Characterizing bioterrorist attacks from a short time series of diagnosed patient data - A Bayesian approach

J. Ray[†,1], Y. M. Marzouk[††], M. Kraus[‡] and P. Fast[§]

[†]Computer Sciences and Information Technologies Center,
Sandia National Laboratories, PO Box 969, Livermore, CA 94550

[††]Biological and Energy Sciences Center,
Sandia National Laboratories, PO Box 969, Livermore, CA 94550

[‡]NORAD-NORTHCOM, 250 S. Peterson Blvd, Peterson AFB, CO 80914-3180

[§]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory,
7000 East Ave, Livermore, CA 94550

Presented at:
Quantitative Methods in Defense and National Security 2007
George Mason University, February 7-8, 2007

Abstract

Terrorist attacks using an aerosolized pathogen preparation have gained credibility as a national security concern after the anthrax attacks of 2001. The ability to characterize such attacks, i.e., to estimate the number of people infected, the time of infection, and the average dose received, is important when planning a medical response. We address this question of characterization by formulating a Bayesian inverse problem predicated on a short time-series of diagnosed patients exhibiting symptoms. To be of relevance to response planning, we limit ourselves to 3–5 days of data. In tests performed with anthrax as the pathogen, we find that these data are usually sufficient, especially if the model of the outbreak used in the inverse problem is an accurate one. We also explore the effect of model error—situations for which the model used in the inverse problem is only a partially accurate representation of the outbreak. We find that while there is a consistent discrepancy between the inferred and the true characterizations, they are also close enough to be of relevance when planning a response.

[1]To whom correspondence should be addressed. Address: MS 9159, PO Box 969, Sandia National Laboratories, Livermore, CA 94550-0969. Email: jairay@somnet.sandia.gov. Phone: 925-294-3638.

1 Introduction

The anthrax attacks of 2001 [1] raised the credibility of aerosolized pathogens being used in bioterrorist attack. Early warning, either in the form of an anomalous increase in syndromes detected by public health monitoring networks [2] or via detection by environmental sensors, holds the highest potential to reduce casualties. However, syndromic surveillance can only provide heightened awareness—it results in neither definite evidence of an attack nor in the identification of the pathogen. Also, the introduction of an aerosolized pathogen into a population may not always be captured on environmental sensors. Examples include small releases that may not travel far, low quality formulations (coarse and heavy particulate matter) which precipitate easily, as well as releases in areas which are not well instrumented. In such a case, the first definitive diagnosis of a patient will be the first intimation of an attack, but by then the disease may have established itself in the population. Being able to infer the characteristics of the release (also referred to as the bioterrorist or BT attack)—i.e., the number N of the people infected, the time τ of infection, and a representative dose D received by the infected people—has important ramifications in planning a response [3]. The inferred characteristics can also serve as initial conditions for various epidemic models that can predict the evolution and spread of the disease in a population [4] and its ramification on society [3, 5].

Inferring the characteristics of the outbreak can be challenging. The observables on which inferences are based consist of the time the diagnosed patients turned symptomatic (typically expressed as a time interval during which they developed symptoms) and the location of their residence and place of work. In case of a very mobile population, e.g., a military force engaged in operations, the location of residence and/or work may be hard to define. The model that relates the time of exhibition of symptoms to the characteristics of the genesis of the outbreak is the incubation period distribution, which in many cases is dependent on the dosage received. To be relevant in an operational, consequence-management sense, these inferences have to be drawn early in the outbreak; a time-series obtained from a 3–5 day observation period may be considered representative. Apart from scarcity of observations, the incubation period distribution used in the inferences may be a poor model for the particular instance of the disease. Thus these inferred characteristics are expected to be rather approximate, and quantifying the uncertainty in the characterization becomes a key requirement of the inference process.

In this paper we will limit ourselves to temporal analysis; we will not take the location of diagnosed patients into consideration. Further, all tests will be performed with anthrax as the pathogen. Broad uniform priors will be used in the inference process. We will operate within a self-imposed limit of a 3–5 day observation period. We present four cases to demonstrate the effect of the size of the outbreak (N) and a representative dose received (D) on the inferences. They will also explore the demonstrate the effect of model mismatch i.e. when the model of anthrax used for the inference is a poor approximation of the actual behavior of anthrax (which produces the observables/data). We conclude with an application of this method to the Sverdlovsk outbreak of 1979 [6]. The results of this study will provide a measure of the accuracy and robustness of this Bayesian method, preparatory to extending this purely temporal analysis into a spatio-temporal one.

2 Previous work

The question of inferring the characteristics of the genesis of an outbreak from a partially observed epidemic has not been extensively studied. Walden and Kaplan [7] developed a Bayesian formulation to estimate the size and time of a bioterrorist attack which they tested on a low-dose anthrax attack corresponding, approximately, to the Sverdlovsk outbreak [6] of 1979, using an incubation period model developed by Brookmeyer [8]. They also demonstrated the use of priors—prior belief regarding the size N of the outbreak—to develop a smooth PDF for N in spite of a small infected population ($N = 100$) and a short time-series (5 days long), with data collected on a daily basis. An alternative approach (maximum likelihood) was used by Brookmeyer and Blades [9] to infer the size of the 2001 anthrax attacks [1], before estimating the reduction of casualties by the timely administration of antibiotics. This inference process was difficult due to the small number of symptomatic patients (11 infectees in 3 separate attacks). They also used the anthrax incubation model in [8]. Both [7] and [8] developed similar expression for the likelihood function, i.e., the probability of observing a time series given an attack at time τ with N infected people. The incubation period distribution was not dose-dependent, and hence no dosages were inferred in the two studies.

