SECURE: Stochastic Enhanced Control of Unstable Regional Environments

Georgiy M. Levchuk[1a], Bruce Skarin[a], Daniel Serfaty[a], and Krishna R. Pattipati[b]

[a]Aptima Inc., 12 Gill Street, Suite 1400,Woburn, MA 01801

[b]Electrical and Computer Engineering Department, University of Connecticut, Storrs, CT 06269

Abstract

In this paper, we present a conceptual integration of pattern classification, dual control under uncertainty, and social dynamics simulation technologies to address the problem of instability management. We conceptually designed a model called Stochastic Enhanced Control of Unstable Regional Environments (SECURE) to provide effective real-time early-warning and decision analysis for monitoring, assessing, forecasting, and preventing the regional conflicts and instability. Our current work is focused on validating the model against real-world and synthetic datasets, and will be reported in future publications.

The SECURE calculates the indicator of a *power balance* in the area of interest based on the *interaction network* that defines the state of and relationships among the groups, organizations, institutions, and individual members of the society. Using this indicator, SECURE tracks the dynamics of the society of interest over time and develops robust dynamic action strategies to maintain stability and prevent crises. SECURE solution is based on the concept of dual control, a judicious integration of actions to *influence the state of the environment* of interest as well as to *gain more knowledge* about the true state of the environment. SECURE is enhanced with the social dynamics simulation models to generate possible dynamics of the society. Such dynamics form the models that are used by the predictive and decision algorithms to recognize and control current and future state of the environment.

1. Motivation: Regional Crises Early Warning

1.1. The Challenge

Today’s world is changing rapidly, generating more asymmetric and unconventional adversaries for the United States. International crises pose more of a challenge to U.S. national security as unstable or failed states fuel regional conflicts, harbor terrorists hostile to the U.S. or are unable to protect themselves against the spread of terrorism. Failed states can also become centers for the trade of illicit drugs and arms, and can form breeding grounds for dangerous diseases. Regional conflicts not only cause humanitarian disasters requiring ever-increasing resources from the international community, but also create political, economic, and social instability in the neighboring countries, a growing number of which possess or are interested in developing weapons of mass destruction.

[1] georgiy@aptima.com; phone 781-935-3966x267; fax 781-935-4385; www.aptima.com

The state failure and large-scale conflicts that ensue in societies do not just happen overnight: the dynamics leading to these events evolve gradually, and display potentially discernable patterns in critical societal relationships. The indicators of state weakness should not be based solely on the problems of state security, but also on the conditions that threaten the status quo, the physical integrity of the state’s borders, the general welfare, and human self-determination. Such conditions and their dynamics need to be assessed and identified early because a truly effective strategy is not one of *reacting* to a disaster, but one that *anticipates* and *prevents* disasters. Interventions should start early, before states begin to fail. However, the feasible preventive actions are severely limited and can only be narrowly focused, as the international community does not have the resources, political will, or know-how to mount comprehensive interventions in every state that gives early signs of failure (Ottaway, and Mair, 2004).

1.2. The Need for Prediction of True Environment Dynamics

The asymmetric threat environment has large information gaps. First, much of the information is missing due to limited resources to collect the data – e.g., open source intelligence (OSINT) only contains information expressed by the actors and is therefore most often an incomplete and misleading statement of their intent. Second, much of the transactional data collected are irrelevant to establishing an accurate description of the environment and hostile activities (“background noise”) – e.g., there are many actors that are not part of the hostile groups and do not influence the onset and evolution of crises. Third, information processing itself can encounter errors – either from automated tools, or from personal biases about intent of actions. Fourth, the groups of interest might possess few members living, planning and operating within a large population (e.g, Al Qeada in Pakistan). And finally, a significant amount of intelligence could be observations of intentionally deceptive actions conducted by the hostile actors.

Due to these sources of uncertainty, actions to prevent crises cannot be based on current observations alone. Instead, we need to use historical observations to track the dynamics of an environment, predict its current true state from partially observed transactions, and forecast future evolutions including potential instabilities.

Lately, U.S. researchers supported by the Department of Defense have conducted a concerted effort to develop early warning indicators (such as failed state index[2]) and decision support tools to forecast regional and state instabilities (O’Brien, 2002; Baker, 2003; King and Zeng, 2001). However, these early warning systems focus on a high-level assessment and do not predict the timing for impending conflicts nor the specific nature of the conflicts. Existing state instability forecasting tools therefore do not provide the insights that short-term planners need – *who is doing what, why, and to whom*, - in order to develop effective countermeasures. Consequently, researchers have begun to dedicate significant attention to analyzing the interactions among individuals, groups, and institutions (Popp et al., 2006; Saunders-Newton, Frank, and Popp, 2005; O’Brien, 2004; Schrodt and Gerner, 2000, 2001). Various models have been applied to study such interactions, including differential equations (Turchin, 2003), interaction-events data analysis (Gerner et al., 2002; O’Brien, 2004), game-theoretic models (Brams and Kilgour, 1988), and others. While these approaches benefit from technologies that can capture entities and events from open sources (e.g., text reports from news media), most of this data is either captured at a national level (O’Brien, 2004) or at a very detailed level of the individual members of the society.

[2] http://www.fundforpeace.org/web/index.php?option=com_content&task=view&id=99&Itemid=140

A large number of models have also been developed using agent-based simulations of societies (Popp et al., 2006). However, the amount of data that is needed to populate these models is enormous, with data and data sources that are usually difficult to acquire. A significant amount of noise events (text parsing errors, misclassifications, missed information, and deceptions) have contributed to misleading forecasts (false alarms and false positives – the recognition of potential threats that have little or no impact) due to the sensitivity of agent-based models to input parameters. In addition, different models worked at different levels of granularity, with no common problem framework developed to integrate model inputs and outputs (Popp et al., 2006). Very few of the models were able to “remove the noise” from the input data, and none of the models were able to work with data sources at different levels of granularity. As the result, the Integrated Crisis Early Warning System (ICEWS) program was initiated at DARPA to develop a comprehensive, integrated, automated, generalizable, and validated system to monitor, assess, and forecast national, sub-national, and international crises. One of the main objectives for this program is to develop predictions that would support decisions on how to allocate resources to mitigate potential crises and instability.

1.3. The Need for Control in the Presence of Uncertainty

One of the needs identified in the ICEWS program is the development of Courses of Action (COAs) that can achieve end-state regional stability objectives in near real-time. Such COAs must be designed in the presence of a very high uncertainty in the data, when no single reliable prediction about the situation can be obtained and instead multiple possibilities are equally likely. In this vein, two types of control actions are possible: on the one hand, we can design **preventive actions** that seek to change the environment towards a desired state while being *robust to the uncertainty* in current forecasts about the environment; on the other hand, we can design **investigative actions** that *improve our knowledge* of the environment to enable better control of the environment in the future.

The preventive actions need to account for the risks associated with potential undesired consequences due to the uncertainty of current environmental state predictions and uncertainty in action outcomes. When the uncertainty in current predictions is large, few, if any, preventive actions are possible that would avoid such consequences while ensuring that the crisis does not occur. Here, injecting investigative actions to gain a better understanding of the environment will lead to better preventive actions in the future.

Investigative actions can be of two types: (i) passive information collection actions; and (ii) active probing actions. Passive information collection actions can extract the relevant information critical to improving the quality of predictions from other (possibly non-government) sources or allocate available sensor resources for future intelligence collection in areas of interest. However, actions limited to passive sensing cannot improve the situational awareness when no revealing transactions have occurred. For example, passive sensing can analyze open-source intelligence from blogs and websites visited by potential leaders of hostile groups in the region; however, no significant information can be gathered if the adversarial actors do not express their views, intent, or do not overtly communicate with each other in any way.

1.4. The Need for Probing when Uncertainty is too High

To discover the hidden state of an environment (e.g., actors, their relationships, their views, intent, and power balance in the society), we can apply actions that are analogous to injecting test signals into physical systems. Such *probing signals* force the system to reveal itself. Mathematically, the main objective of probing actions is to *maximize the information gain* (or minimize the entropy, a measure of uncertainty) of the predictions about the state of the environment. This can be achieved by finding those probing actions that reduce the entropy of the current state of the environment the most. An example of a probing action could be providing information that coerces adversaries to communicate more and hence to reveal their relationships, commit actions that could be observed, or even change some of their plans. The outcome of probing is an improved understanding of the environment and, consequently, a better ability to prevent instability and/or control crises.

