

Global Disease Monitoring and Forecasting with Wikipedia



Nicholas Generous*, Geoffrey Fairchild, Alina Deshpande, Sara Y. Del Valle, Reid Priedhorsky

Defense Systems and Analysis Division, Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America

Abstract

Infectious disease is a leading threat to public health, economic stability, and other key social structures. Efforts to mitigate these impacts depend on accurate and timely monitoring to measure the risk and progress of disease. Traditional, biologically-focused monitoring techniques are accurate but costly and slow; in response, new techniques based on social internet data, such as social media and search queries, are emerging. These efforts are promising, but important challenges in the areas of scientific peer review, breadth of diseases and countries, and forecasting hamper their operational usefulness. We examine a freely available, open data source for this use: access logs from the online encyclopedia Wikipedia. Using linear models, language as a proxy for location, and a systematic yet simple article selection procedure, we tested 14 location-disease combinations and demonstrate that these data feasibly support an approach that overcomes these challenges. Specifically, our proof-of-concept yields models with r^2 up to 0.92, forecasting value up to the 28 days tested, and several pairs of models similar enough to suggest that transferring models from one location to another without re-training is feasible. Based on these preliminary results, we close with a research agenda designed to overcome these challenges and produce a disease monitoring and forecasting system that is significantly more effective, robust, and globally comprehensive than the current state of the art.

Citation: Generous N, Fairchild G, Deshpande A, Del Valle SY, Priedhorsky R (2014) Global Disease Monitoring and Forecasting with Wikipedia. *PLoS Comput Biol* 10(11): e1003892. doi:10.1371/journal.pcbi.1003892

Editor: Marcel Salathé, Pennsylvania State University, United States of America

Received: April 18, 2014; **Accepted:** August 21, 2014; **Published:** November 13, 2014

This is an open-access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the Creative Commons CC0 public domain dedication.

Data Availability: The authors confirm that all data underlying the findings are fully available without restriction. Data are available in the Supplemental Information. Software is available at <http://github.com/reidpr/quac>

Funding: This work is supported in part by NIH/NIGMS/MIDAS under grant U01-GM097658-01 and the Defense Threat Reduction Agency (DTRA), Joint Science and Technology Office for Chemical and Biological Defense under project numbers CB3656 and CB10007. Data collected using QUAC; this functionality was supported by the U.S. Department of Energy through the LANL LDRD Program. Computation used HPC resources provided by the LANL Institutional Computing Program. LANL is operated by Los Alamos National Security, LLC for the Department of Energy under contract DE-AC52-06NA25396. Approved for public release: LA-UR~14-22535. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* Email: generous@lanl.gov

Introduction

Motivation and Overview

Infectious disease remains extremely costly in both human and economic terms. For example, the majority of global child mortality is due to conditions such as acute respiratory infection, measles, diarrhea, malaria, and HIV/AIDS [1]. Even in developed countries, infectious disease has great impact; for example, each influenza season costs the United States between 3,000 and 49,000 lives [2] and an average of \$87 billion in reduced economic output [3].

Effective and timely disease surveillance — that is, detecting, characterizing, and quantifying the incidence of disease — is a critical component of prevention and mitigation strategies that can save lives, reduce suffering, and minimize impact. Traditionally, such monitoring takes the form of patient interviews and/or laboratory tests followed by a bureaucratic reporting chain; while generally considered accurate, this process is costly and introduces a significant lag between observation and reporting.

