SOMA Models of the Behaviors of Stakeholders in the Afghan Drug Economy: A Preliminary Report

Amy Sliva   Maria Vanina Martinez   Gerardo I. Simari   V.S. Subrahmanian
Department of Computer Science and
University of Maryland Institute for Advanced Computer Studies (UMIACS)
University of Maryland College Park
College Park, MD 20742, USA

Abstract

Most cultural reasoning today is done by anthropologists and sociologists who use their detailed knowledge of culture to make predictions about how a given group will respond to a given situation. The main problem with this is that experts in a particular culture or subculture are few and not readily accessible to the many who might suddenly need to tap their expertise (e.g. in cases of war or conflict). In this paper, we briefly describe how the SOMA (Stochastic Opponent Modeling Agents) paradigm proposed by the authors can and has been used to model the behaviors of various stake-holders in the drug trade in Afghanistan.

Introduction

Over the past several decades, Afghanistan has been embroiled in numerous external and internal conflicts—the Soviet invasion in 1979, the struggle for control by rival groups when the Soviets withdrew in 1989, and the subsequent rise of the Islamic fundamentalist Taliban (Chayes 2006). However, in spite of its tumultuous history, Afghanistan has largely resided on the periphery of international politics. Following the terrorist attacks of September 11, 2001, this remote and isolated country has become a major battleground against terrorist groups, and has come to the forefront of foreign policy decision-making.

The international community has placed a large emphasis on improving the stability of the country and the new Afghan government. High among the many challenges to achieving this goal is the Pashtun tribal culture that defines many aspects of Afghan society and its civil institutions. Enacting a unified policy in a society based on tribal divisions and customary laws is very difficult, especially when these tribal customs and behaviors are only understood by experts or actual Afghan citizens. The necessity of constructing policy initiatives that work within the Pashtun culture, and of anticipating the reactions of members of this society, places the knowledge of such experts at a premium.

With a limited number of experts available, computational models of various agents involved in tribal groups or cross-sections of Afghan society—such as those agents engaged in the drug trade—can be used to make this information accessible to policy makers to help them construct culturally relevant and viable policy solutions. Agent models can provide information on the types of behavior that can be expected in various scenarios, so decision-makers would be able to understand, for instance, how Afghan farmers may react to a policy of destroying opium crops in order to interrupt the drug trade.

In this paper we present the Stochastic Opponent Modeling Agents (SOMA) system as a prototype solution for providing a model of cultural behaviors. For this work, we have developed a probabilistic logic framework to represent the likely actions of an agent under certain conditions, and several algorithms for reasoning within this framework to determine the most likely actions an agent will perform in a given scenario. In the following sections, we provide a detailed description of Afghan tribal society and the Afghan drug trade, which comprises an enormous portion of the country's GDP. Using the SOMA framework, we have constructed simple behavioral rules for agents in the Pashtun tribal society, and can demonstrate the system's ability to indicate the most probable actions of these agents given certain situations.

Case Study: the Drug Trade in Afghanistan

In this paper, we will focus on the drug trade in the tribal areas located throughout the Afghan/Pakistani border. This particular geographic region has many features of interest, and a rich history explaining the current state of affairs. For instance, (Goodhand 2005) states that:

> "Afghanistan's present borders were defined by imperial powers in the nineteenth century and successive Afghan rulers attempted to defend, strengthen, and redefine these borders in response to external aggression or internal pressures. Borders are 'political membranes' and markers of the success of the state-building enterprise. As (Scott 1976) argues, borderlands are shadow societies, beyond the reach of the state, often with an 'insurrectionary tradition'. In this ongoing 'conversation' between the state and borderlands, violent conflicts have been defining moments of change, shifting the balance of power back and forth between core and periphery. The contemporary political economy of Afghanistan is a product of this history..."

The reality of the situation in post-Taliban Afghanistan is that of warlords competing against each other, as well as with the central authority, to establish 'mini-states' (Goodhand 2005). A great portion of the country's economy is sustained by the cultivation of poppy, which is later used in the production of opium. Initially, most of the poppy fields were in the provinces of Helmand, Kandahar, Uruzgan, and Nangarhar, which are located in Pashtun belts in the south and east, and divided up among many warlords. Farms tend to be small, but cultivation of poppies can yield about $15,000 per hectare per year, which is about ten times what the cultivation of wheat would yield(Goodhand 2005). The cultivation of poppies has proliferated, and there are now poppy fields in 28 provinces, and an estimated 400 laboratories involved in the production of drugs (Goodhand 2005).

Figure 1: Afghanistan and the neighboring area[1].

Afghanistan is the largest producer of opium in the world, yielding about 3.5 million kilograms per year. The bulk of the production is headed to Europe by means of large cargo planes that fly from Kandahar and Jalalabad to Dubai, and come back with dollars. This drug money has been linked to terrorism, since it is used in part to fund Al-Qaeda and by Pakistani intelligence to support terrorist activities in Kashmir. The tribes in the border areas play a major role in this scheme, since they control much of the activities required in the production and transportation of opium. For instance, various tribes control the Afghan side of the border in the Nangarhar province. A former governor of this province has been identified as a major drug kingpin. Another closely related tribe controls the traffic on the Pakistani side, including customs posts and local administration offices. Because these Pakistani tribes constitute an autonomous tribal agency, the Pakistani central government cannot legally operate without permission from local tribal leaders.

