Trust Assessment Through Continous Behaviour Recognition

Gurleen Kaur
University of Aberdeen
Kings College
Aberdeen, UK
g.kaur@abdn.ac.uk

Timothy J. Norman
University of Aberdeen
Kings College
Aberdeen, UK
t.j.norman@abdn.ac.uk

Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA
katia@cs.cmu.edu

Abstract

Computational models of trust typically assume that an assessment of the trustworthiness of an individual can be formed from learning from the outcomes of a sequence of atomic tasks, as well as other evidence such as reports from third parties. Further, they assume that an agent's trustworthiness can be modelled by a single probability distribution. In this paper we explore alternative mechanisms that allow these assumptions to be relaxed. We propose a trust assessment model based on Markov Switching Regimes, where direct and third-party observations about an agent's behaviour follow interrelated Autoregressive processes. We argue that this offers a richer model of trustworthiness and a means to combine trust assessment with within-task monitoring.

1 Introduction

In dynamic and open systems, diverse autonomous agents interact with their peers to achieve dependent individual, or shared objectives. In such an environment, agents might behave in an untrustworthy manner, delivering unsatisfactory performance, whether this be to perform a task or to deliver an information service. When choosing future partners to rely upon, therefore, agents must consider their likely future behaviour. The uncertainties underpinning these decisions are often captured through computational models of trust. Various forms of evidence have been posited as appropriate to inform trust assessments, including past observations of behaviour given contractual expectations [cZYC09, LD12], assessments from third parties, correlations among agent's behaviour [BNS10, LDRL09], and other contextual factors.

We take a probabilistic approach to modelling trust assessment, and there is an extensive literature on mechanisms of this kind. Such models build primarily upon the Beta model [JI02, TPJL06] and its multivariate generalization, the Dirichlet model [JH07, RPC06] either implicitly or explicitly. These models assume the outcome of a delegated task/goal may be modelled as either a binary variable, representing success or failure, or a multivariate outcome. The underlying Bayesian framework of these models assumes that an agent's behaviour can be approximated by a single, static probability distribution. Based on the outcomes of the interations, the parameters of the prior distribution update over time, thus deriving the posterior distribution. The assumption that an agent's behaviour is static and represented by a single probability distribution thoughout future interations may, however, not be reasonable in many situations. A service provider may modify its behaviour over time

In: R. Cohen, R. Falcone and T. J. Norman (eds.): Proceedings of the 17th International Workshop on Trust in Agent Societies, Paris, France, 05-MAY-2014, published at http://ceur-ws.org

according to changes occurring in its environment. A sensor system, for example, may provide highly trusted target tracking data unless it believes that the target it is asked to track is from a specific organisation.

Rather than looking at an interaction between a consumer and a service provider in a macroscopic way, and considering it as a single entity with an end result of failure or success, we take a more refined view. An interaction between two agents is considered as a sequence of events/sub-tasks, and we assume that the consumer may (partially) monitor progress periodically. After observing and evaluating progress a number of times, it is possible for the agent to build a picture of the service provider's behaviour, and hence predict, to some extent, the likely future progress or result (success/failure) of the delegated task. Behaviour detection helps to provide answers to questions such as whether a delegator/consumer wants to continue with the current interaction/task allocation, or what types of task and when to delegate to this provider in the future. It could also capture any learning over time on the part of the agent that may change/improve its performance.

Consider a scenario where an agent, enters into a contract with two agents capable of tracking objects of interest within an environment. The environment is an area of coastline around a port, and the sensor agents may be unmanned aerial vehicles (UAVs) or ground-based sensor systems. Suppose that two UAVs are delegated the task of identifying, tracking and reporting the location of unauthorised boats within the area. This surveillance task may continue for a substantial period of time with target tracking data (observations) being provided by one or both UAVs during various sub-periods of the on-going task. Similar tasks, initiated by other agents, may be active at the same time. Given this contract will continue over a period of time, there may be opportunities to observe the agents' behaviour such as through correlations between observations reported by the two UAVs. This monitoring could show that the probability that one of the UAVs providing accurate reports has decreased. Options to recruit an additional or replacement sensor agent may then be considered. A traditional trust model can only learn about trust in UAVs identifying and tracking boats by looking at the number of successful unauthorised boats observed, and the numbers of failures. It can't tell you anything about the probability of continued success given some observations of the agent's behaviour while reporting.

