Journal of the Royal Statistical Society Series A: Statistics in Society, 2022, **00**, 1–19
https://doi.org/10.1093/jrsssa/qnac019

ROYAL STATISTICAL SOCIETY DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series A
Statistics in Society

Original Article

A Bayesian decision support system for counteracting activities of terrorist groups

Aditi Shenvi[1], Francis Oliver Bunnin[2] and Jim Q. Smith[1,3]

[1]Department of Statistics, University of Warwick, Coventry, UK
[2]Natwest Markets, London, UK
[3]The Alan Turing Institute, London, UK

Address for correspondence: A. Shenvi, Department of Statistics, University of Warwick, Coventry CV4 7AL, UK.
Email: aditi.shenvi@warwick.ac.uk

Abstract

We present an integrating decision support system designed to aid security analysts' monitoring of terrorist groups. The system comprises of (i) a dynamic network model of the level of bilateral communications between individuals and (ii) dynamic graphical models of those individual's latent threat states. These component models are combined in a statistically coherent manner to provide measures of the imminence of an attack by the terrorist group. Domain knowledge provides the structures of the models, values of parameters and prior distributions over latent variables. Inference of the values is performed using time-series of observed data and the statistical dependencies assumed between said data and model variables. The work draws on social network and graphical models used in sociological, military, and medical fields.

Keywords: Bayesian hierarchical models, counterterrorism, dynamic weighted network models, graphical models, integrating decision support systems, multiregression dynamic models

1 Introduction

The nature of terrorism in the UK has changed over the last two decades. Rather than the large-scale, hierarchically organised terrorist attacks executed by the likes of the Irish Republican Army (IRA) and Al Qaeda, recent years have increasingly seen attacks by independent individual actors and small groups of individuals. In contrast to the established terrorist organisations that prevailed in the twentieth to early twenty-first century, these groups of attackers do not receive orders from terrorist leaders. Often, these groups decide to attack independently after having being radicalised through social media or local inciters of violence. Their use of easily obtainable weapons, such as knives, vehicles, or improvised explosive devices, and the soft nature of their targets, facilitates rapid progression from ideation to planning and execution. The authorities' recent focus has, consequently, been on small terrorist groups (Home Office, 2018; Pantucci, 2016). The UK's strategy for counterterrorism (Home Office, 2018) outlines its framework, built on the four 'P' work strands, Prevent, Pursue, Protect, and Prepare.

The objectives of the 'Pursue' work strand are to detect, understand, investigate, and disrupt terrorist attacks. Counterterrorism authorities can closely monitor the activities of suspected terrorist groups to prevent attacks. Whilst groups intent on performing an act of terrorism may attempt to hide or disguise their intentions and activities to avoid detection and scrutiny, they still need to perform certain preparatory tasks and to communicate in order to plan, organise, and execute a joint terrorist attack. Security analysts monitoring suspects can capture fragmentary data of these activities and communications. The combination of data and domain expertise facilitates inference of the plans and threat state of potential attackers; and from such inference security, analysts make

Received: December 16, 2021. **Revised:** November 21, 2022. **Accepted:** December 12, 2022

decisions on further investigation or interventions. Currently, this process of combining data and expertise is predominantly qualitative, informal, and based on individual analysts' experiences.

The objectives of this line of research are to (i) develop support tools targeted at the current nature of domestic terrorism and (ii) systematically analyse the large and heterogeneous data available to the authorities. In this paper, we present one such tool: a Bayesian integrating decision support system (IDSS) designed to aid monitoring of the threat presented by known or suspected terrorist groups.

1.1 Related work

The statistical use of network data for intelligence use has a long history; going back to at least World War II when *Traffic Analysis* was used by the allied forces. This was defined as 'the study of the external characteristics of signal communications' used for 'drawing deductions and inferences of value as intelligence even in the absence of specific knowledge of the contents' (Cunningham et al., 2015; Departments of the Army and the Air Force, 1948; van Meter, 2002). Since then the statistical aspects of terrorist networks have been researched extensively utilising diverse methods.

Sparrow (1991) emphasised key issues such as 'weak ties' indicating that the most valuable and urgent communication channels are likely to be those that are seldom used, 'fuzzy boundaries' indicating that boundaries of such networks can be quite ambiguous, and 'incompleteness' indicating that data relating to these networks are likely to be incomplete with informative missingness. Toth et al. (2013) used centrality measures to identify key individuals and heterogeneous roles, Ranciati et al. (2020) use Bayesian bipartite graph methods to identify overlapping cells, and multipartite graph methods were used by Campedelli et al. (2019) to cluster similar terrorist groups. For a detailed review of social network analysis of activities of opposition and terrorist forces, see Section 6.4 of Shenvi (2021) and references therein.

The closest research to our own is the adaptive safety analysis and monitoring (ASAM) tool (Allanach et al., 2004; Singh et al., 2004). ASAM is a hierarchical system consisting of a collection of hidden Markov models where each hidden Markov model monitors the probability of a specific type of terrorist attack. ASAM extracts signals of suspicious activities from noisy and partial data to evaluate the overall probability of a terrorist attack.

Our research complements and differs from the majority of terrorist network existing research in four areas: (i) the objective is a real-time support system using incoming data, rather than ex post analyses of data sets; (ii) our synthetic data are based on the actual types of data available to and used by security operatives, (iii) the system is designed and structured to incorporate expert judgement at every level; and (iv) the calculations are in closed form which facilitates speed and, critically in this domain, transparency in how the measures were derived from data and assumptions.

The paper is structured as follows: Section 2 describes the types of data the authorities typically have access to. Section 3 introduces the IDSS: its network and graphical model representing the communications and threat states within a terrorist group, respectively; the assumptions pertaining; and the mode of inference from observable data to latent communication and threat state processes. Section 3.4 demonstrates how indicators of group threat are derived from the IDSS. Section 4 illustrates the model workings using a hypothetical example with synthetic data. We conclude in Section 5 with a discussion of avenues of further research in this and related areas.

2 The nature of data available to the authorities

The data that security analysts use in their investigations are multifarious. They arise from a variety of sources and are heterogeneous, open-ended, and partial. Formats vary from oral reports, handwritten text, physical sightings, to streaming electronic audio and visual data. A non-exhaustive list includes historical police records; information from informants, friends, family; public sighting at events and with other individuals; and public social media posts. The challenge is to use such a wide variety of data in a principled, systematic, and intelligent manner to complement existing expertise. The ever increasing abundance of electronic data made available by new technologies gives increased opportunities for information gathering. As manpower is limited, interest has grown in how to use new technologies to extract intelligence from such large and heterogeneous datasets.

Data on individuals include contents and meta-data from unstructured personal data such as police records, contact with the UK prevent programme (Home Office, 2018), social media posts, internet browsing history, physical or electronic observations, and/or reports of activities, movements, and financial transactions. These types of data are typical of those available to the police and security services (Brain, 2015; R v Ziamani, 2015; Radio4 BBC, 2019).

If the authorities have credible reasons to believe that an individual presents a potential threat, then that individual may be actively investigated, up to and including surveillance. 'Surveillance, …, includes monitoring, observing or listening to persons, their movements, conversations or other activities and communications' (Home Office, 2020). The Regulation of Investigatory Powers Act 2000 governs the methods of information gathering. A warrant to gather information requires permission from the Secretary of State and an independent senior judge. They approve the warrant only when they are convinced that it is necessary and proportionate (Chorley, 2021; UK Public General Acts 2000 c. 23, 2000).

Surveillance under a warrant hence permits interception and continued observation of personal data. These can include a person's mobile phone and private social media communications, their internet browsing activity, and observations of their physical movements and activities. For multiple reasons, including technical limitations, hiding, and disguising, any such observed data are incomplete and the real intent and actions of observees must be inferred. The heterogeneous and fragmentary nature of such data is an obstacle to traditional statistical or AI data processing. Nevertheless, as MI5 director Ken McCallum stated (Chorley, 2021) 'difficult decisions [have to be made] based on fragmentary information'.

The first indication that distinct individuals, that are already under suspicion, may be connected is any form of link between them. There are multiple types of data which may be used to infer ties between suspected terrorists. Here we enumerate five that are known to be used by analysts.

(a) Existing kinship or social links.
(b) Work or other shared affiliations.
(c) Physical meetings (observed directly or through closed circuit television).
(d) Financial connections (e.g., shared accounts, shared asset ownership, bank transfers between accounts).
(e) Bilateral electronic communications (e.g., telephone, email, Whatsapp. etc.).

None of these data necessarily indicate ties that are malign in nature, but in the given context of reasonable suspicion, they may be driven by malicious intent. Moreover pre-existing ties, such as items (a) and (b), facilitate collaboration once other factors have come into play. Concrete real examples of such ties include:

(a) The kinship tie between Saleem and Hashem Abedi—the former being the suicide bomber of the 22 May 2017 Manchester Arena bombing and the latter his brother who was found guilty of aiding Saleem (Parveen & Walker, 2020).
(b) The London Bridge attackers Khuram Butt, Rachid Redouane, and Youssef Zaghba's membership of the Ummah Fitness Centre which was also frequented by the son of Anjem Choudary, a convicted pro-terrorist proselytiser with links to ISIL the Islamic State of Iraq and the Levant (Gardham & Gibbons, 2019).
(c) The sightings of Brusthom Ziamani with members of the proscribed militant group al-Muhajiroun. These occurred both at political demonstrations and at the members' flat into which Ziamani re-located and which was subsequently raided by the police for non-terrorism-related reasons (Counter Extremism Project, 2022).
(d) The use of data from Society for Worldwide Interbank Financial Telecommunication (SWIFT) by the United States Treasury Terrorist Financial Tracking Programme that 'enhances [their] ability to map out terrorist networks, often filling in missing links in an investigative chain' (United States Treasury Department, 2022);

(e) Findings of case study A8/1 in the UK Bulk Powers Review (Anderson, 2016) that the use of new phone numbers 'used by individuals known to be involved in plotting terrorist acts in the UK… [and] further analysis to identify contacts… [explained that] each phone would not necessarily have been identified as suspicious but, when taken as a network, the likely operational nature of the phones was clear to see'. Similarly, see also case study A8/2.

