

A Bayesian Hierarchical Model for Criminal Investigations

F. O. Bunnin^{*,†} and J. Q. Smith^{‡,§}

Abstract. Potential violent criminals will often need to go through a sequence of preparatory steps before they can execute their plans. During this escalation process police have the opportunity to evaluate the threat posed by such people through what they know, observe and learn from intelligence reports about their activities. In this paper we customise a three-level Bayesian hierarchical model to describe this process. This is able to propagate both routine and unexpected evidence in real time. We discuss how to set up such a model so that it calibrates to domain expert judgments. The model illustrations include a hypothetical example based on a potential vehicle based terrorist attack.

Keywords: Chain Event Graphs, hierarchical models, Markov processes, Markov switching models, probabilistic graphical models, decision support systems.

1 Introduction

How to better support police to prevent terrorist attacks continues to be a major political concern due to continued violence perpetrated by extremists Europol (2018); Allen and Dempsey (2018). In contrast to the majority of terrorist incidents in the latter half of the twentieth century which were executed by known organised terrorist groups with substantial planning and sophistication, more recent attacks have often involved individuals or small groups targeting civilians in public places using basic equipment such as vehicles, guns and knives Europol (2018); Lindekilde et al. (2019). Consequentially this entails less sophistication in materials, planning and execution. In terms of analysing how to understand and prevent terrorism, criminologist focus has shifted from “individual qualities (who we think terrorists ‘are’) to . . . what lone-actor terrorists do in the commission of a terrorist attack and how they do it” Gill (2012). Gill, referencing Horgan (2005), notes “it is useful to view each terrorist offence as comprising of a series of stages”.

The case studies of lone-actor terrorists have been analysed extensively both qualitatively and quantitatively for insight into background and preparatory behaviours, vulnerability indicators, radicalisation patterns, and modes of attack planning Bouhana

*The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, OBunnin@turing.ac.uk

†Data Science Institute, Imperial College London, South Kensington Campus, SW7 2AZ, FBunnin@imperial.sc.uk

‡The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, JSmith@turing.ac.uk

§Department of Statistics, University of Warwick, Coventry, CV4 7AL, United Kingdom, J.Q.Smith@warwick.ac.uk

and Wikstrom (2011); Corner et al. (2019); Lindekilde et al. (2019); Bouhana et al. (2016). These studies emphasize that the small number and heterogeneity of cases make rigorous scientific examination of associative and causal relationships extremely difficult. As they indicate it is vital, therefore, to utilise structure from existing domain expertise on the relationships between observable data, preparatory activities, and attack modes in any probabilistic analysis of the progression of an individual to an attack.

Probabilistic models, including Bayesian graphical models, have been used for modelling “comprehension and decision making of law enforcement personnel with respect to terrorism-centric behaviours” Regens et al. (2015), in a terrorist cell actor-event network analysis Ranciati et al. (2017), for “rapid detection of bio-terrorist attacks” Fienberg and Shmueli (2005), for spatio-temporal terrorism analyses Clark and Dixon (2019); Python et al. (2019), and in a “systems analysis approach to setting priorities among countermeasures” against terrorist threat Pat-Cornell and Guikema (2002). Bartolucci et al. (2007) apply a multivariate Latent Markov model to the analysis of criminal trajectories: their focus is on identifying the model structure given longitudinal data on individuals’ criminal convictions and discrete covariates such as gender and age band; the latent states are an individual’s “tendency to commit” certain types of crime.

It is within this context that we present a new class of Bayesian models to dynamically infer the progression of an individual through discrete stages towards a criminal attack. These models have been developed through close discussions over several years with a number of different policing agencies. To our knowledge this approach is novel and complements the existing research.

Overview of the model

A suspect within a subpopulation of interest to the police, $\omega \in \Omega$, is believed to be planning a serious criminal attack against the general public. Typically ω will need to step through various stages of preparation before perpetrating this crime. During this progression police will have the opportunity to observe and evaluate ω ’s status through their record, updated throughout an investigation by sporadic intelligence reports and routine observations of ω ’s activities. A dynamic Bayesian model is uniquely placed to provide decision support for such policing activities. It provides a framework within which to encode criminological theories, domain knowledge available about ω , for example his police record and personal modus operandi, and also draw in evidence from noisy streaming data about ω observed by police. All these features are integrated into a single dynamic probability model. The model we build in this paper tracks the probability ω lies in certain states or makes a transition from one state into another at any given time. These probabilities help to guide interventions and resource allocations.

To be operational such a Bayesian model must be constructed so that current prior information in a given suspect ω ’s record can be quickly updated not only in the light of routine surveillance but also unexpected sources. So, for example, police may well be monitoring the phone log of someone suspected of a serious crime. But within an investigation direct information sporadically comes to light about what ω is doing – unexpected sightings, overheard statements of intent, and so on. It would be unreasonable to assume that this type of information, often critical to a correct appraisal of ω ’s

status, could have been forecast and accommodated into any prior model specification. Any methodology we design for this domain therefore needs to be *open* to manual intervention West and Harrison (1997). Police will then be able to input these unpredicted new sources of information into the system and so improve the probability assessments of the Bayesian model. A three level hierarchy facilitates this openness property.

At the **deepest level** lies a Reduced Dynamic Chain Event Graph (RDCEG). This is a graphically based model drawn from a particular subclass of finite state semi-Markov processes customised to model transition processes in a subpopulation of the general public Shenvi and Smith (2018). This deepest level provides a framework for expressing the probability judgements of police concerning ω 's current threat status.

The **intermediate level** of our hierarchy concerns intelligence police might acquire concerning ω 's enacted intentions. When a suspect is at a particular stage of a criminal pathway, in order to engage in that step of criminality or alternatively to progress to the next step, a set of associated tasks needs to be completed. Intermittent intelligence reports often inform these. Because these are explicit components of the hierarchical model propagating such information corresponds to simple conditioning. A vector of tasks whose components form a signature of various states of criminal intent and capability constitutes the variables that lie on the intermediate layer of the hierarchy.

The **surface layer** of the model then links these tasks to the intensities of certain activities that can be routinely observed by the police if they have the necessary resource and permissions. In the absence of direct information about ω 's engagement in tasks, signals from predesigned filters provide vital information about what ω might be doing. For example suppose the task concerns ω 's intent to travel to a region to learn how to bomb. Then a filter that measures the intensity of the suspect's engagement in searching airline websites would give a noisy signal of his booking a flight. Such information is imperfect: ω may book a flight directly from an airport or to have chosen not to fly to the destination. And of course a high intensity in such activity could be entirely innocent: ω may be booking a vacation, for example. Such measures are nevertheless obviously informative. Appropriately chosen filters of these data streams provide the surface level of our Bayesian hierarchy. The usual Bayesian apparatus then provides a formal and justifiable framework around which police can logically and defensibly propagate information about ω .

Formally describing states by collections of tasks within generic Bayesian models supporting criminal investigations is, to our knowledge, novel. However it is interesting that Ferrara et al. (2016) proposed a similar approach albeit less formally expressed and in a more restricted domain: the discovery of recruiters to radicalisation to extreme violence from Twitter communications. By performing a number of thought experiments with domain experts these authors successfully extracted a collection of tasks that a recruiter would need to engage in to be effective. Although an innocent non-recruiter, such as an academic or journalist, might happen to engage in *some* of the tasks in this collection they would be unlikely to engage in *all* of these tasks simultaneously. The authors then related this vector to various easily extracted meta data signals that could be routinely extracted from an enormous dataset. This provided an analogue of the types of filter of a routinely applied observation vector we discuss later here.

In the next section we describe the RDCEG and demonstrate through some simple examples how it can be used to translate domain experts’ judgements into a latent probability model at the deepest level of a hierarchy. We also illustrate sets of tasks that ω lying in a particular state or transitioning between states might entail. Our core methodology is described in Section 3. We propose a collection of assumptions, elicited from domain experts, about the ways criminal progressions associate to tasks, and how the intensities synthesise various sources of routine measurements of engagement in each of these tasks given background circumstances. Given these assumptions we are able to propagate not only routine indirect but also unexpected direct information about ω ’s current activities to obtain posterior probabilities about ω ’s current position.

The resulting propagation algorithms are straightforward to enact. However the inputs of the model: both the structural prior information and the prior parameter distributions embellishing them need to be carefully specified if the methodology is going to be operationalised. In Section 4 we outline how we do this.

In this setting to explore the methods using known data on given suspects as illustrations is clearly unethical. However it is still possible to demonstrate how the system works in various hypothetical situations whose distributions are informed by publically available data and elicited judgements. Therefore in Section 5 we illustrate the way the system is able to update the state probabilities of an individual under suspicion of a potential attack. We describe the different task sets we have used and the construction of routine filters and test these against a two scenarios. In the concluding section we discuss how we are now extending the methodology to model threatening subpopulations of the general public where estimation and model selection algorithms can also be built to better understand the developing processes.

2 The RDCEG for criminal escalation

2.1 Introduction

Chain event graphs (CEGs) are now an established tool for modelling discrete processes where there is significant asymmetry in the underlying development see e.g. Barclay et al. (2013, 2014); Collazo and Smith (2016); Cowell et al. (2014); G3rgen and Smith (2018); Collazo et al. (2018). Dynamic versions of these processes, using analogous semantics, first appeared in Barclay et al. (2015). However formal extensions of these classes to model open populations have only recently been discovered Smith and Shenvi (2018) and developed Shenvi and Smith (2018); Shenvi and Smith (2019). We briefly review and illustrate the main properties of this class as they apply to the hierarchical model developed here. We refer the reader to the references above for more details.