A commonly-employed class of models applied to the analysis of time series of this kind are those based upon a Hidden Markov Model (HMM) [SS13]. An important drawback to the application of HMMs in trust assessment, however, is the maximum likelihood approach used in the parameter estimation of standard HMMs. This approach is based the 'equally likely' assumption: it assumes each observation in the training data set is of equal importance for a future prediction, no matter how big the training set. This is counter to our intuition that more recent observations should represent more weight of evidence for a trust assessment, although this approach might work well when training sets are relatively small or if the data series studied is not time-sensitive in this manner. Beta, or Dirichlet-based models of trust assessment do not suffer from this limitation, however. They tend to use the principle of exponential decay that discounts past observations, placing more weight on recent evidence.

In this paper we explore the requirements of a trust assessment model where the relationships among time series variables influence an assessment, but also where these statistical relationships are subject to change over time. We discuss a model that integrates an autoregressive process with a Markov switching model [Ham94] to exploit evidence from continuous behavioural monitoring. The autoregressive Markov switching model relaxes the standard HMM conditional independence assumption by allowing observed variables that depend on the current state to also depend on the past output/observation. In this way the autoregressive process weighs more recent observations more highly, thus explicitly modelling some of the dynamic behaviour we are interested in. Our conjecture is that this combination may lead to more accurate predictions.

Before exploring an initial trust assessment model, we formalise the underlying mechanisms that we build upon: Autoregression and Markov switching models.

2 Autoregressive and Markov Switching Models

2.1 Autoregression and Vector Autoregression

Decision makers need to be guided by predictions about how the environment is likely to change. In forcasting the values of important environmental variables, we can assume that the historical behaviour of that variable over time contains information about its future development. This history of behaviour may be records of successes/failures to meet requests, more structured outcomes, or more fine-grained samples of performance as a request is being satisfied. Given the assumption that evidence of past performance can help in estimating future behaviour, we can employ various methods to model, analyse and forecast variables of interest. An autoregressive process [Ham94] is one such tool, which we refer to as $AR(p)$ where p is the length of the window

of past values that we use to predict the next value of some variable.

DEFINITION 1
If we know the parameters of an autoregressive process $(\alpha_1, \ldots, \alpha_p)$, and if we have a sequence of p past observations of the variable of interest, y, the autoregressive equation of order p, can be used to estimate the value of y at time t.

$$y_t = c + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \cdots + \alpha_p y_{t-p} + \varepsilon_t \tag{1}$$

where c is a constant, and ε_t is white noise with zero mean and finite variance.

In the real world, however, the value of one variable not only depends on its own past values but also of the past values of other variables. In order to model these interdependencies, we can extend this *univariate* autoregressive process to model a set of time series variables simultaneously.

The vector autoregression model, $VAR(p)$, [Ham94] is an extension of univariate autoregressive process, $AR(p)$, to model dynamic multivariate time series data, and hence capture the linear interdependencies among multiple time series variables. The only prior knowledge that we require to employ this approach are the variables can affect each other over time.

DEFINITION 2
A vector autoregressive process of order p, $(VAR(p))$ can be written as

$$y_t = c + A_1 y_{t-1} + A_2 y_{t-2} + \cdots + A_p y_{t-p} + u_t \tag{2}$$

where $y_t = (y_{1t}, \cdots, y_{Kt})'$ is a $K \times 1$ vector of time series variables, $c = (c_{1t}, \cdots, c_{Kt})'$ is a $K \times 1$ vector of constants (intercepts), each A_i is a time-invariant $K \times K$ coefficient matrix, $A_i = \begin{bmatrix} a_{11,i} & \cdots & a_{1K,i} \\ \vdots & \ddots & \vdots \\ a_{K1,i} & \cdots & a_{KK,i} \end{bmatrix}$ and $u_t = (u_{1t}, \cdots, u_{Kt})'$ is a $K \times 1$ vector of error terms satisfying:

- $\mathbf{E}(u_t) = 0$, every error term has mean zero.
- $\mathbf{E}(u_t u_t') = \Sigma$, the contemporaneous variance-covariance matrix of error terms is Σ (a $K \times K$ positive-semidefinite matrix)
- $\mathbf{E}(u_t u_s') = 0$, for any $t \neq s$, there is no correlation across time; in particular, no serial correlation in individual error terms.

Rather than being interested in a single variable, y, here each y_t represents a vector of the variables that we assume to depend on each other. Our coefficient matrices, $A_1, \ldots, A_p$, model the extent to which each variable influences each other over the window of length p.