Significantly more effort has been spent in characterizing the incubation period of inhalational anthrax. The bulk of the work has been experimental, with non-human primates being subjected to anthrax challenges [10, 11, 12, 13, 14, 15]. Brookmeyer et al [8] developed a low-dose incubation period model applicable to the Sverdlovsk outbreak; their more recent work, based on a competing risks formulation, includes dose-dependence [16]. A more empirical study, but based on significantly more data, was done recently by Wilkening [17], where he compared four different models called Models A, B, C and D. Model D is a slight modification of Brookmeyer's dose-dependent model described in [16]. While Wilkening's Model A agreed with Model D at the high-dose limit, their low-dose behavior was different. Further, Wilkening developed two variants of his Model A, A1 and A2. A1 is a simpler model but its comparison with experimental results is slightly worse than A2. In this study, we will use Wilkening's Models A2 and D for simulated BT attacks while Model A2 will be used in the inference scheme.

An effort with aims similar to ours is the Bayesian Aerosol Release Detector (BARD) [18]. It poses an inverse problem to infer the location and height of an anthrax release (the approach is general but has only been tested with anthrax), the time of release and the quantity of material released. The observables are the number of respiratory visits to emergency departments collated in 24-hour intervals and by zip code - such information can be obtained from typical syndromic surveillance systems such as RODS [19, 20]. BARD does not calculate the the number of people infected or the dosage - however, given PDFs for the location and quantity, the magnitude of the outbreak and the dosage may be trivially obtained by using the BARD inferences as the initial condition in a Gaussian plume to disperse the aerosol and using Glassman's [21] (or Druett's [22]) model to decide the probability of infection.

3 The inverse problem

Consider an attack at time τ where N people are infected, with each of the N people receiving the same dose of D anthrax spores. The incubation period obeys a dose-dependent dis-

tribution; we refer to its cumulative distribution function (CDF) as $C(t, D)$. For a few days M (say 3–5 days) we can expect (1) a series $t_i,\ i = 0 \ldots M$, of times, the endpoints of 24-hr intervals, when patients' symptoms are observed and (2) the series $n_i,\ i = 0 \ldots M$, of new patients who turned symptomatic between $t_i - \Delta t$ and t_i where $t_i - t_{i-1} = \Delta t, i \geq 0$, and Δt is a constant. We define survival probability as $P_{surv}(t, D) = 1 - C(t, D)$. We can state the problem as such: Given a time-series $(t_i, n_i), i = 0 \ldots M$, of patients showing symptoms over a few days, estimate (N, τ, D) from these data.

Let $L = \sum_{i=0}^{M} n_i$ be the total number of people who have developed symptoms by t_M. Thus $N - L$ infected people are still asymptomatic and the probability of such an event is $\{P_{surv}(t_M - \tau, D)\}^{N-L}$. The probability that n_i people will develop symptoms in the time interval between t_{i-1} and t_i is $\{C(t_i - \tau, D) - C(t_i - \Delta t - \tau, D)\}^{n_i}$. The probability of observing the $\{t_i, n_i\}, i = 0 \ldots M$ time-series given a BT attack characterized by (N, τ, D), or equivalently, the likelihood function $\mathcal{L}$, is

$$\begin{aligned} \mathcal{L}(N, \tau, D) &\equiv p(\{t_i, n_i\}, i = 0 \ldots M | N, \tau, D) \\ &= \frac{N!}{(N-L)! \prod_{i=0}^{M} n_i!} \times \{P_{surv}(t_M - \tau, D)\}^{N-L} \times \\ &\quad \prod_{i=0}^{M} \{C(t_i - \tau, D) - C(t_{i-1} - \tau, D)\}^{n_i}. \end{aligned} \tag{1}$$

Exploiting Bayes rule, we obtain

$$\pi(N, \tau, D | \{t_i, n_i\}, i = 0 \ldots M) = \frac{\mathcal{L}(N, \tau, D) \pi_N(N) \pi_\tau(\tau) \pi_D(D)}{\pi(\{t_i, n_i\}, i = 0 \ldots M)} \tag{2}$$

where π_N, π_τ and π_D are the priors for N, τ and D and $\pi(\{t_i, n_i\}, i = 0 \ldots M)$ is the probability of observing a $\{t_i, n_i\}, i = 0 \ldots M$ time-series in any circumstance. In this study, we use broad uniform distributions as priors. The joint probability distribution $\pi(N, \tau, D | \{t_i, n_i\}, i = 0 \ldots M)$ is marginalized to obtain individual PDFs for N, τ and D. A more detailed derivation can be found in [23].