In dealing with communications data, it is important to differentiate between (i) the content of such communications and (ii) 'secondary data', i.e., meta-data such as the identities of parties and the timing, location, and duration of communications. Often secondary data are available whilst content data are unavailable due to either encryption or limits prescribed by certain interception warrants. Secondary data without content data have nevertheless proven to be extremely useful: '*secondary data* can enable the tracing of contacts, associations, habits and preferences' (Anderson, 2016). Contents data, being in the most part, unstructured, feed into the model through the analysts' choice of priors and the ability to override intermediate model values based on external information. In our model, both the existence of ties and the weight of the ties, representing the latent level of actual bilateral communication, are inferred in the first instance by analysts' a priori knowledge, and subsequently through surveillance data.

3 Model: two components, one decision support system

3.1 IDSSs

An IDSS is a Bayesian unifying and integrating framework that combines component IDSSs—each supporting decision-making about a distinct aspect of a complex system—into a single entity (see Barons et al., 2021; Leonelli & Smith, 2015; Smith et al., 2015 and references therein). The transparent and statistically grounded framework of the IDSS enables a statistician to formally incorporate the judgements and uncertainties of the domain experts and decision-makers. Available data are then fed through the relevant components of the IDSS with full consideration of these judgements and uncertainties. The outputs of all the components are then combined—in a manner that is appropriate for the application—to enable the decision-makers to fully evaluate the effects of any potential policies on their outcomes of interest.

In our IDSS below, we denote the open population of *persons of interest* (POIs) by $\mathcal{P}$, individuals within that population as $p_i \in \mathcal{P}$, and the latent threat state of the individuals p_i as X_i. The latent level of communication between p_i and p_j is the variable φ_{ij} and the observed data informing φ_{ij} are denoted by $\mathbf{s}_{ij}$. $\mathbf{Y}_i$ is the observed data on p_i and ϑ_i are the activities of p_i. These form the main variables of interest of the IDSS, along with parameters and intermediate variables to be introduced in subsequent sections. The flow of inference is from the data $\mathbf{s}_{ij}$, $\mathbf{Y}_i$ to the latent variables φ_{ij}, X_i for each p_i and pair p_i, p_j. The inference from $\mathbf{Y}_i$ to X_i is through the intermediate variables ϑ_i. The flow of causality is the reverse. From the inferred probability distributions of φ_{ij}, X_i indicators of the imminence of an attack are derived. These indicators are denoted by $\{\Lambda_i\}$, $i \in \{0, \ldots, 3\}$.

The primary assumptions of the IDSS are that (i) the two components can be decoupled through the technique of multiregression dynamic models (Queen & Smith, 1993) which is detailed in Online Supplementary Material, Appendix A; (ii) φ_{ijt}, the latent variable φ_{ij} through time, is a Markov process; (iii) the threat state X_{it} through time is a semi-Markov process, and (iv) data on individuals Y_i are conditionally independent of X_i given the activities performed by p_i; i.e., observed data are only of use in inferring the activities ϑ_i of p_i which in turn are used for inference of X_i. The mathematical formulations of these assumptions are provided in the expositions in the remainder of this section and in Online Supplementary Material, Appendices A and B.

3.2 IDSS for activities of terrorist groups

Our IDSS models a terrorist group being monitored by security analysts. The IDSS contains two components: a network model of the level of communications between individuals of the group and a hierarchical graphical model of the threat state of those individuals. The structure of the IDSS is given in Figures 1 and 3.

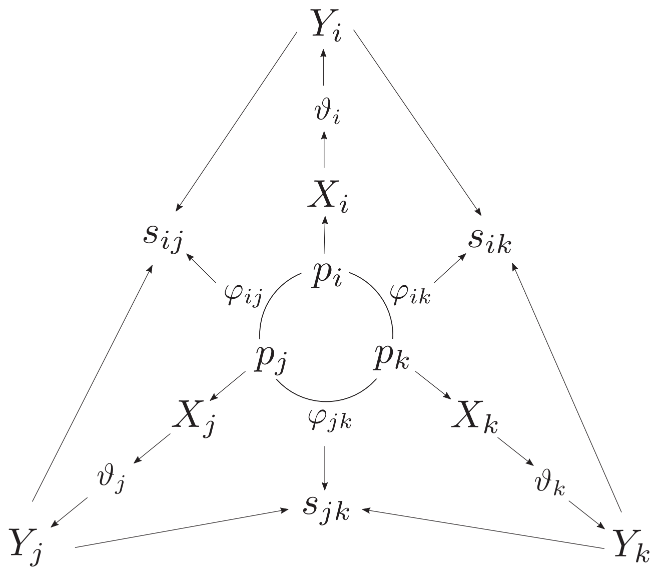

Figure 1. The graph of the overall system: the network of individuals p, edge weights φ, communications data s, threat state X, task vector ϑ, and data on individuals Y. The directed edges indicate the flow of causation, e.g., a particular threat state, say preparing, causes certain activities to be done, which in turn give rise to certain observable data. Inference can be performed in the reverse direction: from data to activities to threat state.

Component 1: the terrorist network model

The terrorist network is an undirected, dynamic, and weighted network model. At each time $t \geq 0$, the network vertices consist of the POIs $\mathcal{P}_t$ whom the counterterrorism authorities choose to observe at time t; and the network edges indicate known or potential ties between these POIs.

Ties between the suspects are informed by observable data as described in Section 2. During each time period, new leads are discovered. From among these leads, new investigative cases are opened for those that pass a triage process (see, for example, details in Radio4 BBC, 2019). The triage process thus gives rise to a set of newly identified individuals $\mathcal{P}_t^+$ at each time interval t. Over the same interval, a set $\mathcal{P}_t^-$ are lost from $\mathcal{P}_t$ for a variety of reasons such as arrest or evidence of innocence. For simplicity, assume that $\mathcal{P}_t^+$ joins the set $\mathcal{P}_t$ at the start of the time period t and existing individuals $\mathcal{P}_t^-$ are lost at the end of t. We then have that

$$\mathcal{P}_t = \{\mathcal{P}_{t-1} \backslash \mathcal{P}_{t-1}^-\} \cup \mathcal{P}_t^+ \tag{1}$$

An undirected network $\mathcal{N}_t = (V(\mathcal{N}_t), E(\mathcal{N}_t))$ is then created at each time t where $V(\mathcal{N}_t) = \mathcal{P}_t$ are the vertices and $E(\mathcal{N}_t)$ are the edges of the network. An edge $e_{ij} \in E(\mathcal{N}_t)$ exists between two individuals p_i and p_j if there is a tie between them. Once an edge is created in $\mathcal{N}_t$ between some $p_i, p_j \in \mathcal{P}_t$, this edge endures for all $\mathcal{N}_t$' where $t' \geq t$ as long as $p_i, p_j \in \mathcal{P}_t$'. Denote by φ_{ijt} the latent random variable measuring the pairwise communications shared between p_i and p_j at time t. Thus, φ_{ijt} acts as a quantitative measure of the information directly exchanged between p_i and p_j and models the edge weight on the edge e_{ij} in $\mathcal{N}_t$. Denote by Φ_t a $|\mathcal{P}_t| \times |\mathcal{P}_t|$ symmetric matrix with its (i, j)th entry given by φ_{ijt}. By convention, we set $\varphi_{ijt} = 0$ if $i = j$ or $e_{ij} \notin E(\mathcal{N}_t)$ for $i \neq j$. The observable pairwise communications data are used to estimate φ_{ijt}. Note that the granularity of the time steps (e.g., hourly, daily, weekly) is chosen to suit the observation process. Online Supplementary Material, Appendix C, gives an illustration of network creation for a simple hypothesised terrorist group.

The counterterrorism authorities are likely to receive data and information from multiple sources. Suppose that there are K such *information channels*. The data from each channel are condensed into a summary measure in the terrorist network. The summary measure used for each channel depends on factors such as the type of data, the data source, the frequency of the observations, and the required granularity of that specific type of data. For instance, for an information channel informing the duration of phone calls or number of text messages exchanged between a pair of suspects, the summary measure may simply be the sum of the observations, whereas for bank transfers between the suspects, the amount of money exchanged might be a suitable summary measure. However, note that these summary measures for the different information channels may be on very different scales of measurement, e.g., x hours of a phone call and £x of money exchanged, and hence, might have a disproportionate effect on the edge weight variable φ_{ijt}. To balance the effect of data relating to different channels on φ_{ijt}, the data obtained through the different channels must be on a comparable scale. This can be achieved through any of the standard methods of scaling or standardisation (see, e.g., Jahan & Edwards, 2015).