We can now apply this kind of model to explore domains in which there are multiple, interdependent variables. If, for example, we anticipate that variables y_1 and y_2 are interrelated, we can use domain data to learn the matrices c (intercepts) and $A_1, \ldots, A_p$ (coefficients). From this model, we can then use this to anticipate how variables y_1 and y_2 change together. For example, when we try to estimate the trustworthiness of a target agent, we generally use two main sources of evidence, direct experiences and third party reputational reports. Suppose that y_1 represents our direct observation of an agent's behaviour and y_2 is some aggregation of third party reports. Using this model, we can exploit the interdependencies between these two variables to predict the future values of y_1.

2.2 Switching regime models

In the vector autoregression model, $VAR(p)$, it is assumed that the parameters of the model, capturing the interrelationships among the variables of interest, are fixed; i.e. the vectors c (intercepts) and $A_1, \ldots, A_p$ (coefficient matrices) do not vary. There may, however, be periods in which the interrelationships among variables change significantly. Switching regime models are used to capture the dynamics of the observed variables of interest. For example, the dynamic behaviour of an agent may fluctuate between trustworthy and untrustworthy states, depending on its environment. The agent in question may be trustworthy when it has a medium or low load due to other commitments, but untrustworthy when it has a high load. These unobserved states (low, medium

or high load) may be reflected in the patterns of behaviour, that are observed either by direct experience or in reputational reports. Using these observations, and clustering techniques, the fluctuations in behaviour can be modelled and their underlying unobserved state/regime can be detected with some level of confidence.

Parameters of the observed series will be time varying (they will take on different values in each different, predetermined number, of regimes or states of the environment) and so fitting a linear model for each regime may be used to approximate the non linear data. The regime at any point in time is an unobserved variable but the stochastic process that determines the unobserved regime variables is known. In this work we consider a vector autoregression time series model with changes in regime and the unobserved regime variable is generated by a discrete-state homogeneous ergodic Markov chain.

Markov Switching Vector Autoregressive Model

DEFINITION 3
A M-state Markov switching, p-lag vector autoregressive process, *MS(M)-VAR(p)* [Kro00] is given by

$$y_t = \nu(s_t) + A_1(s_t)y_{t-1} + \cdots + A_p(s_t)y_{t-p} + u_t \tag{3}$$

where,

- y_t is a K-dimensional time series vector, i.e.; $y_t = (y_{1t}, \cdots y_{Kt})'$.
- M is a finite number of predetermined, feasible regimes or states.
- The unobserved regime variable at time t is s_t, and $s_t \in \{1, 2, \cdots, M\}$ follows a discrete, M-state homogeneous ergodic Markov chain.
- The K-dimensional intercept vector, $\nu(s_t)$, autoregressive parameter matrices, $A_1(s_t), \cdots, A_p(s_t)$, and the variance-covariance matrix, $\Sigma(s_t)$, vary according to the regime (state) of the environment, which is controlled by the unobserved regime variable s_t at time t.
- $u_t \sim NID(0, \Sigma(s_t))$ is a variance-covariance matrix, $\Sigma(s_t)$ of the error terms u_t, which also depends on the unobserved regime variable s_t.

The two main components of the Markov Switching Vector Autoregressive Model model are, therefore:

1. A Markov chain as the regime generating process for unobserved state s_t.
2. A Gaussian vector autoregression as the data generating process of the observed variable y_t, which is conditional on an unobserved regime s_t.

The parameters of $VAR(p)$ will, therefore, be time varying but the process is time invariant, conditional on an unobserved state s_t.

Regime Generation Process

The unobserved regime s_t in a Markov switching model is assumed to be generated by an ergodic Markov chain with a finite number of predetermined feasible states, say M, $s_t \in \{1, 2, \cdots, M\}$, which is defined by the transition probabilities p_{ij}:

$$p_{ij} = Pr(s_{t+1} = j \mid s_t = i), \sum_{j=1}^{M}, \forall i, j \in \{1, 2, \cdots, M\} \tag{4}$$

We collect all the transition probabilities between the states in the transition matrix, P.

$$P = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1M} \\ p_{21} & p_{22} & \cdots & p_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ p_{M1} & p_{M2} & \cdots & p_{MM} \end{bmatrix}$$

Using this law for the regime generating process, the evolution of the unobserved regime can be inferred from the observed time series data using clustering techniques. Let ξ_t be the vector representation of the unobserved regime variable $s_t \in \{1, 2, \cdots, M\}$ at time t. If $s_t = j$, then the unobserved regime vector ξ_t is the j^{th} column of an $M \times M$ identity matrix. The M-dimensional vector ξ_t can also be written as $\xi_t = (\mathbf{I}(s_t = 1), \cdots, \mathbf{I}(s_t = M))'$ where $\mathbf{I}$ is the indicator function.