The CDF $C(t, D)$ in Eq. 1 can be either that of Wilkening's Model A2 or D. These can be found in [23, 17]. The parameters in these models were obtained by fitting to the incubation periods observed in experiments with non-human primates (performed by Henderson et al [10] and Friedlander et al [13]) and the data from the Sverdlovsk outbreak. However, the average dose during the Sverdlovsk outbreak had to be inferred from atmospheric dispersion models and the probability of exhibiting symptoms (in infinite time) given a dose of D spores. This was done by Wilkening [17]. If one uses Glassman's model [21] for the probability of infection, one obtains an average dose of 2.4 spores. Alternatively, if one uses Druett's model [22] one obtains a dose of 300 spores. Wilkening retained both the possibilities and incorporated them into separate models. Model D is based on a dose of 300 spores at Sverdlovsk while A2 assumes 2.4 spores.

In Fig. 1, we plot the median incubation period as predicted by Model A2 and D, as a function of dosage D. The dosage at Sverdlovsk, inferred as 2.4 spores (represented by $\bullet$) is used to derive the parameters for Model A2 (solid line); the alternative inference of 300 spores (represented by a filled ∇) is used for Model D (dashed line). Studies by Henderson [10] with 2.1×10^5, 3.9×10^5 and 7.6×10^5 spores (represented as filled $\diamond$) and Friedlander with 3.5×10^5 spores (represented by filled $\triangle$) were used to derive the parameters of both the models. Studies by Ivins et al [14] (unfilled $\triangle$) and Gleiser et

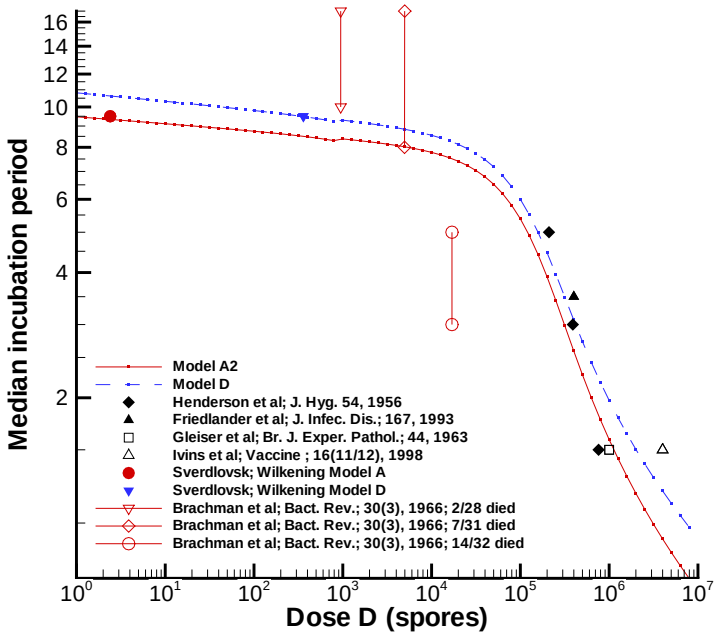

Figure 1: The median incubation period for anthrax as a function of dosage D. The solid line is Model A2 which assumes a dose of 2.4 spores at Sverdlovsk, while the dashed line is Model D, which assumes 300 spores. The solid symbols are median incubation periods which were obtained from experimental investigations or from the data from the Sverdlovsk outbreak. Symbols which are not filled denote experiments where the population of primates were too small to draw statistically meaningful results. The experiments by Brachman et al [15] are shown by vertical lines between symbols. In these tests, only the lower and upper bounds of the incubation period were provided. These were not used for determining model parameters and are only provided for reference.

al [12] (unfilled □) were conducted with very few primates and consequently are plotted only for reference. Brachman [15] conducted studies where he tried to simulate the effect of a low dose, received regularly over an extended period of time, as might be the case in a contaminated wool-sorting mill. The primates went through extended periods when they received no spores at all. The dose was calculated as the total number of spores breathed in and was generally low, between 1000 and 10,000 spores. We plot the ranges of incubation periods observed (only the range was provided) for various dosages for reference.

We see that the tests by Gleiser et al and Ivins et al agree with both the models, which in turn agree with each other, except at the low dose limit. Brachman's tests show median incubation periods which are at odds with the models' predictions; however the mode of infection, that approximating a continuous, low-level infection process spread over days or months was very different from the quick (timescale of an hour) challenge one would expect in a BT attack. Both the models show a kink at $D = 10^3$; this is because they are

evaluated with a lower value of λ ($= 1.3 \times 10^{-6}$ day^{-1}) corresponding to a primate ID_{50} of 55,000 spores for comparison with primate anthrax challenge results at the high dose limit, while the low dose predictions were developed with a human ID_{50} of 8600 spores for comparison with the inferences of the dose received at Sverdlovsk. To the best of the authors' knowledge, this is the sum total of experimental data obtained from anthrax challenges of non-human primates where incubation times were measured. We have omitted a study by Klein et al [24] in which an incubation period increase was observed with increasing doses, because only one primate was subjected to a given dose, making the behavior statistically unreliable.