Denote by s_{ijkt} the scaled or standardised summary measure of the data observed between the pair $p_i, p_j \in \mathcal{P}_t$ from channel k at time t. We assume that the following independence relationship holds:

$$\perp\!\!\!\perp_{k \in \{1,\dots,K\}} s_{ijkt} \tag{2}$$

which implies that the data and information obtained from the different information channels for a given pair $\{p_i, p_j\} \in \mathcal{P}_t \times \mathcal{P}_t$ at time t are mutually independent. This simplifying assumption is conservative and guided by the supporting role in counterterrorism intended for our models: it enables us to ensure that the inference is tractable and can be performed in real time. To account for any correlation in data from the information channels, a multivariate Gamma–Poisson mixture setting can be considered (Andreassen, 2013; Choo & Walker, 2008) instead of the Gamma–Poisson setting introduced below; although inference will be analytically intractable. Denote by S_t the $|\mathcal{P}_t| \times |\mathcal{P}_t|$ *observations array* at time t whose elements are the vectors $\mathbf{s}_{ijt}$ such that $\mathbf{s}_{ijt} = \{s_{ij1t}, \dots, s_{ijKt}\}$. Notice that S_t is symmetric with $\mathbf{s}_{ijt} = \mathbf{s}_{jit}$ due to the nature of the pairwise communications data. We use the convention that $\mathbf{s}_{ijt}$ is a K-dimensional zero vector whenever (i) $i = j$, (ii) $e_{ij} \notin E(\mathcal{N}_t)$, or (iii) whenever no information is observed between two individuals. To indicate the difference in the quality or reliability of data obtained from the different channels, we define a parameter $\xi_k \in (0, 2]$ which denotes the *efficiency* of the intelligence obtained from channel k, for $k = 1, \dots, K$. This efficiency parameter indicates the loss of information expected from a specific information channel. A value closer to 2 represents minimal loss of information (e.g., bank transactions data), whereas a value closer to 0 indicates that the actual observations are likely to be much higher than what has been conveyed to the authorities (e.g., patchy or poor source of secondary data). See Online Supplementary Material, Appendix E, for an illustration of how to scale the observation data and how to set the efficiency parameters.

In order to maintain transparency in the model, interpretability of its parameters, and to enable quick and efficient inference, we use a Gamma–Poisson conjugate setting for updating the distributions of the φ_{ijt}—the random variables modelling the edge weights, for $p_i, p_j \in \mathcal{P}_t$ and $t \geq 0$. Note that unlike the Poisson distribution, Gamma–Poisson compound distributions—which are equivalent to negative binomial distribution—can handle overdispersed data (Schein et al., 2016). We adopt the approach of using discount factors to transform the posterior at time t into the prior at time $t + 1$ as described in West and Harrison (1997) and Smith (1979). The discount factor δ_t is a value in $(0, 1]$ that represents the decay of information from time $t - 1$ to time t. We take φ_{ijt} to be a Markov process and assume that s_{ijt}, the observed communications between p_i and p_j at time t, depends only on φ_{ijt}, the latent edge weight which represents the actual (unobserved) level of communication between p_i and p_j at time t. This is formally stated as

$$\varphi_{ijt} \perp\!\!\!\perp \mathcal{F}_t \mid \varphi_{ij,t-1} \tag{3}$$

$$s_{ijt} \perp\!\!\!\perp (\Phi_t, S_t, \mathcal{F}_t) \mid \varphi_{ijt} \tag{4}$$

where $\mathcal{F}_t$ denotes all past data and edge weight random variables up to but not including time t, i.e., $S_{t'}$ and $\Phi_{t'}$ for $t' < t$. This formalisation enables us to update the edge weight variables φ_{ijt} using observational data $\mathbf{s}_{ijt}$ for each pair p_i and p_j independently, see Online Supplementary Material, Appendix D.

In practice, the efficiency parameters ξ_k for the $k = 1, \ldots, K$ information channels would be determined a priori by the authorities in collaboration with channel-specific experts. The discount factors δ_{ijt} would be set by the authorities based on empirical evidence or estimated from the dataset. One method of estimation of discount factors is the grid search approach used in Barons et al. (2021). Furthermore, both the efficiency parameters and the discount factors can be adjusted at any time to reflect changes in the quality of the incoming data and the rate of decay of past information, respectively.

We now describe the forward filtering equations for each pair $\{p_i, p_j\} \in \mathcal{P}_t \times \mathcal{P}_t$ in the terrorist network:

Initialisation. Set the prior φ_{ijt_0} as follows:

$$\varphi_{ijt_0} \sim \text{Gamma}(\alpha_{ijt_0}, \beta_{ijt_0}) \tag{5}$$

where t_0 is the first time step of the time-series. The parameters α_{ijt_0} and β_{ijt_0} are determined by existing case knowledge. For example, if $e_{ij} \in E(\mathcal{N}_{t_0})$ exists only due to a social relation, then α_{ijt_0} and β_{ijt_0} may be set such that the mean and variance of φ_{ijt_0} are both relatively low. On the other hand, if $\mathcal{P}_i$ and $\mathcal{P}_j$ have a previous joint conviction, then these parameters can be set such that the φ_{ijt_0} has a high mean and lower variance.

Posterior at time $t-1$. Let the posterior of $\varphi_{ij,t-1}$ after observing $\mathbf{s}_{ij,t-1}$ and $\mathcal{F}_{t-1}$ be given by

$$\varphi_{ij,t-1} \mid \mathbf{s}_{ij,t-1}, \mathcal{F}_{t-1} \sim \text{Gamma}(\alpha_{ij,t-1}, \beta_{ij,t-1}) \tag{6}$$

Prior at time t. Using the discount factor $\delta_{ijt} \in (0, 1]$, the posterior at time $t-1$ evolves to the prior at time t as

$$\varphi_{ijt} \mid \mathcal{F}_t \sim \text{Gamma}(\delta_{ijt}\alpha_{ij,t-1}, \delta_{ijt}\beta_{ij,t-1}) \tag{7}$$

Under this posterior-to-prior evolution, the mean of the distribution remains unaffected while the variance either remains the same (when $\delta_{ijt} = 1$) or increases (when $0 < \delta_{ijt} < 1$). Thus, a lower value of δ_{ijt} indicates a reduced confidence in the posterior at the previous time step as the variance increases. This is also associated with a decay of information from the previous time step depending on how much the situation is likely to have evolved since then.

Data generation at time t. The observations from the different information channels are modelled independently as

$$s_{ijkt} \mid \varphi_{ijt}, \mathcal{F}_t \sim \text{Poisson}(\xi_k \varphi_{ijt}), \quad k = 1, \ldots, K \tag{8}$$

Posterior at time t. The posterior when the observation vector $\mathbf{s}_{ijt}$ has at least one non-zero element is given by

$$\begin{aligned} p(\varphi_{ijt} \mid \mathbf{s}_{ijt}, \mathcal{F}_t) &\propto \prod_{k=1}^{K} p(s_{ijkt} \mid \varphi_{ijt}, \mathcal{F}_t) p(\varphi_{ijt} \mid \mathcal{F}_t) \\ &= \varphi_{ijt}^{\sum_k s_{ijkt} + \delta_{ijt}\alpha_{ij,t-1} - 1} \exp\left(-\left(\sum_k \xi_k + \delta_{ijt}\beta_{ij,t-1}\right)\varphi_{ijt}\right) \\ \varphi_{ijt} \mid \mathbf{s}_{ijt}, \mathcal{F}_t &\sim \text{Gamma}(\alpha_{ijt}, \beta_{ijt}) \end{aligned} \tag{9}$$

where $\alpha_{ijt} = \delta_{ijt}\alpha_{ij,t-1} + \sum_k s_{ijkt}$ and $\beta_{ijt} = \delta_{ijt}\beta_{ij,t-1} + \sum_k \xi_k$. For the same value of $\sum_k s_{ijkt}$, a lower overall efficiency of the observations given by $\sum_k \xi_k$ results in a higher mean and larger variance—indicating the associated increase in uncertainty—of φ_{ijt} compared to when the overall efficiency is higher.

The distribution of φ_{ijt} for a pair $\{p_i, p_j\}$ can hence be periodically updated over the evolution of time t in closed form using the above recurrences across the terrorist network given sequential incoming observational data. See Online Supplementary Material, Appendices D and F, for more details and an illustration. The dynamic nature of the open population is easily incorporated in our model by introducing vertices, edges, and priors for immigrants (new entrants) and removing them for emigrants (leavers) at the appropriate time. Finally, we note here that in a policing and counterterrorism setting, it is essential to differentiate between the following cases:

1. $\sum_k s_{ijkt} = 0$ because p_i and p_j were monitored but did not communicate in any way during time t;
2. $\sum_k s_{ijkt} = 0$ because p_i and p_j were not closely monitored during time t.

In the first case, the posterior update is carried out as described above as we have *observed* zero communications. Whereas in the second case, we do not update the posterior. Thus, the posterior mean at time t is the same as the posterior mean at time $t-1$ and the posterior variance is further diffused from time $t-1$ to t. Notice that if no new information is observed through $\mathbf{s}_{ijs}$, $s \geq t$ then the variance of φ_{ijs}, $s \geq t$ will keep increasing. To prevent this and to reflect that we expect a baseline amount of information flow to continue between a pair of suspects p_i and p_j who share an edge between them—until we observe information indicating otherwise—we can set the discount factor as $\delta_{ijt} = d_{ij} + (1 - d_{ij}) \exp(-\sum_k s_{ijk,t-1}\xi_k)$ as detailed in Chen et al. (2018). Here d_{ij} is the baseline discount factor for pair $\{p_i, p_j\}$. This is particularly useful if we expect to have large consecutive gaps of time when we do not expect to observe good quality data on the pairs. When we observe very low levels of quality information in the previous time, the discount factor is closer to 1 and when good quality information is observed, the discount factor will be closer to d_{ij}. This setting allows us to set pair-specific discount factors if required.