These problems have motivated new surveillance techniques based upon internet data sources such as search queries and social media posts. Essentially, these methods use large-scale data mining techniques to identify health-related activity traces within the data

streams, extract them, and transform them into some useful metric. The basic approach is to train a statistical estimation model against ground truth data, such as ministry of health disease incidence records, and then apply the model to generate estimates when the true data are not available, e.g., when forecasting or when the true data have not yet been published. This has proven effective and has spawned operational systems such as Google Flu Trends (<http://www.google.org/flutrends/>). However, four key challenges remain before internet-based disease surveillance models can be reliably integrated into an decision-making toolkit:

C1. Openness. Models should afford review, replication, improvement, and deployment by third parties. This guarantees a high-quality scientific basis, continuity of operations, and broad applicability. These requirements imply that model algorithms — in the form of source code, not research papers — must be generally available, and they also imply that complete input data must be available. The latter is the key obstacle, as terms are dictated by the data owner rather than the data user; this motivated our exploration of Wikipedia access logs. To our knowledge, no models exist that use both open data and open algorithms.

C2. Breadth. Dozens of diseases in hundreds of countries have sufficient impact to merit surveillance; however, adapting a

Author Summary

Even in developed countries, infectious disease has significant impact; for example, flu seasons in the United States take between 3,000 and 49,000 lives. Disease surveillance, traditionally based on patient visits to health providers and laboratory tests, can reduce these impacts. Motivated by cost and timeliness, surveillance methods based on internet data have recently emerged, but are not yet reliable for several reasons, including weak scientific peer review, breadth of diseases and countries covered, and underdeveloped forecasting capabilities. We argue that these challenges can be overcome by using a freely available data source: aggregated access logs from the online encyclopedia Wikipedia. Using simple statistical techniques, our proof-of-concept experiments suggest that these data are effective for predicting the present, as well as forecasting up to the 28-day limit of our tests. Our results also suggest that these models can be used even in places with no official data upon which to build models. In short, this paper establishes the utility of Wikipedia as a broadly effective data source for disease information, and we outline a path to a reliable, scientifically sound, operational, and global disease surveillance system that overcomes key gaps in existing traditional and internet-based techniques.

model from one *disease-location context* to another can be costly, and resources are often, if not usually, unavailable to do so. Thus, models should be cheaply adaptable, ideally by simply entering new incidence data for training. While most published models afford this flexibility in principle, few have been expressly tested for this purpose.

C3. Transferability. Many contexts have insufficient reliable incidence data to train a model (for example, the relevant ministry of health might not track the disease of interest), and in fact these are the contexts where new approaches are of the greatest urgency. Thus, trained models should be translatable to new contexts using alternate, non-incidence data such as a bilingual dictionary or census demographics. To our knowledge, no such models exist.

C4. Forecasting. Effective disease response depends not only on the current state of an outbreak but also its future course. That is, models should provide not only estimates of the current state of the world — *nowcasts* — but also *forecasts* of its future state.

While recent work in disease forecasting has made significant strides in accuracy, forecasting the future of an outbreak is still a complex affair that is sharply limited in contexts with insufficient data or insufficient understanding of the biological processes and parameters underpinning the outbreak. In these contexts, a simpler statistical approach based on leading indicators in internet data streams may improve forecast availability, quality, and time horizon. Prior evaluations of such approaches have yielded conflicting results and to our knowledge have not been performed at time granularity finer than one week.

In order to address these challenges, we propose a new approach based on freely available Wikipedia article access logs. In the current proof of concept, we use language as a proxy for location, but we hope that access data explicitly aggregated by geography will become available in the future. (Our implementation is available as open source software: <http://github.com/reidpr/quac>.) To demonstrate the feasibility of techniques built upon this data stream, we built linear models mapping daily access counts of encyclopedia articles to case counts for 7 diseases in 9 countries, for a total of 14 contexts. Even a simple article selection

method was successful in 8 of the 14 contexts, yielding models with r^2 up to 0.89 in nowcasting and 0.92 in forecasting, with most of the successful contexts having forecast value up to the tested limit of 28 days. Specifically, we argue that approaches based on this data source can overcome the four challenges as follows:

C1. Anyone with relatively modest computing resources can download the complete Wikipedia dataset and keep it up to date. The data can also be freely shared with others.