In [23], Eq. 2 was used to infer "idealized" attacks where N people were infected identically with D spores each. Wilkening's Model A2 was used to evolve the disease in each person and was also used to infer the attack - thus the only source of uncertainty in the inferences was the incompleteness (3–5 days) of the observables. It was noticed that the time of the attack τ was generally easy to infer regardless of the size of the attack. The dose D was virtually impossible for small ($N \approx 100$) attacks and in general large attacks were easier to infer accurately than smaller ones. In some cases, the observed data supported multiple characterizations (and sometimes the wrong characterization more than the correct one) but with the availability of data (and consequently time) the correct characterization was always recovered.

4 Results

4.1 Tests with simulated attacks

We conduct four tests with simulated anthrax BT attacks. The infected population of size N receives a range of doses reflecting atmospheric dispersion of an aerosolized pathogen. We assume a general population density distribution over a 10 km square domain, over which we release 10^{13} spores. Assuming a 4 m/s wind and a Pasquill stability class of "B", we use a Gaussian plume to develop dose contours and expose the population. For Case A and B, we use Glassman's model [21] to determine the probability of infection of a person exposed to a dose D while for Cases C and D, we use Wilkening's Model D. Quantiles of the dose distribution for the four cases are in Table 1. Note that 80% of the infected people receive a range of doses spanning an order of magnitude. Details can be found in [23].

In Cases A and B, the evolution of anthrax in the infected people is governed by Wilkening's Model A2; this model is also used in the inference process. Thus uncertainty in the inference is due to the incomplete nature of the observations and the errors incurred by fitting a constant dose model (Eq. 2) to variable dose data. In Cases C and D, we use Model D for the evolution of anthrax in the infected people; however, Model A2 is used in the inference process. Table 1 contains the time-series for the four attacks, as well as the correct values of N, τ and $\overline{D}$, the average dose received by the N infected people.

In Fig. 2 we plot the probability density functions (PDFs) for Cases A and B for N, τ and $\log_{10}(D)$. In Table 2 we tabulate their maximum a priori (MAP) estimates and 90 % credibility intervals (CIs) developed for the data available on Day 5. These are compared with the true characterizations from Table 1. Since the infected population receives a range of doses, we compare the logarithm of the median dose $\log_{10}(D_{50})$ to the logarithm of the "representative" dose $\log_{10}(D)$ inferred by fitting a constant dose model to variable dose data. In Case B we see multimodal PDFs for N and τ with Day 5 data, though by Day 7 (see [25] for data beyond Day 5) the PDFs are narrow and unimodal. The 90% CIs for N,

Time	Case A	Case B	Case C	Case D
0	1	9	1	3
1	5	73	5	208
2	15	317	7	478
3	29	522	11	565
4	40	628	15	490
5	31	556	19	410
N	318	4537	161	4453
τ	-1.5	-1.25	-0.75	-0.5
$\bar{D}$	2912.8	13,150.4	3606.5	16,532
D_1	5.3×10^1	1.32×10^2	3.41×10^2	3.0×10^2
D_{25}	1.23×10^3	3.47×10^3	1.99×10^3	9.45×10^3
D_{50}	2.91×10^3	1.24×10^4	3.34×10^3	1.57×10^4
D_{75}	4.12×10^3	1.87×10^4	4.79×10^3	2.07×10^4
D_{99}	8.28×10^3	5.91×10^4	9.17×10^3	6.51×10^4

Table 1: The daily number of new symptomatic patients for Cases A, B, C and D. The middle 3 rows of the table provide the true characteristics of the attacks. $\bar{D}$ is the mean of the dose distribution, in units of spores. τ, the time of the attack/infection, is measured from the moment the first patient exhibited symptoms and is therefore always negative. The values D_x are quantiles of the dose distribution; x% of the infected population receives a dose of D_x spores or less. Time is measured in days.

τ and $\log_{10}(D)$ bracket the true values. The time of the attack is easily inferred, though withing the finite time resolution of the observed data. The dose is harder to infer in the smaller Case A. Thus while the errors introduced by violating the assumption of constant dose in Eq. 2 are not negligible, the current formulation provides a reasonable and useful characterization of the attack.

We now proceed to Cases C and D. These include a systematic difference between the actual evolution of the attack and the model used to interpret the data. Table 1 contains the observations and the correct characterization of the attack. Since Model A2 (used in the inference) generally predicts a smaller shorter incubation periods compared to Model D (see Fig. 1), the rise in the epidemic curve, as simulated with Model D, will be slower than that predicted by Model A2. When this data is interpreted using Model A2, the posterior distribution compensated for the slower growth by underestimating N i.e. by suggesting a smaller attack. In Fig. 3 we plot PDFs for Cases C and D. We see that the PDF for τ for Case D is bimodal when only 3 days of data are available; with more data, the PDF is unimodal. The PDF for $\log_{10}(D)$ for Case C ($N = 161$) is too broad to be very informative. The MAP estimate for N, τ and $\log_{10}(D)$ are tabulated in Table 2. The MAP estimate of N is very close to the correct value for Case C and lower than the true number for Case D (for reasons explained above). τ is inferred about a day too late. The inferred dose (MAP estimate) is within an half an order of magnitude of the true figure.