Component 2: the individual terrorist model

The individual terrorist model is a hierarchical Bayesian graphical model of an individual $p_i \in \mathcal{P}$ introduced in Bunnin and Smith (2021). A concise bottom-up description of the three levels of the individual terrorist model for a suspect $p \in \mathcal{P}_t$ is given below.

Latent level. This level consists of a discrete time graphical model (more precisely, a reduced dynamic chain event graph, see Bunnin & Smith, 2021; Shenvi, 2021).

The vertices of the graph are *threat states*. These represent an individual POI's current stage of progress towards a potential terrorist attack. Smith and Shenvi (2018) provide several types of categorisations for a wide range of criminal behaviours which can be used to inform the threat states of the graphical model. Alternatively, these states can be more generically defined (e.g., 'Mobilised', 'Preparing', 'Training', and 'Active/Threatening'). In both cases, the model also includes a 'Neutral' state. This is an absorbing state representing that the suspect no longer presents a threat to the general public. Denote by X_t the latent random variable indicating the threat state occupied by a suspect p at time $t \geq 0$. The sample space of X_t is given by the vertices $\{x_0, x_1, \ldots, x_n\}$ of the graph. Let $\boldsymbol{\pi}_t = \{\pi_{t0}, \pi_{t1}, \ldots, \pi_{tn}\}$ where π_{ti} indicates the probability of the suspects being in threat state x_i at time t for $i \in \{0, 1, \ldots, n\}$. We assume X_t follows a semi-Markov process over the graph whose dynamics are determined by its semi-Markov transition matrix which defines the transition probabilities and the holding time distributions (Çinlar, 1975). See Online Supplementary Material, Appendix B.1, Equation (3) for more details.

At any time t, X_t occupies exactly one of the vertices of the graph. The probability vector $\boldsymbol{\pi}_t$ over the vertices represents the security analysts' level of uncertainty over the actual position of X_t; that is, it represents their *imperfect* knowledge: the actual position can only be known with *perfect* knowledge.

Intermediate level. At this level, we define a collection of R tasks associated with the threat states of the graphical model. Each task is an activity that enables progression along the vertices of the graph. At any time $t \geq 0$, denote the task vector by $\boldsymbol{\vartheta}_t = \{\vartheta_{t1}, \vartheta_{t2}, \ldots, \vartheta_{tR}\}$ where each ϑ_{tj} is an indicator variable such that $\vartheta_{tj} = 1$ if p is enacting task j at time t, for $j \in \{1, 2, \ldots, R\}$. Each task can be associated with one or more threat states of the graphical model. The purpose of the task vector is to enable the counterterrorism authorities to estimate how far along the suspect is in their progression towards a specified or unspecified terrorist attack.

Surface level. This level consists of the observable data $\{\mathbf{Y}_t\}_{t\geq 0}$ relating to task activities of the suspect p. For each task ϑ_{tj} in the intermediate level, we can associate a subset $Y_{tj} \subseteq \mathbf{Y}_t$ of the data stream observed which informs whether p is engaged in task ϑ_{tj}, for $j \in \{1, 2, \ldots, R\}$ at time t. If the data are noisy, a *filter function* (i.e., any suitable function $\tau_j(\cdot)$ of the data Y_{tj}; see Bunnin & Smith, 2021) may be used to obtain some viable scalar signal Z_{tj} from the noisy data subset Y_{tj}. Denote the vector of signals $(Z_{t1}, Z_{t2}, \ldots, Z_{tR})$ at time t by $\mathbf{Z}_t$.

Inference. The inferential recurrences associated with the progressions in the individual terrorist model are described in Online Supplementary Material, Appendix B.1, Equations 3–5. At each time $t \geq 0$, the model takes observed data on the individual to infer which tasks said individual is doing or has completed. The probabilities over tasks are used to infer probabilities over threat states: i.e., the output is the posterior probability vector $\boldsymbol{\pi}_t$ associated with the suspect occupying one of threat states in the underlying model. Concrete examples of the joint and conditional distributions of the data, tasks, and threat states are given in Bunnin and Smith (2021). A general template for a individual terrorist attack can be represented by the graph in Figure 2. The threat states of this model are represented by the vertices of this graph and the edges represent the possible transitions between the threat states. An example of the threat states, tasks and observable data for a gun attack are shown in Table 1.

3.3 Decoupling and coupling of component models

The decoupling methodology of multiregression dynamic models described in Online Supplementary Material, Appendix A, enables us to formally decouple the terrorist network and the individual terrorist models of each $p \in \mathcal{P}_t$ for each time $t \geq 0$ and then recombine them within a modular IDSS. The properties of this methodology rely only on the initial independencies

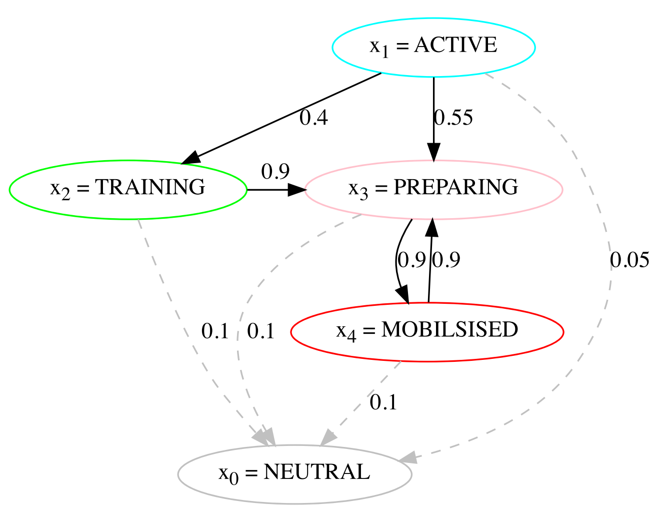

Figure 2. The graph of the state space of X_t latent level of the hierarchical individual terrorist model. Threat states and transition probabilities are shown; state probabilities and holding time distributions are omitted to avoid visual clutter. The relations between this level and the intermediate task ϑ_t and observable surface $\mathbf{Y}_t$ levels of the hierarchical model are shown in Figure 1.

Table 1. Examples of the threat states, tasks, and observable data for a gun attack

X: threat states	ϑ tasks	Y observable data
Active	Engagement with radicalisers	Physical meets with radicals
Training	Make personal threats	Personal threats made
Preparing	Learn to drive	Obtained driving license
Mobilised	Obtain vehicle	Rented car
Neutral	Engage in public threats	Public threats made on social media
	Obtain financial resources	Sold assets
	Learn how to use a gun	Visited shooting ranges
	Acquire a gun	Been seen with a gun
	Acquire ammunition	Met with gun and ammunition dealer
	Reconnoitre targets	Visits to target location made
		Financial transactions
		Contacts with family
		Meetings with trained radicals

Note. The threat states define the sample space $\{x_1, x_2, \ldots, x_n\}$ (with associated probabilities $\boldsymbol{\pi}_t = \{\pi_{t0}, \pi_{t1}, \ldots, \pi_{tn}\}$) of the suspect's state at time t given by X_t. The tasks have associated indicator task vector $\boldsymbol{\vartheta}_t = \{\vartheta_{t1}, \vartheta_{t2}, \ldots, \vartheta_{tR}\}$. Subset Y_{tj} of observable data $\{\mathbf{Y}_t\}_{t\geq 0}$ and its associated signal Z_{tj} inform the suspect's engagement in task ϑ_{tj}.

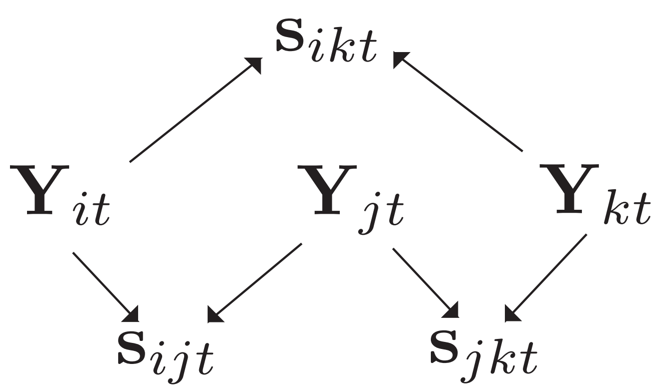

Figure 3. The directed acyclic graph (DAG) of the integrating decision support system (IDSS) after combining the individual terrorist and terrorist network models.

set through the prior parameters and on the directed acyclic graph (DAG) structure linking the components of the time-series.

We first briefly review the required notation. $\mathbf{Y}_t$ refers to the data relating to the activities of a suspect p at time $t \geq 0$ in their individual terrorist model. To generalise this notation to a population of suspects $\mathcal{P}_t$, let $\mathbf{Y}_{it}$ denote the data relating to the activities of suspect $p_i \in \mathcal{P}_t$ at time $t \geq 0$. Furthermore, $\mathbf{s}_{ijt}$ is the K-dimensional vector containing summary measures of the information shared between individuals p_i and p_j through the K information channels at time $t \geq 0$. The terrorist network model $\mathcal{N}_t$ for population $\mathcal{P}_t$ can now be coupled with the $|\mathcal{P}_t|$ individual terrorist models—one for each $p \in \mathcal{P}_t$—through a DAG which contains edges from $\mathbf{Y}_{it}$ and $\mathbf{Y}_{jt}$ to $\mathbf{s}_{ijt}$ for each pair $\{p_i, p_j\} \in \mathcal{P}_t \times \mathcal{P}_t$, and no other edges. For instance, consider $\mathcal{P}_t = \{p_i, p_j, p_k\}$. The DAG combining the individual individual terrorist models for p_i, p_j, and p_k, and the terrorist network model $\mathcal{N}_t$ at time $t \geq 0$, is given in Figure 3.