The RDCEG we use in this paper is a particular family of semi-Markov process that can be expressed by a single graph. Each represented state is called a *position*. In our domain there is an absorbing state – called the *neutral* state that ω enters when presenting no future threat of perpetrating the given crime. The practical challenge is to find a way to systematically construct the set of positions so that the embedded Markov assumptions are faithful to expert judgements. In Barclay et al. (2015) and Shenvi and

Smith (2018) we describe how this can be done. We take a natural language description from domain experts and re-express this as a potentially infinite tree. We then translate this tree into an equivalent graph \mathcal{C} . For the purposes above we will henceforth assume this to have a finite number of vertices.

A position w is connected by a directed edge into another w' in the graph of an RDCEG iff there is a positive probability that the next transition from w will be into w' . Typically although the transition probabilities are fairly stable, the time it takes to make a transition is not. Therefore we need to express this expert judgement as a semi-Markov rather than Markov process. The graph \mathcal{C} is one that on the one hand is often found to be transparent and natural to our users but on the other has a formal Bayesian interpretation. So this elicited graph provides a vehicle to move seamlessly from an expert elicitation into a more formal family of stochastic processes.

The RDCEG developed in Shenvi and Smith (2019) was designed to be applied to public health processes where \mathcal{C} could often be observed directly. For criminal processes this is not usually possible. Therefore for crime modelling an RDCEG process typically remains latent and any prior to posterior analysis of the suspect's positions needs a little more sophistication. The hierarchical structure we define in the next section provides the framework for this update. We give some simplified illustrations below of such RDCEGs.

2.2 A criminal RDCEG and its tasks

We describe an RDCEG, Figure 1, for a politically motivated murder plot illustrating the relationship between its states and tasks: the lowest and intermediate levels in our hierarchical model.

Example 2.1. Electronic posts directly observed by the police suggest woman S is plotting to kill a certain political figure by shooting them. At any time S could lie in a number of positions. In positions w_3 and w_4 she is trained to shoot (T) in w_1 and w_2 not (T^c) and will own a gun (G) – position w_2 and w_4 or not (G^c) when in positions w_1 and w_3 . Each edge and each vertex in the RDCEG \mathcal{C}_1 below can be associated to




Figure 1: The RDCEG \mathcal{C}_1 for a murder plot.

States	Description / State Tasks	Edges	Edge Tasks
$w_0 = N$	plot ends		
$w_1 = T^c, G^c$	can't shoot, no gun	$w_1 \rightarrow w_2$	acquire gun θ_1
		$w_1 \rightarrow w_3$	train to shoot θ_2
$w_2 = T^c, G$	can't shoot, has gun	$w_2 \rightarrow w_4$	train to shoot θ_2
		$w_2 \rightarrow w_1$	lose gun θ_3
$w_3 = T, G^c$	trained to shoot, no gun	$w_3 \rightarrow w_4$	acquire gun θ_1
$w_4 = T, G$	trained to shoot, has gun	$w_4 \rightarrow w_5$	locate and approach target θ_4, θ_5
		$w_4 \rightarrow w_3$	lose gun θ_3
$w_5 = O$	attempt murder θ_6	$w_5 \rightarrow w_4$	fail and escape θ_7

Table 1: States, edges and tasks for example 2.1.

different collections of tasks. For example the vertex $w_5 = O$ is associated with the task “attempt murder”; the edges $w_1 \rightarrow w_2$ and $w_3 \rightarrow w_4$ are associated with the task “acquire gun”. If she cannot shoot she could next choose to learn how next: from a state where she currently owns a gun (T^c, G) or not (T^c, G^c). Alternatively if she currently has no gun then she could next try to acquire one, either when trained to shoot or not. At any point in this process she may enter the neutral state $w_0 = N$: for example the target may die through other natural or unnatural circumstances, S may change her intention, she may be arrested the police having gained enough evidence to charge her. Note that only once she has a gun and can shoot – state (T, G) – can she attempt the murder O by locating and then approaching the target. Implicitly as there is no state “commit murder”, once in the “attempt murder” state w_5 she either enters the neutral state w_0 or re-enters the “trained to shoot, has gun” state w_4 : entering w_4 implies she has failed that particular attempt and may try again; entering w_0 implies that either she failed and cannot make any further attempts or she succeeded and poses no threat to any other individual. The relevant RDCEG \mathcal{C}_1 and a table describing the positions, edges, and tasks are given in Figure 1 and Table 1.

The RDCEG \mathcal{C}_1 does not have any positions such that there exists two or more edges between them and thus it is simply a subgraph of the state transition graph of a semi-Markov process defining the dynamic where the absorbing state N and all edges into it are removed. All states in the process other than w_0 appear as vertices. The absorbing state w_0 is not depicted for three reasons:

- By definition an RDCEG contains the absorbing or “drop-out” state with edges from any position leading to it so depicting would be informationally redundant
- Not depicting it and the multiple edges into it reduces visual clutter on the graph
- Also by definition once the individual enters w_0 they are no longer of interest: hence we focus attention on the active positions by eliding it from the visual depiction.

Tasks can be associated to one or more edges or states: here “acquire gun” is associated with the edges $w_1 \rightarrow w_2$ and $w_3 \rightarrow w_4$. The structure of the transition matrix, M_1

with states w_0, w_1, \dots, w_5 of a semi-Markov process with a configuration of zeros given below. The starred entries represent probabilities that need to be added to complete the matrix

$$\begin{array}{c} w_0 \\ w_1 \\ w_2 \\ w_3 \\ w_4 \\ w_5 \end{array} \begin{pmatrix} w_0 & w_1 & w_2 & w_3 & w_4 & w_5 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ * & 0 & * & * & 0 & 0 \\ * & * & 0 & 0 & * & 0 \\ * & 0 & 0 & 0 & * & 0 \\ * & 0 & 0 & * & 0 & * \\ * & 0 & 0 & 0 & * & 0 \end{pmatrix}.$$

Note this RDCEG has translated the verbal police description above into a semi-Markov process. It can be elaborated into a full semi-Markov model by eliciting or estimating the probabilities in M_1 and the holding times. In the above, because of the sum to one condition, we have eight functionally independent transition probabilities. To complete the specification of this stochastic process it is necessary to define the holding time distributions associated with the active states i.e. how long we believe the suspect will stay in their current state before transitioning into another. These probabilities and the parameters of the holding time distributions may well themselves be uncertain. However of course within a Bayesian analysis their distributions can be elicited or estimated in standard ways O'Hagan et al. (2006).

The implicit Markov hypotheses of a given RDCEG are of course typically substantive. One critical issue is that its positions/states define the *only* aspects of the history of the suspect that are asserted relevant to predicting her or his future acts. The art of the modeler is to elicit positions in such a way that these Markov assumptions are faithful to the expert judgements being expressed. However because the methodology is fully Bayesian, experts can be interrogated as to the integrity of these assumptions just as they can be for other graphical models Smith (2010). In this way we can iterate towards a model which is requisite Phillips (1984). The process of query, critique, elaboration and adjustment is precisely why this bespoke graphical representation is such a powerful tool. In particular a faithful structural model of the domain information can be discovered Smith (2010) *before* numerical probabilities are elicited or estimated: see e.g. Wilkerson and Smith (2018); Collazo et al. (2018).

2.3 From tasks to routinely observed behaviour

Sometimes solid police intelligence, for example from an informant, will confirm that ω is engaged in a particular task. But at other times only echoes of a task will be seen by police. Suppose ω is suspected of being at a stage where they need to accomplish the task of selecting a target location for a bombing or vehicle attack. They might be observed travelling to what the police assess might be a potential target to check timings, the density of people at the venue and its defences activities. This visitation by ω might have been recorded on closed-circuit television (CCTV). In addition or instead, ω might inspect Google maps of the attack area and the route to it or contact

like-minded collaborators by phone or electronic media for advice. So indirect evidence ω engages in such a task can come from a variety of media and platforms.

Even such incomplete and disguised signals can be usefully filtered from complex incoming data about the suspect, albeit with considerable associated uncertainty. The further a suspect currently is from the main focus of an investigation the more indirect information will be. However even then the composition of a collection of weak signals the police are allowed to see may still provide enough information to significantly revise the evaluation of the threat posed by a particular individual.

The hierarchical structure we describe below enables us to draw together all these different types of evidence. If direct information about the tasks a suspect is engaging in is available then, because such tasks are explicitly represented within our model, we can simply condition on this information and so refine our judgements. The Bayesian hierarchical model simply discards the weaker indirect information to focus on what is known. Otherwise the model uses filters of the indirect signals police can see to infer what tasks ω might be engaging in to help inform police of ω 's position.

3 The structure of the hierarchical task model

3.1 Introduction

Henceforth assume that the RDCEG \mathcal{C} correctly specifies the underlying process concerning ω . To build the propagation algorithms we first define our notation.

Let W_t be the random variable taking as possible values the states $\{w_0, w_1, w_2, \dots, w_m\}$ of $\omega \in \Omega$ – the subpopulation of interest – at time $t > 0$ where $\{w_1, w_2, \dots, w_m\}$ are the vertices/positions/active states of \mathcal{C} and w_0 the inactive/neutral state. At any time t , ω might enact one or more of R tasks associated to one or more of the positions $\{w_1, w_2, \dots, w_m\}$ or alternatively to a transition from one position w_- into another w_+ . So let

$$\{\boldsymbol{\theta}_t = (\theta_{t1}, \theta_{t2}, \dots, \theta_{tR}) : t \leq T\}$$

denote the *task vector* $\boldsymbol{\theta}_t$: a vector of binary random variables where $\theta_{ti} = 1$, $i \in \{1, 2, \dots, R\}$, indicates that ω is enacting task i at time t .