2. The SECURE System: Concept Description

2.1. System Workflow

The workflow of the SECURE system is shown in Figure 1. The observations and events from the environment are gathered by the *Entity/Relationship Extraction* component. Such data can be obtained from existing technologies developed under various programs (e.g. ICEWS), and is therefore outside the scope of this research. This data is fed into the *Environment Dynamics Tracking* component, which generates predictions of the environment dynamics and its current and future states. This component is using the set of hypothetical environment dynamics models that are generated by *Social Dynamics Simulation*. If the reliability of predictions is high (low entropy), the *Control* component is activated to find actions that seek to prevent the crises and influence the environment. If the reliability is low, the *Investigative* component is activated that invokes the information collection and/or probing actions to increase our understanding of the environment.

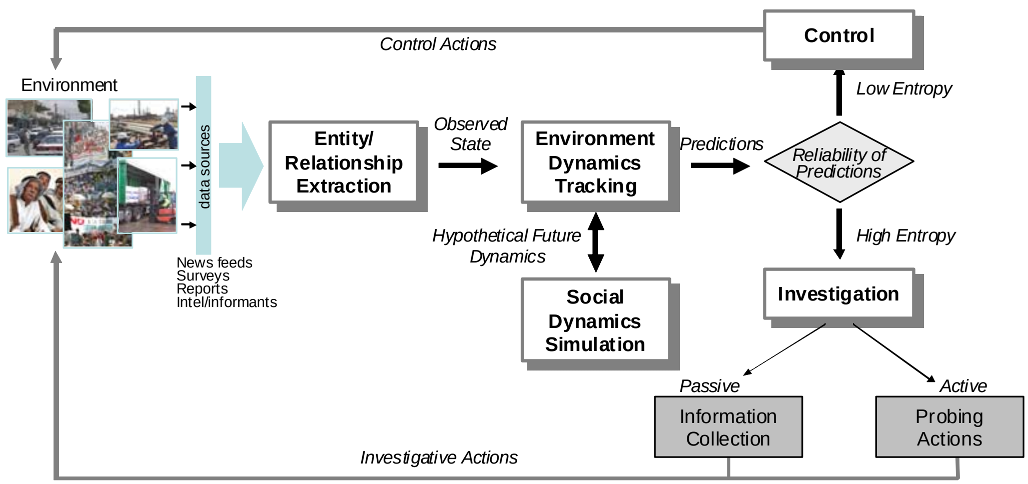

Figure 1: The Workflow of SECURE Model

2.2. Instability Indicator as a Power Balance Network

SECURE calculates the indicator of regional/state instability as a *power balance network* (PBN) in the area of interest. The PBN indicator is based on the *interaction network* that defines the state and relationships among groups, organizations, institutions, and individual members of society (Figure 2). The nodes in the network represent groups, organizations, government, states, and individuals, with node attributes defining size/membership of the groups, beliefs of people, economic power of organizations, social identities of individuals and groups, and other economic and social features. The links in the network represent *who does what to whom* in the area of interest – social interactions and influences, economic transactions, behavioral interactions, political events and activities, etc., with link attributes defining the frequency and types of interactions among the nodes.

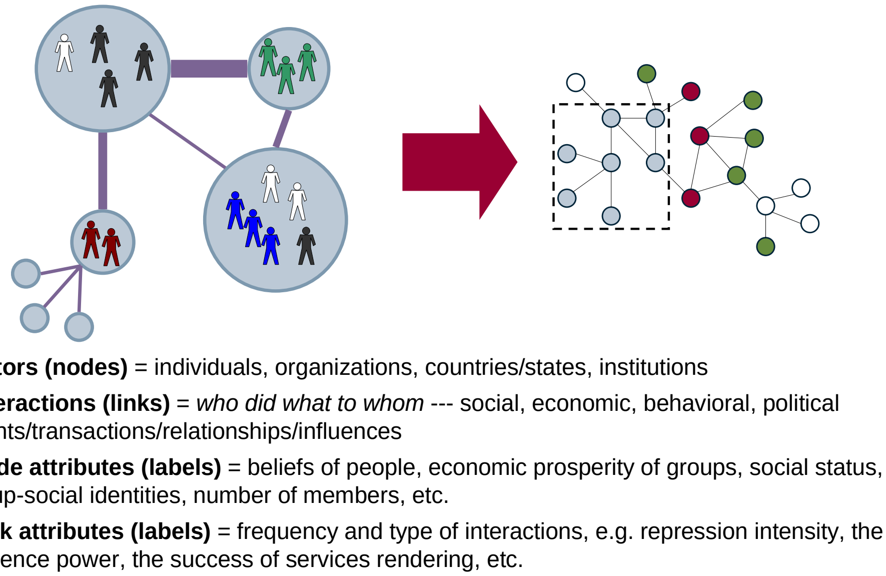

Figure 2: Power Balance Network as a Labeled Graph

There are large information gaps in the data that can be gathered about a society of interest due to factors including:

- **Large numbers of actors/interactions/attributes that are irrelevant:** for example, normal/green only events cannot be filtered out because the same event might be in normal and abnormal behavioral patterns.
- **Large numbers of actors/interactions/attributes that are missing:** for example, it is impossible to know the financial status, opinions, or # of members of every group in the society.
- **Large numbers of actors/interactions/attributes that are uncertain:** for example, we can only approximate the beliefs of every person/group by using “representatives” and drawing conclusions about the group as a whole.
- **Large numbers of actors/interactions/attributes that are deceptive/erroneous:** for example, false information in surveys, data intentionally misrepresented by members of the group, misclassified memberships or events (e.g., errors in text entity extraction from open source data or bias in Subject-Matter Expert’s (SME) input).

As a result, the true state of the power balance is not known (Figure 3). Instead, the obtained observed state of PBN is obtained, which would not have a perfect match with the true state of PBN due to the data gaps and classification issues. The objective of the environment dynamics tracking component is then to recognize the true state of the PBN and find the mapping between the observed and true PBN states that produces the best match.

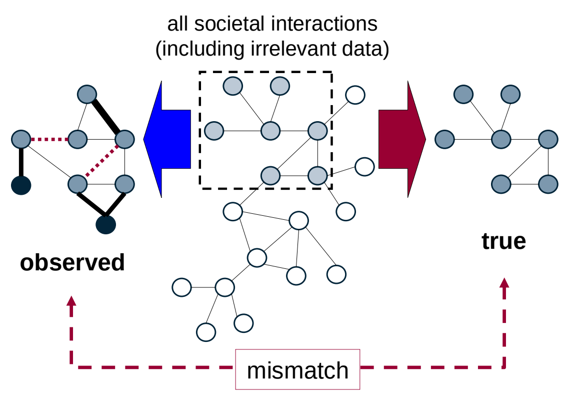

Figure 3: Observed versus True State of PBN

2.3. PBN State Identification via Pattern Matching

To estimate the state of the environment, SECURE performs the *probabilistic network pattern matching* of the observed data about the state of actors and their interactions in the society against a set of hypothetical power balance patterns (Figure 4). The outcome of network matching is the state of the power balance that best explains/matches the observations, and the forecast of the instability associated with the highest-scored mapped pattern(s).

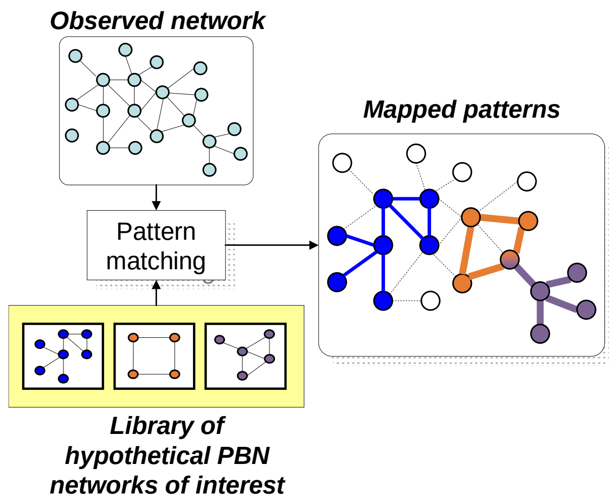

Figure 4: SECURE PBN Inference via Network Pattern Matching

More formally, we define a power balance network as an attributed graph $G_m^M = (V_m^M, E_m^M, A_m^M)$, where for a hypothesis m, often called *model network*, V_m^M is a set of societal actors/nodes, E_m^M is a set of links, and A_m^M is a set of attributes on nodes and links. Similarly, the observed interaction pattern, often called *data network*, is defined as $G^D = (V^D, E^D, A^D)$.