Since $\mathbf{s}_{ijt}$ contains all the observed information about the pairwise communications needed to estimate the edge weight modelled by random variable φ_{ijt} in the terrorist network, and typically

$\mathbf{s}_{ijt} \subset \mathbf{Y}_{it}$ and $\mathbf{s}_{ijt} \subset \mathbf{Y}_{jt}$, the estimation of φ_{ijt} can be performed independently of $\mathbf{Y}_{it}$ and $\mathbf{Y}_{jt}$ when $\mathbf{s}_{ijt}$ is given. Stated formally, we have $\varphi_{ijt} \perp\!\!\!\perp \mathbf{Y}_{it}, \mathbf{Y}_{jt} \mid \mathbf{s}_{ijt}$.

3.4 Indicators of a terrorist attack

The above described network component uses prior knowledge and partial data on bilateral communications $\mathbf{s}_{ijt}$ to infer (i) the existence of connections and (ii) the levels of communications φ_{ijt} along such connections. Likewise the individual terrorist component uses partial data $\mathbf{Y}_{it}$ arising from activities $\boldsymbol{\vartheta}_{jt}$ to infer threat states X_{it} of individuals. The combination of these two components' outputs, namely, φ_{ijt} for each $e_{ijt} \in E(\mathcal{N}_t)$ and X_{it} for each $p_i \in V(\mathcal{N}_t) = \mathcal{P}_t$, facilitates analysis on group entities. These correspond to investigative cases on subjects that may be groups of POIs as well as individual POIs.

The manner of integration of these two components is achieved via the construction of early warning indicators. These give quantitative measures of the imminence of threat posed by groups within $\mathcal{P}_t$. They are designed to facilitate pre-emptive action to frustrate potential attacks through flagging the activities of a group of connected individuals for increased monitoring and possible interventions. We describe below how such indicators might be constructed and how they could be utilised to forewarn the authorities. We shall, hereafter, refer to suspected or known terrorist groups as *cells*.

In our model, we define a cell $C \subset \mathcal{P}_t$ as a group of individuals who induce a connected subgraph in the terrorist network $\mathcal{N}_t$ at time $t \geq 0$.

Thus for any cell C we have $C = \{p_i \in \mathcal{P}_t : e_{ij} \in E(\mathcal{N}_t), p_j \in C\}$. We define unconnected individuals in $\mathcal{P}$ as *trivial cells*, so the set of all cells forms a partition of $\mathcal{N}_t$ which corresponds with the totality of cases being observed by the analysts. Denote the size of a particular cell as $N_c = |C|$.

Collective progress. We construct a *Terrorist Cell* model for modelling the progress of a cell C, as a separate entity, towards a terrorist attack. Let X_t^C and $\boldsymbol{\pi}_t^C$ be defined analogous to X_t and $\boldsymbol{\pi}_t$ in Section 3.2. That is $X_t^C \in \{x_{c0}, x_{c1} \ldots x_{cn}\}$ is the threat state of the cell C and $\boldsymbol{\pi}_t^C$ be the probability function over the cell threat states. Within a collaborative unit such as a cell, there will be some tasks that need only be done by a subset of the members of the cell; for example figuring out the logistics or developing certain skills. Thus, the filtered data $\mathbf{Z}_{ct}$ obtained from the collective data on the cell $\mathbf{Y}_{ct}$ must be set against these requirements to indicate whether the tasks are being sufficiently completed. Let $\mathbf{T}^C$ be the subset of the state space of X_t^C that indicates the set of states considered to be most indicative of an imminent attack by the authorities. One possible measure of collective progress m_1 of the cell C can then be obtained as

$$m_1 = \sum_{x_i^C \in \mathbf{T}^C} \pi_{ti}^C \tag{10}$$

Individual threat. As discussed above, within a cell, not all tasks need to be performed by each and every member of the cell. Ideally we would like to be able to identify, for each member of a cell C, the role that they play within the cell. However, this is not always possible as it requires detailed understanding of the cell's dynamics—intelligence which is extremely sensitive and difficult to gather (Duijn et al., 2014). An alternative is to evaluate the threat status of the individuals in C based on their progress on the tasks $\boldsymbol{\vartheta}_t^* \subset \boldsymbol{\vartheta}_{ct}$ that *most of the members* of C are expected to have the skills to do. The states for each individual's terrorist model can be adapted in line with this to obtain the product of measures of individual threat m_2 for each member of C as

$$m_2 = \prod_{p \in C} \left\{ \sum_{x_i \in T} \pi_{ti}^p \right\} \tag{11}$$

where $\boldsymbol{T}$ denotes the set of most dangerous threat states in the individual terrorist models.

Latent collaboration. In any cell, we may not expect each pair to be communicating with each other. However, for any successful collaboration, a certain amount of connectivity is expected between each communicating pair and overall in the cell. Hence we set up two different measures of latent collaboration. For each communicating pair $\{p_i, p_j\}$ in C, we measure pairwise cohesion m_3^* as

$$m_3^* = p(\varphi_{ijt} > \ell) \tag{12}$$

where ℓ is the lower limit of how much we expect each pair to be communicating for the terrorist attack to be enacted. A cell-level measure of pairwise cohesion m_3 can be obtained as

$$m_3 = \prod_{\{p_i,p_j\}\in\mathcal{P}_t\times\mathcal{P}_t} p(\varphi_{ijt} > \ell) \tag{13}$$

Size of the cell. Although collaborative efforts benefit from sharing resources and skills, a large cell can be unwieldy and increases the risk of the exposure of that cell. As a proxy measure for the size of the cell, we can devise a measure of the level of cohesion with the cell. One such measure would be through the subnetwork density m_4 of C given as

$$m_4 = \frac{k}{(N_c 2)} \tag{14}$$

where $k = |E(C_t)|$ represents the number of ties shared by the members of cell C in the network model $\mathcal{N}_t$ at time $t \geq 0$, and, as before, $N_c = |C_t|$ is the size of the cell C and thus $(N_c 2)$ is the number of possible ties in C.

The above measures are illustrative but as shown in Section 4, can still be powerful. Each of the above measures takes values in [0, 1] with a higher value signalling a greater level of threat based on that measure. In practice, these measures and the way in which they are combined together to create early warning indicators would need to be customised by the authorities, see, e.g., Xu et al. (2004) and Yang et al. (2006). We describe below one possible way in which these measures m_i for $i = \{1, 2, \ldots, 4\}$ could be combined. A cell is most threatening when $m_1 = m_2 = m_3 = m_4 = 1$. We can obtain an ordered set of indicators $\{\Lambda_C(i)\}$, $i \in \{0, \ldots, 3\}$, as

$$\Lambda_C(i) = \prod_{j=1}^{4-i} m'_j$$
$$\{m'_j\}_{j=1,\ldots,4} = \sigma(\{m_i\}_{i=1,\ldots 4}) \tag{15}$$

where σ is a permutation of elements such that for $i = 1, \ldots 4$, we have $0 \leq m'_{i+1} \leq m'_i \leq 1$ and hence for $i = 0, \ldots 3$, we have $0 \leq \Lambda_C(i) \leq \Lambda_C(i+1) \leq 1$. This ordered set is used to check whether the values of one or more measures are overly affecting the base $\Lambda_C(0)$ score. Each of these indicators has the property that a higher value of $\Lambda_C(i)$ indicates a greater imminence and danger of the threat posed by the cell C. Thus, several key factors may be combined to obtain transparent indicators of threat which can guide the counterterrorism authorities to prioritise and de-prioritise cases. These indicators can be plotted against time to analyse how the threat posed by the cell develops dynamically.