Let χ_I denote an indicator on a subset $I \subseteq \{1, 2, \dots, R\}$. Then the tasks a suspect $\omega \in \Omega$ engages in at a given time t can be represented by events of the form $\{\boldsymbol{\theta}_t = \chi_I : 0 \leq t \leq T\}$.

Direct *positive* evidence that ω lies at position w_i is provided by tasks whose indices lie in $I = I(w_i)$ and when ω is transitioning along the edge $e(w_-, w_i)$ by tasks whose indices lie in $I = I(w_-, w_i)$, where $I(w_i)$ are the indices of tasks associated with the state w_i and $I(w_-, w_i)$ are the indices of tasks associated with edges into state w_i and thus $I(w_i)$ and $I(w_-, w_i)$ are both subsets of index set $\{1, 2, \dots, R\}$ of tasks.¹ Note that the $m + 1$ sets $\{I(w_i), I(w_-, w_i) : w_i, w_-, w_i \in \{w_1, w_2, \dots, w_m\}\}$, $i = 0 \dots m$ typically

¹For example in \mathcal{C}_1 of Section 2.2, $I(w_5) = \{6\}$ where θ_6 is “attempt murder” and $I(w_-, w_5)$ is $\{4, 5\}$ where θ_4 is “locate target” and θ_5 is “approach target”.

do not form a partition of $\{1, 2, \dots, R\}$: tasks can be simultaneously suggestive that ω lies in one of a number of different active positions.

Occasionally police may also acquire *negative* evidence from learning that a suspect – thought to have just before lain in position w_i or be transitioning along $e(w_-, w_+)$ – ceases to perform any of the associated tasks. From observing the absence of these tasks they might then infer that ω might have transitioned either to w_0 or a different active state adjacent to w_i in \mathcal{C} . Similar negative inferences might also be made indirectly from learning that ω stops engaging in all tasks associated with an edge emanating from w_i .

With these issues in mind therefore, let $I^*(w_i) \triangleq I_-^*(w_i) \cup I_+^*(w_i)$ where

$$I_+^*(w_i) \triangleq I(w_i) \cup \bigcup_{e(w_-, w_i) \in E(\mathcal{C})} I(w_-, w_i), \quad (1)$$

$$I_-^*(w_i) \triangleq \bigcup_{e(w_-, w_i) \in E(\mathcal{C})} I(w_-) \cup \bigcup_{e(w_i, w_+) \in E(\mathcal{C})} I(w_i, w_+). \quad (2)$$

Thus $I_+^*(w_i)$ is the set of tasks which can positively discriminate w_i from w_0 when the corresponding components take the value 1. The set of tasks in $I_-^*(w_i)$ can negatively discriminate: when taking the value 0 they indicate that ω has ceased to engage in tasks associated with preceding positions and is not engaging in any tasks suggestive of leaving w_i . The set $I^*(w_i)$ is then the set of indices of all tasks in any way relevant to w_i .

For each of the component tasks θ_{tk} at time t we associate a vector of observations of a set of related actions: $\mathbf{Y}_{tk} \in \mathbb{R}^{d_k}$. Let $\mathbf{Y}_t \in \mathbb{R}^d$ be all the routinely observable data on ω so each \mathbf{Y}_{tk} is a projection from \mathbb{R}^d to \mathbb{R}^{d_k} , $d_k \leq d$, $k = 1 \dots R$. It will usually be necessary to work with a filter² of these data streams. So let $Z_{tk} = \tau_k(\mathbf{Y}_{tk})$ denote real functions of these processes and set $\mathbf{Z}_t = (Z_{t1} \dots Z_{tR})$.

Modelling hidden or disguised data

One issue in modelling serious crime is that data concerning a suspect is often hidden, lost, disguised or even be the result of the use of a decoy. This means that the data streams are often intentionally corrupted. However, in contrast to models that describe the data streams directly, our state space model can conceptually accommodate such disruptions: see West and Harrison (1997). Guided by police expert judgement, we can explicitly model the processes designed to disguise or deceive through an appropriate choice of sample distribution of observations given each task.

Informed missingness using CEGs has already been successfully applied in a public health study Barclay et al. (2014). Binary variables were introduced indicating the missingness of readings on mental disability and visual ability for each individual in the Mersey cerebral palsy cohort. The data set including these missingness variables were used to find the best-fitting structural CEG model and from this context specific inferences were made on whether the data were MAR, MCAR or MNAR.³

²Filter in the sense of a function attempting to identify a signal from noisy data.

³Missing at Random, Missing Completely at Random, or Missing Not at Random see Rubin (1976).

In our application we could similarly apply binary variables for the missingness of any of the routinely observable data in \mathbf{Y}_t , and moreover, as also discussed in Barclay et al. (2014), introduce categorical variables for the possible reason for missingness: such as hidden, lost, disguised. The presence of certain patterns of data along with the absence of other data could then influence the probability that certain tasks θ_t were being done despite being hidden or disguised which then would inform the latent state W_t . Alternatively or additionally we could explicitly include *deception* tasks for the hiding of or disguising data and use the above mentioned patterns of data and missing data to perform inference on the probabilities that these deception tasks were being done. This is all, however, beyond the scope of this paper.

3.2 The hierarchical model

The conditional independence structure defining the hierarchy

Because by definition and through the process defined above we would like perfect task information to override all such indirect information we will henceforth assume *task sufficiency*. This states that for all time t

$$W_t \perp\!\!\!\perp \mathbf{Y}_t | \boldsymbol{\theta}_t, \mathcal{F}_{t-}, \quad (3)$$

where \mathcal{F}_{t-} represents the filtration of the past data until but not including time t . This clearly implies that for all time t .

$$W_t \perp\!\!\!\perp \mathbf{Z}_t | \boldsymbol{\theta}_t, \mathcal{F}_{t-}. \quad (4)$$

Ideally we would prefer the filter $\{\mathbf{Z}_t\}_{t \geq 0}$ we use to be sufficient for $\{\boldsymbol{\theta}_t\}_{t \geq 0}$ too i.e. that for all time t

$$\boldsymbol{\theta}_t \perp\!\!\!\perp \mathbf{Y}_t | \mathbf{Z}_t, \mathcal{F}_{t-}. \quad (5)$$

Then there would be no loss in discarding information in \mathbf{Y}_t not expressed in \mathbf{Z}_t . In what we henceforth present, since we develop recurrences only concerning $\{\mathbf{Z}_t\}_{t \geq 0}$ and not $\{\mathbf{Y}_t\}_{t \geq 0}$ we implicitly assume condition 5.

Although condition 5 is a heroic one, in our examples a well-chosen one dimensional time series of intensities Z_{tk} , performs well even when these are chosen to be linear in the records of the component signals \mathbf{Y}_{tk} . One advantage of this simplicity is that the role of the filter can be explained and if necessary adapted by the user, perhaps even customising this filter to their own personal modus operandi and judgements.

Again for simplicity we henceforth assume that any filter $\{\mathbf{Z}_t\}_{t \geq 0}$ will be a *Markov task filter* i.e. that for all $t > 0$

$$\perp\!\!\!\perp \{\mathbf{Z}_{t'}\}_{t' \geq t} | \boldsymbol{\theta}_t, \mathcal{F}_{t-}. \quad (6)$$

This assumption is a familiar one made for dynamic models; see e.g. West and Harrison (1997). It assumes that once the task is known, no further past information about past $\{\mathbf{Z}_t\}_{t \geq 0}$ will add anything further useful for predicting the future. This assumption

enables us, for particular choices of sample distributions, to use all the established recurrences for dynamic state space models – in particular those from dynamic switching models so excellently summarised in Frühwirth-Schnatter (2008). Here our RDCEG probability model specifies such a switching mechanism.

Defining tasks to be fit for purpose

Our interpretation of $I^*(w)$ requires that if suspect ω is known to be either neutral or in any active state w_i then the only components in θ_t that helpfully discriminates between these two possibilities must lie in $\theta_{I^*(w_i),t}$. The assumption *Task set integrity* demands $I^*(w_i)$ is defined so that for all $i = 1, 2, \dots, m$, $0 \leq t \leq T$,

$$W_t \perp\!\!\!\perp \theta_t | W_t \in \{w_0, w_i\}, \theta_{I^*(w_i),t}. \quad (7)$$

This is equivalent to requiring that

$$\frac{p_t(w_i | \theta_t, \mathcal{F}_{t-})}{p_t(w_0 | \theta_t, \mathcal{F}_{t-})} = \frac{p_t(w_i | \theta_{I^*(w_i),t}, \theta_{\widehat{I}^*(w_i),t}, \mathcal{F}_{t-})}{p_t(w_0 | \theta_{I^*(w_i),t}, \theta_{\widehat{I}^*(w_i),t}, \mathcal{F}_{t-})}$$

is a function only of $\theta_{I^*(w_i),t}$ where $\widehat{I}^*(w)$ denotes the set of indices not in $I^*(w)$. Task set integrity is always satisfied by setting $I^*(w) = \{1, 2, \dots, R\}$, $i = 1, 2, \dots, m$ but of course for transparency and computational efficiency ideally $I^*(w)$ is chosen to be a small subset of $\{1, 2, \dots, R\}$. Providing the divisor is not zero, task set integrity holds whenever

$$\widehat{\lambda}_i(\theta_{I^*(w_i),t}) \triangleq \log p_t(\theta_{\widehat{I}^*(w_i),t} | \theta_{I^*(w_i),t}, w_i, \mathcal{F}_{t-}) - \log p_t(\theta_{\widehat{I}^*(w_i),t} | \theta_{I^*(w_i),t}, w_0, \mathcal{F}_{t-})$$

is a function of θ_t only through $\theta_{I^*(w_i),t}$. So by writing the prior and posterior log -odds as

$$\begin{aligned} \rho_{it} &\triangleq \log p_t(w_i | \mathcal{F}_{t-}) - \log p_t(w_0 | \mathcal{F}_{t-}), \\ \rho_{it}^* &\triangleq \log p_t(w_i | \theta_t, \mathcal{F}_{t-}) - \log p_t(w_0 | \theta_t, \mathcal{F}_{t-}) \end{aligned}$$

and the loglikelihood ratio of task vector $\lambda_i(\theta_{I^*(w_i),t})$,

$$\lambda_i(\theta_{I^*(w_i),t}) \triangleq \log p_t(\theta_{I^*(w_i),t} | w_i, \mathcal{F}_{t-}) - \log p_t(\theta_{I^*(w_i),t} | w_0, \mathcal{F}_{t-})$$

then a little rearrangement gives us an adaptation of the usual Bayesian linear updating equation linking posterior and prior odds viz:

$$\rho_{it}^* = \rho_{it} + \lambda_i(\theta_{I^*(w_i),t}) + \widehat{\lambda}_i(\theta_{I^*(w_i),t}). \quad (8)$$