To find the matching between these two networks, we define the mapping of the data network to model network, i.e. the mapping of the nodes in the observed data, such as people, groups, institutions, states, governments, etc., to the nodes in the hypotheses PBN – the power holders and their profiles/roles. The model is based on the NetSTAR formulation of network pattern classification originally developed to recognize adversarial organizational networks (Levchuk et al., 2006; 2007).

The mapping between observed and hypothesized networks is represented as a matrix $S_m = \{s_{ij}^m\}$, where $s_{ij}^m = 1$ if and only if the observed node i is mapped to a node j in the hypothesis. This mapping is found by maximizing the likelihood that the observed data has been generated by the model: $S_m = \arg\max_S P(G^D \mid G_m^M, S)$. We can approximate the negative of log-likelihood function using the quadratic polynomial objective function (this formulation is an alternative to the structural method described in (Levchuk et al. 2006)): $Q(S) = \frac{1}{2}\sum_{kimj} s_{ki} s_{mj} c_{ki;mj} + \sum_{ki} s_{ki} c_{ki}$, where the parameters $c_{ki;mj}$ and c_{ki} are correspondingly the score of mismatch of the links and the score of mismatch of the nodes between the observed and hypothesized PBNs. In our research, we have experimented with several definitions of the mismatch coefficients, including Euclidean norm, cosine similarity, sigmoid, and logistic measures.

The problem (1) then becomes a quadratic assignment problem, which can be solved by several algorithms. One of the good low-computational complexity algorithms is a graduated assignment algorithm developed by Rangarajan, Vuille, and Mjolsness, 1999. This algorithm relaxes the integrality constraint on the objective function and iteratively computes the continuous mapping approximation matrix which gradually converges to a 0-1 solution.

2.4. Identification of Environment Dynamics and Temporal State Tracking

The dynamics of an environment can be represented as stochastic evolution of power balance networks over time. This evolution corresponds to changes in the PBN pattern, where a PBN instance represents the current state of the environment. The same environment dynamic may result in different sequences of PBNs due to uncertainties in the actions and interactions of members of the society. This allows accounting for the diversity of the possible states of the environment and enables predictions of instability that are robust to normal deviations of population behavior and uncertainties in the environment. Different environment dynamics will result in sequences of PBNs that significantly differ from one another. In addition to deviations in environment state, another important challenge is that the knowledge of an environment will always be incomplete. As a result, we need to consider that the state of the environment is a process that is only partially observable over time.

To represent the environment dynamics, we use the Hidden Markov Model (HMM) formalism. We chose to use HMMs because they constitute a principled method for modeling partially observed stochastic processes that have temporal structure. Each HMM can be viewed as a detailed, stochastic time–evolution of a particular system. An HMM can sequentially process new information (a window of data) each time an observed event occurs. The window of observations could contain a single or a batch of observations and activities to improve the efficiency of a solution. The premise behind an HMM is that the true underlying process (defined as a Markov chain representing the evolution of the activities as a function of time) is not directly observable (hidden), yet it can still be probabilistically inferred through another set of stochastic processes (observed events about interactions in the society, for example). In order to account for specific decision points and control in the environment, this model can also be extended to Partially Observable Markov Decision Processes (POMDP).

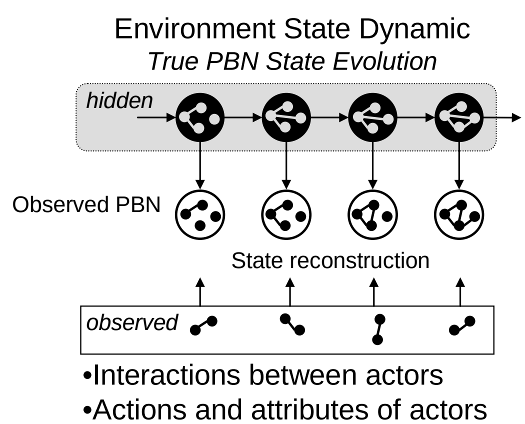

Figure 5: HMM-based Society Dynamics

A single HMM indicates a certain hypothesis about the society dynamics --- the evolution of environment states that can be indicative of a pattern or precursors to instability. HMMs can be constructed by SMEs or learned from data about environment changes in the past. Figure 5 shows a typical HMM. The black colored circles represent true states of the environment (the shaded area). This true process is a “hidden” environment dynamic with series of true state transactions describing the evolution of a particular environment. This true “hidden pattern” is observed through a noisy process represented by a series of observed events (the white rectangle in Figure 5). Our objective here is to detect hidden “true” societal evolution dynamic and its hidden state evolution which is a sequence of activity network states (shown inside black circles) via the observed process (white circles). We can infer the existence of a true pattern based upon a set of observations as shown in Figure 4 because HMM states are statistically related to a noisy observation process. Each state of the HMM corresponds to a true state of PBN, while each observation state is the observed state of

PBN – actors in the society, their attributes, actions and interactions. An HMM is defined using three parameters: the prior probability of a true state of a PBN $\pi_i = \mathrm{P}\left(G_m^M[0] = x_i\right)$ (where $\{x_i, i = 1,..., N\}$ is a set of feasible PBN states and brackets in the definition of the PBN state indicate the time at which the state is considered), the transition probability of moving from one PBN state to another $a_{ij} = \mathrm{P}\left(G_m^M[t+1] = x_j \mid G_m^M[t] = x_i\right)$, and the probability of observing a PBN state given a true state. The observation probability does not have to be specified in advance, but can be calculated as the observations are received. We will calculate observation probability as a likelihood score from the pattern matching step described in the previous section: $b_i(k) = P\left(G^D[t] = o_k \mid G_m^M[t] = x_i\right) \cong \arg\max_S P\left(G^D = o_k \mid G_m^M = x_i, S\right)$.

In the context of predicting and preventing societal instability, we will use HMMs to address the following problems:

- *Problem 1 – Identifying the environment dynamics*: There are multiple HMM representations for different types of environment evolutions, some of which indicate various levels of instability. HMMs allow finding the model of societal evolution dynamics that best explains observations by maximizing the likelihood probability $\mathrm{P}(O \mid \lambda)$, where $O = \{o_1,..., o_T\}$ is a sequence of observed PBN states and λ is a HMM. The maximization is then achieved by using a forward-backward algorithm (Rabiner, 1989), where the forward algorithm is used to update the likelihood scores for all HMMs over time.
- *Problem 2 – Recognizing PBN evolution*: Often, we are interested in finding how the state of an environment has been changing over time, i.e. the most probable sequence of states that resulted in the observations obtained. We can do this by using HMMs to find the sequence of states $Q = \{q_1, q_2,..., q_T\} \subset \{x_i, i = 1,..., N\}$ that maximizes the likelihood probability $\mathrm{P}(Q, O \mid \lambda)$. This can be achieved efficiently using a Viterbi algorithm (Rabiner, 1989).
- *Problem 3 – Prediction of the future state of a society:* To develop effective preventive policies before a potentially avoidable crisis develops we need to anticipate the future states of the environment that can cause instability. Using HMMs, we can calculate the likelihood of an environment to be in a specific state $G_m^M[T] = x_i$ at some future time T given the observations up until time t. This can be obtained by finding $\mathrm{P}\left(G_m^M[T] = x_i, o_1, o_2,..., o_t \mid \lambda\right)$ for each state s_i as a score of *attainability* for the state.
- *Problem 4 – Learning the model of environment dynamics*: When historic data about environment is available, we can find the parameters of HMM models to represent the corresponding state dynamics of the society. This amounts to finding the state transition structure and probabilities for model $\lambda = \{a_{ij}, b_i(k), \pi_i\}$ that maximizes the likelihood probability $\mathrm{P}(O \mid \lambda)$, and can be achieved using Expectation-Maximization and other related algorithms (Cappe, Moulines, Ryden, 2005).

