation of Multi-Stage Decisions for Operation of a Fractionated Satellite Mission," IEEE Aerospace Conference, Big Sky, MT, 2011

PRESENTATIONS - OUTREACH EVENTS

Paté-Cornell, M. Elisabeth – Stanford University

1. Paté-Cornell, M.E., "Risk and Intelligence Actors: Two Cases from the Space World and the Oil Industry," Draper Laboratory and MIT Department of Engineering System Dynamics, Cambridge, MA, 2010

Blum, David – Stanford University

2. Blum, D., "Probabilistic Warning Systems for National Security Crises," School of International Relations, St. Andrews University, St. Andrews, UK, 2010

3. Blum, D., "Indications & Warning: A Bayesian Approach", presentation to the Office of the Director of National Intelligence, McLean, VA, 2010

4. Blum, D., "Indications & Warning: A Bayesian Approach", presentation to IARPA, Washington, DC, 2010

5.2. Models, Databases, and Software Tools and Products

SOFTWARE PRODUCTS							
CREATE Project Leader(s)	Name of Product	Type of Product	Application Area	Intended Users and/or Clients	Specific Users and/or Clients	Status	Other Comments
Paté-Cornell	Warning POMDP	Matlab code	Risk analysis	US intelligence analysts	Stanford University	under development	Initial version used in Pearl Harbor retrospective case study

Matlab code implements the POMDP generalized model. Current version is specific to Pearl Harbor retrospective case study, but a generalized version is under development.

6. Education and Outreach Products

Education and Outreach Initiatives (Please detail below)	**#**
# of students supported (funded by CREATE)	0
# of students involved (funded by CREATE + any other programs)	3
# of students graduated	1
# of contacts with DHS, other Federal agencies, or State/Local (committees)	7
# of existing courses modified with new material	2
# of new courses developed	1
# of new certificate programs developed	0
# of new degree programs developed	0

School of Engineering)

CREATE STUDENTS								
	Last Name	First Name	University	School	Department	Degree	Research Area	Funded
1.	Blum	David	Stanford	Engineering	Management Science & Engineering	PhD	Risk Analysis	
2.	Caswell	David	Stanford	Engineering	Management Science & Engineering	PhD	Risk Analysis	
3.	Daniels	Matthew	Stanford	Engineering	Management Science & Engineering	PhD	Risk Analysis	

CREATE RELATED COURSES			
	Instructor	**University**	**Course Title**
1.	Paté-Cornell	Stanford	Engineering Risk Analysis
2.	Paté-Cornell	Stanford	Project Course in Engineering Risk Analysis
3.	Paté-Cornell & Hecker	Stanford	Seminar in Quantitative Analysis of Issues in National Security

MEMBERSHIP IN MAJOR DHS-RELATED COMMITTEES		
Committee	**Institution**	**Time Period**
M. Elisabeth Paté-Cornell		
Committee to Analysis of the Causes of the Deepwater Horizon Explosion	National Research Council	2010
Board of Trustees	InQtel	2006-present

MEETINGS WITH EXTERNAL ORGANIZATIONS

- CIA
- DIA
- Draper Laboratory
- IARPA
- NASA Ames Research Center
- NSA
- ODNI