Note that (7) holds in particular whenever it is a *simple task vector*; i.e. has the property that for any time $t > 0$, $\omega \in \Omega$ and $i = 1, 2, \dots, m$

$$p_t(\theta_{\widehat{I}^*(w_i),t} | \theta_{I^*(w_i),t}, w_i, \mathcal{F}_{t-}) = p_{tj}(\theta_{\widehat{I}^*(w_i),t} | \theta_{I^*(w_i),t}, w_0, \mathcal{F}_{t-}). \quad (9)$$

Property 9 holds whenever the other tasks are useless for discriminating any threat position w_i from w_0 : the probability ω engages in these tasks does not depend on these other tasks. In this case the term $\widehat{\lambda}_i(\boldsymbol{\theta}_{I^*(w_i)t})$ vanishes and

$$\rho_{it}^* = \rho_{it} + \lambda_i(\boldsymbol{\theta}_{I^*(w_i)t}). \quad (10)$$

For practical reasons we have often found it convenient to decompose $\lambda_i(\boldsymbol{\theta}_{I^*(w_i),t,j})$ into functions of components in $I_-^*(w_i)$ and $I_+^*(w_i)$ respectively. When these are disjoint and conditionally independent given w_i – as in practice we find is often a plausible assumption then

$$\lambda_i(\boldsymbol{\theta}_{I^*(w_i)t}) = \lambda_{-i}(\boldsymbol{\theta}_{I_-^*(w_i)t}) + \lambda_{+i}(\boldsymbol{\theta}_{I_+^*(w_i)t}), \quad (11)$$

where

$$\begin{aligned} \lambda_{-i}(\boldsymbol{\theta}_{I_-^*(w_i)t}) &\triangleq \log p_t(\boldsymbol{\theta}_{I_-^*(w_i)t}|w_i, \mathcal{F}_{t-}) - \log p_t(\boldsymbol{\theta}_{I_-^*(w_i)t}|w_0, \mathcal{F}_{t-}), \\ \lambda_{+i}(\boldsymbol{\theta}_{I_+^*(w_i)t}) &\triangleq \log p_t(\boldsymbol{\theta}_{I_+^*(w_i)t}|w_i, \mathcal{F}_{t-}) - \log p_t(\boldsymbol{\theta}_{I_+^*(w_i)t}|w_0, \mathcal{F}_{t-}). \end{aligned}$$

Note here that, by definition, $\lambda_{-i}(\boldsymbol{\theta}_{I_-^*(w_i)t})$, takes its maximum value when $\boldsymbol{\theta}_{I_-^*(w_i)t} = \mathbf{0}$ and $\lambda_{+i}(\boldsymbol{\theta}_{I_+^*(w_i)t})$ takes its maximum value when $\boldsymbol{\theta}_{I_+^*(w_i)t} = \mathbf{1}$, $i = 1, 2, \dots, m$.

The equation (10) are now sufficient to calculate the probability that ω is in each of the positions w_1, w_2, \dots, w_m given our evidence, using the familiar invertible function from log odds to probability: see Smith (2010).

Model assumptions concerning routine observations

For our chosen filtered sequence $\mathbf{Z}_t \triangleq (Z_{1t}, Z_{2t}, \dots, Z_{Rt})$ designed to pick up the different tasks associated with a criminal process let

$$\mathbf{Z}_{\widehat{kt}} \triangleq (Z_{1t}, Z_{2t}, \dots, Z_{(k-1)t}, Z_{(k+1)t}, \dots, Z_{Rt}).$$

Then a simple but bold type of Naive Bayes assumption is to assume that filter \mathbf{Z}_t is *pure*: i.e. that for any set $\boldsymbol{\theta}_{At}$ containing $\boldsymbol{\theta}_{kt}$ as a component

$$\mathbf{Z}_{kt} \perp\!\!\!\perp \mathbf{Z}_{\widehat{kt}} | \boldsymbol{\theta}_{At}, \mathcal{F}_{t-}. \quad (12)$$

Then Bayes Rule and task set integrity implies that within tasks

$$p_t(\mathbf{Z}_{I^*(w_i)t} | \boldsymbol{\theta}_{I^*(w_i)t}, \mathcal{F}_{t-}) = \prod_{k \in I^*(w_i)} p_t(\mathbf{Z}_{kt} | \boldsymbol{\theta}_{kt}, \mathcal{F}_{t-})$$

whilst across tasks

$$p_t(w_{it} | \mathbf{Z}_t, \mathcal{F}_{t-}) = \sum_{\boldsymbol{\theta}_{I^*(w_i)t} \in \Theta_{I^*(w_i)}} p_t(w_{it} | \boldsymbol{\theta}_{I^*(w_i)t}, \mathcal{F}_{t-}) p_t(\boldsymbol{\theta}_{I^*(w_i)t} | \mathbf{Z}_{I^*(w_i)t}, \mathcal{F}_{t-}), \quad (13)$$

where

$$p(\boldsymbol{\theta}_{I^*(w_i)t} | \mathbf{Z}_{I^*(w_i)t}, \mathcal{F}_{t-}) = \frac{p(\mathbf{Z}_{I^*(w_i)t} | \boldsymbol{\theta}_{I^*(w_i)t}, \mathcal{F}_{t-}) p(\boldsymbol{\theta}_{I^*(w_i)t} | \mathcal{F}_{t-})}{p(\mathbf{Z}_{I^*(w_i)t,j} | \mathcal{F}_{t-,j})}.$$

Let the prior and posterior odds be respectively denoted by

$$\begin{aligned} \phi_{i't}^{iA} &\triangleq \log p_t(\boldsymbol{\theta}_{I^*(w_i)t} = \chi_A | \mathcal{F}_{t-}) - \log p_t(\boldsymbol{\theta}_{I^*(w_i)t} = \chi_\emptyset | \mathcal{F}_{t-}), \\ \phi_{i't}^{*iA} &\triangleq \log p_t(\boldsymbol{\theta}_{I^*(w_i)t} = \chi_A | \mathbf{Y}_{I^*(w_i)t}, \mathcal{F}_{t-}) - \log p_t(\boldsymbol{\theta}_{I^*(w_i)t} = \chi_\emptyset | \mathbf{Y}_{I^*(w_i)t}, \mathcal{F}_{t-}), \\ \gamma_{kt} &\triangleq \log p_t(\mathbf{Z}_{kt} | \theta_{kt} = 1, \mathcal{F}_{t-}) - \log p_t(\mathbf{Z}_{kt} | \theta_{kt} = 0, \mathcal{F}_{t-}) \end{aligned}$$

and let $\gamma_t^A \triangleq \sum_{k \in A} \gamma_{kt}$. Then

$$\log p_t(\mathbf{Z}_{I^*(w_i)t} | \boldsymbol{\theta}_{I^*(w_i)t} = \chi_A, \mathcal{F}_{t-}) - \log p_t(\mathbf{Z}_{I^*(w_i)t} | \boldsymbol{\theta}_{I^*(w_i)t} = \chi_\emptyset, \mathcal{F}_{t-}) = \gamma_t^A$$

and from the above

$$\phi_{i't}^{*iA} = \phi_{i't}^{iA} + \gamma_{i't}^A. \quad (14)$$

For any set A we can therefore calculate

$$p_t(\boldsymbol{\theta}_{I^*(w_i)t} = \chi_A | \mathbf{Z}_{I^*(w_i)t}, \mathcal{F}_{t-}), \quad (15)$$

where the Law of Total Probability implies that for $i = 0, 1, 2, \dots, m$

$$p_t(W_t = w_i | \mathbf{Z}_t, \mathcal{F}_{t-}) = \sum_{\boldsymbol{\theta} \in \Theta_{I^*(w_i)}} p_t(w_i | \boldsymbol{\theta}, \mathcal{F}_{t-}) p_t(\boldsymbol{\theta} | \mathbf{Z}_{I^*(w_i)t}, \mathcal{F}_{t-}). \quad (16)$$

Here the position probabilities over tasks calculated from (14) are averaged over the different tasks possibly explaining the data, weighting using the posterior probabilities given in (15). Note that it is easy to check that these indirect observations provide less discriminatory power than when tasks are observed directly. The assumptions above therefore provide us with a formally justifiable propagation algorithm for updating the probabilities of a suspect's likely criminal status. We next turn to how we might calibrate the model to the expert judgements we might elicit from criminologists, police and technicians about the probable relationships between criminal status, what they might try to accomplish and how this endeavour might be reflected through how they communicate.