4.2 The Sverdlovsk incident of 1979

It is suspected that on April 2$^{\text{nd}}$, 1979, there was an accidental release of a high-grade anthrax formulation from a military facility in Sverdlovsk, Russia [6]. 70 people are be-

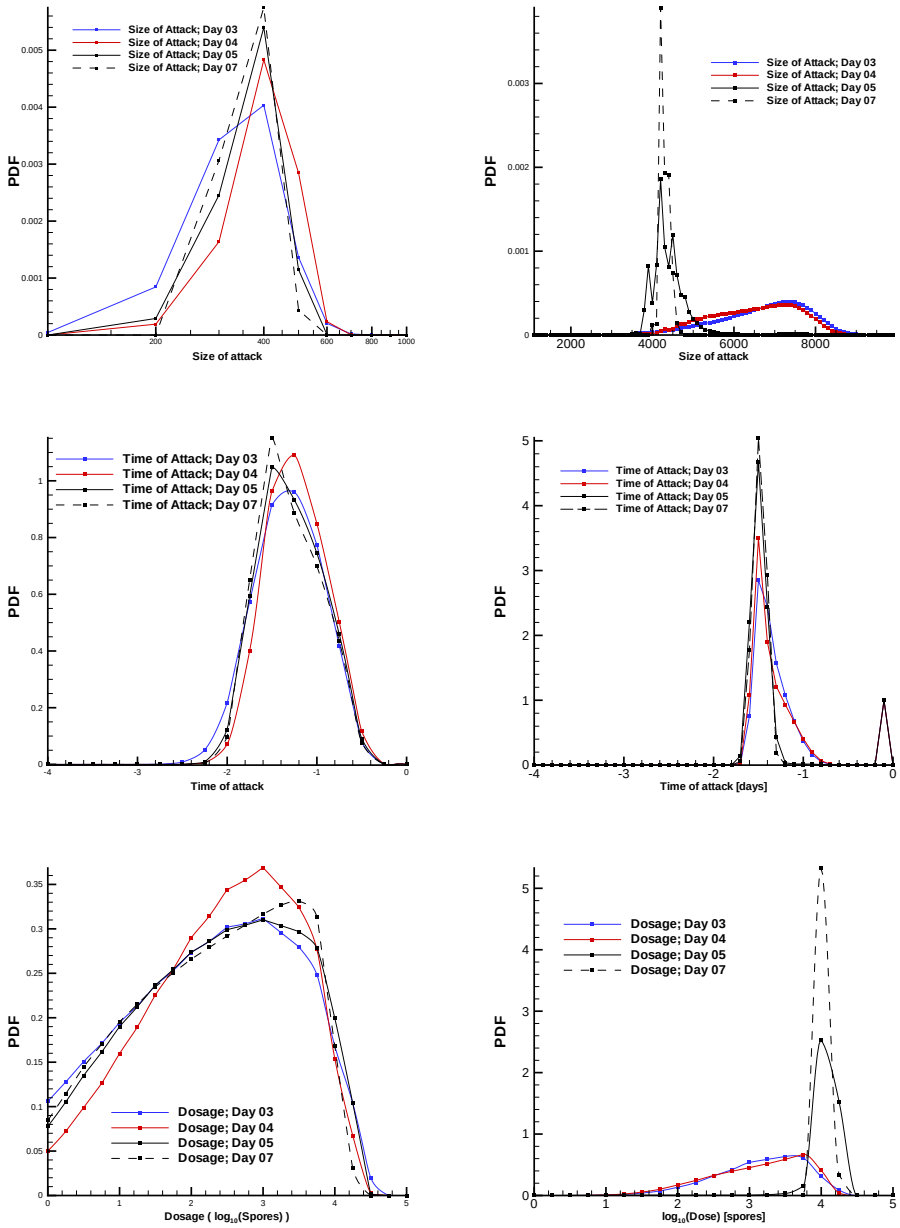

Figure 2: PDFs for N (top), τ and $\log_{10}(D)$ for Cases A (left) and B (right) as developed from the data tabulated in Table 1. The inferences are based on 3, 4 and 5 days of data (blue, red and black lines respectively).