Data Generating Process

The time series process of the observed variable y_t at time t is governed by the underlying hidden regime of the environment, ξ_t. Therefore, for a given regime ξ_t and previous values of the observed variables up to time $t-1$, $Y_{t-1} = (y'_{t-1}, y'_{t-2}, \cdots, y'_1, y'_0, y'_{-1}, \cdots, y'_{1-p})'$, the conditional probability density function of y_t is given by $p(y_t \mid \xi_t, Y_{t-1})$. In the definition of the *MS(M)-VAR(p)* process we assumed that for each regime s_t at time t, the error terms u_t are normally distributed with mean 0 and variance that depend on the regime. This implies that the conditional probability density function of y_t, given an unobserved regime ξ_t, will also be normally distributed. We collect all these Gaussian conditional densities of y_t in an M-dimensional vector, η_t.

$$\eta_t = p(y_t \mid \xi_t, Y_{t-1}) = (p(y_t \mid \xi_t = \iota_1, Y_{t-1}), p(y_t \mid \xi_t = \iota_2, Y_{t-1}), \cdots, p(y_t \mid \xi_t = \iota_M, Y_{t-1}))'$$

2.2.1 Parameter Estimation

The parameters of the *MS(M)-VAR(p)* are estimated using the Expectation Maximisation algorithm introduced by Dempster *et al.* [DLR77]. This is an iterative technique used to obtain maximum likelihood estimates of the model's parameters, where the observed time series data depends on some unobserved or hidden variable.

This two-step algorithm involves an expectation step, in which the optimal inference of the unobserved regime sequence is determined, and a maximization step, in which the parameters of the model are updated by using the maximum likelihood approach.

Expectation Step (E step)

Suppose full observation data up to time T is known and let λ be the parameter vector (to be estimated). During each iteration the unobserved states ξ_t are estimated by their smoothed probabilitites, $\widehat{\xi}_{t|T} = Pr(\xi_t \mid Y_T, \lambda^{j-1})$. These conditional probabilities are calculated using a forward recursive filter and backward recursive smoothing algorithms (see below).

The filter probability is the conditional probability of the hidden regime ξ_t given the observed sample data $Y_t = (y'_t, y'_{t-1}, \cdots, y'_{1-p})'$ up to time t, and the model parameters, $\widehat{\xi}_{t|t} = Pr(\xi_t \mid Y_t)$ [Ham94].

$$\widehat{\xi}_{t|t} = \frac{\eta_t \bigodot \widehat{\xi}_{t|t-1}}{\mathbf{1}'_M(\eta_t \bigodot \widehat{\xi}_{t|t-1})} \tag{5}$$

The forward recursive filter can be used to infer the hidden regime for time $t' \geq t$ given the observed data set up to time t. The optimal m-period forecast of ξ_{t+m} is given by $\widehat{\xi}_{t+m|t} = (P')^m \widehat{\xi}_{t|t}$, where P is the transition matrix.

Similarly, the smoothed probability is the conditional probability of the hidden regime ξ_t given the observed sample data $Y_T = (y'_T, y'_{T-1}, \cdots, y'_{1-p})'$ up to time T, and the model parameters. Smoothing is, therefore, a backward recursive process that infers unobserved states by including the sample information previously neglected in filtering [Kro00].[1]

$$\widehat{\xi}_{t|T} = Pr(\xi_t \mid Y_T) = \left(P\left[\widehat{\xi}_{t+1|T} \oslash \widehat{\xi}_{t+1|t}\right]\right) \odot \widehat{\xi}_{t|t} \tag{6}$$

Maximization Step (M step)

In the E step, the parameter vector, λ, was taken to fixed and known. Within the M step we compute the maximum likelihood estimate for our model parameters. The parameter vector λ contains *VAR* parameters (i.e. intercept, autoregressive matrices and error variance) and the initial and transition probabilities of the underlying hidden Markov chain.

The log likelihood function is given by $L(\lambda \mid Y, \xi) := p(Y_T \mid \lambda, \xi)$. To maximise the value of this function, the latent/hidden variable will be substituted by its expectated value, $\widehat{\xi}_{t|T}$. This means that the conditional regime probability $Pr(\xi_t \mid Y_T, \lambda)$ will be replaced by smoothed probabilities, calculated in the previous expectation step, thus eliminating non-linearities. The parameters for this function are derived from solving the first-order conditions of a constrained log likelihood function (see Krolzig [Kro00] for detailed analytical solution).

[1] Functions $\oslash$ and $\odot$ are element-wise division and multiplication respectively.

