



Bayesian modeling of temporal properties of infectious disease in a college student population

Zhengming Xing, Bradley Nicholson, Monica Jimenez, Timothy Veldman, Lori Hudson, Joseph Lucas, David Dunson, Aimee K. Zaas, Christopher W. Woods, Geoffrey S. Ginsburg & Lawrence Carin

To cite this article: Zhengming Xing, Bradley Nicholson, Monica Jimenez, Timothy Veldman, Lori Hudson, Joseph Lucas, David Dunson, Aimee K. Zaas, Christopher W. Woods, Geoffrey S. Ginsburg & Lawrence Carin (2014) Bayesian modeling of temporal properties of infectious disease in a college student population, Journal of Applied Statistics, 41:6, 1358-1382, DOI: [10.1080/02664763.2013.870138](https://doi.org/10.1080/02664763.2013.870138)

To link to this article: <https://doi.org/10.1080/02664763.2013.870138>



© 2013 The Author(s). Published by Taylor & Francis.



[View supplementary material](#)



Published online: 23 Dec 2013.



[Submit your article to this journal](#)



Article views: 1348



[View related articles](#)



[View Crossmark data](#)



Citing articles: 2 [View citing articles](#)

Bayesian modeling of temporal properties of infectious disease in a college student population

Zhengming Xing^a, Bradley Nicholson^b, Monica Jimenez^b, Timothy Veldman^c, Lori Hudson^c, Joseph Lucas^{c,d}, David Dunson^d, Aimee K. Zaas^{c,e}, Christopher W. Woods^{b,c,e}, Geoffrey S. Ginsburg^{a,c} and Lawrence Carin^{a,c*}

^aElectrical & Computer Engineering Department, Duke University, Durham, NC, USA; ^bDurham Veterans Affairs Medical Center, Durham, NC, USA; ^cDuke Institute for Genomic Sciences & Policy, Durham, NC, USA; ^dStatistical Sciences Department, Duke University, Durham, NC, USA; ^eDuke University Medical Center, Durham, NC, USA

(Received 25 March 2013; accepted 25 November 2013)

A Bayesian statistical model is developed for analysis of the time-evolving properties of infectious disease, with a particular focus on viruses. The model employs a latent semi-Markovian state process, and the state-transition statistics are driven by three terms: (i) a general time-evolving trend of the overall population, (ii) a semi-periodic term that accounts for effects caused by the days of the week, and (iii) a regression term that relates the probability of infection to covariates (here, specifically, to the Google Flu Trends data). Computations are performed using Markov Chain Monte Carlo sampling. Results are presented using a novel data set: daily self-reported symptom scores from hundreds of Duke University undergraduate students, collected over three academic years. The illnesses associated with these students are (imperfectly) labeled using real-time (RT) polymerase chain reaction (PCR) testing for several viruses, and gene-expression data were also analyzed. The statistical analysis is performed on the daily, self-reported symptom scores, and the RT PCR and gene-expression data are employed for analysis and interpretation of the model results.

Keywords: infectious disease; Bayesian; semi-Markov

1. Introduction

There has been significant interest in the analysis of community-to-individual and individual-to-individual transfer of virus [1,3,7,16,20,26]. Many of these studies have been concerned with infection transfer in households [3,7,20], in confined spaces like elementary schools [26], as well as transfer among domestic animal populations [16]. The modeling may take different forms,

*Corresponding author. Email: lcarin@duke.edu

depending upon the data considered and questions being asked. There has been a significant interest in influenza, and in such studies the interest is typically on influenza-like illness (ILI). For example, based on the incidences of ILI, one may be interested in the analysis of large-scale dynamics of epidemic propagation [1,10,14], where in this case the data may be counts of space–time ILI events. There are other studies for which the modeling is performed at or near the level of the symptoms or biomarkers, which are noisy and often imperfect [16]. The studies are complicated by missing and incomplete data, and an unknown number of competing pathogens [26].

Infection dynamics are complex, and therefore the power and flexibility of Bayesian models are attractive [1,3,7,26]; we employ such a modeling approach in this study. Most of these models assume that a given individual is in a particular state of health, such as susceptible (S), exposed (E), infective (I) and recovered (R); an individual in state I is infectious, in that they are capable of transmitting the virus. The sequence of states considered in such a model defines its character, for example susceptible–infective–recovered models are widely considered [14], and once in the R state individuals are often assumed removed from the population from the standpoint of infection transfer (because of acquired immunity, or because of death; in some settings R represents ‘removed’ rather than recovered).

When dealing with many competing pathogens [26], such as distinct viruses characteristic of common/typical colds, even after an individual recovers from one virus, he or she may soon be susceptible to another. That is the case considered in this paper, which motivates an susceptible–infective–susceptible (SIS) model. We effectively ignore the short time period that may exist between state R and the return to S ; during this short time period there may be some cross-immunity between pathogens [26]. Additionally, we do not explicitly model the distinction between states E and I , as these states are not distinguishable with the data considered. With the observed data under consideration, only symptoms allow distinction of states, and therefore we assume state S is one characterized by no or minimal symptoms, and state I is one in which symptoms are observed (there are complications with this definition of state I , as discussed below).

There are potential pitfalls associated with attempting to model person-to-person transfer, when this mechanism is tied to symptoms. It has been observed that individuals may be infected with a virus but display no symptoms [21]; additionally, for those who do ultimately have symptoms, pathogen transfer may occur in the presymptomatic state [21]. In other words, even though the absence of symptoms from an individual may indicate that she is in state S , she in fact may be in state I , and she may transfer/shed virus. The effectiveness of pathogen transfer from asymptomatic shedders is not well understood. Additionally, the data of interest for person-to-person transfer may be incomplete, in that it only accounts for a subset of close contacts. In [3] the authors modeled infection transfer within elementary schools, but not within the households of the students; in [26] the authors modeled person-to-person transfer within households, but not among other close contacts outside the households. For these reasons, and because of the characteristics of the data considered here (detailed in Section 2), we only model community-to-person transfer. This allows for the possible transfer of pathogen from an asymptomatic (but virus shedding) member of the community to a member of our study. However, an asymptomatic member of our study will still be deemed in state S by our model, based on symptoms, even though they may be in state I and asymptotically shedding virus; this issue is examined in detail when presenting results.

Within the proposed SIS model, we assume that the probability of transferring from state S to state I is time-dependent. Further, we assume that different individuals have distinct degrees of susceptibility to common viruses, and this is modeled as well (i.e. there is a person-dependent character to the degree of susceptibility, and hence to the characteristics of state S). A unique aspect of this study is that the modeling is performed directly at the level of observed symptoms, rather than using pre-specified means of defining whether one is in the S or I state. Specifically, in most of the above studies the state S/I of the individual was assumed observed, and the goal

was to infer the statistics of the state dynamics (e.g. the probability of transiting from S to I , and the duration of being in state I). In this study the symptoms are the observed data, and the state S/I is treated as being *latent*, and to be inferred. As discussed in Section 2, we also have access to (imperfect) labels on the health of the individual at a given time, based upon real-time (RT) polymerase chain reaction (PCR) testing and gene-expression data. We compare the model-inferred state of the individual to the state based upon RT PCR and gene expression. This comparison provides insights into such mechanisms as the aforementioned asymptomatic virus shedders, as well as individuals who are symptomatic but not in state I as defined by RT PCR.

The time dependence in the probability of transiting from S to I captures time variation in the viruses present at a given time, as well as time-dependent dynamics of human interaction (e.g. the mixing of a new set of people may increase virus transfer [17,23]). A unique characteristic of the data considered here is that it is collected daily, for an entire university academic year; further, we have data from three full academic years. By comparison, the data in [3] only existed for 15 days after a household index case. The long time scale, and the daily sampling, introduce interesting phenomena that have not been investigated previously, to our knowledge (in [1] weekly sampling of symptoms was performed). Specifically, how one reports symptoms may be linked to their mood, which may vary with the day of the week. For example, it has been demonstrated that the way in which individuals rate music is linked to the day (and even time) of reporting [15,22]. Since we are analyzing symptom data, we must consider biases in data reporting that may occur based upon the day of reporting. The degree of missingness tends to also be linked to the day of the week. This therefore motivates employing a semi-periodic, or weekly effect in the probability of transiting from state S to I . This is a novel characteristic of the model developed here, in which we generalize the use of models that employ seasonal terms [25] (here they become weekly, and they are modeling a latent process).

Covariates may be available that can be employed to impact the probability of transiting from state S to I . Given the very long longitudinal length of our data, we may consider new forms of covariates, becoming available from a web-centric world. We consider the Google Flu Trends data [13]. These covariates are constituted by region-specific web searches of words linked to ILI, and specifically here we employ the Google Flu Trends for Durham, NC, the city in which Duke University resides. Other recent statistical analyses have modeled the space-time properties of such Google data alone [11,12], but here we are focused on observed symptoms and the time-dependent Google data is a covariate. In [11] the focus is on modeling an epidemic like influenza, and therefore they employ a susceptible-exposed-infected-recovered (SEIR) model; once in the recovered state, an individual is effectively removed from the pandemic dynamics because of immunity. Here we are interested in modeling long-term illness dynamics from common viruses (producing ILI), in addition to influenza, and therefore removing an individual from the population upon recovery is not appropriate (in the data we observe some people with repeated ILI). This may motivate an SEIS model, but for simplicity we consider an SIS model [14].

2. Data and motivating questions

2.1 Self-reported daily symptom data

Self-reported symptom-score data were collected from undergraduate students at Duke University, following guidelines specified by the Duke Institutional Review Board. Data were collected daily during the 2009–2010, 2010–2011 and 2011–2012 academic school years; in each year, data were collected from the beginning of September until May, using a web-based tool. For each of these collection periods, respectively 246, 378 and 242 students participated. The total number of days in which data were recorded were, respectively, 222, 214, and 227 over each collection period. The 2009–2010 collection period coincided with the novel H1N1 pandemic [24].