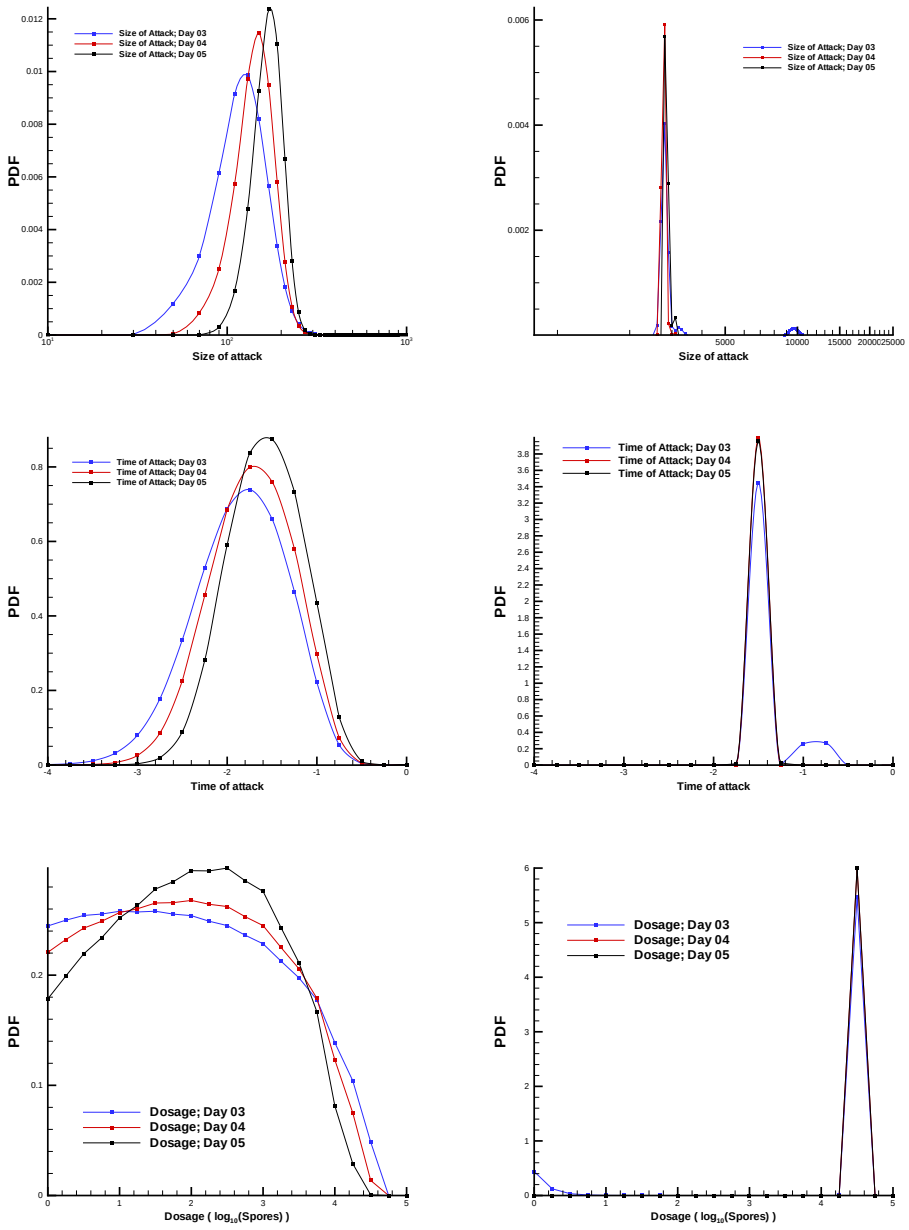

Figure 3: PDFs for N (top), τ and $\log_{10}(D)$ for Cases C (left) and D (right) as developed from the data tabulated in Table 1. The inferences are based on 3, 4 and 5 days of data (blue, red and black lines respectively).

Case	N	τ	$\log_{10}(D)$
A	400, {318}	-1.4=5, {-1.5}	3.0, {3.46}
	[225.8, 580.2]	[-1.97, -0.73]	[0.48, 4.12]
B	4200, {4437}	-1.5, {-1.25}	4.0, {4.09}
	[3894, 5143]	[-1.67, -1.33]	[3.8, 4.7]
C	170, {161}	-1.5, {-0.75}	2.5, {3.52}
	[124.9, 238.9]	[-2.4, -0.94]	[0.263, 3.87]
D	2800, {4453}	-1.5, {-0.75}	4.5, {4.20}
	[2726, 2998]	[-1.97, -1.3]	[4.272, 4.725]
Sverd-	50, {75–80}	-2, {-2}	1.3, { ? }
lovsk	[41.15, 66.49]	[-3.22, -1.38]	[0.18, 3.5]

Table 2: The MAP estimate and the 90% credibility intervals (in square brackets) for N, τ and $\log_{10}(D)$ as calculated from data in Table 1. Data from Day 5 (Day 9 for Sverdlovsk) are used. Correct values for N and τ are in { }. The "correct" representative dose is taken to be $\log_{10}(D_{50})$, also in { }.

lieved to have died [6, 8] and it has been estimated that 80 were infected [8]. This estimate was obtained under the assumption that all the fatalities were due to inhalational anthrax. The Sverdlovsk outbreak provides a good real-world test case for our inference procedure. Wilkening [17] estimates that the average dosage was either around 2-3 spores, based on his Model A, or around 300 spores based on his Model D, which is similar to the competing risks model of Brookmeyer [16]. Meselson [6] estimates 100-2000 spores as the likely dosage.

The Sverdlovsk case presents significant challenges. It was a low-dose attack infecting fewer than a hundred people. The first patient was detected on April 4^{th}, 1979. The time-series of symptom onset is available on a day-by-day basis in [26]. Around April 12^{th}, tetracycline was administered around Sverdlovsk; around the middle of April people were vaccinated. These prophylactic measures probably cured a few and increased the incubation period in others. Further, the data we work with almost certainly contains some recording errors. Noisiness of the data, the effect of prophylaxis (which is not modeled in our inference process), and the small size of the infected population are expected to stress our inference process.