4 The elicitation process

4.1 Introduction

Copying the standard protocols for the elicitation of a Bayesian Network: see e.g. Korb and Nicholson (2010), as in e.g. Wilkerson and Smith (2018) our process begins with the elicitation of structure. We perform a sequence of three structural elicitations for each of the three levels. These can proceed almost entirely using natural language descriptions of

the process. Because the representation of the structure of each of the levels is formal and compatible with a probability model the structural elicitation can take place before the model is quantified. This is extremely helpful because structural information is typically much easier to elicit faithfully than quantitative judgments. The RDCEG defining this structure, the list of tasks and how these might interrelate and then the choice of filter of the routine observations follows:

1. First the decision analyst elicits the positions of the process via the careful conversion of natural language expressions within the domain experts' description of the process into the topology of a RDCEG – often somewhat more nuanced than the ones we discussed above and in the example below.
2. Second the positions and edges of this RDCEG are then associated to elicited portfolios of tasks.
3. Finally each task is associated with the way domain experts and police believe ω might behave in order to carry out these tasks, including how they might choose to disguise these actions and so what signals might be visible when ω enacts a task.

We now briefly outline each of these steps in turn in a little more detail.

4.2 Choosing an appropriate RDCEG

Firstly, when eliciting a RDCEG we aim to keep the number of positions as small as possible within the constraint that they are sufficient to distinguish relevant states. The choice of topology should reflect what is known about the development of the modelled criminal behaviour. Positions may depend on the history, environmental and personality profile covariates exhibited by a suspect. Relevant population studies of criminal behaviours are often helpful here. Based on historical cases, criminologists' analyses Gill (2012) and discussions with practitioners, we have found that the coarsest type of model – an illustration of such given in the next section – of different types of attack and concerning different people – are often generic.

Secondly positions need to be well defined enough to pass the Clarity Test Howard (1988); Smith (2010). This is achieved by demanding that the suspect could, if they were so minded, place themselves in a particular position. Such categories are often syntheses of standard scales used by social workers and probation workers across the world. Many examples of these types of categorisations, based on fusing various publicly available categorisations – for example those found in the training manuals of social workers in detecting people threatening to eventually perpetrate acts of severe violence – are given in Smith and Shenvi (2018).

Thirdly positions must be defined such that for each position there is a collection of tasks associated with it that jointly informs whether the suspect is in that position or transitioning from that position to another position. A stylised example of this association was given earlier in this paper and we will illustrate the process in more detail in the next section with a deeper illustration.

Once the RDCEG has been drawn its embedded assumptions can be queried by automatically generating logical deductions concerning the implications of the model. If such deductions appear implausible to relevant domain experts then positions need to be redefined and graphs redrawn until they are. The ways of iterating until a model is requisite and the nature of these deductions is beyond the scope of this paper but are discussed in Collazo et al. (2018).

The final step in the elicitation of the RDCEG will be the prior conditional probabilities associated with the positions and the hyperparameters of the holding times. Suitable generic methods for this elicitation are now very well established: see e.g. O’Hagan et al. (2006); Smith (2010) and these need little adaptation to be applied. Note, in particular, that the methods described for the elicitation of the position probabilities in the RDCEG are essentially identical to those for the CEG as for example discussed in a chapter of Collazo et al. (2018).

4.3 Elicitation of portfolios of tasks

Clustering the tasks

The next elicitation process is to take each position in turn and a list of associated tasks conditional on ω being known to lie in that position. The questions we might ask would be something like “Now suppose that you happen to learn that ω lies in position w_i . What behaviours/tasks would you expect them to perform that would be different from what they would typically do were they neutral?” We try to ensure that either ω ’s engagement in such tasks could be learned through intelligence or alternatively be indicated through certain filters.

Typically in well designed models we specify tasks so that they are as specific to only a small proportion of ω ’s active states. This makes them as discriminatory as possible. Note that each component of θ_t must be defined sufficiently precisely – i.e. pass the clarity test for ω to be able to divulge its value if so inclined; see Smith and Shenvi (2018).

We have found it useful to toggle between specifying positions and specifying tasks: sometimes aggregating positions if they appear associated with the same sets of tasks or splitting a position into a set of new ones if a finer definition can discriminate between one position and another. It is also sometimes helpful to readjust the definition of tasks once we have elicited possible signals.

Once the task sets are requisite Phillips (1984); Smith (2010) we need to specify the various odds ratios against the neutral state. We illustrate this in the next section.

Simplifying assumptions that can ease task probability elicitation

Although the log score updating formulae ((8), (10)) are simple ones, to evaluate the log odd scores above can demand a great many probabilities, both of each task given its associated position and of seeing that task performed if ω were neutral, to be elicited or estimated. This can destabilise the system unless various simplifying assumptions are made.

Conditional on an active position w_i we recommend that the first probability to be elicited is when ω is engaging in *all* the tasks in the portfolio of tasks associated with w . We then use this elicitation to benchmark the probability that ω is engaging in a subset of these tasks. To calculate the odds of the portfolio against the neutral suspect, a default assumption that is sometimes appropriate is simply to assume that people will engage in these probabilities independently – a *naive Bayes Assumption* Smith (2010) for the neutral suspect. In this case for any time $t > 0$, $\omega \in \Omega$ and $i = 1, 2, \dots, m$

$$p_t(\boldsymbol{\theta}_t | w_0, \mathcal{F}_t) = \prod_{k=1}^R p_t(\theta_{tk} | w_0, \mathcal{F}_t)$$

It is easy to check that when a portfolio contains more than one task, and when such an assumption is valid, it can provide the basis of a very powerful discriminatory tool. This is because the divisor in the relevant odds reduces exponentially with the number of tasks whilst in the denominator does not: see e.g. Smith and Shenvi (2018) for an example of this.

Of course in some instances this naive Bayes assumption may not be appropriate. It will then need to be substituted. When population statistics associated with public engagement in different task activities are available these can be used to verify this assumption or form the basis of constructively replacing it.

4.4 Choosing an appropriate filter of routine data streams

Typically we would like the components of $\{\mathbf{Z}_t\}_{t \geq 0}$ to measure an intensity of activity related to a position or edge task. In this sense we would therefore like to factor out all signals that might be considered typical of ω 's innocent activities so that we can focus on the incriminating signals. The full data stream $\{\mathbf{Y}_t\}_{t \geq 0}$ collected on ω tends to be a highly non-stationary multivariate time series. However we strive to construct the filter $\{\mathbf{Z}_t\}_{t \geq 0}$ so that the stochastic dependence it exhibits is explained solely by ω 's engagement in certain tasks (see (6)). This filter is clearly dependent both on population level signals and what we know about ω 's personality. We therefore usually need expert judgments to choose $\{\mathbf{Z}_t\}_{t \geq 0}$ so that it is fit for purpose.

There are some generic features that are worth introducing at this stage. First in the case of edge tasks we typically observe something different than before as ω begins to enact a new task in order to make a transition. So some components of \mathbf{Z}_t will be defined as first differences of derived series. Secondly indicative observations may also need to be smoothed from the past – either because what we see may forewarn a task is about to be enacted, or simply because short term averages – for example any measure of intensity of communication – will often be better represented by an average over the recent past rather than by an instantaneous measure.

To construct our hierarchical model we typically loop around the bullets below:

1. Reflect on what functions of the vector of observable data available to the police might help indicate that a suspect really does lie in a particular task rather than

other related innocent activities, $k = 1, 2, \dots, R$. This choice should be informed by the ease at which such signals can be filtered but also how easy it might be for a criminal to disguise that signal were they to learn that the chosen candidate filter was being used.

- Using expert judgments and any survey data available, reflect on what the distribution of \mathbf{Z}_{kt} might be were the suspect actually engaged in the particular task and if they were not. Thus specify

$$p(\mathbf{Z}_{kt}|\theta_{kt} = 1) \text{ and } p(\mathbf{Z}_{kt}|\theta_{kt} = 0).$$

- Check that these two distributions are not close to one another. If not return to the first step.

Finally in many instances of such police work routine measurements concerning suspects are typically recorded and reported over fixed periods of time. This means that the filtered observation sequence is a discrete time filter. For modelling purposes it has been necessary to define the deep stochastic process as semi-Markov. However the semi-Markov structure with holding times and transition probabilities specified will retain the Markov structure over the fixed time points. This means once appropriate transformations are applied standard updating rules associated with Markov switching models are then valid Frühwirth-Schnatter (2008): see Appendix B in the supplementary material Bunnin and Smith (2019).

5 A vehicle attacker example

5.1 States

We now give a more detailed example of a vehicle attacker that illustrates the three level hierarchical model of latent states, tasks, and routinely observable data. We specify the states of the RDCEG to be:

$$W = \{N, A, T, P, M\}, \tag{17}$$

where N is “Neutral”, A is “ActiveConvert”, T is “Training”, P is “Preparing”, and M is “Mobilised”. Based on existing information about the suspect we assign prior probabilities to each state as shown in Figure 2; implicitly the prior probability for the elided “Neutral” state is 0.05. Based on knowledge about such attacks we hypothesize that the suspect may transition from A to T or P, from T to P, from P to M, and from M back to P. These transitions are indicated by the directed edges between the vertices on said figure. The weights labelling the transitions are the probabilities of transitions from the source vertex to the destination vertex conditional on a transition having occurred (i.e. the entries labelled m_{w_i, w_j} in Table 1 in Appendix B of the supplementary material). The probability of transition into the “Neutral” state from any represented state is implied by the sum of all the emanating edges’ probabilities summing to one. This is in contrast to Shenvi and Smith (2018) where the transition probabilities are conditional on not moving to the absorbing state. This prior RDCEG is used in all the examples in this section.




Figure 2: RDCEG for Vehicle Attacker.