The students' reported symptoms were routinely monitored by Duke University health professionals. When a student was deemed – from reported symptoms – to likely be sick with an infectious disease (e.g. virus), the student was contacted and nasal and blood samples were taken. These are termed index cases. Further, each student provided a list of close contacts (other students they interacted with frequently). Blood samples were then collected daily for a week on these close contacts, with the hope that we may observe the transfer of infectious disease (and to analyze that in the context of the blood samples).

2.2 Virus identification and gene-expression data

For the students from whom samples were collected, RT PCR testing was done for a set of viruses. The particular viruses for which a RT PCR test was available were: Rhinovirus, Coxsackie, Echovirus, Coronavirus (229E, HKU1, NL63, OC43), Parainfluenzavirus (1, 2, 3, 4), RSV A/B, Influenza A, Influenza B, Metapneumovirus (A & B), and Adenovirus E & B, and Bocavirus (platform used: Qiagen ResPlex II V2.0). Therefore, *if* one of these viruses was responsible for the student's illness, and *if* the virus was present in the collected sample, and *if* the RT PCR test worked properly, then the virus type responsible for illness can be detected. However, a negative RT PCR result does not necessarily imply that the student is not sick with a virus, as there may have been a poor sample (in which markers of the virus were not present), a non-tested (within the RT PCR library) virus may have been responsible for illness, and the RT PCR test is itself imperfect. In the context of this study, across the three years, 897 viral etiology results based on RT PCR were constituted.

In addition to the aforementioned RT PCR tests, we also used the available blood samples to perform gene-expression analysis. Let $\mathbf{x}_q \in \mathbb{R}^G$ represent the expression data for subject q , for G genes. We performed sparse Bayesian factor analysis on the set of data $\mathbf{X} \in \mathbb{R}^{G \times Q}$, where column q of \mathbf{X} corresponds to \mathbf{x}_q . Details on the factor analysis method may be found in [2,4,6]. In [4,6,28] it was demonstrated that one of the factors in such an analysis may be linked to the host response to virus, and in multiple experiments this signature has been found invariant to the particular type of virus studied. Therefore, for the Q samples defining \mathbf{X} , we used this factor analysis to define which of these subjects appear to be infected by a virus (by the presence of an elevated form of a specific factor [4,6,28]). We emphasize that this test is also imperfect, but it provides more generality than the RT PCR tests, which are constrained to specific types of viruses. In these experiments $Q = 34$ and $G = 22, 277$.

Based upon the RT PCR and gene-expression tests outlined above, we can (imperfectly) label each of the subjects for whom samples were collected as being sick with a virus or not, at a given point in time; these tests are used to assess the quality of the model developed in Section 3 to detect sickness based on the symptom scores alone (with state of health defined by the inferred latent state, S or I). It is important to emphasize that the labels we will use to assess performance are imperfect, in the sense that the RT PCR and gene-expression tests are only testing for the presence of a virus (and even these tests are not perfect). It is possible that an individual may be sick for another reason (e.g. due to allergies or bacteria); in this case the virus-driven labels may indicate that the student is not sick, while the symptoms indicate otherwise (our symptom-based declaration of health or sick may indicate state I , while the virus-driven labels may indicate S). These issues will be revisited when presenting results.

2.3 Impact of form of data on the developed model

This paper is principally directed toward analyzing the self-reported symptom-score data, with a focus on community-to-person transmission of pathogens. In each student dorm (living facility), roughly 10–20% of the students participated in the study, and therefore the close contacts are

very sparsely sampled. Further, many of the students spend most of their time outside the dorm, and interact infrequently with many members of the same dorm (SD). It was therefore deemed inappropriate to try to model person-to-person pathogen transfer. We also considered developing dorm-dependent models for pathogen transfer, but the dynamics across the different dorms (e.g. fraction of students sick at any given time) did not vary substantially, and therefore it was deemed most appropriate to develop a single community-to-person model for all students who participated in the study.

However, as detailed below, the close contacts constitute a separate set of data (with associated ‘truth’ for the presence/absence of virus), within the context of the imperfections of the RT PCR and gene-expression data. We therefore use these data for model testing.

2.4 Questions to be examined in this study

- We have access to data over three academic years, with one year corresponding to the presence of novel H1N1 virus. During that year there was heightened awareness on campus about non-pharmacological ways to reduce virus transmission, with many highly visible reminders (e.g. students were prominently reminded about hand washing, use of disinfectants, not touching eyes and nose, etc.). Disinfectant soap was widely accessible throughout the campus, at locations in which people congregate. We wish to examine how this heightened awareness affected the time-dependent hazard of community-to-person transmission of pathogens, relative to the other two years of the study, in which pathogen transmission was far less of a focus.
- The students who participated in this study primarily lived on a separate campus dedicated for first-year students. Therefore, most of the students were Freshman, and at the beginning of the academic year most of these students were coming together, and living in close proximity (in dorms), for the first time. We wish to examine the impact of this new mixing of people on the time-dependent hazard of community-to-person transmission of pathogens.
- The Duke University campus resides within the surrounding city of Durham, NC. We wish to examine how the time-dependent hazard of community-to-person transmission of pathogens of Duke students relates to such metrics as Google Flu Trends. Specifically, we wish to examine the extent to which Google Flu Trends for Durham, NC predicts the hazard of pathogen transmission on the Duke campus.
- We have access to which dorm room each student lived in. Based upon the symptom scores, the model predicts whether each student is infected at a given time. While we do not explicitly model person-to-person transmission within the model (for reasons stated above), we may use model predictions on the state of health to examine whether someone getting infected at time t within a given dorm raises (or lowers) the incidence of infection of other students in the study who lived in the SD (and more specifically, on the SD floor). For example, an infected neighbor in a dorm may *heighten* awareness of the danger of pathogen transfer, yielding phenomenon like that associated with the exposure to novel H1N1 (heightened awareness, and hence precautions). We examine this issue in detail, as a function of the type of virus associated with each index case (with virus type imperfectly determined via RT PCR).
- We examine and analyze real-world characteristics of daily self-reported symptom data. This includes day-of-the-week dependent phenomenon in the data (weekly, semi-periodic effects), and connections to data missingness.
- We examine the utility of using symptoms alone for classification of the latent state S/I , with comparisons to RT PCR. This is of clinical relevance, as clinicians typically make a diagnosis based directly on symptoms. We also examine the presence of non-symptomatic individuals who are shedding the virus. Further, we examine cases for which symptoms are clear, but extensive RT PCR testing is negative.

3. Methods

3.1 Observed symptoms and the latent state of health

Assume access to self-reported data from N individuals, provided daily over multiple months. The data correspond to the strength of various infectious-disease-related symptoms, reported separately by each of the N students. Eight symptoms are recorded: nasal discharge, nasal congestion, sneezing, cough, malaise, throat discomfort, fever and headache. Each of the eight symptoms is reported on an ordinal scale, from 0 to 4, with 0 being no symptoms, and 4 ‘maximum’ symptoms. Before the study each of the students is instructed on how to connect perceived symptoms to this scale. Nevertheless, there is clearly subjectivity to the mapping from perceived symptoms to ordinal data, and this subjectivity should be accounted for in the statistical analysis. We note that such subjectivity is always present when individuals report symptom severity to a doctor or nurse.

Let $\mathbf{y}_{nt} \in \{0, \dots, M\}^J$ represent the J symptom scores reported by individual $n \in \{1, \dots, N\}$ on day t , where for our study $J = 8$ and $M = 4$; we use generalized notation because the basic modeling strategy may be applied to other types of related data. It is assumed that, at a given time, individual n is either in an infective state I or in susceptible state S . When in state S , the student is not currently sick from a virus, and therefore does not display ILI symptoms; however, the student is assumed susceptible to virus infection. When in state S , different individuals may have distinct levels of susceptibility to virus-borne illness, and this is accounted for in the model. We have tied state I to symptoms, as is common [1,3,7,16,20,26]. However, there are asymptomatic individuals who shed virus [21] and hence are in the infective state I ; these individuals are identified and discussed when presenting results.

We employ an ordinal probit model to link the J observed symptoms to the latent state. Specifically, consider $\mathbf{r}_{nt} \in \mathbb{R}^J$, drawn conditioned on the latent state as

$$\mathbf{r}_{nt} \mid z_{nt} \sim \mathcal{N}(\boldsymbol{\mu}_{z_{nt}}, \boldsymbol{\Sigma}_{z_{nt}}^{-1}), \tag{1}$$

where $z_{nt} = S$ or $z_{nt} = I$. Let r_{njt} represent the j th component (symptom) of \mathbf{r}_{nt} , and y_{njt} similarly represent the j th component of \mathbf{y}_{nt} . The mapping from real r_{njt} to ordinal y_{njt} is manifested via a traditional probit model as

$$y_{njt} = m \quad \text{if } \tau_{j,m-1} < r_{njt} \leq \tau_{j,m}, \tag{2}$$

where each $\tau_{j,m} \in \mathbb{R}$, $\tau_{j,m-1} < \tau_{j,m}$, $\tau_{j,-1} = -\infty$, $\tau_{j,0} = 0$, and $\tau_{j,M} = \infty$. We wish to infer $\{\tau_{j,1}, \dots, \tau_{j,M-1}\}$, with this performed by considering an improper uniform prior on $\tau_{j,1} < \dots < \tau_{j,M-1}$ [19]. Improper uniform priors result in proper posterior distributions under mild conditions, as detailed in [19], yielding practically useful sufficient conditions that are met in our study. As discussed in Appendix A, for identifiability purposes the covariance matrices $\boldsymbol{\Sigma}_S^{-1}$ and $\boldsymbol{\Sigma}_I^{-1}$ are restricted to correspond to correlation matrices [5], with diagonal elements all equal to one.

3.2 Semi-Markov latent-state dynamics

The probability of individual n transiting from a state of susceptibility at time $t - 1$, $z_{n,t-1} = S$, to a state of illness/infection at time t , $z_{n,t} = I$, is modeled as

$$p(z_{n,t} = I \mid z_{n,t-1} = S) = \Phi(\gamma_{nt}), \tag{3}$$

where $\Phi(x) = \int_{-\infty}^x d\eta \mathcal{N}(\eta; 0, 1)$ is a cumulative distribution function, with $\mathcal{N}(\eta; 0, 1)$ a normal distribution function for variable η , characterized by zero mean and unit variance (probit transition statistics). The model of the time-evolving variable $\gamma_{nt} \in \mathbb{R}$ is discussed in Section 3.3.