In Fig. 4 we plot the inferences for N, τ and $\log_{10}(D)$, based on the data in [26]. Model A2 is used for inference. The data was collected on a daily basis for 42 days, the duration of the outbreak. The time of release τ was easy to infer. The PDFs for dosage (omitted here) are indeterminate. The inference for N centers around 50 by Day 9 (April 13^{th}), though the earlier inferences underestimate N. Nevertheless, we are certainly within a factor of two of the correct value of N even with 9 days of data. By Day 9 it is also clear that the the size of the outbreak would almost certainly be less than 100. However, medical measures during the outbreak affected about 59,000 people in the Chkalovskiy *raion*, of which about 80% were vaccinated at least once [6]. In Table 2, we summarize the MAP estimates and confidence intervals developed from 9 days of data.

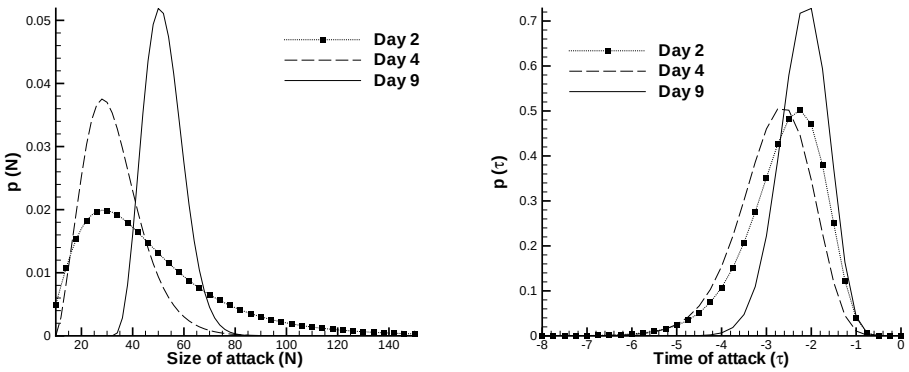

Figure 4: PDFs for the N and τ for the Sverdlovsk incident. The MAP values from 9 days of data are compared with the true values in Table 2.

5 Conclusions

We have developed a Bayesian approach to characterize BT attacks from a time series of diagnosed patients. Our tests with anthrax show that an observation period of 3–5 days may be sufficient to estimate the number of asymptomatic infected people, the time of infection, and a representative dose, and to provide quantified uncertainty intervals around these estimates; in the absence of an accurate disease model, we may arrive within a factor of two of the size of the attack. The resolution of the time series of diagnosed patients has a small impact if the disease model is accurate; otherwise, model errors dominate.

This Bayesian approach is amenable to extension and improvement in many ways. Informative prior distributions for N and τ, drawn from syndromic surveillance data, may increase the efficiency of the inference process. The ability to "fuse" disparate sources of data via prior distributions contributes significantly to the robustness of Bayesian inference in data-starved environments. One could also incorporate atmospheric transport processes into the likelihood function, thus using the spatial locations of diagnosed patients to guide posterior estimates, though for urban terrains this could lead to very involved computations. Also, the present approach can immediately be applied to other noncontagious diseases, as well as to contagious diseases with long incubation periods, such as smallpox, where secondary cases do not appear in the early time series of patient data.

The importance of quantitatively characterizing a BT attack was explicitly identified in the "Dark Winter" exercise [3]. Participants sought the ability "...to immediately predict the likely size of the epidemic on the basis of the initial cases; to know how many people were exposed." Thus the primary utility of our inference procedure is within the framework of a response plan. Response to a BT attack would typically involve confirmatory testing and logistics (the transport of medical materiel and personnel), both of which could be better targeted by a quantitative characterization of the attack. The probabilistic characterizations developed here, along with resource hedging for risk-mitigation, support a *measured* approach to addressing BT attacks. In addition to being more sustainable, measured responses may introduce fewer undesirable side effects and be less susceptible to feints.

References

[1] John A. Jernigan et al. Bioterrorism-related innhalational anthrax: The first 10 cases reported in the United States. *Emerging Infectious Diseases*, 7(6):933–944, 2001.

[2] Anna Goldenberg, Galit Shmueli, Richard A. Caruana, and Stephen E. Fienberg. Early statistical detection of anthrax outbreaks by tracking over-the-counter medication sales. *Proceedings of the National Academy of Sciences*, 99:5237–5240, 2002.

[3] Tara O'Toole, Michael Mair, and Thomas V. Inglesby. Shining light on Dark Winter. *Clinical Infectious Diseases*, 34:972–983, 2002.

[4] Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltan Toroczkai, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.

[5] Lynn Yang, Ben Wu, Dawn Manley, Julie Freutel, Marilyn Hawley, Keith Vanderveen, Jaideep Ray, Zach Heath, Christine Yang, Donna Djordjevich, Han Lin, John Jungels, Luis Hernandez, Bill Wilcox, Andrew Rothfuss, Michael Chen, Stephen Mueller, Howard Hirano, and Michael Johnson. BioNet Systems Modeling and Analysis. Technical Report SAND2005-7842, Sandia National Laboratories, Albuquerque, NM, 2005. Unclassified limited release.

[6] Matthew Meselson, Jeanne Guillemin, Martin Hugh-Jones, Alexander Langmuir, Ilone Popova, Alexis Shelokov, and Olga Yampolskaya. The Sverdlovsk anthrax outbreak of 1979. *Science*, 266:1202–1208, 1994.