C2. In cases where estimation is practical, our approach can be adapted to a new context by simply supplying a reliable incidence time series and selecting input articles. We demonstrate this by computing effective models for several different contexts even with a simple article selection procedure. Future, more powerful article selection procedures will increase the adaptability of the approach.

C3. In several instances, our models for the same disease in different locations are very similar; i.e., correlations between different language versions of the same article and the corresponding local disease incidence are similar. This suggests that simple techniques based on inter-language article mappings or other readily available data can be used to translate models from one context to another without re-training.

C4. Even our simple models show usefully high r^2 when forecasting a few days or weeks into the future. This suggests that the general approach can be used to build short-term forecasts with reasonably tight confidence intervals.

In short, this paper offers two key arguments. First, we evaluate the potential of an emerging data source, Wikipedia access logs, for global disease surveillance and forecasting in more detail than is previously available, and we argue that the openness and other properties of these data have important scientific and operational benefits. Second, using simple proof-of-concept experiments, we demonstrate that statistical techniques effective for estimating disease incidence using previous internet data are likely to also be effective using Wikipedia access logs.

We turn next to a more thorough discussion of prior work, both to set the stage for the current work as well as outline in greater detail the state of the art's relationship to the challenges above. Following that, we cover our methods and data sources, results, and a discussion of implications and future work.

Related Work

Our paper draws upon prior scholarly and practical work in three areas: traditional patient- and laboratory-based disease surveillance, Wikipedia-based measurement of the real world, and internet-based disease surveillance.

Traditional disease surveillance. Traditional forms of disease surveillance are based upon direct patient contact or biological tests taking place in clinics, hospitals, and laboratories. The majority of current systems rely on syndromic surveillance data (i.e., about symptoms) including clinical diagnoses, chief complaints, school and work absenteeism, illness-related 911 calls, and emergency room admissions [4].

For example, a well-established measure for influenza surveillance is the *fraction of patients with influenza-like illness*, abbreviated simply ILI. A network of outpatient providers report the total number of patients seen and the number who present with symptoms consistent with influenza that have no other identifiable cause [5]. Similarly, other electronic resources have emerged, such as the Electronic Surveillance System for the Early Notification of Community Based Epidemics (ESSENCE), based

on real-time data from the Department of Defense Military Health System [6] and BioSense, based on data from the Department of Veterans Affairs, the Department of Defense, retail pharmacies, and Laboratory Corporation of America [7]. These systems are designed to facilitate early detection of disease outbreaks as well as response to harmful health effects, exposure to disease, or related hazardous conditions.

Clinical labs play a critical role in surveillance of infectious diseases. For example, the Laboratory Response Network (LRN), consisting of over 120 biological laboratories, provides active surveillance of a number of disease agents in humans ranging from mild (e.g., non-pathogenic *E. coli* and *Staphylococcus aureus*) to severe (e.g., Ebola and Marburg), based on clinical or environmental samples [4]. Other systems monitor non-traditional public health indicators such as school absenteeism rates, over-the-counter medication sales, 911 calls, veterinary data, and ambulance run data. For example, the Early Aberration Reporting System (EARS) provides national, state, and local health departments alternative detection approaches for syndromic surveillance [8].

The main value of these systems is their accuracy. However, they have a number of disadvantages, notably cost and timeliness: for example, each ILI datum requires a practitioner visit, and ILI data are published only after a delay of 1–2 weeks [5].

Wikipedia. Wikipedia is an online encyclopedia that has, since its founding in 2001, grown to contain approximately 30 million articles in 287 languages [9]. In recent years, it has consistently ranked as a top-10 website; as of this writing, it is the 6th most visited website in the world and the most visited site that is not a search engine or social network [10], serving roughly 850 million article requests per day [11]. For numerous search engine queries, a Wikipedia article is the top result.