2.2.2 Forecasting

Despite being a non-linear model, the attractive feature of Markov switching vector autoregression, is its simplicity of forecasting. To obtain the optimal h-step forecast, the mean squared prediction error (MSPE) criterion may be used (i.e. we minimise the squares of the forecast errors).

$$\widehat{y}_{t+h|t} := \arg\min_{\widehat{y}} \mathbf{E}\left[(y_{t+h} - \widehat{y})^2 \mid Y_t\right]$$

Given the information Y_t up to time t, therefore, the optimal h-step forecast of the observed time series is given by the conditional mean:

$$\widehat{y}_{t+h|t} = \mathbf{E}\left[y_{t+h} \mid Y_t\right]$$

Since the data generating process is nonlinear, the MSPE optimal forecast is a not a linear predictor of the observed temporal data (see Krolzig [Kro00] for further details).

3 A simple Markov switching trust assessment model

We may now demonstrate how this general model, *MS(M)-VAR(p)*, may be applied to trust assessment. Consider a system with n agents, where $\boldsymbol{A} = \{a_1, a_2, \ldots, a_n\}$ is the set of all agents. Agents interact with one another and work together to acomplish various tasks. Direct experiences from these transactions can aid both parties, say a_i and a_j, in forming opinions about the trustwortiness of each other. We consider these experiences to be *both* the results of monitoring actions during a transaction and the transaction outcomes.

The rating that an agent gives to an element of an interaction (assessed through monitoring) may belong to a discrete set of values or from a continous range, such as $[0, 1]$. Over time, interactions between agents produce a history of direct evaluations and ratings, thus forming a time series of observations of that variable.

We may discretise the series of observations made by agent a_i of agent a_j such that for each time period $[t-1, t]$, a_i may make a direct trust evaluation of a_j. A simple method would be to compute the average of the outcomes of observations made during that time period, but other aggregation methods are possible. We refer to the direct observation made by agent a_i of agent a_j at time t as $Y_{ij}(t)$.

Evaluating the trustworthiness of an agent is time-varying process; more recent behaviour should have greater influence on a trust assessment. An autoregressive process offers a means to model this. The output variable of this process at time t depends on its own previous values, thus capturing any positive or negative effect of the observed data. The dynamic, unobserved behaviour of a service provider may also change with time; changes that may be manifest from observations acquired through monitoring. If positive observations are made then the agent is more likely to be behaving in a trustworthy manner and vise versa. These observations may, however, change dramatically away from long-run mean. This volatility could indicate a change in the regime of the observed temporal process. The series is, therefore, assumed to be dependent on an unobserved stochastic process. This unobserved process models the actual state of the agent's behaviour and the evaluations of their trustworthiness at time t are the observed data determined by this hidden variable at that time.

If two agents have interacted with each other within the society, then they will have a temporal data set reflecting their observations of the encounter. Based on these direct experiences, an agent may build a model to predict likely future behaviour.

Other forms of evidence may be exploited during trust assessment. Opinions, derived from behavioural observations, about the target of a trust assessment may be acquired from third parties. Taking into consideration such indirect evaluation about a service provider's behaviour may improve the accuracy of a trust assessment. Although useful evidence, the use of third party reports is not without its risks. It is possible that a recommender can provide misleading or biased feedback about other agents in the society unintentionally or otherwise. These reputation reports may undervalue the true behaviour of a peer or represent an unjustifiably positive opinion.

It is notoriously difficult to detect misleading recomendation reports, but failing to consider all feedback from peers has its own risks. Ideally we would want to weigh the recomendations received from different agents to mitigate the influence of biases and misleading reports, if not eliminate their effects entirely. We now discuss a simple means to take into account such evidence within an *MS(M)-VAR(p)* process.

At time t, if the agent a_i wants to assess the behaviour of another agent a_j then it may seek opinions about a_j from all the other agents in the system. An aggregation of this feedback may then be integrated with its own

view of the target agent, a_j. Let $R_{ij}(t)$, denote agent a_i's estimate of the community opinion of agent a_j at time t obtained by aggregating the recomendation reports from its peers in the environment. A simple method of obtaining $R_{ij}(t)$ is to compute the weighted average of all the reputation reports received.

$$R_{ij}(t) = \frac{1}{\sum\limits_{\substack{a_x \in \boldsymbol{A} \\ a_x \neq (a_i, a_j)}} W_{ix}(t)} \sum_{\substack{a_x \in \boldsymbol{A} \\ a_x \neq (a_i, a_j)}} W_{ix}(t) Y_{xj}(t) \tag{7}$$

Here, $W_{ix}(t)$ is the weight given by agent a_i to agent a_x's reputation report and $Y_{xj}(t)$ is the reported behavioural observation of agent a_j by a_x at time t. These aggregated reports may then be exploited as a second variable within the *MS(M)-VAR(p)* process that the predicted trustworthiness of the target agent depends upon.