Related forms of time and covariate dependent probabilities of community-to-person transfer have been considered in [26]; however, in that work, and much of the literature, the state S or I was assumed observed, where here the state is latent and is to be inferred upon the observed symptoms.

If individual n transits from $z_{n,t-1} = S$ to $z_{n,t} = I$, then it is assumed that $z_{n,t+d} = I$ for $0 \leq d \leq D_{nt}$, where D_{nt} is a random variable defining the number of days of infection. We employ the model

$$D_{nt} = c + \hat{D}_{nt}, \quad \hat{D}_{nt} \sim \text{Pois}(\lambda_n), \quad (4)$$

where $c > 0$ is a minimum number of days infected, and the rate parameter λ_n is assumed drawn from a gamma distribution. We discuss setting c when presenting experimental results; the imposition of a lower bound c on the number of days of being infected (i.e. in the state I) helps distinguish isolated days when one may not feel well, for various reasons, from actual extended periods of infection.

In [1] the length of time D_{nt} in state I was a real random variable, and was drawn from a gamma distribution. Here we observe discrete temporal data (days), and employ Poisson random variables for the length of time in state S ; the lower bound c assures that we are not undermined by draws from $\text{Pois}(\lambda_n)$ that could be equal to zero.

We apply the following approach to help the model distinguish between states S and I , imposing a strong identifiability condition. Based upon the model construction above, the only way states S and I are distinguished is via a requirement that an individual remain in the I state for a minimum of c days. Based on expertise of the infectious disease medical doctors who are co-authors on this study, we set $c = 3$ days, consistent with the minimum length of time one would be anticipated to manifest symptoms due to infectious disease of the type associated with common viruses. In addition, recall that a subset of the subjects within the study were confirmed via RT PCR and/or gene-expression analysis to be ill due to a virus. A small subset of these infective individuals had their data removed from the subsequent analysis, and the correlation between the symptoms of this subset of confirmed cases were used to set the hyperparameters in the prior for Σ_I . This setting of the model parameters significantly distinguished the S and I states, yielding interpretable results.

Note that we assume that the statistics of the symptoms for the infective individuals, characterized by $\mathcal{N}(\mu_I, \Sigma_I^{-1})$, are independent of the length of time in which the individual has been in state I . This is a modeling simplification, and one may also link the symptom statistics to the length of time the subject has been in the infective state. The variability in the symptom scores within a given state, characterized by μ_I and Σ_I , account for variability in how a given individual maps perceived symptom strength to ordinal values. Further, Σ_I accounts for variability in symptom strength across an extended period of infection (typically from weak, to strong, and back to weak symptoms over the period of infection).

3.3 Modeling the time-dependent probability of becoming infected

The time-evolving parameter γ_{nt} , in concert with a probit link function, defines the probability with which one transits from a susceptible to infective state. We model this time-evolving parameter via four terms:

$$\gamma_{nt} = a_n + \sum_{i=1}^3 \gamma_t^{(i)} \quad (5)$$

with $\gamma_t^{(1)}$ modeling the general trend within the population to become infected, $\gamma_t^{(2)}$ is associated with periodic (weekly) effects characterizing unique aspects of the day of the week, and

$\gamma_t^{(3)}$ is a regression term. Concerning $\gamma_t^{(3)}$, we specifically perform regression to the Google Flu Trends data. Note that $\{\gamma_t^{(i)}\}_{i=1,3}$ are independent of the individual index n , and are therefore shared across the population. The term $a_n \in \mathbb{R}$ is an individual-dependent tendency to get infected, which we place a normal prior on (when a_n is large and positive the n th individual has a heightened susceptibility toward illness, with the opposite true when a_n is negative with large magnitude).

The model in [26] also imposed covariate-dependent state-transition statistics. However, the length of the data considered in that study, and in most of the literature, precluded the need to consider semi-periodic terms. Further, most such models are not performed at the level of symptoms, and therefore they do not have to address semi-periodic missing data phenomenon, and other characteristics of the symptoms.

3.3.1 General-trend term

An autoregressive model is employed for $\gamma_t^{(1)}$:

$$\gamma_t^{(1)} \sim \mathcal{N}(\omega\gamma_{t-1}^{(1)}, \beta^{-1}), \tag{6}$$

where a gamma prior is placed on β and $\omega \in (0, 1)$ is drawn from a truncated normal distribution, $\omega \sim \mathcal{N}_{(0,1)}(\mu_\omega, \sigma_\omega)$. This imposes that the time dependence of the general trend toward illness varies smoothly.

3.3.2 Weekly or periodic term

The observed data are characterized by clear dependencies on the day of the week on which symptoms are reported. The day of the week may impact general feelings of well being (Monday vs. Friday), and certain portions of the week may be characterized by heightened student workload, stress, and lack of sleep/exercise. A 7-day semi-periodic term is therefore employed to model $\gamma_t^{(2)}$. This term builds upon modeling strategies discussed in [25, Chapter 8.6]; for completeness, we here provide some details.

Using notation from [25], the unique terms of a periodic function may be represented in terms of φ_j , for $j = 0, \dots, p - 1$, where p is the period of the weekly/repeating term (for our problem $p = 7$, for the days of the week). Basic Fourier analysis dictates that

$$\varphi_j = a_0 + \sum_{r=1}^h [a_r \cos(\alpha r j) + b_r \sin(\alpha r j)] = a_0 + \sum_{r=1}^h S_r(j), \tag{7}$$

where h is the largest integer not exceeding $p/2$, $\alpha = 2\pi/p$, and $a_r \in \mathbb{R}$ and $b_r \in \mathbb{R}$ are Fourier components. One may readily demonstrate that

$$S_r(j) = \mathbf{e}^T \boldsymbol{\theta}_r(j), \quad \boldsymbol{\theta}_r(j) = \mathbf{J}(\alpha r) \boldsymbol{\theta}_r(j - 1), \tag{8}$$

where $\boldsymbol{\theta}_r(0) = (a_r \ b_r)^T$, $\mathbf{e} = (1 \ 0)^T$ and $\mathbf{J}(w) = \begin{pmatrix} \cos(w) & \sin(w) \\ -\sin(w) & \cos(w) \end{pmatrix}$.

Generalizing Equation (8) to the stochastic case, and motivated by a dynamic linear model [25], the time dependence of the r th Fourier component within a particular time period is modeled as

$$S_r(j) = \mathbf{e}^T \boldsymbol{\theta}_r(j) + v_r(j), \quad \boldsymbol{\theta}_r(j) = \mathbf{J}(\alpha r) \boldsymbol{\theta}_r(j - 1) + \boldsymbol{\epsilon}_r(j), \tag{9}$$

where the components of the two-dimensional vector $\boldsymbol{\epsilon}_r(j)$ are drawn $\boldsymbol{\epsilon}_r(j) \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{\theta_r}^{-1})$, and $v_r(j) \sim \mathcal{N}(0, \zeta_{S_r}^{-1})$, with a gamma prior placed on ζ_{S_r} and a Wishart prior on $\boldsymbol{\Sigma}_{\theta_r}$. The term

$\epsilon_r(j)$ models noise in the Fourier components over a given time period, and $v_r(j)$ represents measurement noise.

A prior is placed on $\theta_r(0)$, corresponding to the Fourier components over the first week of data, and then Equation (9) is repeated cyclically over the multiple weeks, through sequential draws of $\{v_r(j)\}$ and $\{\epsilon_r(j)\}$. Note that with the zero mean priors on $\{v_r(j)\}$ and $\{\epsilon_r(j)\}$, conditioned on $\theta_r(0)$, the expectation of Equation (9) corresponds to Equation (8). A zero-mean normal prior is placed on $\theta_r(0)$, for each r . With $S_r(j)$ so drawn, one may superpose the Fourier components to constitute $\gamma_t^{(2)}$; for the weekly data under consideration, there are $h = 3$ Fourier components, in addition to the mean a_0 .

It is important to emphasize that the day-dependent term is stochastic, and it is not the same for each week (all the different weeks share the same basic Fourier model [25], but there is statistical variation from week-to-week). We also considered a fixed day-of-the-week dependent term, in which $\gamma_t^{(2)}$ was modeled as a constant 7-dimensional vector for each week. This was found to not be as effective as the above model, which allows for differences in the characteristics of weeks dependent upon the time of the semester (e.g. exam periods are more stressful, and characterized often by less sleep and rest than other parts of the semester).

3.3.3 Regression term

Assume that we have access to a time-dependent covariate f_t , which in our problem corresponds to the Google Flu Trends [13] data for the region in which the individuals under study reside. The regression term is modeled as

$$\gamma_t^{(3)} \sim \mathcal{N}(\xi f_t, \alpha_f^{-1}), \tag{10}$$

where a zero-mean normal prior is placed on ξ and a gamma prior is placed on α_f .

3.4 Clustering tendency toward infection, and length of infection

In the previous section it was assumed that the parameters λ_n and a_n were drawn i.i.d., with the former controlling the length of time individual n tends to be in an infective state, and the latter controlling the tendency of individual n to get infected. The parameter a_n has the impact of controlling the degree to which one is susceptible to virus, and hence to transition from state S to I (large a_n implies higher susceptibility).

It is anticipated that individuals may cluster in terms of their (e.g. genetic or behavioral) tendency to get infected, and in the length with which they stay infected. It is desirable to account for this in the model (it allows sharing of statistical strength between individuals). Additionally, for the data set that motivates this paper, we have access to the residence location of each student, and therefore it is possible to use this spatial information as a covariate. For example, one may consider the spatial location of each student when modeling the time-dependent tendency to get infected, via including spatial information in $\gamma_t^{(1)}$, for example. Other modeling issues discussed below include consideration of missing data, and the joint modeling of data from multiple years.