Wikipedia contrasts with traditional encyclopedias on two key dimensions: it is free of charge to read, and anyone can make changes that are published immediately — review is performed by the community *after* publication. (This is true for the vast majority of articles. Particularly controversial articles, such as “George W. Bush” or “Abortion”, have varying levels of edit protection.) While this surprising inversion of the traditional review-publish cycle would seem to invite all manner of abuse and misinformation, Wikipedia has developed effective measures to deal with these problems and is of similar accuracy to traditional encyclopedias such as *Britannica* [12].

Wikipedia article access logs have been used for a modest variety of research. The most common application is detection and measurement of popular news topics or events [13–17]. The data have also been used to study the dynamics of Wikipedia itself [18–20]. Social applications include evaluating toponym importance in order to make type size decisions for maps [21], measuring the flow of concepts across the world [22], and estimating the popularity of politicians and political parties [23]. Finally, economic applications include attempts to forecast movie ticket sales [24] and stock prices [25]. The latter two applications are of particular interest because they include a forecasting component, as the present work does.

In the context of health information, the most prominent research direction focuses on assessing the quality of Wikipedia as a health information source for the public, e.g., with respect to cancer [26,27], carpal tunnel syndrome [28], drug information [29], and kidney conditions [30]. To our knowledge, only four health studies exist that make use of Wikipedia access logs. Tausczik et al. examined public “anxiety and information seeking” during the 2009 H1N1 pandemic, in part by measuring traffic to

H1N1-related Wikipedia articles [31]. Laurent and Vickers evaluated Wikipedia article traffic for disease-related seasonality and in relation to news coverage of health issues, finding significant effects in both cases [32]. Aitken et al. found a correlation between drug sales and Wikipedia traffic for a selection of approximately 5,000 health-related articles [33]. None of these propose a time-series model mapping article traffic to disease metrics.

The fourth study is a recent article by McIver & Brownstein, which uses statistical techniques to estimate the influenza rate in the United States from Wikipedia access logs [34]. In the next section, we compare and contrast this article with the present work in the context of a broader discussion of such techniques.

In summary, use of Wikipedia access logs to measure real-world quantities is beginning to emerge, as is interest in Wikipedia for health purposes. However, to our knowledge, use of the encyclopedia for quantitative disease surveillance remains at the earliest stages.

Internet-based disease surveillance. Recently, new forms of surveillance based upon the social internet have emerged; these data streams are appealing in large part because of their real-time nature and the low cost of information extraction, properties complementary to traditional methods. The basic insight is that people leave traces of their online activity related to health observations, and these traces can be captured and used to derive actionable information. Two main classes of trace exist: *sharing* such as social media mentions of face mask use [35] and *health-seeking behavior* such as web searches for health-related topics [36]. (In fact, there is evidence that the volume of internet-based health-seeking behavior dwarfs traditional avenues [37,38].)

In this section, we focus on the surveillance work most closely related to our efforts, specifically, that which uses existing single-source internet data feeds to estimate some scalar disease-related metric. For example, we exclude from detailed analysis work that provides only alerts [39,40], measures public perception of a disease [41], includes disease dynamics in its model [42], evaluates a third-party method [43], uses non-single-source data feeds [39,44], or crowd-sources health-related data (*participatory disease surveillance*) [45,46]. We also focus on work that estimates biologically-rooted metrics. For example, we exclude metrics based on seasonality [47,48] and over-the-counter drug sales volume, itself a proxy [49].

These activity traces are embedded in search queries [36,50–76], social media messages [77–92], and web server access logs [34,72,93]. At a basic level, traces are extracted by counting query strings, words or phrases, or web page URLs that are related to the metric of interest, forming a time series of occurrences for each item. A statistical model is then created that maps these input time series to a time series estimating the metric’s changing value. This model is trained on time period(s) when both the internet data and the true metric values are available and then applied to estimate the metric value over time period(s) when it is not available, i.e., *forecasting* the future, *nowcasting* the present, and *anti-forecasting* the past (the latter two being useful in cases where true metric availability lags real time).