2.5. Control of Environment Dynamics

To develop Courses of Action (COAs) that can achieve end-state regional stability objectives in near real-time, we must deal with the presence of very high uncertainty in the data. That is, where no single reliable prediction about the environment and the situation can be obtained and instead multiple possibilities are equally likely. In this vein, two types of control actions are possible: in one case, we can design **preventive actions** that seek to change the environment towards a desired state while being *robust to the uncertainty* in current forecasts about the environment; in the other case, we can design **investigative actions** that *improve our knowledge* of the environment to enable better control of the environment in the future. Preventive actions need to account for the risks associated with potentially undesired consequences that can result from uncertainty in the current environmental state predictions and action outcomes.

Regional control to achieve stability objectives can be modeled using Partially Observable Markov Decision Process (POMDP). A POMDP model adds the concept of actions and state rewards to the HMM (Figure 6). Actions affect the state transition probabilities, so we redefine the environment/PBN state transition probabilities as dependent on the action $a[t]$ taken at the current time: $\mathrm{P}\left(G_m^M[t+1]=x_j \mid G_m^M[t]=x_i, a[t]\right)$. A reward can be positive or negative; it is positive for stable states and negative for states where conflicts exist. Rewards can also incorporate the cost of conducting preventive actions. For simplicity, we assume that the environment has finite number of states and the control options are finite as well. All the derivations extend to the case of infinite state space representations *mutatis mutandis*.

If the state of the environment is fully observable (that is, for each environment state x_i there is an observation o_k and $P\left(o_k \mid x_i\right)=1$), then the problem simplifies to a Markov Decision Process (MDP). The MDP solution is significantly easier to obtain, with algorithms running in time $O(|N|^2)$ per iteration or $O(|N|^3)$ for a closed-form solution (Bellman, 1957). The POMDP solution is significantly more complex, and the problem of finding optimal policies is PSPACE-complete (Papadimitriou and Tsitsiklis, 1987). There exist two main classes of approximate algorithms for POMDP: value-function methods that seek to approximate the value of *belief states* – probability distributions over the environment states (Hauskrecht, 2000), and policy-based methods that search for a good policy within some restricted class of policies. The approach we undertake in SECURE system belongs to the second class due, in large part, to the large number of possible states of the environment and the intractability of efficiently storing the continuous belief state distributions.

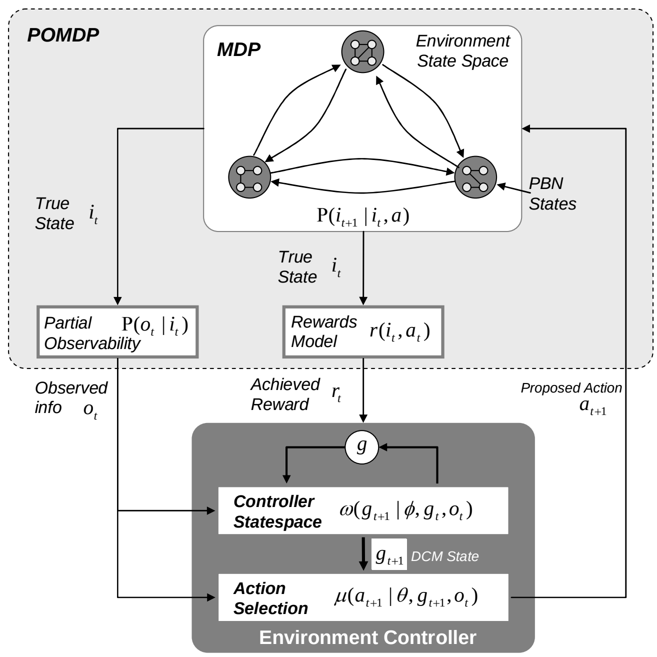

Figure 6: The POMDP Model and Environment Controller

SECURE identifies the preventive action policy using a finite state environment controller (Figure 6). This setup allows us to trade-off optimality with complexity and to train the controller over time. As the new observation $o \in O$ is received, the controller adapts using a parameterized stochastic function $\omega(h \mid \phi, g, o)$ equal to the probability of moving the controller from state $g \in \Omega$ to state $h \in \Omega$, where ϕ is a set of parameters, and $\Omega = \{1,...,|\Omega|\}$ is the set of internal-states of the *SECURE stochastic controller*. The plan for courses of action is defined via a stochastic parameterized policy $\mu(a \mid \theta, g, o)$ equal to the probability of taking action $a \in A$ given observation $o \in O$, $g \in \Omega$ is the state of the SECURE controller, and θ is the set of parameters. To optimize the COA, we need to find the parameters ϕ and θ that would optimize a state stability benefit expressed as the *long-term average reward*: $\eta(\phi,\theta) = \lim_{T\to\infty} \frac{1}{T} E_{\phi,\theta}\left[\sum_{t=0}^{T} r(i_t)\right]$, where $E_{\phi,\theta}[\cdot]$ is the expectation over all possible environment-controller state trajectories $\{(i_0, g_0),...,(i_T, g_T)\}$ when the controller's parameters are ϕ and θ.

The most simple and least complex solution assumes the controller to have a single state, in which case the action selection function $\mu(a \mid \theta, o)$ is a purely reactive policy. This policy can be trained from previously monitored and analyzed situations for the environment of interest, but it then chooses the control actions based only on the current observations, disregarding the *recent history* of environment states and control actions. In the presence of observation noise, such purely reactive policies are severely suboptimal. Instead, to achieve better conflict prevention and to steer the environment to a desired state (or to prevent it from moving into an undesired state), the decision model must remember the features of the history. For example, if the current observation is a political rally, this observation alone cannot tell the model what the true situation in the region of interest is and whether conflicts between the groups and members of the society can occur. Instead, knowing the causes of the rally, the position of the participants in the society, and the strengths of the political system would provide a better understanding of the societal dynamics, prediction of possible futures, and identification of actions that could be used to control a fragile situation.

A more efficient approach would be to store all the history of observations and actions $H_T = \{(o_0, a_0), (o_1, a_1),...,(o_T, a_T)\}$ as the state in the controller (e.g., *utile distinction trees* (McCallum, 1996) and *prediction suffix trees* (Ron et al., 1994)). However, this will require the amount of memory to represent policies to grow exponentially with the number of events (observations and actions). While such models are intractable in situations that involve a large number of temporal observations-actions, simpler approaches exist that can code the state representations more efficiently.

The POMDP solution can maintain a state space in the form of *beliefs about the true environment states*. These beliefs are probability distributions over the environment states and provide sufficient statistics to act optimally. Even though the set of belief states is infinite, the structure of the POMDP problem allows the efficient clustering of all beliefs into a limited set of states, and several popular algorithms have been developed for finding the value-function using value iteration (Sondik, 1971; Kaelbling et al., 1996; Cassandra, 1998). However, when the number of environment states is very large, finding an optimal solution and representation of the exact value-function over belief states is computationally prohibitive. Only problems with tens to hundreds of states can be solved using these approaches. The optimal solution was proven to be PSPACE-complete, and as a result, researchers have recently looked at new methods that restrict the set of states that the POMDP controller can have (Sutton et al., 2000; Poupart et al., 2001; Aberdeen and Baxter, 2002). One such solution using an internal-state policy-gradient algorithm (Aberdeen,

2003) was shown to solve the problem with tens of thousands of possible environment states in reasonable time (30 minutes). This is a solution used in SECURE system to find control actions.

2.6. The SECURE Investigation Model

The investigative actions can also be incorporated into the POMDP model. However, this would increase its complexity significantly, making the solution intractable when the number of intelligence collection actions and environment states is large. Instead, we can decouple the environment control into investigation and prevention models and conduct investigative activities when certain conditions in the prevention model are not met. The investigative policy will then define the information collection or probing action to find the information element (feature) of the environment (e.g., adversarial organization, intent, plans, materials flow, individual actor profiles, etc.) that is the most relevant or critical for predictions at the current time.