A natural means of clustering λ_n and a_n is to employ a Dirichlet process [18], with which the number of clusters may be inferred nonparametrically. Specifically, we draw

$$\lambda_n \sim G_\lambda, \quad G_\lambda \sim \text{DP}(\alpha_{0\lambda} G_{0\lambda}), \tag{11}$$

$$a_n \sim G_a, \quad G_a \sim \text{DP}(\alpha_{0a} G_{0a}), \tag{12}$$

where the base measures $G_{0\lambda}$ and G_{0a} correspond, respectively, to gamma and normal distributions. Gamma priors are placed on the DP parameters $\alpha_{0\lambda}$ and α_{0a} .

3.5 Spatial covariates

For the motivating data, we have knowledge of the residence location (dorm room) of each individual (student), and therefore it is possible to exploit spatial information when modeling the general trend toward being infected, reflected in $\gamma_t^{(1)}$. One could also consider utilizing spatial information when modeling the weekly (semi-periodic) term $\gamma_t^{(2)}$ and the regression term $\gamma_t^{(3)}$, but spatial dependencies for these terms are less well motivated.

In our numerical experiments, we considered assigning a separate $\gamma_t^{(1)}$ for each floor of a dorm. In this case all students on a given floor shared the same floor-dependent variant of $\gamma_t^{(1)}$ (i.e. rather than sharing a single $\gamma_t^{(1)}$ across all students, a separate such term was employed for each door floor). We also considered assigning a separate term of the form $\gamma_t^{(1)}$ to each dorm (i.e. all residents in a given dorm, independent of floor, shared the same $\gamma_t^{(1)}$). In our experiments, we found that such added modeling complexity did not improve the predictive performance of the model, and in some cases reduced performance (since the students were spatially segregated in these tests, fewer students were associated with a particular floor/dorm-dependent $\gamma_t^{(1)}$, and therefore statistical strength was diffused). There did not appear to be clear situations for which a given dorm or specific dorm floor had a greater tendency toward health (state S) or sickness (state I) than the general population. A potential reason for this is that students spend a significant portion of their time away from their dorm, in classes and other activities, mixing with the general population.

For these reasons, for the results below we do not explicitly leverage spatial covariates for student dorm rooms within the model. However, when presenting results we will examine some of the inferred parameters in the context of student residency location.

3.6 Missing data

There is a substantial quantity of missing data in self-reported studies, and it is anticipated that the missingness is *not* at random. It is likely that individuals are less likely to pay attention to reporting symptoms when they are feeling well, with greater attention paid during the time of actual illness. If data are missing from individual n on day t , the ‘observations’ are denoted $y_{nt} = \emptyset$. The probability of the null observation in states S is defined as $\eta \in (0, 1)$, and the probability of a null observation in state I is $\rho \in (0, 1)$. We now consider the case of missing data a null observation, and the observation probability for symptom j , individual n and day t is generalized to [27]

$$y_{njt} | S \sim \left[\eta \delta_{\emptyset} + (1 - \eta) \sum_{k=0}^M p(y_{njt} = k | S) \delta_k \right], \tag{13}$$

$$y_{njt} | I \sim \left[\rho \delta_{\emptyset} + (1 - \rho) \sum_{k=0}^M p(y_{njt} = k | I) \delta_k \right], \tag{14}$$

where $p(y_{njt} = k | S)$ and $p(y_{njt} = k | I)$ are the observation probabilities from Section 3.1 (assuming symptoms are not missing); the symbol δ_k is a unit measure concentrated at the point k . It is assumed that η and ρ are drawn from uniform priors over $[0,1]$.

3.7 Modeling multiple years of data

The experiments detailed in Sections 2 and 4 correspond to (ideally) daily student recording of symptom scores, over an entire academic year; imperfections in this process naturally manifest

missing data. Data of this type were collected over three academic years. It is desirable to analyze all of these data jointly, to achieve maximal statistical strength in the results. However, because of the influx of new (freshman) students, and the exit/graduation of others (seniors), the explicit set of students considered on consecutive years is largely distinct. Additionally, each year is characterized (for example) by a distinct respiratory viral illness season, and this must be accounted for when deciding which components of the model to share between multiple years. For example, one of the years during which we collected data corresponded to the presence of an unusual (and potentially dangerous) novel H1N1 flu, which had characteristics (e.g. time of arrival) distinct from typical flu seasons.

So motivated, in the experiments that follow, the explicit $\{\gamma_t^{(1)}, \gamma_t^{(2)}, \gamma_t^{(3)}\}$ are modeled as being distinct among the three years of data. However, the priors on parameters with which these time-dependent functions are constituted are shared across years. Specifically, for the AR(1) model of $\gamma_t^{(1)}$, the priors for ω and β are shared across the multiple years of data. For $\gamma_t^{(2)}$, the parameters Σ_{θ_r} and ζ_{S_r} are shared across years, as is the prior on $\theta_r(0)$. Finally, for $\gamma_t^{(3)}$, the priors for ξ and α_f are shared across the multiple years.

Concerning λ_n and a_n , the DP-drawn priors G_λ and G_a are shared across the multiple years, and therefore the clustering of types of people (by susceptibility toward illness, and length of illness) is performed jointly across the multiple years. Finally, concerning the observed symptoms, the parameters $\{\mu_S, \mu_I, \Sigma_S, \Sigma_I\}$ are shared across the multiple years, as are the ordinal probit cut points $\{\tau_{j,1}, \dots, \tau_{j,M-1}\}_{j=1..J}$, and η and ρ (for missing data).

3.8 Model summary

The overall modeling framework has multiple components, and therefore to aid understanding, in Figure 1 is provided a summary of the semi-Markov SIS structure.




Figure 1. Summary of the semi-Markov SIS model. The underlying states, S and I , are latent, and for individual n we observe a vector of symptom scores y_{nt} at time t . If the individual is in state S at time $t - 1$, there is a probability $\Phi(\gamma_{nt})$ of transiting to state I at time t , and probability $1 - \Phi(\gamma_{nt})$ of staying in state S , where $\Phi(\cdot)$ is a probit distribution function, with parameters γ_{nt} . If the individual transits into state I at time t , she stays in that state for $D_{nt} = c + \hat{D}_{nt}$ days, where $\hat{D}_{nt} \sim \text{Pois}(\lambda_n)$. After being in state I , the individual transits back to state S . The observations for symptom j , y_{njt} , are conditionally dependent on the latent state S or I , and the probability of missing data is tied to this state. When in state S there is a probability $\eta \in [0, 1]$ of missing symptoms, and in state I a probability $\rho \in [0, 1]$. The function $\gamma_{nt} = a_n + \sum_{i=1}^3 \gamma_t^{(i)}$, where term $i = 1$ is a community-wide susceptibility toward illness, term $i = 2$ characterizes weekly semi-periodic effects, and terms $i = 3$ is a regression to the Google Flu Trends data. The term a_n characterizes the general susceptibility of person n to illness. Finally, the parameters λ_n and a_n are clustered among the N people in the study, via Dirichlet process models.

4. Results

The modeling software was implemented in MATLABTM. On a laptop with a 2.7 GHz dual core CPU, each Gibbs iteration takes about 6 s, to process all three years of data. We consider 20,000 MCMC samples, with the first 2000 samples discard as burn-in. After a few hundred iterations, the log posterior reached and sustained a high probability. We also performed the convergence diagnosis tests reviewed in [8] for model parameters; the results show reasonable convergence and mixing.

4.1 Example student trajectories and inferred diagnoses

For each individual in the study, as a function of time (day), we infer the probability that the student is in state I (i.e. that they are sick). To demonstrate this, and to give a sense of the reported symptom scores, in Figure 2 self-reported data and the inferred probability of being in state I are depicted for four example individuals. In Figure 2, based upon averaging across all MCMC collection samples, we plot the probability the individual is in the infective state I , for each day. The results in Figure 2 are based upon a joint analysis of all self-reported symptom-score data, across all three years.

4.2 Characteristics of missing data

In Table 1 we show data on the characteristics of student reporting and associated symptom scores. In this table is depicted the percentage of days students did not report data, and the number of such students in each class of missingness. In Figure 3 we show the fraction of students who do not report symptoms (fraction of missing data), as a function of day for each of the three years of the study. The clear a weekly semi-periodic effect, which has motivated the term $\gamma_i^{(2)}$ in the model.

4.3 Virus infection probability over time

We examine the probability of whether a student is in state I , relative to that student's living conditions with respect to another student who had a positive RT PCR test (RT PCR test results are available for 897 samples). Specifically, assume that a given student has a positive RT PCR test. Over a period of a week after that positive RT PCR test, we examine the probability of being in an infective state for all students who shared a dorm with the student confirmed by RT PCR as being in state I . We also examined the probability of being in state I for all students on the SD floor (not just the SD) of a student confirmed by RT PCR as being infected, again for a week after RT PCR confirmation. When multiple instances overlap in time (multiple positive RT PCR tests), average results are presented across those multiple instances.

To be precise, let \mathcal{X} represent a particular set of students (e.g. a set of students in the SD of a RT PCR-confirmed infective student, or a set of students on the SD floor of a RT PCR-confirmed infective student). Let $|\mathcal{X}|$ represent the number of individuals in this set. Then we are interested in computing $S_{\mathcal{X}} = \frac{1}{|\mathcal{X}|} \sum_{n \in \mathcal{X}} p(z_{nt} = I | \mathbf{y}_{nt})$, where $p(z_{nt} = I | \mathbf{y}_{nt})$ is computed from our model. This provides a means of examining the inferred degree of enhanced probability of becoming ill with virus, given a nearby confirmed case (recognizing the imperfections in our $p(z_{nt} = I | \mathbf{y}_{nt})$, most notably that one may become sick for other reasons than virus transfer). Such that we have enough individuals in a given set to make this investigation meaningful, we only consider cases for which $|\mathcal{X}| \geq 5$; e.g. when examining propagation of infectious disease on a dorm floor, we only consider cases for which five or more students within the study live on the same floor.