Typically, this model is linear, e.g.:

$$M = \sum_{j=1}^J \alpha_j x_j \quad (1)$$

where x_j is the count of some item, J is the total number of possible items (i.e., vocabulary size), M is the estimated metric

value, and α_j are selected by linear regression or similar methods. When appropriately trained, these methods can be quite accurate; for example, many of the cited models can produce near real-time estimates of case counts with correlations upwards of $r = 0.95$.

The collection of disease surveillance work cited above has estimated incidence for a wide variety of infectious and non-infectious conditions: avian influenza [52], cancer [55], chicken pox [67], cholera [81], dengue [50,53,84], dysentery [76], gastroenteritis [56,61,67], gonorrhea [64], hand foot and mouth disease (HFMD) [72], HIV/AIDS [75,76], influenza [34,36,54,57,59,62,63,65,67,68,71,74,77–80,82,83,85–93], kidney stones [51], listeriosis [70], malaria [66], methicillin-resistant *Staphylococcus aureus* (MRSA) [58], pertussis [90], pneumonia [68], respiratory syncytial virus (RSV) [52], scarlet fever [76], stroke [69], suicide [60,73], tuberculosis [76], and West Nile virus [52].

Closely related to the present work is an independent, simultaneous effort by McIver & Brownstein to measure influenza in the United States using Wikipedia access logs [34]. This study used Poisson models fitted with LASSO regression to estimate ILI over a 5-year period. The results, Pearson's r of 0.94 to 0.99 against official data, depending on model variation, compare quite favorably to prior work that tries to replicate official data. More generally, this article's statistical methods are more sophisticated than those employed in the present study. However, we offer several key improvements:

- We evaluate 14 location-disease contexts around the globe, rather than just one. In doing so, we test the use of language as a location proxy, which was noted briefly as future work in McIver & Brownstein. (However, as we detail below, we suspect this is not a reliable geo-location method for the long term.)
- We test our models for forecasting value, which was again mentioned briefly as future work in McIver & Brownstein.
- We evaluate models for translatability from one location to another.
- We present negative results and use these to begin exploring when internet-based disease surveillance methods might and might not work.
- We offer a systematic, well-specified, and simple procedure to select articles for model inclusion.
- We normalize article access counts by actual total language traffic rather than using a few specific articles as a proxy for total traffic.
- Our software is open source and has only freely available dependencies, while the McIver & Brownstein code is not available and depends on proprietary components (Stata).

Finally, the goals of the two studies differ. McIver & Brownstein wanted to “develop a statistical model to provide near-time estimates of ILI activity in the US using freely available data gathered from the online encyclopedia Wikipedia” [34, p. 2]. Our goals are to assess the applicability of these data to global disease surveillance for operational public health purposes and to lay out a research agenda for achieving this end.

These methods are the basis for at least one deployed, widely used surveillance system. Based upon search query data, Google Flu Trends offers near-real-time estimates of influenza activity in 29 countries across the world (15 at the province level); another facet of the same system, Google Dengue Trends (<http://www.google.org/denguetrends/>) estimates dengue activity in 9 countries (2 at the province level) in Asia and Latin America.

Having laid out the space of quantitative internet disease surveillance as it exists to the best of our knowledge, we now consider this prior work in the context of our four challenges:

C1. Openness. Deep access to search queries from Baidu, a Chinese-language search engine serving mostly the Chinese market (<http://www.baidu.com>) [64,74,76]; Google [36,50–54,56–60,65–67,69–73,75]; Yahoo [55,68]; and Yandex, a search engine serving mostly Russia and Slavic countries in Russian (<http://www.yandex.ru>), English (<http://www.yandex.com>), and Turkish (<http://www.yandex.com.tr>) [75], as well as purpose-built health website search queries [61–63] and access logs [72,93] are available only to those within the organizations, upon payment of an often-substantial fee, or by some other special arrangement. While tools such as Baidu Index (<http://index.baidu.com>), Google Trends (<http://www.google.com/trends/>), Google Correlate (<http://www.google.com/trends/correlate/>), and Yandex's WordStat (<http://wordstat.yandex.com>) provide a limited view into specific search queries and/or time periods, as do occasional lower-level data dumps offered for research, neither affords the large-scale, broad data analysis that drives the most effective models.