There are many actors involved in the Afghan opium economy among which we can mention *farmers, itinerant laborers, members of regional militia, landowners, traffickers, warlords,* and *government figures*, not only local but also at the national level. The motives and methods used by each group vary based on their geographic location, economic circumstances, relationships with ethnic groups and external parties, and prevailing political conditions.

For our case study we focus on a small set of actors that are in some way or another involved in the drug economy of a typical village in the Afghan/Pakistani border region. The following characterization of actors was synthesized from (Goodhand 2005; Kakar 2005).

- **Malik:** Maliks are the intermediaries and representatives between the village community and central power/government, and are in charge of solving communal disputes and maintaining communal property.
- **Khan:** In Arabic, the word "khan" denotes a ruler or central authority figure. Of course, Arabic is not spoken in Afghanistan - there, the word khan colloquially denotes a large feudal landowner who controls many resources in the community along with providing jobs to laborers and land to sharecroppers. Khans may also arbitrate conflicts.
- **Ulema:** The Ulema shura is a group of religious leaders who lead prayers, give sermons, and have the power of moral judgement in the community. They are also involved in resolving conflicts from the point of view of Shariah (Islamic law).
- **Member of Shura:** The Shura is a tribal council that meets only as problems arise in order to solve them. Such imminent problems range from personal disputes to maintenance of communal property.
- **Warlord:** Warlords are the large, regionally based commanders (remnants of past conflicts) which were able to utilize ties of solidarity based on ethnicity or regional allegiances to build up enough support among local commanders to control the area. They collect taxes, control borders, control local resources, produce and sell drugs, arms, etc.
- **Farmer:** Farmers are the workers of the land, who may own the land or rent it from khans or warlords. Cultivating poppies has, in many places, become the main way for farmers to gain access to land or to seasonal employment, giving them access to land on which they can also cultivate food crops to support themselves and the community. In addition, through the tribal salaam system of lending, farmers can gain access to credit; farmers are given an advance payment on a fixed amount of future production, and opium, with its reliable rates of return, is favored by tribal money lenders.
- **Village Citizens:** This class includes general men, children, and women from a village.

[1]"Afghanistan" [map]. Visual Scale. The World Factbook: Afghanistan. http://www.cia.gov/cia/publications/factbook/geos/af.html. (2007)

While the micro-economy of each village or region has it's own unique properties, poppy cultivation and the opium trade has an impact on the lives and behaviors of the above agents throughout much of the tribal Afghan/Pakistani border region.

There is a plethora of methods that have been attempted to ameliorate the problems that come with the drug economy. In order to provide a more comprehensive illustration of the scenario, we will enumerate some of these methods here.

- *Destroy the poppy fields.* This has been tried extensively without much success. The market reacts by treating this as a drop in production (supply) with no corresponding drop in demand. Hence, prices go up. Moreover, fields have relocated, penalizing farmers instead of distributers.
- *Destroy the production laboratories.* Even though there is a relatively low number of labs in Afghanistan (numbering in the 100s), this method proved ineffective due to warnings of raids by corrupt officials. As in the case where poppy fields were destroyed, this only resulted in raised retail prices and a penalization for the farmers.
- *Invest money in basic infrastructure.* An apparent decline in poppy production in the Nihag valley in Pakistan was due to efforts to enhance irrigation, road and terrace construction, other crops, livestock, and electrification. However, all this tactic actually did was push production into Afghanistan and to tribal areas (Khyber, Bajaur, North and South Waziristan). (Blanchard 2005)
- *Legalize opium production.* This proposal would set up a legal agency to license, produce and buy opium from farmers at a fixed price (around $52/kg; current Afghan price = $113; legalized price in India about $28) and then export opium under UN licenses to pharmaceutical companies. Another option is to manufacture morphine/codeine in Afghanistan itself and then export the manufactured morpheine/codeine products under UN export licenses. The key expected pitfalls in this case include (i) resistance from distributors, and (ii) the difficulty of implementing such a system in Afghanistan in the presence of corrupt key officials overseeing the licensing process. In addition, there could be political opposition in the US to legalizing drug production.
- *Enlist help of fundamental Islamic clerics and Islamic law for enforcement.* Strict fundamentalist Islam forbids consumption of opium,. Moreover, the lowest historic opium production levels in Afghanistan coincided with the rule of the Taliban who strongly suppressed the trade. Therefore, enlisting help (with compensation) from fundamentalist clerics to send an anti-opium message could be helpful. The same can be done with fundamentalist forces to enforce drug laws. However, this legitimizes past enemies and places influence in their hands. Furthermore, distributors are cut out of the loop and can be expected to both violently rebel against such a plan and influence the current Afghan administration to deny political support.
- *Take distributors out of the loop.* It is conceivable to make the current opium distributors the distributors of legally exported opium/morphine/codeine. This has many advantages, such as bringing them into the legal regulatory system, guaranteeing income to them for lost drug income. Regulatory control also ensures a fair deal for farmers vs. distributors. Finally, potential immunity for past crimes may be desirable, as they can now spend part of their previously ill-gotten gains legally.

  However, there is always the danger of distributors siphoning off parts of the opium to the illegal market at higher prices. Furthermore, will revenues be adequate to bind them?