4 Analysis of a hypothetical terrorist group

Context

We present a hypothetical case to illustrate the functionality, inference, and potential output of the IDSS. Four individuals in close proximity, p_1, p_2, p_3, and p_4, have been observed to have posted pro-terrorist material on social media and have been triaged into $\mathcal{P}_{t_1}$, the observed subpopulation

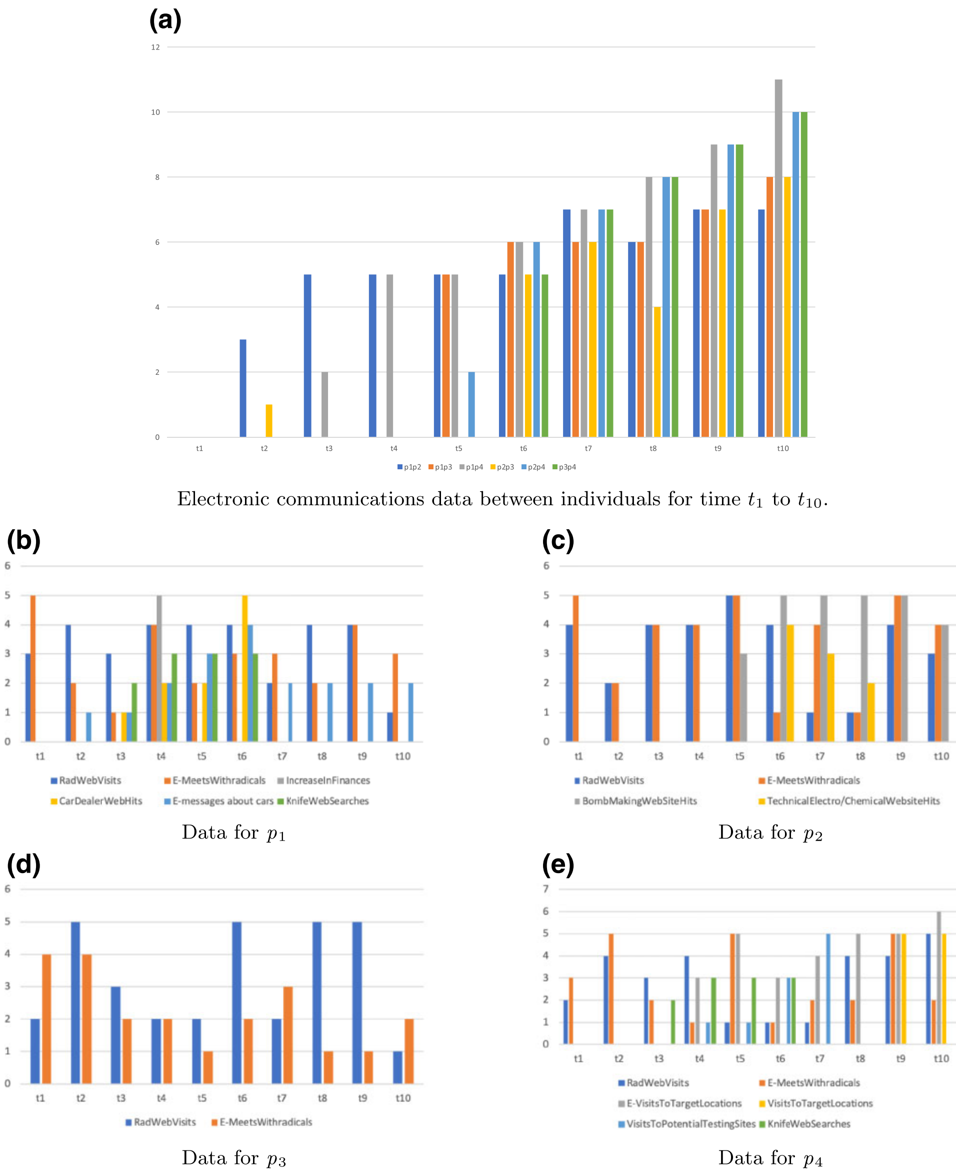

Electronic communications data between individuals for time t_1 to t_{10}.

Data for p_1

Data for p_2

Data for p_3

Data for p_4

Figure 4. Observed data on the four suspects over the observed time period of ten weeks.

at time t_1. A time step here corresponds to one week. A preliminary investigation revealed that p_1 and p_2 attended the same secondary school and are the same age, and that p_2 and p_3 attend the same gym and are frequently seen together.

Synthetic dataset

As in Section 2, the data that are fed to the model consist of variables on individuals and variables on connections between the individuals; the vertices and edges of the network respectively. Section 2 covered the types of data that can be obtained to inform the existence and strength of connections between the individuals. In this example, we take the weekly total duration of mobile phone calls as a proxy for all bilateral communications. Mobile phone data are typical of that used by the police and security analysts in identifying criminal networks as described in the cases in Anderson (2016) and Kennedy (2021). Figure 4 charts electronic data on the group observed over a period of ten weeks and Online Supplementary Material, Appendix F, tables the values of these data.

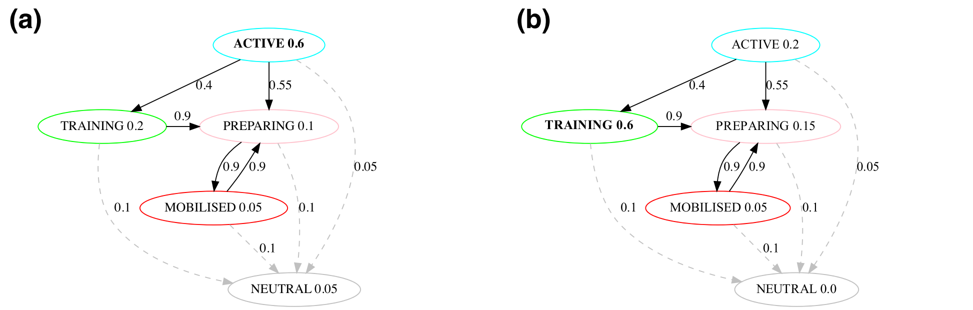

Graph of individual terrorist model for p_1, p_2 and p_3. Graph of individual terrorist model for p_4.

Figure 5. In both figures, the vertex labels include the prior state probability, and edge labels denote the conditional transition probability at time t_1.

As in Section 3.2, the space of threat states is {'Active', 'Training', 'Preparing', 'Mobilised', 'Neutral'}. Using the criminal profiles of these suspects and based on their past and current activities, the prior probabilities of the state X_{i,t_1} occupied by these individuals at time t_1 are shown in Figure 5. Suspect p_4 is believed to have received training by pro-terrorist groups and hence has probability weighted toward the 'Training' state, whereas the others have only stated their views and intentions but there is no indication otherwise of them training or preparing, hence they are weighted toward the 'Active' state.

Over the following weeks, it is observed that suspect p_1's internet activities include repeated visits to websites of car dealers and car rentals, as well as knife retailers. Their bank account also shows a large influx of funds from an overseas bank account. The internet activity of suspect p_2 includes visits to illegal bomb making websites, and repeated visits to and comments on extremist radical forums. Suspect p_4's internet activity includes searches for online maps and blueprints of government buildings and densely populated commercial areas of the town. Suspect p_4 is also observed to have physically visited potential bomb testing sites.

Inference

Using the data on individuals in Figure 4, the posterior probabilities of X_{i,t_k} are updated in the individual terrorist model over the ten weeks for each p_i as shown in Figure 6 for $i = 1, 2, 3, 4$ and time $1 \leq k \leq 10$.

The phone call data, which inform the connections between the suspects, are summarised as the sum of the phone calls in hours between the pair observed over the week. Based on this, ties are revealed represented as edge $e_{1,4}$ at time t_3, edges $e_{1,3}$ and $e_{2,4}$ at time t_5, and edge $e_{3,4}$ in the network at time t_6. Hence at time t_6, the network becomes a complete graph. The total time of calls increase from weeks t_7 to t_{10}.

The edge weight distributions φ_{ij,t_k} for $i, j = 1, 2, 3, 4$, $i \neq j$ and $1 \leq k \leq 10$, for the terrorist network model can be estimated as follows. The prior distributions for φ_{ij,t_k} are set by specifying the α and β parameters of the prior Gamma distributions. For instance, based on the prior knowledge the policing authority has on the suspects, they believe that the extent of information shared between p_1 and p_2 and between p_2 and p_3 is relatively low, with some uncertainty at time t_0. Hence, the α and β parameters are set as 1 and 2, respectively, for $\varphi_{1,2,t_0}$ and $\varphi_{2,3,t_0}$. The setting of the α and β parameters for φ_{ij,t_k} for all pairs over the ten week period are given in Online Supplementary Material. Online Supplementary Material, Appendix F, Figure 3, show the evolution of φ_{ij,t_k} through the posterior densities from time t_3 to t_6. The discount factor δ_{ij,t_k} is set to $0.9512 = \exp -0.05$ across all pairs and for the entire ten week duration. With this setting, information has a half-life of approximately 14 weeks.

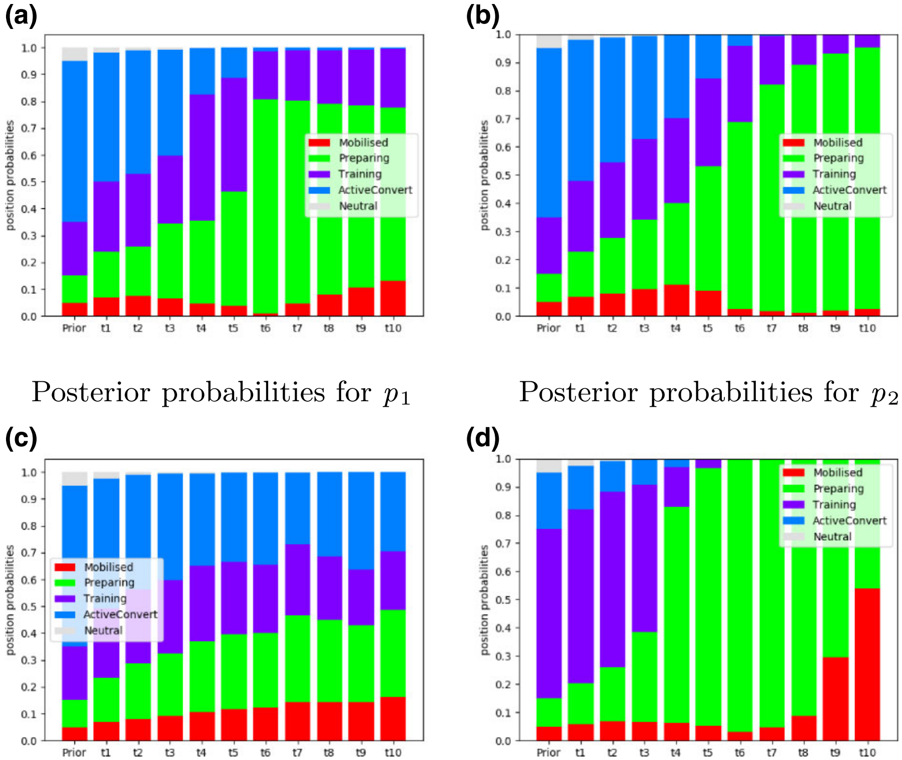

Posterior probabilities for p_1 Posterior probabilities for p_2

Posterior probabilities for p_3 Posterior probabilities for p_4

Figure 6. Posterior threat state probabilities from the individual terrorist models of the suspects over the ten weeks.