At each time step, we can calculate the likelihood distribution of the beliefs of a subset of environment states that are most likely. We can do this by looking at the environment-controller state (i,g) as a state of the hidden network and using the Hidden Markov Model forward algorithms to estimate the forward probabilities for each state, given previous observations. We will then select the most likely subset of states at every iteration of the forward algorithm to reduce the complexity of the solution. Let’s denote the probability of being in state $i \in \mathrm{N}$ at the current time as $p_t(i)$. To distinguish the states and determine the information collection and probing actions, we define the vector of features $\varsigma^i = \{\varsigma_1^i,...,\varsigma_N^i\}$ for each feasible state of the environment i. For simplicity of notation, we can assume that each feature ς_f^i can be collected using some action, and that the observation ξ_f will be obtained with the probability $p_f(\xi \mid \varsigma) = \mathrm{P}(\xi_f = \xi \mid \varsigma_f^i = \varsigma)$. The actions to collect the feature information can involve both types of investigative actions described above. For the purposes of the mathematical formulation, we ignore the differences between passive information collection and probing, but we note that the main differences in real-world settings would involve the probability of collecting correct intelligence, cost of the actions, and the action specifics.

Without loss of generality, we assume that all environment states are among the current potentials. We use the entropy as a score of confidence in current predictions, as it characterizes “how much uncertainty is there in predicting the true state of the environment given the data already collected?”. When the entropy is high (i.e., there is a high uncertainty that the current predictions can achieve a correct result, so $H_t = -\sum_i p_t(i) \log p_t(i)$ is close to $\log|\mathrm{N}|$), the likelihood estimates of the environment state predictions are similar. If these states carry significantly different projections in terms of planned preventive actions, i.e. the best policies for each state are significantly different and none is satisfactory for all states, we cannot rely with confidence on the policy derived from the current state prediction. Since significant uncertainty in predictions is often due to missing data, we can instead attempt to identify the features that are critical to prediction, so that the collection of these features would achieve the largest reduction in the prediction’s ambiguity.

First, we perform feature extraction by selecting the subset of all features $F \subset \{1,...,N\}$, where for each selected feature $f \in F$ there exist at least two states i,j such that $\varsigma_f^i \neq \varsigma_f^j$. Second, we order these features to achieve the most distinguishability among state predictions while satisfying resource constraints. The measure of the benefit for conducting the investigative action for feature $f \in F$ is then the information gain:

$$gain(f) = H(\mathrm{N} \mid O) - H(\mathrm{N} \mid O, f) = -\sum_i p_t(i)\log p_t(i) + \sum_{x,y} \frac{\left|k \in S : \{\varsigma_f^k = \varsigma\}\right|}{|\mathrm{N}|} p_f(\xi \mid \varsigma) \sum_{i \in \mathrm{N}:\{x_f^i = x\}} p_t(i \mid \xi_f = \xi)\log p_t(i \mid \xi_f = \xi)$$

In the above, $p_t(i \mid \xi_f = \xi) = \dfrac{p_f\left(\xi \mid \varsigma^i\right)p_t(i)}{\sum_{j \in S} p_f\left(\xi \mid \varsigma^j\right)p_t(j)}$

Finally, we can construct the sequence of investigative actions by defining a decision tree (Figure 7), where each internal node corresponds to the collection action (probe), and the links out of the nodes correspond to the action outcomes (collected information). The leaf nodes of the collection tree correspond to the environment states or groups of states. In the latter case, the leaf node is a belief about being in each of the states of the set, and we can compute the probability of each of those states.

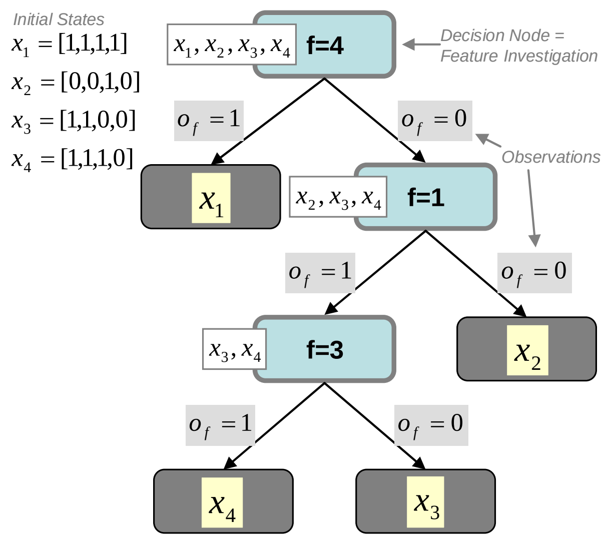

Figure 7: Example of an Investigative Action Plan

The probing plan tree can be constructed using several approaches. One of the simplest algorithms computationally, selects the probes greedily based on the maximum reduction in the entropy. Another approach can consider the outcomes of the probes in the future, so that the policy that maximizes the long-term reduction in the entropy is used. We can use a combination of the roll-out and greedy methods to design such a probing action tree.

Execution of the probing plan is accomplished by following the decision tree and the outcomes of the investigative actions. Each decision to collect the data splits the set of current hypotheses (environment states) into two, and results in a reduction in the entropy (or increased information gain).

3. Generation of Hypothetical Future Environment Dynamics

The knowledge of potential interaction networks and their dynamics indicative of impending conflict relies on ontology of power balance patterns. This ontology can be constructed and updated over time using agent-based simulations of the socio-cultural dynamics for a geographic region of interest. In these simulations, the support and membership for different identity groups is driven by changes in opinions and attitudes. These changes are the result of events in the environment and the exchange of information across the agents' social networks. The model, called SCIPR (Simulation of Cultural Identities for Prediction of Reactions), is a simulation tool that is currently being used to predict the reactions of culturally diverse groups, such as insurgents, political factions, or civilian populations in response to U.S.-influenced events or adversarial actions (Figure 8). By modeling responses and behaviors to "what-if" scenarios such as reconstruction or military intervention, SCIPR helps to gauge the effects of alternative courses of

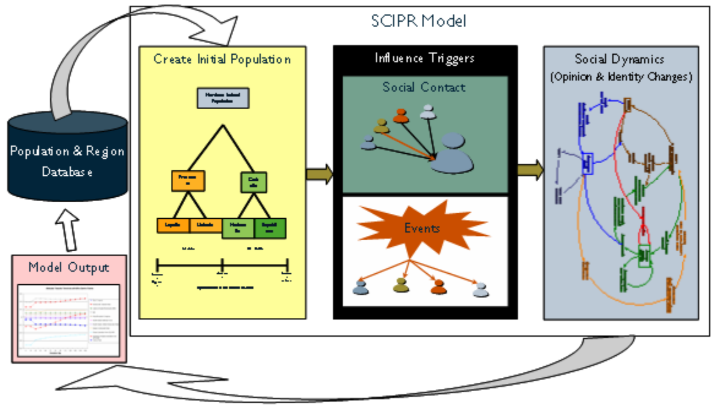

Figure 8: SCIPR Overview

action on measures of interest (attitudes and opinions, such as support for the government or militant groups) and membership of groups (such as political parties and activist groups).

SCIPR represents the changes in identities and opinions via social interactions and regional events. SCIPR is used in SECURE to generate the set of possible futures and construct the dynamic evolution of the state of power balance networks. Each future is generated using a set of parameters and distributions of the attributes and the relationships among the groups in the society. As this data is only partially available, we can apply the orthogonal array or importance sampling methods to determine a subset of attributes (with instantiations for the missing data) that would provide the most comprehensive set of possible futures/PBN networks.

3.1 Population Representation

SCIPR produces these futures by first generating a representative population from census, polling, ethnographic, and SME data on the following categories:

- Primary Identity Groups: Gender, religion, political affiliation, tribe, etc.
    - Identities: Male/Female, Catholic/Protestant/Muslim, Republican/Democrat, etc.
    - Identity relationships: Given an identity (male), what is the distribution of identities in another group (political affiliation)
    - Identity connectivity: Given an identity group (religion) what is the average percentage of contacts that are of the same identity.
- Regions: Country(s), states, counties, cities, etc.
    - Region details: Population, identity composition, connectivity distances, media, etc.
- Opinions and Attitudes: Support for government, stability, foreign presence, etc.
    - Identity to Opinion Ranges: Given and identity (Democrat) distributions of attitudes (for/against or believe/don’t believe) on a given opinion (stability)
- Events: Courses of action, hostile events, natural events, etc.
    - Event reactions: Given an event (attack), an identity (Male), and opinion (stability), what is the expected the reaction (decrease/increase in belief)

This data is then used to create agents with the following attributes:

- Identities: Male, Catholic, Democrat, etc.
- Connection network: Number of contacts, distance range and identity diversity of contacts, access to phone and media, etc.
- Region: Home state, county, city, etc.
- Initial Opinions and Attitudes: Support government, believe there is stability, don’t support foreign presence, etc.