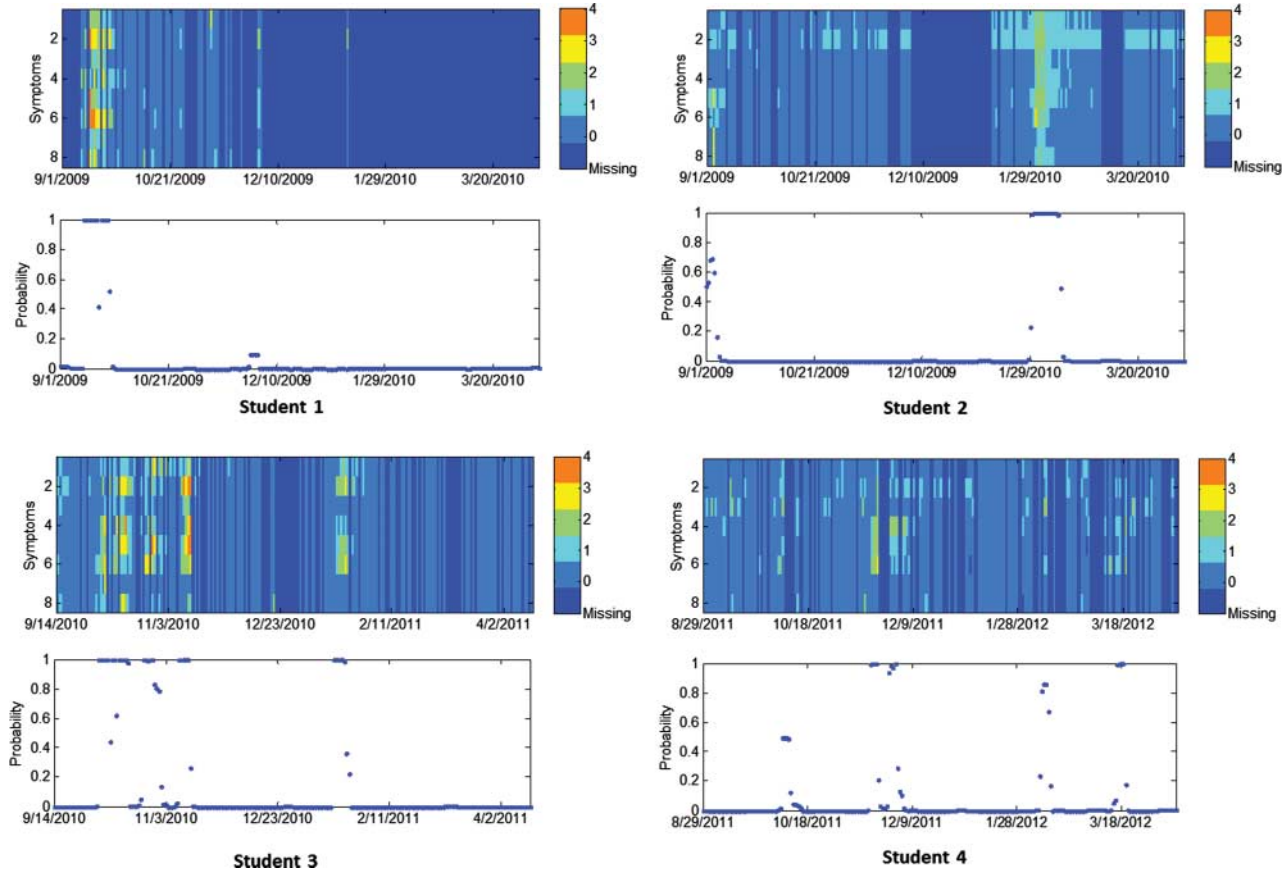


Figure 2. State of health of four example students. For each student, self-reported symptom scores are shown in the top figure. Different colors denote different scores (missing, 0, 1, 2, 3, 4). The probability that a student is in an infective state I at a given time is presented in the bottom subfigure for each of the four students. Each interval in horizontal axis refers to 50 days.

Table 1. Summary on properties of student reporting frequency and associated reported symptom scores.

Missingness	0–20%	20–40%	40–60%	60–80%	80–100%
# Students	22	158	303	178	205
Avg. symp. score	2.67	2.22	2.64	3.27	4.33
Avg. sick prob.	0.12	0.11	0.15	0.2	0.32

Note: The average symptom score reported is the average of the *sum* of the scores for eight symptoms.




Figure 3. Fraction of missing data over days. From top to bottom are the results for academic years 2009–2010, 2010–2011 and 2011–2012. The gray bars reflect, from left-to-right, Thanksgiving break, Winter/Christmas break, and Spring break. Each interval in horizontal axis refers to 50 days.

In Figures 4–6 are shown results for, respectively, academic years 2009–2010, 2010–2011, 2011–2012. On the bottom of each figure, a black bar represents the presence of a RT PCR-confirmed virus of noted type. At the top of each figure is shown the average probability of being in the infective state, under three circumstances. In red are shown results for all students and all times, and therefore in this case \mathcal{X} denotes the set of all students. The blue curve corresponds to the case for which \mathcal{X} corresponds to the set of students from the same dorm and same floor (SDSF) of a RT PCR-confirmed case. Finally, for the black curve, \mathcal{X} corresponds to the set of students in the SD of a RT PCR-confirmed case. Unlike the case of all students (red), for the cases of SDSF and SD the curves are not shown at all times, because we only look within a window of 7 days after a RT PCR-based detection, and in some cases there are no data (e.g. a given RT PCR-confirmed student doesn't have a sufficient number of students in the study on the same floor, or there were no RT PCR-confirmed detections in the last 7 days).

In Figures 7–9 we show results like discussed above, but now for the SD and SDSF cases we only consider situations in which there was PCR-confirmed Rhinovirus-induced illness. In




Figure 4. Top figure: probability of being in the infective state I on a given day, for academic year 2009–2010. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm and same floor with infective individuals. ‘SD’ refers to the average of students living in the same dorm with infective individuals. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time certain type of virus was detected.

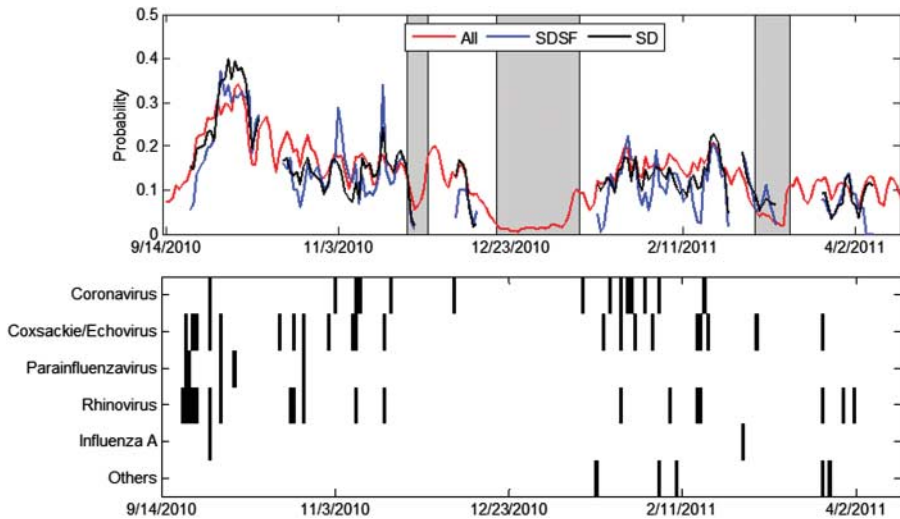


Figure 5. Top figure: probability of being in the infective state I on a given day, for academic year 2010–2011. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm and same floor with infective individuals. ‘SD’ refers to the average of students living in the same dorm with infective individuals. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time certain type of virus was detected.




Figure 6. Top figure: probability of being in the infective state I on a given day, for academic year 2011–2012. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm and same floor with infective individuals. ‘SD’ refers to the average of students living in the same dorm with infective individuals. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time certain type of virus was detected.




Figure 7. Top figure: probability of being in the infective state I on a given day, for academic year 2009–2010. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm same floor with an infective individual. ‘SD’ refers to the average of students living in the same dorm with an infective individual. The SD and SDSF cases are only for confirmed cases of Rhinovirus. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time Rhinovirus was detected.

Figures 10 and 11 similar results are shown as above, but now only Influenza A is considered for the SD and SDSF cases.

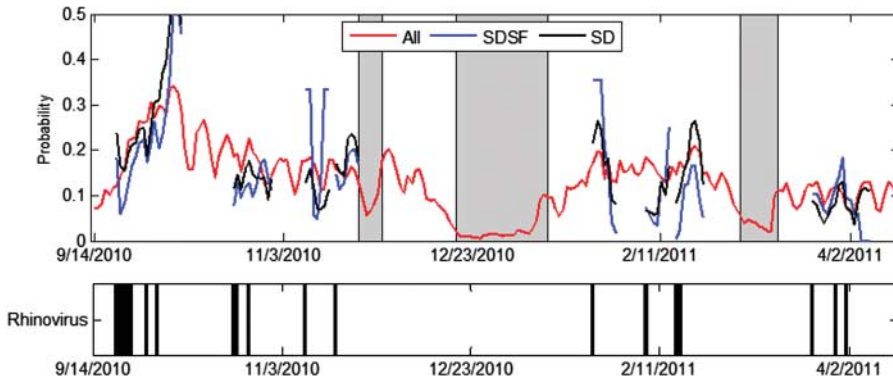


Figure 8. Top figure: probability of being in the infective state I on a given day, for academic year 2010–2011. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm same floor with an infective individual. ‘SD’ refers to the average of students living in the same dorm with an infective individual. The SD and SDSF cases are only for confirmed cases of Rhinovirus. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time Rhinovirus was detected.

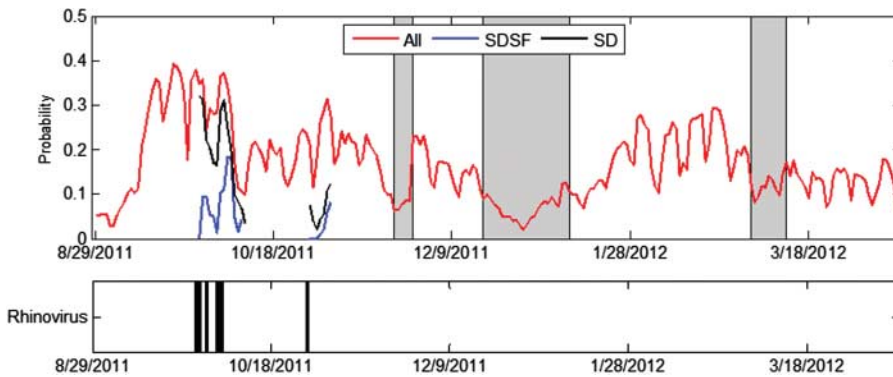


Figure 9. Top figure: probability of being in the infective state I on a given day, for academic year 2011–2012. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm same floor with an infective individual. ‘SD’ refers to the average of students living in the same dorm with an infective individual. The SD and SDSF cases are only for confirmed cases of Rhinovirus. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time Rhinovirus was detected.