The situation is only somewhat better for surveillance efforts based upon Twitter [77–92]. While a small portion of the real-time message stream (1%, or 10% for certain grandfathered users) is available outside the company at no cost, terms of use prohibit sharing historical data needed for calibration between researchers. Access rules are similar or significantly more restrictive for other social media platforms such as Facebook and Sina Weibo, the leading Chinese microblogging site (<http://weibo.com>). Consistent with this, we were able to find no research meeting our inclusion criteria based on either of these extremely popular systems.

We identified only one prior effort making use of open data, McIver & Brownstein with Wikipedia access logs [34]. Open algorithms in this field of inquiry are also very limited. Of the works cited above, again only one, Althouse et al. [50], claims general availability of their algorithms in the form of open source code.

Finally, we highlight the quite successful Google Flu and Dengue Trends as a case study in the problems of closed data and algorithms. First, because their data and algorithms are proprietary, there is little opportunity for the wider community of expertise to offer peer review or improvements (for example, the list of search terms used by Dengue Trends has never been published, even in summary form); the importance of these opportunities is highlighted by the system's well-publicized estimation failures during the 2012–2013 flu season [94] as well as more comprehensive scholarly criticisms [43]. Second, only Google can choose the level of resources to spend on Trends, and no one else, regardless of their available resources, can add new contexts or take on operational responsibility should Google choose to discontinue the project.

C2. Breadth. While in principle these surveillance approaches are highly generalizable, nearly all extant efforts address a small set of diseases in a small set of countries, without testing specific methods to expand these sets.

The key exception is Paul & Dredze [91], which proposes a content-based method, *ailment topic aspect model* (ATAM), to automatically discover a theoretically unbounded set of medical conditions mentioned in Twitter messages. This unsupervised machine learning algorithm, similarly to latent Dirichlet allocation (LDA) [95], accumulates co-occurring words into probabilistic

topics. Lists of health-related lay keywords, as well as the text of health articles written for a lay audience, are used to ensure that the algorithm builds topics related to medical issues. A test of the method discovered 15 coherent condition topics including infectious diseases such as influenza, non-infectious diseases such as cancer, and non-specific conditions such as aches and pains. The influenza topic's time series correlated very well with ILI data in the United States.

However, we identify three drawbacks of this approach. First, significant curated text input data in the target language are required; second, output topics require expert interpretation; and third, the ATAM algorithm has several parameters that require expert tuning. That is, in order to adapt the algorithm to a new location and/or language, expertise in both machine learning as well as the target language are required.

In summary, to our knowledge, no disease measurement algorithms have been proposed that are extensible to new disease-location contexts solely by adding examples of desired output. We propose a path to such algorithms.

C3. Transferability. To our knowledge, no prior work offers trained models that can be translated from one context to another. We propose using the inter-language article links provided in Wikipedia to accomplish this translation.

C4. Forecasting. A substantial minority of the efforts in this space test some kind of forecasting method. (Note that many papers use the term *predict*, and some even misuse *forecast*, to indicate nowcasting.) In addition to forecasting models that incorporate disease dynamics (recall that these are out of scope for the current paper), two basic classes of forecasting exist: *lag analysis*, where the internet data are simply time-shifted in order to capture leading signals, and statistical forecast models such as linear regression.