As we will explain in the next sections, all of this knowledge of the scenario, including facts about how the world is organized, actions taken in the past, etc., can be used to build a model of the major players, which is useful both in understanding and predicting future actions. In the next section, we will present the SOMA system for cultural and behavioral modeling.

The SOMA System

The Stochastic Opponent Modeling Agents (SOMA) system we have developed is a prototype system for reasoning with behavioral and cultural models and determining the most probable actions that an agent will take in a given situation. The current system allows a user to select an agent model and an initial state of the world. Users are then able to calculate the k most probable courses of action (see Definition 4)—or sets of actions—that the agent might take in that state. Successive versions of the SOMA system will provide more user features, such as the ability to create new agent models or scenarios.

The SOMA Language

In this section, we will present a brief introduction to the SOMA language, which is the foundation of the SOMA system. SOMA programs are an immediate variant of the probabilistic logic programs introduced in (Ng & Subrahmanian 1991; 1992). We assume the existence of a logical alphabet that consists of a finite set $\mathcal{L}_{cons}$ of constant symbols, a finite set $\mathcal{L}_{pred}$ of predicate symbols (each with an associated arity) and an infinite set $\mathcal{V}$ of variable symbols. Function symbols are not allowed in our language. Terms and atoms are defined in the usual way (Lloyd 1987). We assume that a subset $\mathcal{L}_{act}$ of $\mathcal{L}_{Pred}$ are designated as *action symbols* - these are symbols that denote some action. Thus, an atom $p(t_1, \ldots, t_n)$, where $p \in \mathcal{L}_{act}$, is an *action atom.* Every atom (resp. action) is a wff (resp. action wff). If F, G are wffs (resp. action wffs), then $(F \wedge G), (F \vee G)$ and $\neg G$ are all wffs (resp. action wffs).

Definition 1 *If F is a wff (resp. action wff) and $\mu = [\alpha, \beta] \subseteq [0, 1]$, then $F : \mu$ is called a p-annotated (resp. ap-annotated—short for "action probabilistic" annotated) wff. μ is called the p-annotation (resp. ap-annotation) of F.*

Without loss of generality, we assume that F is in conjunctive normal form (i.e. it is written as a conjunction of disjunctions).

support(X,Y): $[0.75, 0.9]$	$\leftarrow$	*is_warlord(X)* $\wedge$ *profit_from_alliance(X,L)* $\wedge$ *is_leader(L,Y)* $\wedge$ *soc_or_pol_religious_regime(Y).*
buy_war_materials: $[0.9, 1]$	$\leftarrow$	*is_warlord(X)* $\wedge$ *needs(X,show_strength).*

Figure 2: A simple example of a SOMA program.

Definition 2 (SOMA-rules) *If F is an action formula, $A_1, A_2, ..., A_m$ are action atoms, $B_1, \ldots, B_n$ are non-action atoms, and $\mu, \mu_1, ..., \mu_m$ are ap-annotations, then $F : \mu \leftarrow A_1 : \mu_1 \wedge A_2 : \mu_2 \wedge \ldots \wedge A_m : \mu_m \wedge B_1 \wedge \ldots B_m$ is called an $SOMA$-*rule*. If this rule is named c, then $Head(c)$ denotes $F : \mu$; $Body^{act}(c)$ denotes $A_1 : \mu_1 \wedge A_2 : \mu_2 \wedge \ldots \wedge A_m : \mu_m$ and $Body^{state}(c)$ denotes $B_1 \wedge \ldots B_n$.*

Intuitively, the above SOMA-rule says that an unnamed entity (e.g. a group g, a person p etc.) *will take action F with probability in the range μ if $B_1, \ldots, B_n$ are true in the current state (we will define this term shortly) and if the unnamed entity will take each action A_i with a probability in the interval μ_i for $1 \leq i \leq n$.*

Definition 3 (SOMA-program) *A* SOMA program *is a finite set of $SOMA$-rules.*

Figure 2 presents a small SOMA program consisting of two rules. For example, the first rule states that a warlord will support a social or political religious regime if he can profit from an alliance with its leader.

Definition 4 (COA/state) *A* course of action *(COA) is any set of ground action atoms. A* state *is any finite set of ground non-action atoms.*

Note that both COAs and states are just ordinary Herbrand interpretations. As such, it is clear what it means for a state to satisfy $Body^{state}$.

Definition 5 *Let Π be a $SOMA$-program and s a state. The* reduction of Π w.r.t. s, *denoted by Π_s is $\{F : \mu \leftarrow Body^{act} \mid s$ satisfies $Body^{state}$ and $F : \mu \leftarrow Body^{act} \wedge Body^{state}$ is a ground instance of a rule in $\Pi\}$.*

Note that Π_s never has any non-action atoms in it.

In order to apply our reasoning algorithms to a SOMA-program, we want to assign a probability range to each formula in the heads of the SOMA-rules. To accomplish this, we can associate a fixpoint operator T_{Π_s} with a SOMA-program Π and a state s which maps sets of ground ap-annotated wffs to sets of ground ap-annotated wffs as follows.