4.4 Classification performance based on symptoms

We examine the accuracy of the prediction of infection, relative to an objective ‘truth’. To do this, we considered all 897 individuals for whom RT PCR-based virus-identification was performed (for a subset of these, for which the RT PCR test was negative, confirmation gene-expression analysis was also performed). The receiver operating characteristic (ROC) curve is manifested by thresholding $p(z_{nt} = I \mid -)$, and is shown in the left column of Figure 12. The model achieves a 90% detection rate at a false-alarm rate of 15%.

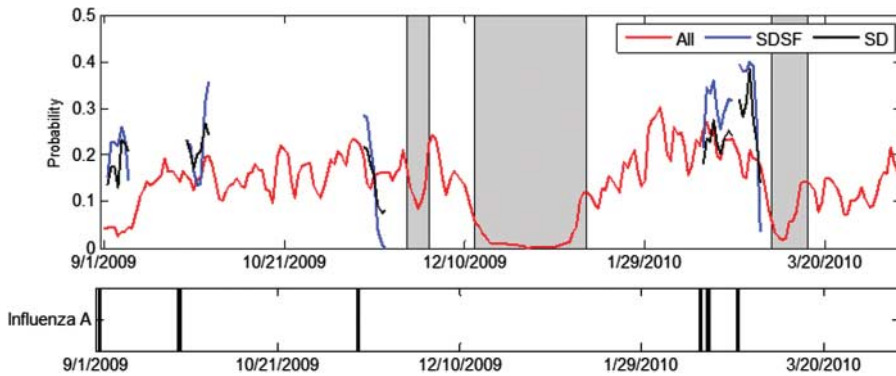


Figure 10. Top figure: probability of being in the infective state on a given day, for academic year 2009–2010. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm and same floor with an infective individual. ‘SD’ refers to the average of students living in the same dorm with an infective individual. The SD and SDSF cases are only for confirmed cases of Influenza A. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time Influenza A was detected.

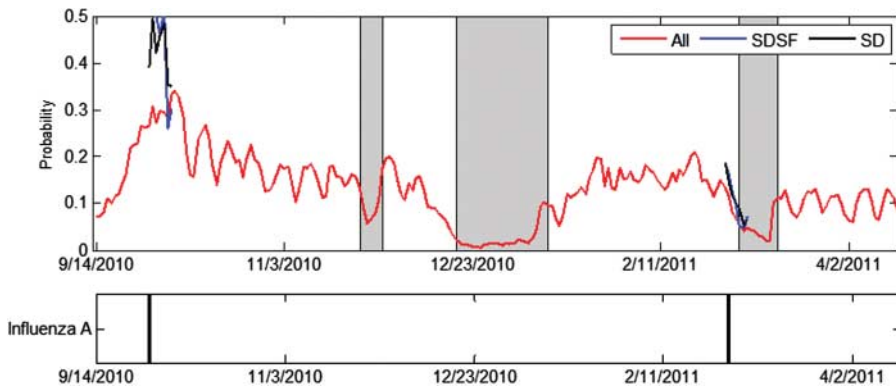


Figure 11. Top figure: probability of being in the infective state on a given day, for academic year 2010–2011. ‘All’ refers to the average across all the students. ‘SDSF’ refers to the average of students living in the same dorm and same floor with an infective individual. ‘SD’ refers to the average of students living in the same dorm with an infective individual. The SD and SDSF cases are only for confirmed cases of Influenza A. The vertical gray bars represent, from left-to-right, Thanksgiving break, inter-semester (Winter/Christmas) break, and Spring break. Each interval in horizontal axis refers to 50 days. Bottom figure: RT PCR test results, black line denotes at that time Influenza A was detected.

4.5 Online prediction of health

The results in Figure 12 on predicting the state of health were based on *all* of the self-reported data, at all times for which data were reported. A clinician must predict the state of health only based upon symptoms up to the point at which a diagnosis is made. It is desirable to predictive probability that a particular student is in state I on day $t + 1$, based on symptom scores up to day t .

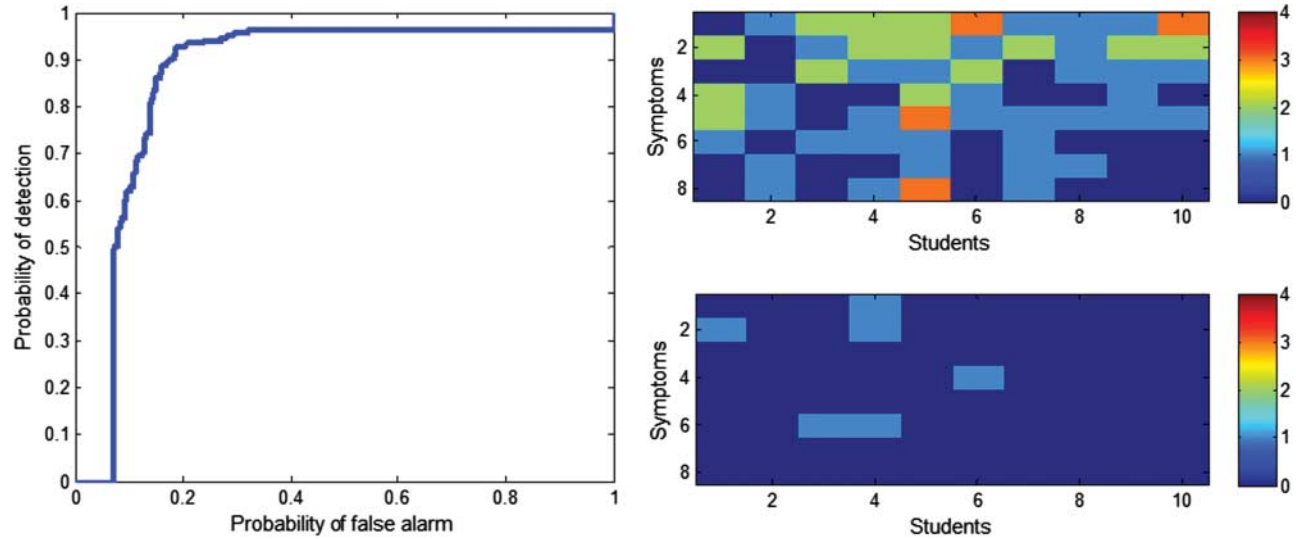


Figure 12. Left column: ROC curve. Right column: Top figure is the symptom scores of students who are healthy (in state S) but labeled infective (in state I) with high probability by the model. The bottom figure shows the symptom scores of students who are infective (in state I) but labeled healthy (in state S) by the model.

Let $\mathbf{y}_1^t = \{y_1, \dots, y_t\}$ represent the symptom scores up to time t for the student in question. The probability that the student is in state I on day $t + 1$ may be expressed as

$$p(z_{t+1} = I \mid \mathbf{y}_1^t, \Omega) = \frac{\sum_{d_{t+1}=1}^{D_{\max}} p(z_{t+1} = I, d_{t+1}, \mathbf{y}_1^t \mid \Omega)}{\sum_{d_{t+1}=1}^{D_{\max}} p(z_{t+1} = I, d_{t+1}, \mathbf{y}_1^t \mid \Omega) + p(z_{t+1} = S, \mathbf{y}_1^t \mid \Omega)}, \quad (15)$$

where $p(z_{t+1} = I, d_{t+1}, \mathbf{y}_1^t \mid \Omega)$ represents the joint probability of data \mathbf{y}_1^t , that the student is in state I on day $t + 1$, and that they are in day d_{t+1} of being infected; $d_{t+1} \in \{1, \dots, D_{\max}\}$ is the number of days left in infective state at time $t + 1$. $p(z_{t+1} = S, \mathbf{y}_1^t \mid \Omega)$ represents the joint probability of the data and being in the susceptible state S . In both cases, Ω represents model parameters learned from data up to day t . The details for calculating $p(z_{t+1} = I, d_{t+1} \mid \mathbf{y}_1^t, \Omega)$ and $p(z_{t+1} = S, \mathbf{y}_1^t \mid \Omega)$ are provided in Appendix B.

In this experiment, the first two years of data, and the third year of data up to day $t = 140$ are employed to learn Ω (these are typical results for many values of t selected in Year 3). In Figure 13 are shown the model predictions for all students in Year 3 (2011–2012), where in Figure 13 the students are ordered from left to right from the most to least probable of being in state I on day $t + 1$. At the top in Figure 13 are shown the symptoms reported on day $t + 1$ (for those for whom scores were provided), and it is evident that the individuals who are deemed most likely to be in state I on day $t + 1$ (based on data up to day t) tend to have the strongest symptoms on that day. In the middle in Figure 13 is shown the probability of being in an infective state on day $t + 1$, based on data up to day t . Finally, the bottom part of Figure 13 shows the probability of being in an infective state on day $t + 1$ based on all of the data. There is generally good agreement (middle and bottom figures) on which students are most likely to be in the infective state on day $t + 1$.

5. Discussion

5.1 Characteristics of missing data

In Section 3.6 we proposed a model for the missing data. Specifically, it was assumed that if a student is in the susceptible state, S , they do not report symptoms (which are likely negligible) with probability η , thereby manifesting missing data. By contrast, when in the infective state I , it is anticipated that one may be more likely to report symptom scores (which are non-negligible, by definition); the probability of not reporting when in state I is represented by ρ (see Section 3.6). Within the analysis, we inferred a mean $\eta = 0.65$, with standard deviation of η equal to 0.01 (reflecting the uncertainty in this parameter from the approximate posterior); the inferred mean for ρ was 0.28, with standard deviation 0.03. Hence, the model infers that when a student is in state S (healthy), a student doesn't report any symptoms 65% of the time, while when in state I (sick) the students don't report symptoms 28% of the time.