Lag analysis has shown mixed results in prior work. Johnson et al. [93], Pelat et al. [67], and Jia-xing et al. [64] identified no reliable leading signals. On the other hand, Polgreen et al. [68] used lag analysis with a shift granularity of one week to forecast positive influenza cultures as well as influenza and pneumonia mortality with a horizon of 5 weeks or more (though these indicators may trail the onset of symptoms significantly). Similarly, Xu et al. [72] reported evidence that lag analysis may be able to forecast HFMD by up to two months, and Yang et al. [73] used lag analysis with a granularity of one month to identify search queries that lead suicide incidence by up to two months.

The more complex method of statistical forecast models appears potentially fruitful as well. Dugas et al. tested several statistical methods using positive influenza tests and Google Flu Trends to make 1-week forecasts [57], and Kim et al. used linear regression to forecast influenza on a horizon of 1 month [86].

In summary, while forecasts based upon models that include disease dynamics are clearly useful, sometimes this is not possible because important disease parameters are insufficiently known. Therefore, it is still important to pursue simple methods. The simplest is lag analysis; our contribution is to evaluate leading information more quantitatively than previously attempted. Specifically, we are unaware of previous analysis with shift granularity less than one week; in contrast, our analysis tests daily shifting even if official data are less granular, and each shift is an independently computed model; thus, our ± 28 -day evaluation results in 57 separate models for each context.

In summary, significant gaps remain with respect to the challenges blocking a path to an open, deployable, quantitative internet-based disease surveillance system. In this paper, we

propose a path to overcoming these challenges and offer evidence demonstrating that this path is plausible.

Methods

We used two data sources, Wikipedia article access logs and official disease incidence reports, and built linear models to analyze approximately 3 years of data for each of 14 disease-location contexts. This section details the nature, acquisition, and processing of these data as well as how we computed the estimation models and evaluated their output.

Wikipedia Article Access Logs

Access logs for all Wikipedia articles are available in summary form to anyone who wishes to use them. We used the complete logs available at <http://dumps.wikimedia.org/other/pagecounts-raw/>. Web interfaces offering a limited view into the logs, such as <http://stats.grok.se>, are also available. These data are referred to using a variety of terms, including *article views*, *article visits*, *pagecount files*, *page views*, *pageviews*, *page view logs*, and *request logs*.

These summary files contain, for each hour from December 9, 2007 to present and updated in real time, a compressed text file listing the number of requests for every article in every language, except that articles with no requests are omitted. (This request count differs from the true number of human views due to automated requests, proxies, pre-fetching, people not reading the article they loaded, and other factors. However, this commonly used proxy for human views is the best available.) We analyzed data from March 7, 2010 through February 1, 2014 inclusive, a total of 1,428 days. This dataset contains roughly 34,000 data files totaling 2.7TB. 266 hours or 0.8% of the data are missing, with the largest gap being 85 hours. These missing data were treated as zero; because they were few, this has minimal effect on our analyses.

We normalized these request counts by language. This yielded, for each article, a time series containing the number of requests for that article during each hour, expressed as a fraction of the hour's total requests for articles in the language. This normalization also compensates for periods of request undercounting, when up to 20% fewer requests were counted than served [96]. Finally, we transposed the data using Map-Reduce [97] to produce files from which the request count time series of any article can be retrieved efficiently.

Disease Incidence Data

Our goal was to evaluate a broad selection of diseases in a variety of countries across the world, in order to test the global applicability and disease agnosticism of our proposed technique. For example, we sought diseases with diverse modes of transmission (e.g., airborne droplet, vector, sexual, and fecal-oral), biology (virus, bacteria, protozoa), types of symptoms, length of incubation period, seasonality, and prevalence. Similarly, we sought locations in both the developed and developing world in various climates. Finally, we wanted to test each disease in multiple countries, to provide an opportunity for comparison.

These comprehensive desiderata were tempered by the realities of data availability. First, we needed reliable data establishing incidence ground truth for specific diseases in specific countries and at high temporal granularity; such official data are frequently not available for locations and diseases of interest. We used official epidemiological reports available on websites of government public health agencies as well as the World Health Organization (WHO).