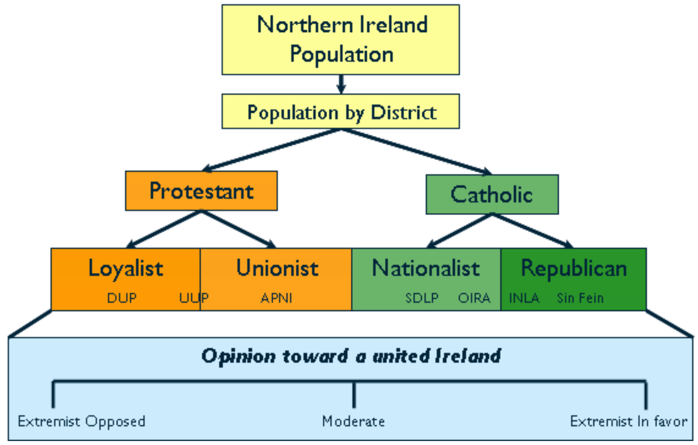

Figure 9: SCIPR Population Generation Example

An example of how an ideally structured population dataset is used to create an agent population is shown in Figure 9. In this example the population of Northern Ireland is distributed by district. For each district, the number of Catholics and Protestants is known. After some analysis it is clear that given that a person is Catholic or Protestant he or she is likely to be one of only a subset of all the political parties (for Protestant it is Loyalist or Unionist parties, for Catholic it is Nationalist or Republican parties). In addition, for all regions it was estimated that 40% of a person's contacts would be of the same religion. Finally given a person's political affiliation, an initial stance on his or her opinion toward having a united Ireland can be derived.

3.2 Social Dynamics

The generated population is then simulated over time allowing agents to exchange opinions over their social networks and to respond to events that occur within their region. The theoretical basis for SCIPR's algorithms controlling the responses to opinion exchanges and events comes from the integration of social identity theory and theories of social influence. Social identity theory was originally developed by Henri Tajfel and John Turner (Tajfel, 1978; Tajfel & Turner, 1979). Many other scholars have continued to develop and test the hypotheses of social identity theory, notably Dominic Abrams and Michael Hogg (2004). Currently, social identity theory is the most well-developed and well-tested theory of cultural change. Social identity theory is highly compatible with theories of social influence, most notably Friedkin's network-oriented structural theory of social influence (Friedkin, 1999, Salzarulo, 2006) as well as other classic research into aspects of social influence such as conformity (Asch, 1955) and group conflict (Sherif et al, 1988; Brown, 2000). Below is a description of all the most important causal mechanisms in social identity theory and theories of social influence.

The primary variables related to social identity theory (Tajfel & Turner, 1979) are identity (sometimes called membership) and opinion (sometimes called attitude). The most relevant definitions of the four words are shown in Table 1, taken from the Merriam-Webster Online Dictionary (2006). While the definitions of identity and membership are very similar as are the definitions of opinion and attitude, the terms identity and opinion will be largely preferred because they are more commonly used in social identity and social influence literature.

Table 1. Important definitions.

Identity: the relation established by psychological identification
Membership: the state or status of being a member
Opinion: a view judgment, or appraisal formed in the mind about a particular matter
Attitude: a mental position with regard to a fact or state; a feeling or emotion toward a fact or state

In social identity theory, people may have multiple identities to which they subscribe at any one time. Minimally, a person has a unique individual identity that determines opinions, perceptions, and actions. In addition, almost all people identify themselves as members of groups. Categories of groups may be at any level of analysis and tend to be hierarchical. For example, some groups may be: gender, age, race, religion, political view, political affiliation, university, hair-combing style, MBTI personality type, etc… Depending on the group and the person, a person may hold multiple identities with regard to the same category. For example someone may consider himself as both a Republican and a Democrat. Social identity theory is concerned most with this *perception* of identity and the *actions* that arise from this perception, rather than institutional membership in a group (being a registered member of the Republican Party or Democratic Party).

As noted above the second concept of importance in social identity theory is opinion. The most useful variables to describe opinion come from models of opinion dynamics (Deffuant, 2006; Deffuant et al, 2002; Hegselmann & Krause, 2002; Salzarulo, 2006). These variables are ***opinion*** and ***certainty***. Opinion is the name of the feeling/judgment about something in the world. Certainty is the strength with which the opinion is held.

Social identity theory (Tajfel & Turner, 1979; Hogg & Grieve, 1999; Hogg et al, 2006) and theories of social influence (Asch, 1955; Festinger, 1954; Friedkin, 1999; Milgram, 1974) suggest three main reasons that people change their identities and opinions: improvement of ***self esteem,*** increase of ***certainty*** (decrease of uncertainty) about the world, and ***conformity*** to social pressure. The self esteem motivation (Tajfel & Turner, 1979) says that a person can improve his/her self esteem by identifying with a group and thinking about how his/her group is good in some way (better than other groups, improving over time, better than some benchmark, etc…). A person can improve his/her certainty about the world by identifying with groups and taking on their opinions (Hogg & Grieve, 1999; Hogg et al, 2006) and by communicating with other people to find out their opinions (Festinger, 1954). A person can also be motivated by the desire to belong, the fear of physical punishment, and the fear of social stigma to change identities and attitudes to conform to the opinions of other people (Asch, 1955; Milgram, 1974).

The implementations of these social dynamics in SCIPR are used to assess how agents change their opinions and how they determine their cognitive membership to any of the available identities based on initial conditions and events in the environment (including actions by interested sides). The simulations therefore provide forecasts of how different COAs might lead to desired or undesired changes in opinions of interest and material support for hostile identities.

4. Conclusions and Future Research

The concept of controlling instability using integration of pattern classification, dual control under uncertainty, and social dynamics simulation technologies promises to result in an effective decision support system for the strategic and operational planners. The SECURE model described in this paper has component technologies that have been validated in previous studies and achieved high performance ratings. We are currently working on obtaining a data set and prototyping component integration for the full-range SECURE system validation study. In the following, we describe the types data useful for secure system, propose types of analyses that SECURE system will be used for, and outline a design for SECURE technology validation. This constitutes our current research in the instability forecasting and prevention domain.

4.1. Example of the Real-world Datasets for Validating the SECURE System

Performing modeling and simulation research on regional instability requires robust and comprehensive data. However, these data are difficult to find and often inappropriate for the development of robust models. Furthermore, relying on less-than-ideal data can undermine the success of modeling and simulation projects. As the result, we decided to leverage four datasets that complement each other to develop valuable output and validate SECURE models.

This data covers the events in the Balkans from April 1989 through July 2003, and most of it is publicly available. We propose to use Balkans data sets to validate SECURE-based conflict prevention technology for the following reasons:

- **Availability and comprehensiveness:** The events in the Balkans have been well covered by the media, and there is a large body of data due to the conflict’s duration. As a result, our data is unclassified, and most of it has been collected from open sources.
- **After-the-fact analysis:** The conflict in former Yugoslavia has been well researched, with many methods and statistical analysis applied to predict various events of interest and analyze the hypotheses about the conflict. In addition, the researchers were able to develop data sets about various socio-cultural-economic variables in this region, and understand the causes of some of the conflicts in various case studies.
- **Varied conflict prevention actions:** Various actors have been engaged in mediation and prevention of the conflict, including U.S., U.N., E.U., Russia, etc. Consequently, the ground truth is available on the preventive and reactive actions performed by the parties involved, and these actions can be used in analyzing the ability of SECURE to propose similar or better actions.