From Table 1 note that the largest group of students, with 303 members, did not report symptoms on 40–60% of the days. For each class of missingness, we also report the average reported symptom score, recalling that the values were 0 to 4, with 4 the largest/strongest symptom (eight different symptoms are considered). Note that the average symptom score is particularly large for those students who report data infrequently, and the probability that students are in state I when reporting is heightened for the group that rarely reports. The data in Table 1 motivates the model in Section 3.6, in which the degree of missingness is assumed to be linked to the latent state of health.

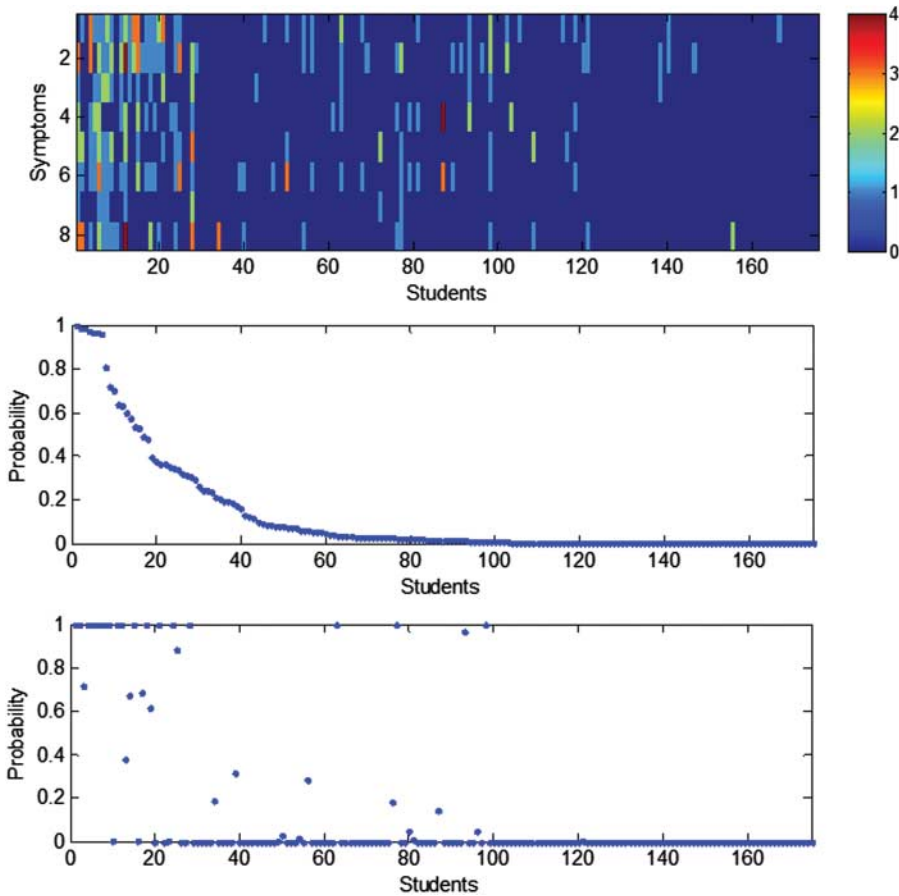


Figure 13. The top figure is the symptoms scores for students at time $t + 1$ (can be considered as ‘truth’). The middle figure is $p(z_{nt+1} = I \mid \mathbf{y}_{n1}^t, -)$, the predictive probability that a given student is in the infective state at $t + 1$. The bottom figure is the probability that students stay in infected at $t + 1$ given all the data.

5.2 Location dependence to probability of getting infected

From Figure 4, we typically see the following trend. If a student gets sick (is in state I) within a given dorm, from one of the specified viruses, then over the preceding 7 days the average probability of students within the SD will have a heightened probability of being in the infective state as compared to the general student population. Moreover, if a student gets infected, within a week students on the SDSF typically have, on average, an even higher probability of being in the infective state. This is not always true, but it seems to be a fairly common situation.

In Figures 5 and 6 results are shown in this same format as in Figure 4, for the 2010–2011 and 2011–2012 academic years, respectively. Given the relatively small number of RT PCR-confirmed cases relative to the total population size, it is difficult to make strong conclusions from Figures 4–6. There are periods in which being on the SD as an infective student clearly manifests increased probability of being in the infective state, over a subsequent 7 day period, but during other times this trend is not evident (e.g. during the 2011–2012 academic year). One interpretation is that the presence of dorm colleagues who are sick may heighten attention to protecting oneself, through hand washing, etc. Therefore, from this perspective, the presence of a sick student may actually encourage more-healthy behavior in others.

To examine this issue from a finer perspective, we examine these same types of curves, but for two specific viruses: Rhinovirus and Influenza A. Rhinovirus is associated with the ‘common cold’, and therefore it is a virus that all students will come in contact with, in and outside their dorm. Therefore, in the case of Rhinovirus, the connection to the dorm, and who is infected there at a given time, may be more tenuous (students will come in contact with Rhinovirus and associated infective students in their classes, and other activities outside their dorm). Influenza A occurs more rarely, and therefore if someone is confirmed as infected with Influenza A, it is anticipated that students within the SD and SDSF may be at higher risk of infection.

For the case of Rhinovirus (Figures 7–9), we generally observe that if in the SD or on the SDSF, when a given student is infected his/her dorm neighbors have a heightened probability of being in the infective state over the next week. However, there are cases for which this is not the case, which indicates that for Rhinovirus transmission activities outside the dorm may be as or more important than the degree of infection within the dorm.

There are fewer Influenza A cases than Rhinovirus, so conclusions must be drawn with care. Nevertheless, for the case of Influenza A (Figures 10 and 11), the SD probability of being infected within a week of a confirmed Influenza A case is heightened relative to the general population, and the SDSF is generally further heightened. Note that in many cases, after roughly 5 days from a confirmed Influenza A case, the SD/SDSF probability of being infected is *less than* that of the general population; right after the confirmed Influenza A case the SD/SDSF probability of being infected increases, but then it diminishes relative to the general student population (e.g. see the case in November 2009, in Figure 10). This phenomenon may be attributed to acquired immunity, after being infected.

5.3 Dynamic effects on the probability of being infected

An interesting phenomenon is observed in Figures 7–9, when considering the probability of being in an infective state for all of the students (red curve). When the students come together at the beginning of the school year, and after the long Winter break, a general increase in the probability of being in an infective state is observed. Note that at the beginning of the school year, this is particularly evident in the 2010–2011 year (Figure 8), and in 2011–2012 (Figure 9). Therefore, in Figures 8 and 9 the students are coming together for the first time at the beginning of the school year, from all over the USA, and from many other countries across the world. This phenomenon of increased probability of infection as students come together for the first time, or after extended break, may be associated with the general spread of infectious disease caused by a new mix of people, as been observed previously in the literature [17,23].

Figure 7 for 2009–2010 has temporal dependence (red curve) that is distinct from 2010–2011 and 2011–2012 (respectively, Figures 7 and 8). This may be attributed to the fact that the 2009–2010 academic year was the year of the novel H1N1 virus, and significantly heightened on-campus attention to protecting oneself from virus transfer. These results seem to indicate that the heightened attention to viral transfer manifested by the novel H1N1 virus had a significant impact in reducing the probability of students transiting from the *S* to *I* state (not just from H1N1 virus, but from all viruses), when compared to two years in which such attention to viruses was far more muted on campus.

5.4 Classification performance

The quality of the detection-performance ROC is undermined by imperfections in the definition of ‘truth’. People who are sick as a result of illnesses other than virus will be deemed as healthy in the truth (negative RT PCR test), but in reality they are sick. Another source of errors are manifested by positive RT PCR tests for the presence of virus, but the individual shows

no symptoms – these are termed ‘shedders’ in the medical community [9]. These individuals are carrying the virus, and shedding the virus, but they do not show any symptoms. The RT PCR test will deem these individuals as being in the infective state I , but from the standpoint of symptoms, which is what our analysis considers, these people are not infected (there are no symptoms present that would allow one to declare they are infected, based on symptoms alone).

In our study there were two reasons a given student could have a RT PCR test performed: (i) based upon their self-reported data, a doctor in (near) RT determined that they were infected, and therefore they were contacted for acquisition of a sample; (ii) a given individual was a close contact of a person who was sampled in the case of (i). Therefore, for the close contacts, the student may not be in an infective state at the time of sampling, either because there was no disease transfer, or because the onset of illness was manifested at a later time.

Based upon the symptoms, and the doctor-based diagnosis, all individuals in case (i) above are defined as being in the infective state, essentially by definition (these students were only contacted because their symptoms were deemed above a threshold of illness). For the close contacts, all individuals for whom the RT PCR test was negative will be deemed as healthy (in susceptible, S) state, and all others were deemed to be in the infective state (the RT PCR test may miss some infected people, which the gene-expression analysis can pick up, and this issue is discussed when presenting results).

In Figure 12, left, note (a) the presence of many false alarms before any detections are achieved (left-most part of the ROC), and (b) after a probability of false alarms of about 0.35, the detection probability is stuck at around 0.95, until the very rightmost part of the ROC. Concerning (a), on the top-right of Figure 12, we show the symptom scores for the 10 students who characterize the individuals detected as being infected, but RT PCR deems as being healthy, or in state S (the false alarms at the beginning of the ROC). Based upon the symptoms (right in Figure 12), these students are almost certainly sick due to some cause other than virus, or because of limitations of the RT PCR test (e.g. poor samples, or because the illness was caused by a virus other than that tested by the RT PCR).

At bottom-right in Figure 12 is shown the symptom scores of students who were deemed infective via RT PCR, but our model deemed as healthy, based upon the symptoms. We see that the symptoms of these students are indeed very mild, or absent. These individuals were likely carrying a virus that was tested, and that was detected via RT PCR. However, these individuals were likely recently infected with the virus, and still carrying it, but no longer infected. Alternatively, these individuals may have been asymptomatic shedders.