5.2 Tasks

We hypothesize the tasks relevant for these positions, i.e. the tasks that are related to which position the suspect is in, are:

$$\theta = \{\theta_j\}_{j=1}^{10},$$

where:

θ_1 is Engaging with Radicals	θ_2 is Engaging in Public Threats
θ_3 is Making Personal Threats	θ_4 is Fewer Public Engagements in Radicalisation
θ_5 is Fewer Contacts with Family and Friends	θ_6 is Securing Monetary Resources
θ_7 is Learning to Drive Large Vehicle	θ_8 is Obtaining Vehicle
θ_9 is Reconnaissance of Target Locations	θ_{10} is Moving to Target Location

For each position w_i , a particular subset of the above tasks are taken to be indicators that the suspect is there. $I^*(w_i)$ is the index set for this subset and we need to specify the distribution $p(\theta_{I^*(w_i)}|w_i)$ for each position. Appendix A of the supplementary material details a methodology for this specification that makes the model discriminatory and Table 4 shows the resulting probabilities used.

5.3 Routinely observable data

The data used to estimate the probabilities that the suspect is engaging in any particular task or tasks are varied and various and may change as technologies and data gathering methods change. In addition as new evidence is gained and the threat level of a suspect increases the authorities may decide to increase monitoring and hence gain more and new types of data. Therefore having the tasks θ intermediate between the positions W and the observed data Y is desirable for both model structural reasons and practical data abstraction purposes. We denote the observable data as a d -dimensional vector

process in discrete time $\mathbf{Y}_t = (Y_{i,t})_{i=1}^d$. In general we assume $\mathbf{Y}_t \in \mathbb{R}^d$. In this example, however, several of the components are count data, such as the number of times such events are observed in a given period, so that: $Y_{i,t} \in \mathbb{Z}^+$. We set here:

Y_1 : radical website visits	Y_2 : physical meetings with known radicals
Y_3 : electronic meetings with known radicals	Y_4 : meetings with trained radicals
Y_5 : meetings with known cell members	Y_6 : seen at radical demonstrations
Y_7 : contacts with non-radicals	Y_8 : public threats made
Y_9 : personal threats made	Y_{10} : increase in known financial resources
Y_{11} : decrease in known financial resources	Y_{12} : obtaining large vehicle driving licence
Y_{13} : vehicle dealer or rental website visits	Y_{14} : vehicle dealer or rental physical visits
Y_{15} : E-visits to target locations	Y_{16} : physical visits to target locations
Y_{17} : statements of intent	Y_{18} : legacy statements

As described in Section 3 we assume that for each task θ_j we can construct a filter Z_j of the relevant data: Here we define the function τ :

$$\begin{aligned} \tau : \mathbb{R}^d &\mapsto \mathbb{R}^R, Z_{j,t} = \tau_j(\mathbf{Y}_t), \\ \tau_j(\mathbf{Y}_t) &= \frac{1}{|I_{\theta_j}|} \sum_{i \in I_{\theta_j}} \tilde{y}_{i,t}, \end{aligned}$$

where \tilde{y}_i is a standardisation⁴ of Y_i and $|I_{\theta_j}|$ is the cardinality of the index set of components of Y dependent on the j^{th} task. We could also set additional components of \mathbf{Y} to be changes over time of other components of \mathbf{Y} and thus monitor drops or spikes in, for example, communication levels with known radicalisers, or with family and friends. We specify the relationship between the observable data and the tasks in Table 2 where the (Y_i, θ_j) entry indicates whether the i th variable is relevant data for the j th task.

5.4 Specifications of the distributions of the task set given position and task set likelihood

For each position we set the probability that all the tasks in that position’s task set are being done to 0.4. The individual probabilities that each task is being done given the suspect is in the neutral state are specified in the column labelled “Neutral” in Table 3. The log odds interpolation methodology detailed in Appendix A in the supplementary material is then used to construct the probabilities that no tasks or less than all the tasks were being done. For example, as shown in the Table 3, for the “Mobilised” position, the tasks “Engaging in public threats”, “Making personal threats”, “Reconnaissance of target locations”, and “Moving to target location” are relevant. Table 4 has the resulting probabilities for each point of $\theta_{I_{w_i}^*} \in \{0, 1\}^{|I_{w_i}^*|}$. For simplicity we used shifted asymmetric logistic functions to construct the density functions $p(Z_{I^*(w_i)} | \theta_{I^*(w_i)})$ for given values of $\theta_{I^*(w_i)}$: see (18). The shift parameters $x_{0,j}$ and the growth rate parameters $k_{0,j}$, $k_{1,j}$ were used to construct functions that were relatively unresponsive when

⁴We subtract a pre-defined mean and divide by a pre-defined standard deviation estimated by investigators’ judgement, experience and historical data.

Observable	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	θ_9	θ_{10}
RadWebVisits	1	0	0	1	1	0	0	0	0	0
PhysicalMeetsWithRadicals	1	0	0	0	1	0	0	0	0	0
E-MeetsWithradicals	1	0	0	1	1	0	0	0	0	0
MeetTrainedRadicals	1	0	0	0	1	0	0	0	0	0
MeetCellMembers	1	0	0	0	1	0	0	0	0	0
SeenAtRadicalDemonstrations	1	0	0	0	1	0	0	0	0	0
ContactsWithNonRadicals	0	0	0	1	0	0	0	0	0	0
PublicThreatsMade	0	1	0	0	0	0	0	0	0	0
PersonalThreatMade	0	0	1	0	0	0	0	0	0	0
IncreaseInFinances	0	0	0	0	0	1	1	1	0	0
DecreaseInFinances	0	0	0	0	0	0	1	1	0	0
ObtainLGVLicence	0	0	0	0	0	0	1	0	0	0
CarDealerWebHits	0	0	0	0	0	0	0	1	0	0
CarDealerPhysicalVisits	0	0	0	0	0	0	0	1	0	0
E-VisitsToTargetLocations	0	0	0	0	0	0	0	0	1	0
VisitsToTargetLocations	0	0	0	0	0	0	0	0	1	0
LegacyStatements	0	0	0	0	0	0	0	0	0	1
StatementOfIntent	0	1	1	0	0	0	0	0	0	1

Table 2: Routine Observation versus Task dependency structure.

	ActiveConvert	Training	Preparing	Mobilised	Neutral
State_Task_Index_Sets					
EngageWithRadicalisers	1	0	0	0	0.020
EngageInPublicThreats	0	0	1	1	0.001
MakePersonalThreats	0	0	1	1	0.001
RedPubEngInRad	1	1	0	0	0.600
RedCntctWthFmlyFrnds	1	0	0	0	0.300
ObtainResources	1	1	0	0	0.300
LearnToDrive	0	1	0	0	0.300
ObtainVehicle	0	1	1	0	0.200
ReconnoitreTargets	0	0	1	1	0.100
MoveToTarget	0	0	0	1	0.200
Cardinality	4	4	4	4	
p+	0.400	0.400	0.400	0.400	
p0	0.001	0.011	0.00000002	0.00000002	
ξ	1.051	2.076	0.332	0.331	

Table 3: Task/position dependencies; probability of task given Neutral state; p+ is probability of all tasks being done, p0 that of none being done; ξ_i is solved for to make the probabilities sum to one for each position.

$Z_j < x_{0,j}$ but sharply responsive when $Z_j \geq x_{0,j}$. An illustration of the form of these functions is provided in Figure 3 for the one and two-dimensional cases i.e. when there are one or two tasks in the task set $I^*(w_i)$; this is purely for ease of plotting: as shown

State_Task_Index_Sets	ActiveConvert	Training	Preparing	Mobilised
np_0	0.00108	0.01080	0.00000002	0.00000002
np_0001	0.00475	0.01341	0.00112	0.00112
np_0010	0.00475	0.01341	0.00112	0.00112
np_0011	0.02319	0.02742	0.01860	0.01860
np_0100	0.00475	0.01341	0.00112	0.00112
np_0101	0.02319	0.02742	0.01860	0.01860
np_0110	0.02319	0.02742	0.01860	0.01860
np_0111	0.11019	0.09277	0.12098	0.12098
np_1000	0.00475	0.01341	0.00112	0.00112
np_1001	0.02319	0.02742	0.01860	0.01860
np_1010	0.02319	0.02742	0.01860	0.01860
np_1011	0.11019	0.09277	0.12098	0.12098
np_1100	0.02319	0.02742	0.01860	0.01860
np_1101	0.11019	0.09277	0.12098	0.12098
np_1110	0.11019	0.09277	0.12098	0.12098
np_1111	0.40000	0.40000	0.40000	0.40000
ntp	1.00000	1.00000	1.00000	1.00000

Table 4: Probabilities of task sets given each position using method in Appendix A.


in Table 3 four dimensional task sets were used in the example scenarios

$$\begin{aligned}
p(Z_{I^*(w_i)}|\theta_{I^*(w_i)}) &= |\theta_{I^*(w_i)}|^{-1} \sum_{j \in I^*(w_i)} g(Z_j|\theta_j, x_{0,j}, k_{0,j}, k_{1,j}), \\
g(x|\theta, x_0, k_0, k_1) &= \frac{1}{1 + \exp(-k_0(x - x_0))} \mathcal{X}_{\{x < x_0\}} + \frac{1}{1 + \exp(-k_1(x - x_0))} \mathcal{X}_{\{x \geq x_0\}}.
\end{aligned} \tag{18}$$


5.5 Scenarios

We illustrate the propagation of probabilities through the model based on scenarios of simulated data. We use the same framework as above and manually set the routinely observed data \mathbf{Y}_t through 24 weekly time steps to examine how the currently parameterised model behaves under each scenario.