We propose the use of four data sets for our analysis. The first set is event-based data generated by Drs. Gerner and Schrodt using CAMEO (Conflict and Mediation Event Observations) coding scheme from open sources, such as the Reuters News Feeds. The Balkans data set[3] contains events for the major actors (including ethnic groups) involved in conflicts in the former Yugoslavia from April 1989 through July 2003. The coding scheme used to extract the event data from open text sources has been specifically optimized for the study of mediation, and contains a number of tertiary sub-categories specific to third-party mediation in international and inter-ethnic conflict. We consider mediation as one of the conflict prevention-related actions, as it is a specific type of political activity that highlights the role of a third party in facilitating a negotiating process, while not imposing any solution on the parties involved. In addition to the dataset, we will use its analysis (Gerner and Schrodt, 2002) [4] that studied the impact of mediation on the level of hostilities in the region. These results were obtained by analyzing post-mediation events. We will compare these results with predictions of the intervention and mediation actions that the SECURE tool will produce.

The second data source will be a collection of dyadic events generated by Dr. King and VRA[5]. The data available here includes almost 10 million individual events, each coded to the exact day they occur or become known. Each event is summarized in the data as “Actor A does something to Actor B”. Actors A and B record about 450 entities – both countries and within-country groups. The statements of “does something to” are coded in ontology of about 200 types of actions. The data were coded by computer from millions of Reuters news reports. We will use a portion of this data that relates to the Balkans conflict; this would allow us to have a richer of the state of the environment in former Yugoslavia, especially interactions between different actors (including ethnic groups) in the region, and correspondingly increase the environment state prediction

[3] Available at http://web.ku.edu/keds/data.dir/cameo.html
[4] http://web.ku.edu/keds/papers.dir/gerner02.pdf
[5] http://gking.harvard.edu/events/

accuracy. This data also contains information on intervention actions other than mediation, and will be helpful in comparing the ability of SECURE to correctly assess the influence of an action.

The third data source will be from the collection of case-studies conducted by CDA Collaborative Learning Projects[6]. In one of their analysis reports, CDA studied how certain ethnic communities in Bosnia were able to exempt themselves from the violence, particularly ethnic cleansing, that surrounded them, in an effort to identify patterns and to draw lessons about conflict prevention. The reactions of the local communities and the actions they perform to control the crises are essential in the analysis of preventive actions and their impact. In addition to providing data about their coping mechanisms, these case studies provide rich contextual information about the actors of interest in the environment we will analyze. This information, used as attributes of the environment actors, will improve our ability to predict the current state of the environment, but can also be used to analyze the effectiveness of probing actions, where we can assume that this data is not available to the system *a priori* and can only be obtained if a specific intelligence collection activity is performed.

The fourth data source is from the ongoing *Flashpoints* project led by the Center for Emerging Threats and Opportunities (CETO), Marine Corps Warfighting Lab (MCWL), in Quantico, VA. Annual *Flashpoint*s reports identify those nations and regions most likely to experience conflicts in the future that may result in US military intervention. CETO, with assistance from the Potomac Institute for Policy Studies, developed Flashpoints as an index calculated with factors and indicators correlated with instability. The 11 factors are associated with four to six indicators that are derived from existing open-source datasets. *Flashpoints* ranks 158 countries by their relevant indicators and then normalizes the results to allow a ranking from the country most at risk of instability to the country least at risk of instability.[7]

These four, open source data sources will permit a rigorous and comprehensive test of the SECURE model. Furthermore, SECURE’s flexibility permits the inclusion of alternative or complementary data sources, should they become available.

4.2. Types of Analysis Produced by the SECURE System

The SECURE system enables the conduct of the following analyses:

(1) **Learning/Causality:** *Discover Power Balance Patterns Indicating Conflict Tendencies*. SECURE finds which power balance/interaction networks correspond to what incidents using historic data, and infers the causes of the conflicts by looking at the socio-economic power imbalance among specific groups in PBN.
(2) **Inference:** *Discover the power balance that most explains observed data*. Given uncertain observations, the SECURE finds the true state of the power balance /interaction patterns.
(3) **Forecasting:** *Track balance dynamics and predict its evolution*. SECURE forecasts the most likely evolution of the power balance dynamics over time.
(4) **Prevention:** *Act to influence and prevent conflicts*. SECURE finds the cost-effective action strategies to prevent the conflicts and state failures by (i) driving the power balance to desired state; (ii) preventing undesired power balance states; and (iii) collecting the data for increased awareness.

---

[6] http://www.cdainc.com/publications/steps_case_studies.php
[7] Rauch, Dale. Center for Emerging Threats and Opportunities, Marine Corps Warfighting Lab. Telephone Interview. 28 Jan 2008.

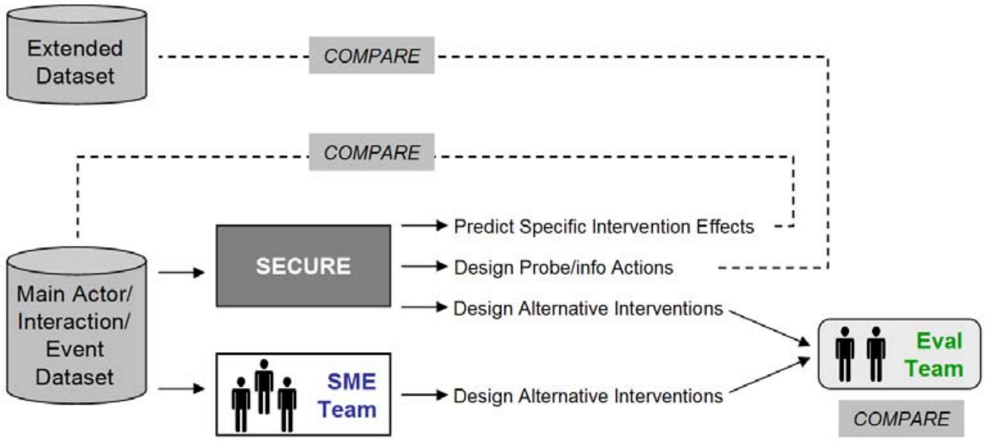

Figure 10: SECURE Validation Experiment Setup

4.3. The Design of the SECURE System Validation Study

In our current research, we are working to validate the SECURE system in three ways against real-world data set:

(a) **Accuracy of predicting the effects of specific interventions:** SECURE system will be used to forecast the impact of intervention actions (e.g., 3-rd party mediation, military action, economic sanctions, etc.) that have been taken in the past. SECURE will only use previous event data to make its predictions about the next state of the environment. These forecasts will then be compared against the real-world event data.

(b) **Effectiveness of probing and info collection actions:** SECURE system will be used to generate probing and info collection actions for the data that will be missing from the main data set (e.g., CDA's data will not be included in the main data set). We will then analyze to see if these probing and passive info collection actions lead to critical info events.

(c) **Effectiveness in designed alternative preventive actions:** SECURE will be used to generate control actions based on the POMDP controller. These actions will oftentimes be different from the actions that occurred in the real-world. Hence, we would need an impartial assessment of the effectiveness of these actions. We propose to conduct an experiment to compare the effectiveness of actions produced by SECURE and by a SME team. All the data on historic environmental dynamics, as well as the intervention policies undertaken in the real-world, will be available to the SME team, which could select the action taken earlier or one of their own. The preventive actions generated by SECURE and the SME team will then be compared blindly by an expert strategic decision making team (Figure 10).

References:

Abrams, D., & Hogg, M. A. (2004). Metatheory: Lessons from Social Identity Research. *Personality and Social Psychology Review, 8*(2), 98-106.

Aberdeen, D., and J. Baxter, "Scalable Internal-State Policy-Gradient Methods for POMDPs", ICML. 2002, pp.3-10

Aberdeen, D., "Policy-Gradient Algorithms for Partially Observable Markov Decision Processes", PhD Thesis, The Australian National University, April 2003

Asch, S. E. (1955). Opinions and Social Pressure. *Scientific American*, pp. 31-35.

Baker, P.H. “Conflict Resolution: A Methodology for Assessing Internal Collapse and Recovery”, in: Pumphrey, Carolyn/Rye Schwartz-Barcott (Hg.), Armed Conflict in Africa. Lanham, MD and Oxford: Triangle Institute for Strategic Studies/The Scarecrow Press, 2003

Bellman, R.. *Dynamic Programming*. Princeton University Press, Princeton, N.J.,, 1957.