Gene expression data were available for six of the individuals considered at right in Figure 12. In all of these cases, the gene-expression analysis was able to confirm the labels inferred by our algorithm based on symptoms.

5.5 Breaking out model components

From the analysis we can break out the three components of the model responsible for the probability of transiting from the susceptible state to the infected state. We do not show the detailed plots of this breakout, for brevity. However, we discuss the findings of that analysis.

We observed that the weekly pattern $\gamma_t^{(2)}$ (impacting the probability of transiting from healthy to infective state) is typically heightened at either Wednesday or Thursday, and tends to be smaller around the weekend. This is possibly reflective of the fact that students are more likely to report symptoms during the school week than they are on the weekend, when they may be distracted by funner activities. Of course, another interpretation is that the probability that the students will *feel* infected/sick is diminished during the weekend, relative to the middle of the week, when they may be under greater stress.

Figure 3 depicts the degree of missingness on average as a function of days. By construction, heightened missingness is deemed associated with health, and weekends tend to be periods of high missingness. Whatever the cause of the weekly effects (student laziness/distraction or actual health), model interpretation may be improved by removing this effect. We consider this below.

It was found that the contribution of the Google Flu Trend term was relatively small (large posterior mass concentrated around zero, particularly for the first two years), which implies that the spread of infectious disease among students on the Duke University campus is a relatively isolated ecosystem, distinct from the city and community of Durham used here for f_i .

We considered the inferred probability of transiting from state S to state I , as a function of day, for each of the three years of the study. To remove the daily oscillations, we removed the periodic term from this analysis. The data were analyzed using all components of the model. However, after this analysis, to remove the effects of the weekly term, we considered the model-inferred probability of transiting from $S \rightarrow I$, and with the weekly term removed. At the beginning of the academic year in 2010–2011 and 2011–2012, there was a clear increased probability of getting infected within the first month or so the students are together, presumably a mixing effect [17,23] caused by interactions of many people who have never met before, coming from all over the USA, and also from outside the USA. It appears that the heightened attention to viruses (from the alarm associated with novel H1N1) dampened this phenomenon in 2009–2010. During the first semester of 2009–2010, when there was so much attention to viruses on campus, there was a noticeable decrease in the probability of transiting from state S to I , after the weekly effects are removed (relative to 2010–2011 and 2011–2012).

6. Conclusions

A statistical model has been developed for analysis of the time-dependent symptom scores provided by a large group of undergraduate college students. Unlike almost all studies of data related to infection transfer, the model has operated directly on the observed symptoms, and the state of the students were assumed to be latent. The community-to-person mechanism for pathogen transfer has been modeled in terms of an SIS analysis, and computations have been performed using Bayesian (MCMC) methods. A detailed characterization of the data and the scientific questions that have motivated this study are discussed in Section 2; a comprehensive answering of these questions with the available data has been provided in Sections 4 and 5. For brevity, we do not repeat these details here. We note that these data are presented here for the first time, and were collected by the authors; all data will be made available to the research community.

There are further questions that may be examined with the collected data, and that are worthy of future study. The identification of the virus responsible for each illness has been (imperfectly) constituted via RT PCR, for a large set of common viruses considered. We have access to the dorm in which each individual resided. A more detailed analysis of pathogen transfer as a function of virus type can be examined. In this paper we have presented results in this direction, but more explicit modeling could be performed (not necessarily at the symptom level, but after the responsible virus has been identified by RT PCR).

The gene expression data from this study have only been employed here in a limited manner, as the focus has been on self-reported symptom scores. However, for the close contacts, we have daily gene expression data for a week. For close contacts who transited from state S to I , we have the opportunity to analyze the time trajectory of the gene expression data as the host responds to the (known) virus. We have performed work of this type for people enrolled in challenge studies (controlled experiments) [6]; the data from this study offers the potential for similar studies on data from individuals who became ill in natural settings. We have preliminary results in this direction on these data, which are encouraging and will be presented elsewhere.

References

- [1] P. Birrell, G. Ketsetzib, N. Gayc, B. Cooperd, A. Presanisa, R. Harrisb, A. Charlettb, X.-S. Zhangb, P. Whiteb, R. Pebody, and D. D. Angelisa, *Bayesian modeling to unmask and predict influenza A/H1N1pdm dynamics in London*, Proc. Nat. Acad. Sci. 108(45) (2011), pp. 18238–18243.
- [2] C. Carvalho, J. Chang, J. Lucas, J. Nevins, Q. Wang, and M. West, *High-dimensional sparse factor modelling: Applications in gene expression genomics*, J. Amer. Statist. Assoc. 103(484) (2008), pp. 1438–1456.
- [3] S. Cauchemez, F. Carrat, C. Viboud, A. Valleron, and P. Boelle, *A Bayesian MCMC approach to study transmission of influenza: Application to household longitudinal data*, Stat. Med. 23(22) (2004), pp. 3469–3487.
- [4] B. Chen, M. Chen, J. Paisley, A. Zaas, C. Woods, G. Ginsburg, A. Hero III, J. Lucas, D. Dunson, and L. Carin, *Bayesian inference of the number of factors in gene-expression analysis: Application to human virus challenge studies*, BMC Bioinformatics 11 (2010), p. 552.
- [5] M. Chen and D. Dey, *Generalized Linear Models: A Bayesian Prospective*, CRC Press, Boca Raton, FL, 2000.
- [6] M. Chen, A. Zaas, C. Woods, G. Ginsburg, J. Lucas, D. Dunson, and L. Carin, *Predicting viral infection from high-dimensional biomarker trajectories*, J. Amer. Statist. Assoc. 106 (2011), pp. 1259–1279.
- [7] D. Clancy and P. O’Neill, *Exact Bayesian inference and model selection for stochastic models of epidemics among a community of households*, Scand. J. Statist. 34(2) (2007), pp. 259–274.
- [8] K. Cowles and B. Carlin, *Markov chain Monte Carlo convergence diagnostics: A comparative review*, J. Amer. Statist. Assoc. 91(434) (1996), pp. 883–904.
- [9] N. Dimmock, A. Easton, and K. Leppard, *Introduction to Modern Virology*, Blackwell, Oxford, 2007.
- [10] W. Dong, A. Pentland, and K. Heller, *Graph-coupled HMMs for modeling the spread of infection*, in *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, N. de Freitas and K.P. Murphy, eds., Catalina Island, CA, USA, August 14–18, AUAI Press, Corvallis, OR, 2012, pp. 227–236.
- [11] V. Dukic, H. Lopes, and N. Polson, *Tracking epidemics with Google Flu trends data and a state-space SEIR model*, J. Amer. Statist. Assoc. 107(500) (2012), pp. 1410–1426.
- [12] E. Fox and D. Dunson, *Bayesian nonparametric covariance regression*, preprint (February 2011). Available at arXiv:1101.2017v2.
- [13] J. Ginsberg, M. Mohebbi, R. Patel, L. Brammer, M.S.Smolinski, and L. Brilliant, *Detecting influenza epidemics using search engine query data*, Nature 457 (2009), pp. 1012–1014.
- [14] H. Hethcote, *The mathematics of infectious diseases*, SIAM Rev. 42(4) (2000), pp. 599–653.
- [15] M. Jahrer and A. Toscher, *Collaborative filtering ensemble*, J. Mach. Learn. Res. 18 (2012); pp. 61–74.
- [16] G. Jones, W. Johnson, W. Vink, and N. French, *A framework for the joint modeling of longitudinal diagnostic outcome data and latent infection status: Application to investigating the temporal relationship between infection and disease*, Biometrics 68(2) (2012), pp. 371–379.
- [17] V. Kak, *Infections in confined spaces: Cruise ships, military barracks, and college dormitories*, Infect. Dis. Clin. North Am. 21(3) (2007), pp. 773–784.
- [18] D. Navarro, T.L. Griffiths, M. Steyvers, and M. Lee, *Modeling individual differences using Dirichlet processes*, J. Math. Psychol. 50(2) (2006), pp. 101–122.
- [19] S. O’Brien and D. Dunson, *Bayesian multivariate logistic regression*, Biometrics 60(3) (2004), pp. 739–746.
- [20] P. O’Neill, D. Balding, N. Becker, M. Eerola, and D. Mollison, *Analyses of infectious disease data from household outbreaks by Markov chain Monte Carlo methods*, Appl. Stat. 49 (2000), pp. 517–542.
- [21] E. Patrozou and L. Mermel, *Does influenza transmission occur from asymptomatic infection or prior to symptom onset?* Public Health Rep. 124(2) (2009) 193–196.
- [22] J. Silva and L. Carin, *Active learning for online Bayesian matrix factorization*, Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, 2012, pp. 325–333.
- [23] Y. Sun, Z. Wang, Y. Zhang, and J. Sundell, *In China, students in crowded dormitories with a low ventilation rate have more common colds: Evidence for airborne transmission*, PLOS One 6(11) (2011), p. e27140.
- [24] V. Trifonov, H. Khiabani, and R. Rabadan, *Geographic dependence, surveillance, and origins of the 2009 Influenza A (H1N1) virus*, N. Engl. J. Med. 361 (2009), pp. 115–119.
- [25] M. West and J. Harrison, *Bayesian Forecasting and Dynamic Models*, Springer, New York, NY, 1989.
- [26] Y. Yang, M. Halloran, M. Daniels, I. Longini, Jr., D. Burke, and D. Cummings, *Modeling competing infectious pathogens from a Bayesian perspective: Application to influenza studies with incomplete laboratory results*, J. Amer. Statist. Assoc. 105(492) (2010), pp. 1310–1322.
- [27] S. Yu and H. Kobayashi, *A hidden semi-Markov model with missing data and multiple observation sequences for mobility tracking*, Signal Process. 83(2) (2003), pp. 235–250.
- [28] A. Zaas, M. Chen, J. Lucas, T. Veldman, A. Hero, J. Varkey, R. Turner, C. Oien, S. Kingsmore, L. Carin, C. Woods, and G. Ginsburg, *Peripheral blood gene expression signatures characterize symptomatic respiratory viral infection*, Cell Host & Microbe 6(3) (2009), pp. 207–217.