Scenario 5.1. In this scenario the suspect increases their web visits to target locations from week 8 and their physical visits to target locations from the week 21; they are in constant communication with known radicals and there is an increase in their finances followed by a decrease in first few weeks, during which time they are seen to be visiting car dealers electronically and physically. They make public and personal threats and in the last weeks of the period the threatening data increases with a legacy statement and a statement of intent. Figures 4a, 5a, 5c show the increase in threat level resulting from this scenario.




(a) Orange line shows $p(Z_{I^*(w_i)} | \theta_{I^*(w_i)} = 1)$; Blue line shows $p(Z_{I^*(w_i)} | \theta_{I^*(w_i)} = 0)$



(b) Two dimensional. Each coloured surface is for a different value of $\theta_{I^*_{w_i}} \in \{0, 1\}^2$; $Z_{I^*_{w_i}}$



(c) $\theta_{I^*_{w_i}} = (0, 0)$



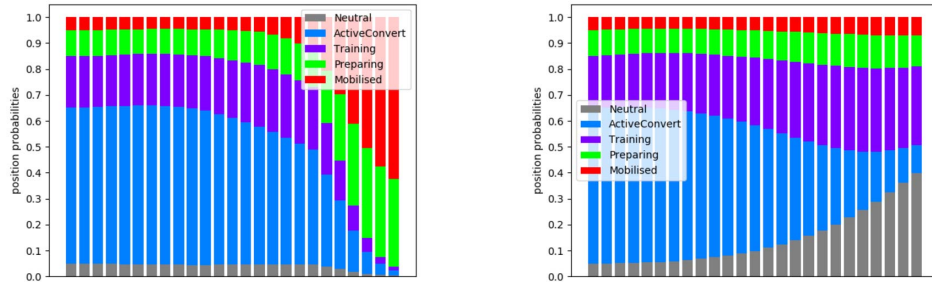
(d) $\theta_{I^*_{w_i}} = (1, 1)$

Figure 3: Illustrative task likelihood functional forms for one and two dimensional task sets.

Scenario 5.2. The suspect's communications and possible training/preparing type data linearly decreases from initial levels similar to scenario 5.1 to zero over the 24 weeks. Moreover there are no threats made during the whole period. Figures 4b, 5b, 5d show the decreasing threat level resulting from this scenario.

5.6 Eventual probability of mobilisation

For any individual suspect or group of suspects the medium to long-term probabilities of mobilisation is of key interest and can aid as a model diagnostic tool. We can estimate this by using the semi-Markov transition matrix to evolve the current probabilities. The RDCEG in this section has the neutral state as the single absorbing state hence asymptotically the probability of this state will go to one; However in the practical medium term we can examine the behaviour of the active positions including the mobilised state. Under the configuration of the priors and the edge probabilities as in



(a) Scenario 5.1: a suspect pursuing activities consistent with preparing for an attack

(b) Scenario 5.2: a suspect reducing activities consistent with no evidence of threat

Figure 4: Posterior state probabilities over time under scenario 5.1 and 5.2.



Figure 2 and with the holding time distribution $\zeta_i(t, t')$ set to a constant 0.01 for the set time period of $t' - t$ equal to one week (as used in the examples above), the qualitative behaviour of the RDCEG's state probabilities can be seen in Figures 6a and 6b. Figures 6c and 6d show the long term behaviour under alternative specifications where the mobilised position is another absorbing state: the suspect once having mobilised and executed the attack cannot transition to any of the other states including the neutral state. The reasoning behind this latter configuration is an assumption that once the individual has mobilised this entails an attack and the end of this particular police case.

6 Model diagnostics

6.1 Robustness to RDCEG structure specification


The graph representation of the RDCEG, that is the set of states and the directed edges between them, form the structure of the RDCEG that is meant to faithfully represent the possible pathways of individuals towards, in this application, acts of terrorism. The actual structure chosen is based on historical cases, existing research by criminologists and discussions with practitioners and is predicated on the assumption that the states are “self-identifiable” that is that the actual individual would be able to place themselves in one of these states at any given time.

Assuming such an approach is valid, without being able to actually look into an individual's mind, we are liable to mis-specify the structure: for example construct states that could meaningfully and usefully be split into finer sub-states, or have a set of states that should be collapsed into one state; or have edges where they should not exist or have edges missing. Whether the set of states are “correct to the individual's mind” is arguably less relevant for our purposes than whether the set of states are a useful discretisation of the individual's potential pathway from the investigators' perspective; and for this we can be directly guided.




(a) Position Score over time based on Scenario 5.1

(b) Position Score over time based on Scenario 5.2




(c) RDCEG with final posterior probabilities based on Scenario 5.1




(d) RDCEG with final posterior probabilities based on Scenario 5.2

Figure 5: Position Score and RDCEG for 5.1 and 5.2.


It is still of use to analyse the sensitivity of behaviour and results to changes in the structure chosen. To this end we examine the effect of using different sets of states and different edges by using the same data sets with different RDCEG structures. We coarsen the RDCEG used in Section 5 by collapsing the “Training” state into the “Preparing” state, and then separately refine the RDCEG by splitting the “Preparing” state in two states: “Preparing for a Vehicle Attack” and “Preparing for a Bomb Attack”, so that we have two new alternative RDCEG structures to compare with the original. We expand the task sets and data (see supplementary material) and run two new data scenarios involving a potential joint vehicle and bomb attack on these two new structures along with the original structure giving in total six sets of results. Results for this analysis along with details of the scenarios are given in Appendix D and Appendix C in the supplementary material. The impacts on the posterior probabilities are as expected: the “Neutral” state’s probability is relatively unchanged and the probability mass of the coarsened state is roughly the sum of the finer states in each example under both scenarios.




(a) Position probabilities over 10 thousand periods based purely on semi-Markov transition matrix



(b) Position probabilities over 2 million periods based purely on semi-Markov transition matrix



(c) Position probabilities with two absorbing states over one thousand periods based purely on semi-Markov transition matrix; holding time transition rate: $\zeta_i(t, t') = 0.01$; edge probabilities $m_{a,n} = m_{t,n} = m_{p,n} = 0.3$



(d) Chart showing the two absorbing states' stationary probabilities varying by directly setting the per-period probability of transition from each transient active state to neutral from 0.01 to 0.99

Figure 6: Figures for long-term probability of mobilisation.

In contrast to our elicited RDCEG structure, Shenvi and Smith (2019) construct an RDCEG structure explicitly from data on individuals within an open population using a hierarchical clustering algorithm and Bayes Factor based model selection in health care settings: recurrent falls in the elderly in the context of care services and drug effects on early epilepsy cases.

6.2 Sensitivity analysis

To improve our understanding of the behaviour and robustness of the model to subjective inputs, and to identify key parameters that merit extra analysis to determine their optimal value, we perform sensitivity analysis against a base scenario. The model is high dimensional in the sense that the number of configurable parameters is large: so for the time being this analysis has focussed on the state prior probabilities and the

holding distributions the latter of which are key to determining the speed of transition.

We tested the sensitivity of the evolution of the state probabilities applying both moderate and large shifts to the state priors. We performed a similar analysis shifting the holding time distribution parameter $\zeta_{i,j}$ which represents the probability of a transition from state w_i to w_j in a given time period (here one week) given that an edge exists from w_i to w_j .

For moderate changes in priors, increasing the prior for a less threatening state has an initial effect but is outweighed by any data that indicates threat; whilst an increase in prior for a threat state accelerates the effect of threat data. For extreme increases in prior for the active states the prior does dominate the evolution; but for the “Neutral” state the initial very high prior is subsequently outweighed by data. See Appendix E in the supplementary material for Figures.

7 Discussion

In this paper we have described a novel three level hierarchical model that utilises at its deepest level an RDCEG modelling the state of a suspect within the stages of a potential attack. We illustrated how such an analysis can synthesise information concerning that suspect, through sets of tasks to produce snap shot summaries of the likely position of this person and the current threat they might present.

Currently, working with various domain experts, we are in the process of constructing a suite of RDCEG templates and their associated tasks Smith and Shenvi (2018). These describe different criminal processes associated with assaults or violence against the general public, indexed by type of crime, that build on existing criminological models. This type of technology has already been well developed for Bayesian Networks (BNs) within the context of forensic science Aitken and Teroni (2004); Mortera and Dawid (2017) and has established frameworks of processes linking activities with evidence. The structure we use here helps in this development because only the top layer of the hierarchy usually needs regular refreshing: the possible positions and associated tasks are fairly stable over time. We hope that within this paper we have illustrated that, just as in forensic science, such methods are both promising and feasible. Indeed the harmonisation of this class of models to forensic analogues means that evidence applied within an investigation can be coherently integrated into case reports associated with criminal proceedings if the suspect does attempt to perpetrate a crime.

In the next phase of this programme, building on these models of individual suspects, we are developing a network model for the stochastic evolution of open populations of violent criminals. This issue is complicated by the fact that many suspects are working in teams and often coordinated. This dependence structure and communications between individuals therefore have to be carefully modelled for such models to be realistic. However this more challenging domain is also a potentially very fertile one – where standard estimation of hyperparameters associated with different units and the Bayesian selection of the most promising models can begin to be applied. The challenge is that

series are frequently disrupted so the systematic estimation of their hyperparameters is hard and needs strong prior information to be effective.

The eventual objective is for the hierarchical model for individuals described in this paper and the network model described in the paragraph above, to be incorporated into a decision theoretic framework that also models the action space and objectives of the authorities. This would aim to aid investigators decide which actions to take; such as which cases to prioritise, or to deprioritise to free up resources, which cases to increase surveillance on, and which cases to make which kind of interventions on; all in order to minimise the expectation of a multi-attribute loss function over incidents, casualties, public terror etc., given the constraints of resources, personal freedom, democratic legitimacy and proportionality.