Indicators of an attack

The connectivity and subsequent completion of the graph at times t_5 and t_6 indicates that the four suspects are working together within a cell. The measures m_i for $i = 2, 3, 4$ described in Section 3.4 are calculated. Note that for the terrorist cell model for measure m_1, the task set and observation data are given by the union of the task sets and observation data for the individual terrorist models of the cell's members. The prior threat state probabilities for the terrorist cell model are taken as the corresponding prior probabilities of the suspect within the cell with the highest prior threat, i.e., p_4. Figure 7a shows the evolution of the threat state probabilities in the terrorist cell model. The posterior probability of the cell being in the 'Preparing' state increases from time t_5 as the communications within the cell and the overall activities of the cell increase. Thereafter around time t_9, the posterior probability of the cell being in the 'Mobilised' state increases sharply. These measures are combined to obtain the indicators of a terrorist attack Λ_C as shown in Figure 7b. If we were to signal a warning when $\Lambda_C(\cdot)$ reaches a certain threshold, say 0.2, then we can see that for $\Lambda_C(0)$ this is not reached till time t_7 whereas for $\Lambda_C(2)$ this is reached by time t_3. In practise, the measures informing these indicators and the chosen thresholds would need to be calibrated using domain experience and judgement.

This simple worked example demonstrates how observed activity data and communications data obtained on monitored suspects, when combined with prior distributions calibrated to the investigator's knowledge, can give real-time indicators of the evolving threat posed by individuals acting in collaboration. In the scenario investigated here, driven by the increase in specific activity data and phone call duration, probability of the suspects forming the cell being in either the 'Preparing' or 'Mobilised' states by week t_{10} increased, and correspondingly, the cell as a whole

(a)

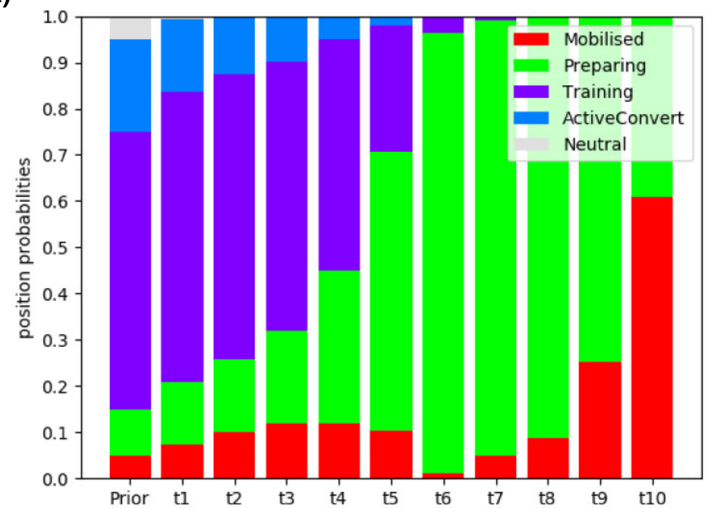

Posterior threat state probabilities for the Terrorist Cell over the ten weeks.

(b)

	Prior	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10
m1	0.15	0.21	0.26	0.32	0.45	0.71	0.96	0.99	1	1	1
m2	0	0	0.01	0.01	0.04	0.09	0.22	0.31	0.32	0.31	0.36
m3	0.14	0.05	0.14	0.04	0.03	0	0.3	1	1	1	1
m4	0.67	0.67	0.67	0.67	0.67	0.67	0.67	0.67	0.67	0.67	0.67
$\Lambda_C(0)$	0	0	0	0	0	0	0.04	0.21	0.21	0.21	0.24
$\Lambda_C(1)$	0.01	0.01	0.02	0.01	0.01	0.04	0.19	0.66	0.67	0.67	0.67
$\Lambda_C(2)$	0.1	0.14	0.17	0.21	0.3	0.48	0.64	0.99	1	1	1
$\Lambda_C(3)$	0.67	0.67	0.67	0.67	0.67	0.71	0.96	1	1	1	1

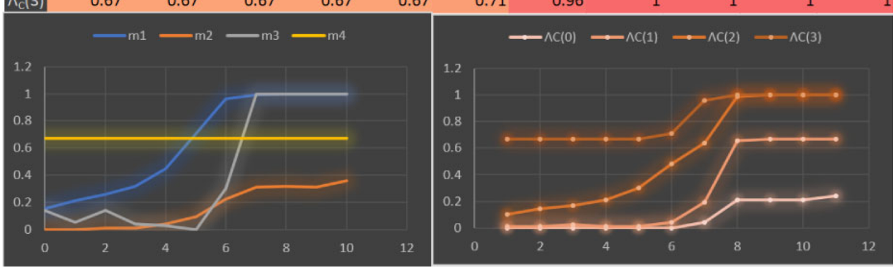

Threat Measures and Indicators for the Terrorist Cell over the ten weeks.

Figure 7. Dashboard threat monitor. In practise we envisage each case being investigated to have such charts giving visualisations of each case's recent and current levels of threat according to the model.

appeared to move from state 'Preparing' occupied since week t_5 to 'Mobilised' by week t_{10}. The indicators of an attack reflected a similar trend. $\Lambda_C(2)$ increases from 0.21 at week t_3 to 0.48 at week t_5 and then saturates at 1 for weeks t_7 to t_{10}. $\Lambda_C(1)$ increases from 0.01 at week t_3 to 0.04 at week t_5 and reaches 0.66–0.67 for weeks t_7 to t_{10}.

Sensitivity analyses for the individual terrorist model were presented in Online Supplementary Material in Bunnin and Smith (2021); the results therein apply also to the cell-level extension of the individual model created to arrive at the m_1 measure. For the network model, results for sensitivity analyses to the Gamma distribution priors $\alpha_{ijt_0}, \beta_{ijt_0}$ and the discount factor and efficiency parameters δ_{ijt}, ξ_k are presented in Online Supplementary Material, Appendix G. Overall the results from these analyses confirm the intuitive meaning of the parameters. The sensitivities can be used to calibrate the priors and parameters to case data.

5 Discussion

In this paper, we proposed a two-part IDSS to support security analysts monitor potential terrorist cells. The IDSS combines the outputs of a terrorist network model and a collection of hierarchical individual terrorist models. It outputs real-time indicators of threat levels. These aim to facilitate pre-emptive action to frustrate attacks. When combined with utility or loss functions elicited from the authorities, the IDSS can be used in a decision-theoretic framework for optimal resource allocation. See Online Supplementary Material, Appendix H, for ideas in this direction that will be presented elsewhere.

There are several avenues of research that can follow from the work presented in this paper. Recall that the edge weight φ_{ijt} along an edge in the terrorist network $\mathcal{N}_t$ is a measure of the pairwise communications shared directly between the suspects p_i and p_j at time t connected by the edge. This definition of the edge weight then leads to the following conditional independence assumption:

$$\perp\!\!\!\perp_{\{p_i,p_j\}\in\mathcal{P}_t\times\mathcal{P}_t}\varphi_{ijt} \mid \mathcal{F}_t \tag{16}$$

Thus, pairwise communications data can be used to estimate the edge weight φ_{ijt}. For an alternative interpretation of the edge weights as measures of the extent of collaboration between the individuals connected by the edge, the above conditional independence statement does not hold. For instance, the extent of collaboration φ_{ijt} between p_i and p_j is also affected by the communications and interactions they both share with a common neighbour, say p_k. This would additionally lead to the violation of the output independence assumption stated in Equation (4) in Section 3.2. Under this independence structure, we would no longer be able to estimate φ_{ijt} for each pair p_i and p_j independently. In this case, the decouple/recouple strategy introduced by Gruber and West (2016) and Zhao et al. (2016) for financial and economic multivariate time-series applications could be explored, although the recoupling is unlikely to be closed form and would need to be numerically estimated. Here, any gains achieved by refining the interpretation of φ_{ijt} to include collaboration in a broader sense would need to be weighed against the loss in transparency and interpretability due to the numerical estimation.

Another avenue of research would be to incorporate link detection within the terrorist network using existing link detection methods (see Section 1.1) to identify potentially hidden ties. This work can further be extended by developing a bespoke clustering algorithm using domain-specific stochastic set functions (as have been used in the criminology literature, see, e.g., Wang et al., 2013) to identify previously unknown cells or monitor the evolution of new cells within the network.

Finally, the generic architecture of an IDSS using the decoupling methodology might be applicable to other domains where there is a requirement to integrate individual time-series with dynamic interactions among individuals, modelled by a network, who collaborate to realise a shared objective. Examples of this include social processes within politics, governments or communities where complex interacting individuals have shared objectives.

Acknowledgments

A.S. and F.O.B. contributed equally.

Supplementary material

Supplementary material is available online at *Journal of the Royal Statistical Society*.

Conflict of interest: The authors have no conflict of interest to disclose.