There are, of course, other means to aggregate third-party opinions. The use of stereotypes [BNS10, LDRL09], for example, may obviate the need to maintain weights for the opinions of other agents. Stereotypes can be used to weigh reports from sources' opinions based on the group to which they belong. Alternatively, we could model each stereotypical group as a variable within the *MS(M)-VAR(p)* process, each of which influencing the variable representing the trustworthiness of the target agent.

The *MS(M)-VAR(p)* process relies on a series of data points within the window of length p to forecast the variable of interest, which, in our case, is the trustworthiness of the target agent. In trust assessment, this is a challenge to the application of the model because there may be significant gaps in direct interaction between agents. Although it is less likely that there are no third-party opinions to exploit, the observed time series data from direct interactions will have missing values. To deal with this challenge, various interpolation techniques may be employed such as using regression or spines [HK10].

Suppose an iteration between two agents in the society ends at time t and starts again at time $t + 4$; we are missing observations for 3 time steps. We first try to predict those missing values by using the data set up to time t, estimate the model parameters, and then forecast the missing observations. Missing data is then replaced with predicted data. We can then use this full time series of direct interactions along with the reputation reports for trust assessment.

4 Illustration

To demonstrate the approach we propose, we simulated a multi-agent system in which reputation reports are exchanged, interactions occur (over periods of time), and agents assess the trustworthiness of potential partners. We simulate ten agents with different behaviour profiles. We investigated the process whereby one agent, a_1, attempts to evaluate the trustworthiness of another, a_2. We simulated relatively long-term interactions happening frequently between agents, to minimise the need for data imputation/interpolation. These transactions between a_1 and a_2 produce time series of direct observations. In parallel a_1 interacts with other agents in the society, producing other observations over time of their behaviour.

At each time step, agent a_1 will query other agents in the society for observations regarding a_2. Agent a_1 aggregates these third party reports using the simple weighted average method described above. This provides a time series of aggregated reputation reports for a_2. Using these two time seires data we will try to evaluate the switches in the unobserved state of agent a_2's behaviour.

For the sake of simplicity, we assume there are two possible hidden states: trustworthy and untrustworthy. The purpose is, therefore, to predict this unobserved state.

To select the vector autoregression lag length, we use the AIC criteria. The table below shows that, in this simple illustration, *VAR(1)* had the smallest AIC value. Given we consider two variables (direct observations, *DO*, and agregated third party reports, *RR*), we use a *MS(2)-VAR(1)* process; i.e. a 2-variable Markov switching autoregressive model with a lag of 1.

AIC and BIC values for VAR		
VAR Lag	AIC Value	BIC Value
1	-2.314199	-2.263234
2	-2.298469	-2.213529
3	-2.289073	-2.170156
4	-2.282472	-2.129579
5	-2.272324	-2.085454

We collect the values acquired for our variables of interest (DO and RR) in a vector $y_t = (DO_t, RR_t)$. The vector autoregression with lag 1 will create two VAR equations, where the temporal variable DO changes and its future development depends on its own past outcomes and on the past outcomes of the RR variable. Thus the marginal change in RR variable will effect DO, hence affecting the unobserved/hidden state of the agent's level of trust.