Moreover the actions of the authorities in terms of preventative measures, defence of targets and pursuit methods will influence the actual decisions and trajectories of the suspects as is documented in case reports (see for example Gill (2012)). These recursive aspects introduce game-theoretic ideas into the models which complicate the inferential process and probability propagation. This work has to be done in conjunction with the authorities in order to properly represent a realistic, constrained action space, to elicit a realistic loss function, and to gain insight on the dynamic interplay between investigator and investigated.

Thus the work we present in this paper is the first phase of a longer programme working with practitioners that eventually aims to provide a theoretically valid, practically useful, and legally defensible system to support the prevention of acts of extreme radical violence.

Supplementary Material

Supplementary material for: A Bayesian Hierarchical Model for Criminal Investigations (DOI: [10.1214/19-BA1192SUPP](https://doi.org/10.1214/19-BA1192SUPP); .pdf).

References

- Aitken, C. and Teroni, F. (2004). *Statistics and the Evaluation of Evidence for Forensic Scientists*. Wiley. 26
- Allen, G. and Dempsey, N. (2018). “Terrorism in Great Britain: the statistics. Briefing Paper Number CBP7613.” Technical report, House of Commons. 1
- Barclay, L. M., Collazo, R. A., Smith, J. Q., Thwaites, P. A., Nicholson, A. E., et al. (2015). “The dynamic chain event graph.” *Electronic Journal of Statistics*, 9(2): 2130–2169. MR3400535. doi: <https://doi.org/10.1214/15-EJS1068>. 4
- Barclay, L. M., Hutton, J. L., and Smith, J. Q. (2013). “Refining a Bayesian Network using a Chain Event Graph.” *International Journal of Approximate Reasoning*, 54: 1300–1309. MR3115418. doi: <https://doi.org/10.1016/j.ijar.2013.05.006>. 4

- Barclay, L. M., Hutton, J. L., and Smith, J. Q. (2014). “Chain Event Graphs for Informed Missingness.” *Bayesian Analysis*, 9(1): 53–76. MR3188299. doi: <https://doi.org/10.1214/13-BA843>. 4, 9, 10
- Bartolucci, F., Pennoni, F., and Francis, B. (2007). “A Latent Markov Model for Detecting Patterns of Criminal Activity.” *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 170(1): 115–132. URL <http://www.jstor.org/stable/4623137>. MR2339101. doi: <https://doi.org/10.1111/j.1467-985X.2006.00440.x>. 2
- Bouhana, N., Thornton, A., Corner, E., Malthaner, S., Lindekilde, L., B., S., and Perry, G. (2016). “EU FP7 PRIME risk analysis framework.” Technical report, Public version. 1
- Bouhana, N. and Wikstrom, P.-O. H. (2011). “Al Qai‘da-influenced radicalisation: A rapid evidence assessment guided by Situational Action Theory.” Rds occasional paper 97, London, England: Home Office Research, Development and Statistics Directorate. 1
- Bunnin, F. O. and Smith, J. Q. (2019). “Supplementary material for: A Bayesian Hierarchical Model for Criminal Investigations.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/19-BA1192SUPP>. 17
- Clark, N. J. and Dixon, P. M. (2019). “Modeling and Estimation for Self-Exciting Spatio-Temporal Models of Terrorist Activity.” *The Annals of Applied Statistics*, 12(1): 633–653. MR3773408. doi: <https://doi.org/10.1214/17-AOAS1112>. 2
- Collazo, R. A., Goergen, C., and Smith, J. Q. (2018). *Chain Event Graphs*. CRC Press. MR3752634. 4, 7, 15
- Collazo, R. A. and Smith, J. Q. (2016). “A New Family of Non-Local Priors for Chain Event Graph Model Selection.” *Bayesian Analysis*, 11(4): 1165–1201. MR3577376. doi: <https://doi.org/10.1214/15-BA981>. 4
- Corner, E., Bouhana, N., and Gill, P. (2019). “The multifinality of vulnerability indicators in lone-actor terrorism.” *Psychology, Crime and Law*, 25(2): 111–132. 1
- Cowell, R. G., Smith, J. Q., et al. (2014). “Causal discovery through MAP selection of stratified chain event graphs.” *Electronic Journal of Statistics*, 8(1): 965–997. MR3263109. doi: <https://doi.org/10.1214/14-EJS917>. 4
- Europol (2018). “Terrorism Situation And Trend Report.” Technical report, European Union Agency for Law Enforcement Cooperation. 1
- Ferrara, E., Wang, W., Varol, O., Flammini, A., and Galstyan, A. (2016). “Predicting online extremism, content adopters, and interaction reciprocity.” *CoRR*, arXiv:1605.00659. 3
- Fienberg, S. E. and Shmueli, G. (2005). “Statistical issues and challenges associated with rapid detection of bio-terrorist attacks.” *Statistics in Medicine*, 12(24): 513–529. MR2134521. doi: <https://doi.org/10.1002/sim.2032>. 2

- Frühwirth-Schnatter, S. (2008). *Finite mixture and Markov switching models*. Springer. MR2265601. 11, 17
- Gill, P. (2012). “Tracing the Motivations and Antecedent Behaviors of Lone-Actor Terrorism A Routine Activity Analysis of Five Lone-Actor Terrorist Events.” Technical report, International Center for the Study of Terrorism, The Pennsylvania State University. 1, 14, 27
- Görge, C. and Smith, J. Q. (2018). “Equivalence classes of staged trees.” *Bernoulli*, 24(4A): 2676–2692. MR3779698. doi: <https://doi.org/10.3150/17-BEJ940>. 4
- Horgan, J. (2005). *The Psychology of Terrorism*. Abingdon: Routledge, first edition. 1
- Howard, R. A. (1988). “Decision Analysis: Practice and Promise.” *Management Science*, 34(6): 679–695. URL <https://doi.org/10.1287/mnsc.34.6.679>. 14
- Korb, K. B. and Nicholson, A. E. (2010). *Bayesian Artificial Intelligence, Second Edition*. Boca Raton, FL, USA: CRC Press, Inc., 2nd edition. MR3100449. 13
- Lindekilde, L., O’Connor, F., and Schuurman, B. (2019). “Radicalization patterns and modes of attack planning and preparation among lone-actor terrorists: an exploratory analysis.” *Behavioral Sciences of Terrorism and Political Aggression*, 11(2): 113–133. 1
- Mortera, J. and Dawid, A. P. (2017). “Graphical Models For Forensic Analysis.” Departmental Working Papers of Economics – University ‘Roma Tre’ 0224, Department of Economics – University Roma Tre. URL <https://ideas.repec.org/p/rtr/wpaper/0224.html>. MR3888008. 26
- O’Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E., and Rakow, T. (2006). *Uncertain Judgements: Eliciting Expert Probabilities*. Chichester: John Wiley. URL <http://oro.open.ac.uk/17948/>. 7, 15
- Pat-Cornell, E. and Guikema, S. (2002). “Probabilistic Modeling of Terrorist Threats: A Systems Analysis Approach to Setting Priorities Among Countermeasures.” *Military Operations Research*, 7: 5–23. 2
- Phillips, L. D. (1984). “A theory of requisite decision models.” *Acta Psychologica*, 56(1): 29–48. URL <http://www.sciencedirect.com/science/article/pii/0001691884900052>. 7, 15
- Python, A., Illian, J. B., Jones-Todd, C. M., and Blangiardo, M. (2019). “A Bayesian approach to modelling subnational spatial dynamics of worldwide non-state terrorism, 2010–2016.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(1): 323–344. MR3902646. 2
- Ranciati, S., Vinciotti, V., and Wit, E. C. (2017). “Identifying overlapping terrorist cells from the Noordin Top actor-event network.” *arXiv e-prints*, arXiv:1710.10319. 2
- Regens, J. L., Mould, N., Jensen, C. J., Graves, M. A., and Edger, D. N. (2015). “Probabilistic Graphical Modeling of Terrorism Threat Recognition Using Bayesian Networks and Monte Carlo Simulation.” *Journal of Cognitive Engineering and Decision Making*, 9(4): 295–311. 2

- Rubin, D. B. (1976). “Inference and missing data.” *Biometrika*, 63(3): 581–592. MR0455196. doi: <https://doi.org/10.1093/biomet/63.3.581>. 9
- Shenvi, A. and Smith, J. Q. (2018). “The Reduced Dynamic Chain Event Graph.” *arXiv e-prints*, arXiv:1811.08872. 3, 4, 17
- Shenvi, A. and Smith, J. Q. (2019). “A Bayesian Dynamic Graphical Model for Recurrent Events in Public Health.” Working paper. 4, 5, 25
- Smith, J. Q. (2010). *Bayesian decision analysis: principles and practice*. Cambridge, UK; New York: Cambridge University Press. URL <http://wrap.warwick.ac.uk/45002/>. MR2828346. doi: <https://doi.org/10.1017/CB09780511779237>. 7, 12, 14, 15, 16
- Smith, J. Q. and Shenvi, A. (2018). “Assault Crime Dynamic Chain Event Graphs.” Working paper. URL <http://wrap.warwick.ac.uk/104824/>. 4, 14, 15, 16, 26
- West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models (2nd Ed.)*. Berlin, Heidelberg: Springer-Verlag. MR1482232. 3, 9, 10
- Wilkerson, R. L. and Smith, J. Q. (2018). “Customised Structural Elicitation.” *arXiv e-prints*, arXiv:1807.03693. 7, 13

Acknowledgments

This work was funded by The Alan Turing Institute Defence and Security Project G027. We are grateful for the comments of two anonymous referees and the associate editor which greatly improved this paper.