Definition 6 *Suppose X is a set of ground action atoms. We first define an intermediate operator $U_{\Pi_s}(X)$ as follows. $U_{\Pi_s}(X) = \{F : \mu \mid F : \mu \leftarrow A_1 : \mu_1 \wedge \cdots \wedge A_m : \mu_m$ is a ground instance of a rule in Π_s and for all $1 \leq j \leq m$, there is an $A_j : \eta_j \in X$ such that $\eta_j \subseteq \mu_j\}$.*

Intuitively, $U_{\Pi_s}(X)$ contains the heads of all rules in Π_s whose bodies are deemed to be "true" if the action wffs in X are true.

In order to assign a probability interval to each ground action atom, we use the same procedure followed in (Ng & Subrahmanian 1991). We use $U_{\Pi_s}(X)$ to set up a linear program $CONS_U(\Pi, s, X)$ as follows. For each COA c_i, let p_i be a variable denoting the probability of c_i being the "real COA" taken by an agent. As each c_i is just a Herbrand interpretation, the notion of satisfaction of an action formula F by a COA c, denoted by $c \mapsto F$, is defined in the usual way. The following constraints are in $CONS_U(\Pi, s, X)$:

1. If $F : [\ell, u] \in U_{\Pi_s}(X)$, then $\ell \leq \Sigma_{w_i \mapsto F}\, p_i \leq u$ is in $CONS_U(\Pi, s, X)$.
2. $\Sigma_{w_i} p_i = 1$ is in $CONS_U(\Pi, s, X)$.

We refer to these as constraints of type (1) and (2), respectively. Our operator $T_{\Pi_s}(X)$ is then defined as follows.

Definition 7 *Suppose Π is a SOMA-program, s is a state, and X is a set of ground ap-wffs. Our operator $T_{\Pi_s}(X)$ is then defined to be $\{F : [\ell(F), u(F)] \mid (\exists \mu)\ F : \mu \in U_{\Pi_s}(X)\} \cup \{A : [\ell(A), u(A)] \mid A$ is a ground action atom $\}$.*

Thus, $T_{\Pi_s}(X)$ works in two phases. It first takes each formula $F : \mu$ that occurs in $U_{\Pi_s}(X)$ and finds $F : [\ell(F), u(F)]$ and puts this in the result. Once all such $F : [\ell(F), u(F)]$'s have been put in the result, it tries to infer the probability bounds of all ground action atoms A from these $F : [\ell(F), u(F)]$'s. The $T_{\Pi_s}(X)$ operator has a least fixpoint, $T^{\omega}_{\Pi_s}$, which contains all of the ground action atoms in X annotated with tight probability intervals.

The final SOMA-program resulting from the application of $T_{\Pi_s}(X)$ to a program Π and a state s is used as input to the SOMA reasoning engine for determining the most probable course of action an agent will take. This is described more thoroughly in the following section.

Reasoning in the SOMA System

We have developed several algorithms to reason about the most probable actions that an agent will take in a given situation; these algorithms find the most probable course of action from a SOMA-program and a state.

Definition 8 (lower/upper probability of a COA) *Suppose Π is a $SOMA$-program and s is a state. The* lower probability, $\mathsf{low}(c_i)$ *of a COA c_i is defined as:* $\mathsf{low}(c_i) = \textbf{minimize}\ p_i$ ***subject to*** $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$. *The* upper probability, *$up(c_i)$ of world w_i is defined as* $up(c_i) = \textbf{maximize}\ p_i$ ***subject to*** $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$.

Thus, the low probability of a COA c_i is the lowest probability that that world can have in any solution to the linear program $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$. Similarly, the upper probability for the same COA represents the highest probability that that world can have. It is important to note that for any COA c, we cannot *exactly* determine a point probability for c. This observation is true even if all rules in Π have a point probability in the head because our framework does not make any simplifying assumptions (e.g. independence) about the probability that certain things will happen.

A *naive* algorithm for finding the most probable world follows directly from the definition of $CONS_U(\Pi, s, X)$ for a SOMA-program Π, a state s, and a set X of ground action atoms:

1. Compute $T^{\omega}_{\Pi_s}$; $Best = NIL$; $Bestval = 0$;
2. For each world c_i,
   (a) Compute $low(c_i)$ by minimizing p_i subject to the set $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$ of constraints.
   (b) If $low(c_i) > Bestval$ then set $Best = c_i$ and $Bestval = low(c_i)$;
3. If $Best = NIL$ then return any COA whatsoever, else return $Best$.

The naive algorithm does a brute force search after computing $T^{\omega}_{\Pi_s}$. It finds the low probability for each COA and chooses the best one. Clearly, we can use it to solve the for the maximal upper probability of a COA by replacing the minimization in step 2(a) by a maximization. Because the number of variables in $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$ is exponential in the size of Π and s, the naive algorithm often takes intractable amounts of time to compute the most probable COAs. We have developed several heuristic approximation algorithms that can significantly reduce the running time of the calculation, while producing satisfactorily accurate results.

Heuristic Approximation Algorithms The goal of the heuristic approximation algorithms is to reduce the number of variables in the linear program $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$ for a SOMA-program Π, state s, and set of ground action literals X.

Random Sampling Heuristic: Suppose we make an *a priori* commitment to only look at some set S_k of k variables from the linear program. In this case, we could eliminate variables not in S_k from any summation in constraints of type(1) as defined in the previous section. We then solve for the most probable COA as in the naive algorithm, minimizing each variable in S_k with respect to the reduced set of constraints, $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$ and returning the variable (and value) with the highest value.