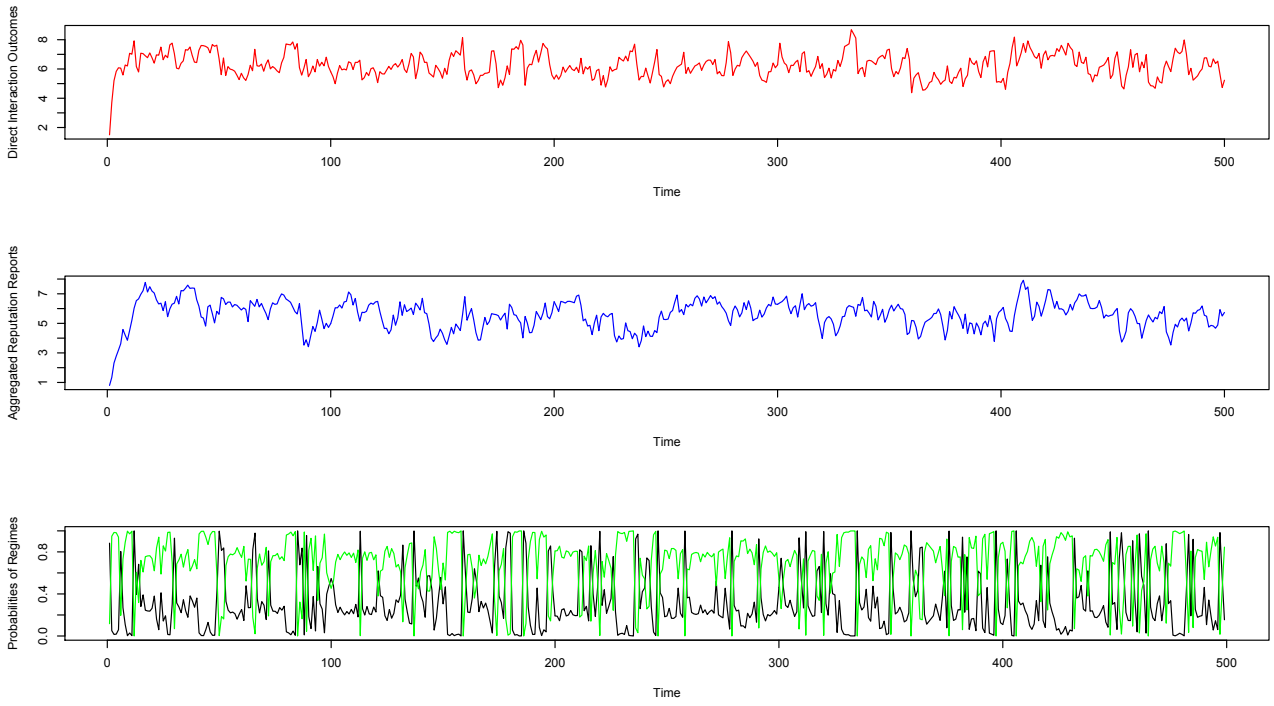

Figure 1: Regime Probabilities

In the above graph, the first time series represents the direct observations of agent a_1 regarding a_2. The second series is the time series of aggregated reputational reports about a_2 received by agent a_1 from the other agents in the society. These reports can be biased, and so we marginalise each report according to agent a_1's level of trust in the report provider.

The third graph shows the switches between the unobserved states of the agent's trust and how these unobserved regimes evolve over time based on the observed data. The black curve represents the probability that service provider a_2 is in the trustworthy state and the green curve represents the probability that a_2 is in the untrustworthy state. It can be seen from graphs that there seems to be a high correlation between direct interaction outcomes (the DO variable) and these hidden states. If the direct interaction ratings are high then the probability of being in a trustworthy state is higher. At the same time, the second series (aggregated reputational reports) also has the ability to pull down or push up the probability of being in a certain unobserved state.

5 Discussion

The most prominent existing research on probabilistic models of trust are grounded upon the Beta reputation models or its multivariate extentions (for interatctions with multiple outcomes). They use either Beta or Dirichlet distributions to represent the probability distribution over interaction outcomes [JI02] [JH07] [MMH02]. These models have also been extended to deal with deception and unfair raitings [WJI05] [TPJL06] [RPC06]. Generally, these models assume that agents' behaviour is static; i.e. represented by a fixed probability distribution. The limitations of this assumption are mitigated by treating recent interaction outcomes as more representative of the likely future behaviour of an agent; e.g. the use of *exponential decay* or *forgetting factor* in Jøsang & Haller [JI02].

Recently, a number of trust assessment models based on Hidden Markov Models (HMM) have been proposed. El Salamouny & Sassone [SS13], for example, propose an HMM-based model of evaluating trust that exploits

direct experiences and reputational reports from an agent's peers. Moe *et al.* [MTK08] propose a trust model that combines an HMM and with reinforcement learning. After learning about the environment from its RL module, the parameters of the HMM module are restimated to detect an agent's behaviour more reliably. Boer *et al.* [SKN07] propose a computational trust model based on HMMs, comparing this with existing probabilistic computational trust models, demonstrating that the non-HMM-based models were unable to deal with dynamic behaviour. A similar study by Moe *et al.* [MHK09] provides results of a comparison of the effectiveness of Beta models with decay factors and an HMM based trust approach; the conclusion being that the later was more realistic and effective in dynamic environments. These HMM-based trust models focus on the interaction history without considering the context of the interaction. Liu & Datta [LD12], however, consider an HMM-based context aware trust model to predict an agent's trustworthiness is dynamic environments. Empirical assessment of this model, which uses multiple discriminent analysis to select appropriate features of the context, demonstrates that it out-performs standard HMM-based models in detecting dynamic behaviour patterns.