Funding

A.S. was supported by the University of Warwick Chancellor's International PhD Scholarship and the Alan Turing Institute's PhD Enrichment Award whilst working on this project during her PhD. J.Q.S. was supported by the Alan Turing Institute and funded by the EPSRC [grant number EP/K03 9628/1]. F.O.B was funded by The Alan Turing Institute Defence and Security Project G027.

Data availability

The code for generating and analysing the synthetic data underlying this article is available as part of the supplementary materials.

References

Allanach J., Tu H., Singh S., Willett P., & Pattipati K. (2004). Detecting, tracking and counteracting terrorist networks via hidden Markov models. In *IEEE aerospace conference proceedings (IEEE Cat. No. 04TH8720)* (Vol. 5, pp. 3246–3257). IEEE.

Anderson D. (2016). *Investigatory powers bill: Bulk powers review*. https://www.gov.uk/government/publications/investigatory-powers-bill-bulk-powers-review

Andreassen C. M. (2013). *Models and inference for correlated count data* [PhD thesis]. Aarhus University, Department of Mathematics.

Barons M. J., Fonseca T. C., Davis A., & Smith J. Q. (2021). A decision support system for addressing food security in the UK. *Journal of the Royal Statistical Society Series A (Statistics in Society)*, *185*(2), 447–470. https://doi.org/10.1111/rssa.12771

Brain J. (2015). Brusthom Ziamani: Teenager guilty of plot to behead soldier. *BBC*. bbc.co.uk/news/uk-31540281

Bunnin F. O., & Smith J. Q. (2021). A Bayesian hierarchical model for criminal investigations. *Bayesian Analysis*, *16*(1), 1–30. https://doi.org/10.1214/19-BA1192

Campedelli G. M., Cruickshank I., & Carley K. M. (2019). A complex networks approach to find latent clusters of terrorist groups. *Applied Network Science*, *4*(1), 1–22. https://doi.org/10.1007/s41109-019-0184-6

Chen X., Irie K., Banks D., Haslinger R., Thomas J., & West M. (2018). Scalable Bayesian modeling, monitoring and analysis of dynamic network flow data. *Journal of the American Statistical Association*, *113*(522), 519–533. https://doi.org/10.1080/01621459.2017.1345742

Choo L., & Walker S. G. (2008). A new approach to investigating spatial variations of disease. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *171*(2), 395–405. https://doi.org/10.1111/j.1467-985X.2007.00503.x

Chorley M. (2021). *Interview with MI5*. https://www.thetimes.co.uk/podcasts/red-box

Çinlar E. (1975). Markov renewal theory: A survey. *Management Science*, *21*(7), 727–752. https://doi.org/10.1287/mnsc.21.7.727

Counter Extremism Project (2022). *Brusthom Ziamani*. Retrieved July 2022, from https://www.counterextremism.com/extremists/brusthom-ziamani

Cunningham D., Everton S. F., & Murphy P. J. (2015). Casting more light on dark networks: A stochastic actor-oriented longitudinal analysis of the Noordin top terrorist network. In Luke M. Gerdes (Ed.), Illuminating dark networks: The study of clandestine groups and organizations (pp. 171–185). Cambridge University Press.

Departments of the Army and the Air Force (1948). *Fundamentals of traffic analysis (radio-telegraph)* (Tech. Rep. US manual number TM 32-250 - AFM 100-80). US Department of Defense.

Duijn P. A., Kashirin V., & Sloot P. M. (2014). The relative ineffectiveness of criminal network disruption. *Scientific Reports*, *4*, 4238. https://doi.org/10.1038/srep04238

Gardham D., & Gibbons K. (2019). London bridge attackers may have intended to target Oxford street, inquest hears. *The Times Newspaper*. https://www.thetimes.co.uk/article/london-bridge-attackers-met-and-trained-in-east-london-gym-hnjvb7r5b

Gruber L., & West M. (2016). GPU-accelerated Bayesian learning and forecasting in simultaneous graphical dynamic linear models. *Bayesian Analysis*, *11*(1), 125–149. https://doi.org/10.1214/15-BA946

Home Office (2018). *Counter-terrorism strategy (CONTEST): The United Kingdom's strategy for countering terrorism* (Tech. Rep.). Government of the United Kingdom. https://www.gov.uk/government/publications/counter-terrorism-strategy-contest-2018

Home Office (2020). *Guidance: Regulation of investigatory powers act 2000 (RIPA)*. https://www.gov.uk/government/publications/regulation-of-investigatory-powers-act-2000-ripa/regulation-of-investigatory-powers-act-2000-ripa

Jahan A., & Edwards K. L. (2015). A state-of-the-art survey on the influence of normalization techniques in ranking: Improving the materials selection process in engineering design. *Materials & Design (1980–2015)*, *65*, 335–342. https://doi.org/10.1016/j.matdes.2014.09.022

Kennedy N. (2021). *Hunting the essex lorry killer*. https://www.bbc.co.uk/iplayer/episode/m0010ldl/hunting-the-essex-lorry-killers

Leonelli M., & Smith J. Q. (2015). Bayesian decision support for complex systems with many distributed experts. *Annals of Operations Research*, *235*(1), 517–542. https://doi.org/10.1007/s10479-015-1957-7

Pantucci R. (2016). *Lone actor terrorism* (Tech. Rep.). Royal United Services Institute. https://rusi.org/explore-our-research/projects/lone-actor-terrorism

Parveen N., & Walker A. (2020). Brother of Manchester Arena bomber found guilty of murder. *The Guardian Newspaper*. https://www.theguardian.com/uk-news/2020/mar/17/brother-of-manchester-arena-bomber-hashem-abedi-guilty-murder

Queen C. M., & Smith J. Q. (1993). Multiregression dynamic models. *Journal of the Royal Statistical Society. Series B (Methodological)*, *55*(4), 849–870. https://doi.org/10.1111/rssb.1993.55.issue-4

Radio4 BBC (2019). *Analysis: Understanding the risks of terrorism*. https://www.bbc.co.uk/programmes/m0006dnh

Ranciati S., Vinciotti V., & Wit E. C. (2020). Identifying overlapping terrorist cells from the Noordin Top actor–event network. *The Annals of Applied Statistics*, *14*(3), 1516–1534. https://doi.org/10.1214/20-AOAS1358

R v Ziamani (2015). *Sentencing remarks of HHJ Pontius*. Central Criminal Court.

Schein A., Wallach H., & Zhou M. (2016). Poisson-Gamma dynamical systems. In *30th conference on neural information processing systems*, 5012–5020.

Shenvi A. (2021). *Non-stratified chain event graphs dynamic variants, inference and applications* [PhD thesis]. University of Warwick.

Singh S., Allanach J., Tu H., Pattipati K., & Willett P. (2004). Stochastic modeling of a terrorist event via the ASAM system. In *IEEE international conference on systems, man and cybernetics (IEEE Cat. No. 04CH37583)* (Vol. 6, pp. 5673–5678). IEEE.

Smith J. Q. (1979). A generalization of the Bayesian steady forecasting model. *Journal of the Royal Statistical Society: Series B (Methodological)*, *41*(3), 375–387.

Smith J. Q., Barons M. J., & Leonelli M. (2015). *Coherent frameworks for statistical inference serving integrating decision support systems*, arXiv:1507.07394, preprint: not peer reviewed.

Smith J. Q., & Shenvi A. (2018). *Assault crime dynamic chain event graphs* (Working Paper). http://wrap.warwick.ac.uk/104824/

Sparrow M. K. (1991). The application of network analysis to criminal intelligence: An assessment of the prospects. *Social Networks*, *13*(3), 251–274. https://doi.org/10.1016/0378-8733(91)90008-H

Toth N., Gulyás L., Legendi R. O., Duijn P., Sloot P. M., & Kampis G. (2013). The importance of centralities in dark network value chains. *The European Physical Journal Special Topics*, *222*(6), 1413–1439. https://doi.org/10.1140/epjst/e2013-01935-7

UK Public General Acts 2000 c. 23 (2000). *Regulation of investigatory powers act 2000 (RIPA)*. https://www.legislation.gov.uk/ukpga/2000/23/contents

United States Treasury Department (2022). *Terrorist finance tracking program*. https://home.treasury.gov/policy-issues/terrorism-and-illicit-finance/terrorist-finance-tracking-program-tftp

van Meter K. M. (2002). Terrorists/liberators: Researching and dealing with adversary social networks. *Connections*, *24*(3), 66–78.

Wang T., Rudin C., Wagner D., & Sevieri R. (2013). Learning to detect patterns of crime. In H. Blockeel, K. Kersting, S. Nijssen, & F. Železný (Eds.), *Joint European conference on machine learning and knowledge discovery in databases* (Vol. 8190, pp. 515–530). Springer.

West M., & Harrison J. (1997). *Bayesian forecasting and dynamic models*. Springer.

Xu J., Marshall B., Kaza S., & Chen H. (2004). Analyzing and visualizing criminal network dynamics: A case study. In *International conference on intelligence and security informatics* (pp. 359–377). Springer.

Yang C. C., Liu N., & Sageman M. (2006). Analyzing the terrorist social networks with visualization tools. In *International conference on intelligence and security informatics* (pp. 331–342). Springer.

Zhao Z. Y., Xie M., & West M. (2016). Dynamic dependence networks: Financial time series forecasting and portfolio decisions. *Applied Stochastic Models in Business and Industry*, *32*(3), 311–332. https://doi.org/10.1002/asmb.v32.3