It is immediately apparent that as all the lower bounds in $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$ are set to ℓ, a solution to this reduced constraint set may or may not exist. Rather than weakening the lower bound from ℓ to 0 (which would guarantee a solution), we wondered how "close" to ℓ one can get while still having a solvable system of equations, yeilding the next heuristic algorithm.

Binary Heuristic: The *binary heuristic* works as follows by *only modifying lower bounds* of a reduced set of constraints $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$. We start with $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$ and see if it is solvable by itself. If so, we return the same solution as the random sampling heuristic. If not, we try to decrease the lower bounds of one or more constraints in $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$ as follows. Suppose $c^{\star}$ is one such type (1) constraint of the form

$$\ell^{\star} \leq \Sigma_{q_i \in S_k} q_i \leq u$$

In this case, we try to replace $\ell^{\star}$ by $\frac{\ell^{\star}}{2}$. If this yields a solvable set of equations, we try to replace $\frac{\ell^{\star}}{2}$ by $\frac{3\times\ell^{\star}}{4}$—if the resulting system of equations is unsolvable, we try to replace it with $\frac{5\times\ell^{\star}}{8}$ and so forth. Effectively, we try to keep the lower bounds of constraints as close to those in the original $\mathsf{CONS}_U(\Pi, s, T^{\omega}_{\Pi_s})$ as possible, while still being solvable when terms not in S_k are eliminated from the type (1) constraints. We will call this the *binary heuristic* due to the fact that it resembles a binary search.

Once we have completed this process of modifying the lower bounds of constraints in $\mathsf{CONS}'_U(\Pi, s, T^{\omega}_{\Pi_s})$ (let the resulting system of constraints be called $\mathsf{CONS}^{\bullet}_U(\Pi, s, T^{\omega}_{\Pi_s})$ we minimize each variable in S_k subject to the constraints in $\mathsf{CONS}^{\bullet}_U(\Pi, s, T^{\omega}_{\Pi_s})$, returning the COA with the highest minimal value (together with its value).

In the next section, we will see how to apply the SOMA language and algorithms to behavioral models of the major actors in the Afghan opium economy, presenting a sample set of rules and demonstrating the functionality of the current SOMA system.

The SOMA System Applied to the Afghan Drug Trade Domain

Based on (Kakar 2005; Goodhand 2005), we were able to draw a conceptual map involving several possible actors in the Afghan drug economy (as given above), and the relationships amongst them that encourage the cultivation of poppies, production and trafficking of drugs, and corruption in the region.

We have developed agent models for each of these actors in the Afghan/Pakistani drug economy scenario. Each of these groups demonstrates patterns of behavior that relate to the production and trade of opium. Based on the dynamics of the drug trade, there are countless actions and situations that we could model; however, to date we have only constructed models for a small portion of what each agent could possibly do. Below we provide a brief summary of the role the various actors play in the drug economy, as well as a description of some of the behaviors we have chosen to model:

- *Warlord:* Warlords and their militias have a great deal of power in the village we are modeling, and control a large portion of the economy. In this type of situation, many of the taxes collected and the revenues obtained through the sale of opium go directly to funding these militias. Because the central government is still relatively weak in the tribal border regions, the warlords also have control over border posts and local military leaders, bribing and threatening them in order to successfully smuggle drugs across the Pakistani border.

  Sample Actions:

  – enforce poppy ban/ban poppy cultivation—One of the possible methods for combating the drug trade has been to ban the cultivation of poppies. However, warlords