We have presented an early exploration of the use of autoregression and Markov switching methods in trust assessment. There are a number of simplifications in how we have applied these techniques to the trust assessment problem such as in the aggregation of third party observations. Existing models propose clever fusion techniques to aggregate these reports such as those used in TRAVOS [TPJL06], or evaluate reports separately and aggregate the results [SS13]. Here, although we use a simple weighted average approach to combine reports, the use of *VAR* enables us to simultaneously model this time series with that from direct experience and model the effects that these temporal variables have on each other. In addition to exploring refinements of our model, we need to thoroughly investigate the accuracy of forecasting future dynamic behaviour of a target agent.

6 Conclusion

We have proposed a novel approach to the development of computational models of trust grounded upon a Markov Switching Regime model (equivalent to an HMM) where the observed data follows an Autoregressive process. The means by which we generate the data that drives the Markov switching model relaxes the assumption used in all HMM-based models of trust: that each observation is of equal importance to the assessment of trust. By considering that the observed data follows an autoregressive process, we place more weight on more recent evidence. The use of a Markov switching model enables us to model non-linear behaviour of a target agent by constructing a vector of linear models of behaviour, given (ideally) distinct behavioural states or regimes, along with a model of how the agents behaviour switches between these states/regimes. This means we are not relying on the assumption made by most non-HMM-based models of trust: that the behaviour of each agent can be modelled by a single, static probability distribution. Further, we do not need to treat interactions (or transactions) between agents as atomic, and use final outcomes as evidence for future assessments. We can exploit the techniques we propose to monitor progress of longer-term delegated tasks to inform interim decisions regarding the dependency between agents.

Acknowledgements

This research was sponsored by Selex ES.

References

[BNS10] C. Burnett, T. J. Norman, and K. Sycara. Bootstrapping trust evaluations through stereotypes. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 241–248, 2010.

[cZYC09] M. Şensoy, J. Zhang, P. Yolum, and R. Cohen. POYRAZ: Context-aware service selection under deception. *Computational Intelligence*, 25(4):335–366, 2009.

[DLR77] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.

[Ham94] J. D. Hamilton. *Time-series analysis*. Princeton Univerity Press, 1994.

[HK10] J. Honaker and G. King. What to do about missing values in time series cross-section data. *American Journal of Political Science*, 54:561–581, 2010.

[JH07] A. Jøsang and J. Haller. Dirichlet reputation systems. In *Proceedings of the International Conference on Availability, Reliability and Security*, pages 112–119, 2007.

[JI02] A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.

[Kro00] H.-M. Krolzig. Predicting markov-switching vector autoregressive processes. Economics Series Working Papers 2000-W31, University of Oxford, Department of Economics, 2000.

[LD12] X. Liu and A. Datta. Modeling context aware dynamic trust using hidden Markov model. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1938–1944, 2012.

[LDRL09] X. Liu, A. Datta, K. Rzadca, and E-P. Lim. StereoTrust: A group based personalized trust model. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, pages 7–16, 2009.

[MHK09] M. E. G. Moe, B. E. Helvik, and S. J. Knapskog. Comparison of the Beta and the hidden Markov models of trust in dynamic environments. In E. Ferrari, N. Li, E. Bertino, and Y. Karabulut, editors, *Trust Management III*, volume 300 of *IFIP Advances in Information and Communication Technology*, pages 283–297. Springer, 2009.

[MMH02] L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, pages 2431–2439, 2002.

[MTK08] M. E. G. Moe, M. Tavakolifard, and S. J. Knapskog. Learning trust in dynamic multiagent environments using HMMs. In *Proceedings of the 13th Nordic Workshop on Secure IT Systems*, 2008.

[RPC06] K. Regan, P. Poupart, and R. Cohen. Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 1206–1212, 2006.

[SKN07] V. Sassone, K. Krukow, and M. Nielsen. Towards a formal framework for computational trust. In F. S. Boer, M. M. Bonsangue, S. Graf, and W.-P. Roever, editors, *Formal Methods for Components and Objects*, volume 4709 of *Lecture Notes in Computer Science*, pages 175–184. Springer, 2007.

[SS13] E. El Salamouny and V. Sassone. An HMM-based reputation model. In A. Awad, A. Hassanien, and K. Baba, editors, *Advances in Security of Information and Communication Networks*, volume 381 of *Communications in Computer and Information Science*, pages 111–121. Springer, 2013.

[TPJL06] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Journal of Autonomous Agents and Multi-Agent Systems*, 12:183–198, 2006.

[WJI05] A. Whitby, A. Jøsang, and J. Indulska. Filtering out unfair ratings in Bayesian reputation systems. *Icfain Journal of Management Research*, 4(2):48–64, 2005.