Brams, S. J., and Kilgour, D. M. (1988). *Game Theory and National Security*. New York: Basil Blackwell.

Brown, R. (2000). Social Identity Theory: Past Achievements, Current Problems, and Future Challenges. *European Journal of Social Psychology, 30*, 745-778.

Cappé, O., E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*, Springer, 2005

Cassandra, A.R., “Exact and Approximate Algorithms for Partially Observable Markov Decision Processes”, Ph.D. Thesis, Brown University, Department of Computer Science, Providence, RI, 1998

Deffuant, G. (2006). Comparing Extremism Propagation Patterns in Continuous Opinion Models. *Journal of Artificial Societies and Social Simulation, 9*(3).

Deffuant, G., Amblard, F., Weisbuch, G., & Faure, T. (2002). How Can Extremism Prevail? A Study on the Relative Agreement Interaction Model. *Journal of Artificial Societies and Social Simulation, 5*(4).

Festinger, L. (1954) A Theory of Social Comparison Processes. *Human Relations 7,* 117-40.

Friedkin, N. E. (1999). Choice Shift and Group Polarization. *American Sociological Review*, *64*(6), 856-875.

Gerner, D.J., P.A. Schrodt, Ö. Yilmaz, and R. Abu-Jabr, “Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for a Post Cold War World”, 2002 Annual Meeting of the American Political Science Association, 29 August - 1 September 2002

Hauskrecht, M. “Value-function approximations for partially observable Markov decision processes”, *Journal of Artificial Intelligence Research*, 13:33-94, August 2000.

Hegselmann, R., & Krause, U. (2002). Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artificial Societies and Social Simulation, 5*(3).

Hogg, M. A., Sherman, D. K., Dierselhuis, J., Maitner, A.T., & Moffitt, G. (2006). Uncertainty, Entitativity, and Group Identification. *Journal of Experimental Social Psychology* (in press).

Hogg, M.A., & Grieve, P. (1999). Social Identity Theory and the Crisis of Confidence in Social Psychology: A Commentary, and Some Research on Uncertainty Reduction. *Asian Journal of Social Psychology, 2,* 79-93.

Kaelbling, L.P., M.L. Littman and A. R. Cassandra, “Planning and Acting in Partially Observable Stochastic Domains”, Tech Report CS-96-08, Brown University, February 1996

King, G., and Zeng. L. “Improving Forecasts of State Failure”, World Politics, Vol. 53, No. 4 (July, 2001): 623-58

Levchuk, G., Y. Levchuk, and K. Pattipati, “Identifying Command, Control and Communication Networks from Interactions and Activities Observations”, *Command and Control Research and Technology Symposium*, 2006, San Diego, CA.

Levchuk, G., F. Yu, H. Tu, K. Pattipati, and Y. Levchuk. “Identifying the Enemy – Part I: Automated Network Identification Model”, 12th ICCRTS, 2007, Newport, RI.

McCallum, R. A. (1996), “Learning to use selective attention and short-term memory in sequential tasks.” In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S. W., editors, From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, Cambridge, MA, pages 315-324. MIT Press, Bradford Books.

Merriam-Webster Online. (2006). Retrieved Oct 19, 2006 from http://www.m-w.com/.

Milgram, S. (1974). Obedience to Authority. New York: Harper and Row.

O’Brien, S.P. 2002. “Anticipating the Good, the Bad, and the Ugly: An Early Warning Approach to Conflict and Instability Analysis.” Journal of Conflict Resolution 46(6): 791-811.

O’Brien, S.P. 2004 “Near-Term Forecasts of Crisis and Instability Using Text-Based Events (NEAR-TERM FORECITE) Date: 04 OCT 2004”, Proceedings of the 72nd MORS Symposium, 22-24 June.

Ottaway, M.S. , and Mair, S. “States at Risk and Failed States”, *Carnegie Endowment for International Peace, Policy Outlook*, September 2004

Papadimitriou, C.H., and J. N. Tsitsiklis. “The complexity of Markov decision ,processes”, *Mathematics of Operations Research*, 12(3):441-450, 1987.

Popp, R., D. Allen and C. Cioffi-Revilla, “Utilizing Information and Social Science Technology to Understand and Counter the 21st Century Strategic Threat,” Emergent Information Technologies and Enabling Policies for Counter-Terrorism, (Eds. R. Popp, J. Yen), Wiley-IEEE Press, pp. 1-27, Jun 2006.

Popp, R. Kaisler, S.H. Allen, D. Cioffi-Revilla, C. Carley, K.M. Azam, M. Russell, A. Choucri, N. Kugler, J. “Assessing nation-state instability and failure”, Aerospace Conference, 4-11 March 2006

Popp, R.,“Exploiting AI, Information and Computational Social Science Technology to Understand the Adversary,” AAAI Spring Symposium on AI Technologies for Homeland Security, Stanford University, CA, March 2005

Poupart. P., L.E. Ortiz, and C.Boutilier, “Value-directed sampling methods for monitoring POMDPs”, Uncertainty in Artificial Intelligence, August 2001, URL www.citeseer.nj.nec.com/445996.html

Rangarajan A, A. Vuille, E. Mjolsness, “Convergence properties of the softassign quadratic assignment algorithm”, Neural Comput. 1999 Aug 15;11(6):1455-74.

Rotberg, R., “The New Nature of Nation-State Failure”, *The Washington Quarterly,* 25 3 (Summer 2002): pp. 85-96.

Rotberg, R., “The Failure and Collapse of Nation-States: Breakdown, Prevention, and Repair,” in When States Fail: Causes and Consequences, (Ed. R. Rotberg), Princeton University Press, pp. 1 – 49, 2004.

Rabiner, L.R., “A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition”, *Proceedings of the IEEE*, 77 (2), p. 257–286, February 1989.

Rangarajan, A., A. Yuille, and E. Mjolsness “Convergence Properties of the Softassign Quadratic Assignment Algorithm”, *Neural Computation*, August 15, 1999, Vol. 11, No. 6, Pages 1455-1474

Ron, D., Singer, Y., and Tishby, N. (1994). “Learning probabilistic automata with variable memory length.” In Aleksander, I. and Taylor, J., editors, Proceedings Computational Learning Theory. ACM Press.

Salzarulo, L. (2006). A Continuous Opinion Dynamics Model Based on the Principle of Metacontrast. *Journal of Artificial Societies and Social Simulation, 9*(1).

Saunders-Newton, D., A. Frank, R. Popp, R. “Avoiding 21st Century Societal and Regional Conflicts Through Advanced Technologies for Understanding, Anticipating and Shaping”, 2005 IEEE Aerospace Conference, 5-12 March, 2005, pp. 1- 7

Schrodt, P.A., “Automated Coding of International Event Data Using Sparse Parsing Techniques”, Fifth International Conference on Social Science Methodology, Cologne, Germany, October 3 - 6, 2000.

Schrodt, P.A., and D.J. Gerner, (2001) Analyzing International Event Data: A Handbook of Computer-Based Techniques, Draft Book, available at http://web.ku.edu/keds/papers.dir/AIED.Preface.pdf

Sherif, M., Harvey, O. J., White, B. J., Hood, W. R., & Sherif, C. W. (1988). *The Robbers Cave experiment: Intergroup conflict and cooperation.* Middletown, CT: Wesleyan University Press.

Sondik, E.J., “The Optimal Control of Partially Observable Markov Processes”, PhD Thesis, Stanford University, Stanford, CA, 1971

Sutton, R.S., D. McAllester, S. Singh, and Y. Mansour, “Policy gradient methods for reinforcement learning with function approximation. Advances in Neural Information Processing Systems, 12:1057-1063, 2000.

Turchin, P. Historical Dynamics: Why States Rise and Fall. Princeton, NJ: Princeton University Press, 2003

Tajfel, H. (1978). *Differentiation between Social Groups*. Cambridge, Cambridge University Press.

Tajfel, H., & Turner, J. C. (1979). "An Integrative Theory of Intergroup Conflict." In S. Worchel, W. G. Austin (Eds), *The Social Psychology of Intergroup Relations*. Monterey, Brooks/Cole.