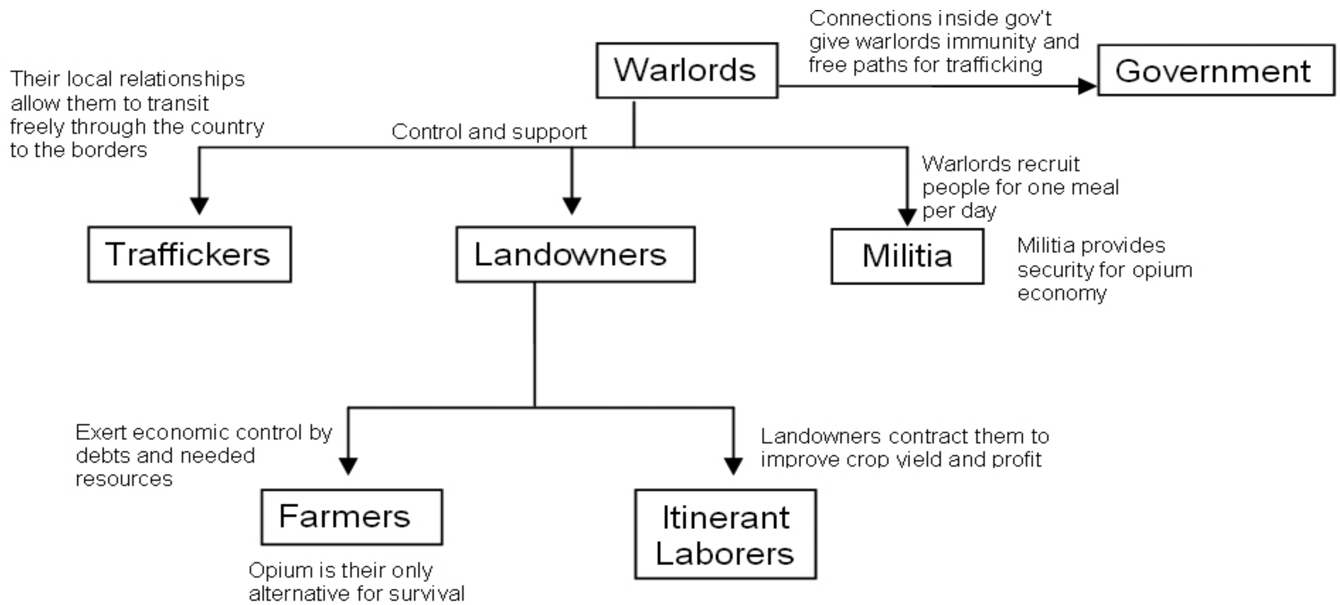

Figure 3: The most important actors involved in the Afghan drug economy.

who have a vested interest in the opium trade will not comply with such a ban, or enforce it in lands that they control.

- purchase arms or other war materials—Warlords who are attempting to seize power from a rival warlord or to gain control of a territory from the central government will purchase weapons in order to increase the capabilities of their backing militia.
- collect taxes—Those warlords who have control over villages on the Afghan-Pakistani border are very likely to collect extra taxes from those traders attempting to smuggle opium through border checkpoints.

- *Khan:* Because of the relative financial advantage of farming poppies as opposed to other crops, khans have a large incentive to force their sharecroppers to cultivate poppies for the drug trade. Opium production has become such an essential part of tribal village economies, that khans determine the price of land rent based on its expected yield of poppies. In addition, powerful warlords and currupt officials who control the trade of opium can threaten landowners to ensure that poppies are grown.

  Sample Actions:

  - switch from poppy cultivation to a legal crop—Policymakers have attempted to encourage the cultivation of legal crops rather than poppies. There is a very low probability, however, that a khan will decide to switch to growing one of the accepted crops on his lands if the poppy continues to yield a greater revenue.
  - enforce poppy ban—Depending on his beliefs, loyalties, and affiliations, a khan might enforce a ban on poppy cultivation with a probability of around 50%. Because of the economic advantages of involvement in opium production, a khan would only comply with a government ban under pressure from a certain powerful group or individual—such as government representatives, a warlord, or a foreign group—with which he desires good relations. Furthermore, if the religious tribal authorities have declared opium to be unislamic, a religious khan will consider enforcing a poppy ban as well.
  - provide financial backing—A khan will provide financial support for a group or individual—such as a militia or warlord—if he expects to gain further financial rewards from such an action. This action leads to the formation of networks of powerful individuals all involved in the opium economy.

- *Farmer:* Afghan farmers have grown to rely increasingly on the cultivation of poppies for their survival and livelihood. Because opium reliably brings greater revenue than legal crops, money lenders give preferential loans to farmers growing poppies, and many farmers' quality of life has vastly improved with the greater income from poppy cultivation. Even farmers who would like to comply with the anti-drug laws are often influenced or intimidated by the powerful khan whose land they farm or warlords who control the transport and trade of drugs; in many cases the cultivation of poppies is not actually a choice for the farmers.

  Sample Actions:

  - cultivate poppy—A farmer will probably cultivate poppies if he is in debt, because loans are easier to secure if poppies are being grown. Similarly, if growing poppies as a sharecropper will allow the farmer to own his own land, he is very likely to choose poppy cultivation.

- switch from poppy cultivation to a legal crop—A farmer growing poppies is very unlikely to switch to a legal crop when the revenue from poppies is higher.

- *Malik, Member of Shura, Ulema:* In spite of the rise of opium production, the traditional civil and religious institutions in the villages often remains completely intact. Maliks still serve as the intermediaries between the tribal community and the central government, and Ulemas (councils) and the Shura are still convened to arbitrate local conflicts and preside over many public works and village-wide issues. However, tensions do exist between the traditional social structure and the new drug economy, as most of the new wealth is controlled by the young men in the village who work for warlords or khans. Similarly, powerful warlords can threaten or bribe civil officials in the village into supporting the drug trade or local militias.

  Sample Actions:

  - support—An ulema is very likely to support a local warlord who has threatened him. This allows warlords to have access to the local councils and further control over the village. Slightly less likely, an ulema will support a social, political, or religious regime—such as the central government or an Islamist group—if he shares beliefs with the regime.
  - provide financial support—provide financial backing—An ulema will provide financial support for a group or individua if he expects to gain further financial rewards or political control from such an action.

In Figure 6, we present a selection of the SOMA-rules for several agents representing their behavior and involvement in the Afghan drug trade; for clarification, each rule is followed by a brief English description. We reiterate that the SOMA system does *not* make any independence assumptions regarding the actions. We also reiterate that the probability annotations in these rules were derived by means of analyses of qualitative (as opposed to quantitative) data and discussions with subject matter experts on the topic. For each agent, we created between 30 and 40 rules about their behavior. This is one reason why we have derived only 30–40 rules per stakeholder group.

We are separately working on mechanisms for automatically inferring these agent models from raw data, thereby increasing both the statistical significance of the rules and possibly discovering relationships and behavioral patterns that humans would not have been able to recognize. However, these results are preliminary and we are not in a position to report on them at this time.