[7] J. Walden and E. H. Kaplan. Estimating time and size of bioterror attack. *Emerging Infectious Diseases*, 10(7):1202–1205, 2004.

[8] R. Brookmeyer, N. Blades, M. Hugh-Jones, and D. A. Henderson. The statistical analysis of truncated data: application to the Sverdlovsk anthrax outbreak. *Biostatistics*, 2:233–247, 2001.

[9] Ron Brookmeyer and Natalie Blades. Statistical models and bioterrorism : Application to the U.S. anthrax attacks. *Journal of the American Statistical Association*, 98(464):781–788, 2003.

[10] D. W. Henderson, S. Peacock, and F. C. Belton. Observations on the prophylaxis of experimental pulmonary anthrax in the monkey. *Journal of Hygiene*, 54:28–36, 1956.

[11] Wilhelm S. Albrink and Robert J. Goodlow. Experimental inhalational anthrax in the chimpanzee. *The American Journal of Pathology*, 35(5):1055–1065, 1959.

[12] C. A. Gleiser, C. C. Berdjis, H. A. Hartman, and W. S. Gochenour. Pathology of experimental respiratory anthrax in *Macaca Mulatta*. *British Journal of Experimental Pathology*, 44:416–426, 1963.

[13] M. M. Friedlander, S. L. Welkos, M. L. Pitt, J. W. Ezzell, P. L. Worsham, K. J. Rose, B. E. Ivins, J. R. Lowe, G. B Howe, P. Mikesell, and Wade B. Lawrence. Postexposure prophylaxis against experimental inhalation anthrax. *Journal of Infectious Disease*, 167(5):1239–1243, 1993.

[14] B. E. Ivins, M. L. M. Pitt, P. F. Fellows, J. W. Farchaus, G. E. Benner, D. M. Waag, S. F. Little, G. W. Anderson, P. H. Gibbs, and A. M. Friedlander. Comparative efficacy of experimental anthrax vaccine candidates against inhalational anthrax in rhesus macaques. *Vaccine*, 16(11/12):1141–1148, 1998.

[15] Philip S. Brachman, Arnold F. Kaufmann, and Frederic G. Dalldorf. Industrial inhalational anthrax. *Bacteriological Reviews*, 30(3):646–657, 1966.

[16] R. Brookmeyer, E. Johnson, and S. Barry. Modelling the incubation period of anthrax. *Statistics in Medicine*, 24:531–542, 2005.

[17] D. Wilkening. Sverdlovsk revisted : Modeling human inhalational anthrax. *Proceedings of the National Academy of Science*, 103(20):7589–7594, May 2006.

[18] W. R. Hogan, G. Cooper, G. L Wallstrom, and M. Wagner. An Improved Bayesian Aerosol Release Detector. Technical report, The RODS Laboratory, 550 Cellomics Building, 100 Technology Drive, Pittsburgh, PA 15219, 2005.

[19] Real-time Outbreak and Disease Surveillance. http://rods.health.pitt.edu/default.htm.

[20] Fu-chang Tsui, Jeremy U. Espino, Virginia M. Dato, Per H. Gesteland, Judith Hutman, and Michael M. Wagner. Technical Description of RODS : A Real-time Public Health Surveillance System. *Journal of the Americal Medical Informatics Association*, 10(5):399–408, 2003.

[21] H. N. Glassman. Discussion on industrial inhalational anthrax. *Bacteriological Review*, 30:657–659, 1966.

[22] H. A. Druett, D. W. Henderson, L. Packman, and S. Peacock. Studies on respiratory infection. I.The influence of particle size on respiratory infection with anthrax spores. *Journal of Hygiene*, 51:359–371, 1953.

[23] J. Ray, Y. M. Marzouk, M. Kraus, and P. Fast. A Bayesian method for characterizing distributed micro-releases: II. Iinference under model uncertainty with short time-series data. SAND Report SAND2006-7568, Sandia National Laboratories, Livermore, CA 94551-0969, January 2007 2006. Unclassified unlimited release.

[24] Frederick Klein, Jerry S. Walker, David F. Fitzpatrick, Ralph E. Lincoln, Bill G. Mahlandt, William L. Jones, James P. Dobbs, and Kenneth J. Hendrix. Pathophysiology of anthrax. *Journal of Infectious Diseases*, 116(2):123–138, 1966.

[25] J. Ray, Y. M. Marzouk, H. N. Najm, M. Kraus, and P. Fast. A Bayesian method for characterizing distributed micro-releases: I. The single-source case for non-contagious diseases. SAND Report SAND2006-1491, Sandia National Laboratories, Livermore, CA 94551-0969, March 2006. Unclassified unlimited release.

[26] Thomas V. Inglesby, Donald A. Henderson, John G. Barteltt, Michael S. Ascher, Edward Eitzen, Arthur M. Friedlander, Jerome Hauer, Joseph McDade, Michael T. Osterholm, Tara O'Toole, Gerald Parker, Trish M. Perl, Philip K. Russel, and Kevin Tomcat. Anthrax as a biological weapon - Medical and public health management. *J. Am. Med. Assoc.*, 281(18):1735–1745, 1999.