The remainder of this section will walk through an example using the SOMA system in the Afghan drug trade scenario. For this example, we will assume that the user has chosen a farmer as the agent model to reason about, and the state given in Figure 4 as the state of the world. When operating the SOMA system, the user must decide how to balance the tradeoff between solution quality—the accuracy of the probability values assigned to the courses of action—and running time of the computation. As the naive algorithm, which yields exact probability results for the k most probable COAs, can take prohibitive amounts of time even for small SOMA-programs, the user can choose from among several approximation algorithms described above. The user can also choose between solving for the lower bound probability of each COA c_i ($\mathsf{low}(c_i)$), the upper bound ($\mathsf{up}(c_i)$), and the average of $\mathsf{low}(c_i)$ and $\mathsf{up}(c_i)$. Choosing to solve for the lower bound is the most conservative approach and gives the most guarantees, as we know that the course of action will be taken with at least a probability of $\mathsf{low}(c_i)$. However, this often yields a probabilty of zero for every course of action; therefore, the upper bound can give us more useful information by giving us a nonzero probability.

$s = \{is\_khan(khan1),$
$drought,$
$is\_farmer(farmer1),$
$has\_debt(farmer1),$
$land\_dispute(farmer1, warlord1),$
$is\_warlord(warlord1),$
$yields\_greater\_revenue(drug\_trade, legal\_business),$
$needs(farmer1, money),$
$increases\_quality\_of\_life(peace, farmer1),$
$is\_malik(malik1, village1),$
$solves\_conflicts\_satisfactorily(malik1),$
$helps\_in(malik1, community\_affairs),$
$respects(malik1, sharia),$
$consults(malik1, ulema1)$,$is\_ulema(ulema1),$
$lives(farmer1, village1)\}$

Figure 4: State of the world used in calculating the most probable course of action for a farmer group (farmer1) in a village (village1) containing a warlord (warlord1), malik (malik1), and ulema (ulema1) using the SOMA System

Here we will investigate the 5 most probable courses of action for a farmer using the binary heuristic algorithm and solving for the upper bound of the COAs. Figure 5 contains the resulting courses of action found by the SOMA system and their respective lower bound probabilities. According to the system, the most likely thing a farmer will do under this state of the world is that a farmer will support a local malik, arrange a salaam (or loan) with a khan in the village, join a jaba (or local militia), and be involved in the drug trade by cultivating poppies; the farmer will take these actions with a probability of about 80%. Based on the state of the world, this result reflects the power of the khans to influence the behavior of village citizens and control the economy. Even though the farmer may support the malik, which indicates a willingness to support, or at least work with the central authority, he is still becoming involved in the drug trade because of pressure from the landowners and the fact that poppy cultivation yields greater revenues than legal crops. Because the farmer is joining a jaba, this also demonstrates the power that extra-governmental forces have on the villagers, as they find it more advantageous to support local warlords or leaders than the centralized goverment.

While the results obtained by the SOMA System may be somewhat simplistic, they indicate the potential for a tool

$COA1 = \{support(farmer1, malik1),$
$arrange\_salaam(farmer1, khan1),$
$support(farmer1, peace),$
$join\_jaba(farmer1),$
$get\_involved\_in\_drug\_trade(farmer1),$
$cultivate\_poppy(farmer1)\}$
Probability: 0.8

$COA2 = \{fight(farmer1, land, warlord1),$
$support(farmer1, malik1),$
$arrange\_salaam(farmer1, khan1),$
$get\_involved\_in\_drug\_trade(farmer1)\}$
Probability: 0.25

$COA3 = \{fight(farmer1, land, warlord1),$
$support(farmer1, malik1),$
$arrange\_salaam(farmer1, khan1),$
$join\_jaba(farmer1),$
$get\_involved\_in\_drug\_trade(farmer1)\}$
Probability: 0.25

$COA4 = \{support(farmer1, malik1),$
$arrange\_salaam(farmer1, khan1),$
$join\_jaba(farmer1)\}$
Probability: 0.225

$COA5 = \{fight(farmer1, land, warlord1),$
$arrange\_salaam(farmer1, khan1),$
$support(farmer1, peace)\}$
Probability: 0.225

Figure 5: The 5 most probable courses of action found by the SOMA System for a farmer in the state of the world given in Figure 4

that will provide policy-makers with valuable insights into the behavioral patterns of a cultural group. For instance, in this example the models indicate that neither farmers, warlords, nor khans would be likely to comply with a government ban on poppy cultivation if this interfiers with their revenues or thier power in the region. Decision-makers could use results such as these as incentive to develop alternate policies that might be more effective, and test these scenarios using the SOMA system.

Conclusions

There is a growing need to be able to reason about cultural groups from around the world. The study of computational models to understand cultural behaviors is rather new (Subrahmanian *et al.* MarchApril 2007) and has a long way to go. In this paper, we have briefly described how the paradigm of SOMA rules has been applied by us to model the behaviors of various stakeholders in the Afghan drug economy.

Our modeling has split the stakeholders in the Afghan drug trade into six categories — *warlords, khans, maliks, shura members, ulema shura members,* and *farmers*. For each of these groups, we have derived a set of approximately 35–40 SOMA rules that describe the conditions under which those groups take certain actions (and the probability with which they do so).

The current system can be used to infer what might happen in a planned or hypothetical situation. By changing the existing state to reflect the planned or hypothetical situation, the SOMA rule interpreter can compute the most probable response that one or more of these group might have to a given situation. This can be a valuable aid in determining what actions to take so as to elicit the desired response(s) from a given group or groups.

Acknowledgements

Researchers funded in part by grant N6133906C0149, ARO grant DAAD190310202, AFOSR grants FA95500610405 and FA95500510298, and NSF grant 0540216.

References

Blanchard, C. M. 2005. Afghanistan: Narcotics and us policy. *CRS Report for Congress*.

Chayes, S. 2006. *The Punishment of Virtue: Inside Afghanistan After the Taliban*. New York, NY, USA: The Penguin Press.

Goodhand, J. 2005. Frontiers and wars: the opium economy in afghanistan. *Journal of Agrarian Change* 5(2):191–216.

Kakar, P. L. 2005. Fine-tuning the nsp: discussions of problems and solutions with facilitating partners. *Afgan women and the National Solidarity Programme*.

Lloyd, J. W. 1987. *Foundations of Logic Programming, Second Edition*. Springer-Verlag.

Ng, R. T., and Subrahmanian, V. S. 1991. A semantical framework for supporting subjective and conditional probabilities in deductive databases. In Furukawa, K., ed., *Proc. of the 8th Int. Conf. on Logic Programming*, 565–580. The MIT Press.

Ng, R. T., and Subrahmanian, V. S. 1992. Probabilistic logic programming. *Information and Computation* 101(2):150–201.

Scott, J. C. 1976. *The moral economy of the peasant: rebellion and subsistence in Southeast Asia*. Boston, MA, USA: Yale University Press.

Subrahmanian, V.; Albanese, M.; Martinez, V.; Nau, D.; Reforgiato, D.; Simari, G.; Sliva, A.; Udrea, O.; and Wilkenfeld, J. March/April 2007. Cara: A cultural adversarial reasoning architecture. *IEEE Intelligent Systems* 12–15.

Farmer

cultivate_poppy(X): [0.7,0.9] ← is_farmer(X) ∧ has_debt(X).
There is a high probability that a farmer will cultivate poppy if he is in debt.

cultivate_poppy(X): [0.8,1.0] ← is_farmer(X) ∧ land_opportunity(X).
There is a high probability that a farmer will cultivate poppy if this gives him a chance to have is own land.

switch_crop(X,poppy,C): [0,0.1] ← traditional_crop(C) ∧ cultivates(X,poppy) ∧ is_farmer(X) ∧ yields_greater_revenue(poppy,C).
There is a very low probability that a farmer will switch crops from poppy to a traditional one if poppy cultivation yields greater revenues.

arrange_salaam(X,K) [0.7,0.9] ← is_farmer(X) ∧ is_khan(K) ∧ banned(poppy_cultivation).
There is a high probability that a farmer will arrange a salaam (a loan) with a khan, if poppy cultivation has been banned in the region.

Khan

switch_crop(X,poppy,Z) [0,0.1] ← is_khan(X) ∧ traditional_crop(C) ∧ cultivates(X,poppy) ∧ yields_greater_revenue(poppy,C).
With a low probability a khan will switch crops from poppy to a traditional crop if poppy yields more revenue than the other crop.

enforce_poppy_ban(X): [0.5,0.65] ← is_khan(X) ∧ supports(Y,poppy_ban) ∧ declared_unislamic(Y,opium) ∧ wants_good_relations(X,Y).
With moderate probability, a khan will enforce a poppy ban if a certain group or individual supports such a ban, opium is declared unislamic, and he desires to have a good relationship with the supporter.

finance(X,Y): [0.6,0.7] ← is_khan(X) ∧ expects_rewards(X,Y).
With moderate probability, a khan will finance a certain group or individual if he expects rewards to come from this kind of support.

Ulema

support(X,Y): [0.80,0.95] ← is_ulema(X) ∧ is_warlord(Y) ∧ threatens(Y,X).
There is a high probability that an ulema will support a warlord who threatens him.

finance(X,Y): [0.6,0.7] ← is_ulema(X) ∧ expects_rewards(X,Y).
There is a moderate probability that an ulema will finance a certain group or individual if he expects rewards to come from this action.

support(X,Y) [0.75, 0.9] ← is_ulema(X) ∧ soc_or_pol_religious_regime(Y) ∧ shares_beliefs(X, Y).
There is a fairly high probability that an ulema will support a social/political/religious regime, if he shares beliefs with the regime.

Warlord

∼ban_poppy_cultivation(X) ∧ ∼enforce_poppy_ban(X): [0.9,1.0] ← is_warlord(X) ∧ is_involved_in_drug_trade(X).
There is a high probability that a warlord will neither ban the cultivation of poppy nor enforce such a ban if he is himself involved in drug trade.

buy_war_materials(X): [0.9,1] ← is_warlord(X) ∧ wants(X,seize_power).
There is a high probability that a warlord will buy war materials if he desires to seize power.

tax(X,Y) [0.7,0.85] ← is_warlord(X) ∧ area(L) ∧ dominates(X,Loc) ∧ uses_route_for_drug_trading(Y,Route) ∧ is_in(Route,Loc) ∧ trader(Y).
There is a moderate to high probability that a warlord will tax a trader who uses routes in his lands for trade drugs.

Figure 6: A sample set of SOMA rules for some of the classes of actors in the Afghan opium economy.