Second, we needed article access counts for specific countries. This information is not present in the Wikipedia article access logs (i.e., request counts are global totals). However, a proxy is sometimes available in that certain languages are mostly limited to one country of interest; for example, a strong majority of Thai speakers are in Thailand, and the only English-speaking country where plague appears is the United States. In contrast, Spanish is spoken all over the world and thus largely unsuitable for this purpose.

Third, the language edition needs to have articles related to the disease of interest that are mature enough to evaluate and generate sufficient traffic to provide a reasonable signal.

With these constraints in mind, we used our professional judgement to select diseases and countries. The resulting list of 14 disease-location contexts, which is designed to be informative rather than comprehensive, is enumerated in Table 1.

These incidence data take two basic forms: (a) tabular files such as spreadsheets mapping days, weeks, or months to new case counts or the total number of infected persons or (b) graphs presenting the same mapping. In the latter case, we used plot digitizing software (Plot Digitizer, <http://plotdigitizer.sourceforge.net>) to extract a tabular form. We then translated these diverse tabular forms to a consistent spreadsheet format, yielding for each disease-location context a time series of disease incidence (these series are available in supplemental data S1).

Article Selection

The goal of our models is to create a linear mapping from the access counts of some set of Wikipedia articles to a scalar disease incidence for some disease-location context. To do so, a procedure for selecting these articles is needed; for the current proof-of-concept work, we used the following:

1. Examine the English-language Wikipedia article for the disease and enumerate the linked articles. Select for analysis the disease article itself along with linked articles on relevant symptoms, syndromes, pathogens, conditions, treatments, biological pro-

cesses, and epidemiology. For example, the articles selected for influenza include “Influenza”, “Amantadine”, and “Swine influenza”, but not “2009 flu pandemic”.

2. Identify the corresponding article in each target language by following the inter-language wiki link; these appear at the lower left of Wikipedia articles under the heading “Languages”. For example, the Polish articles selected for influenza include “Grypa”, “Amantadyna”, and “Świńska grypa”, but not “Pandemia grypy A/H1N1 w latach 2009–2010”, respectively.
3. Translate each article title into the form that appears in the logs. Specifically, encode the article’s Unicode title using UTF-8, percent-encode the result, and replace spaces with underscores. For example, the Polish article “Choroby zakaźne” becomes `Choroby_zaka%C5%Bane`. This procedure is accomplished by simply copying the article’s URL from the web browser address bar.

This procedure has two potential complications. First, an article may not exist in the target language; in this case, we simply omit it. Second, Wikipedia contains null articles called *redirects* that merely point to another article, called the *target* of the redirect. These are created to catch synonyms or common misspellings of an article. For example, in English, the article “Flu” is a redirect to “Influenza”. When a user visits <http://en.wikipedia.org/wiki/Flu>, the content served by Wikipedia is actually that of the “Influenza” article; the server does not issue an HTTP 301 response nor require the reader to manually click through to the redirect target.

This complicates our analysis because this arrangement causes the redirect itself (“Flu”), not the target (“Influenza”), to appear in the access log. While in principle we could sum redirect requests into the target article’s total, reliably mapping redirects to targets is a non-trivial problem because this mapping changes over time, and in fact Wikipedia’s history for redirect changes is not complete [98]. Therefore, we have elected to leave this issue for future work; this choice is supported by our observation below that when target and redirect are reversed, traffic to “Dengue fever” in Thai follows the target.

Table 1. Diseases-location contexts analyzed, with data sources.

Disease	Country	Language	Start	End	Resolution	Sources
Cholera	Haiti	French	2010-12-05	2013-12-05	daily	[99]
Dengue	Brazil	Portuguese	2010-03-07	2013-03-16	weekly	[100]
	Thailand	Thai	2011-01-01	2014-01-31	monthly	[101]
Ebola	Uganda/DRC	English	2011-01-01	2013-12-31	daily	[102–104]
HIV/AIDS	China (PRC)	Chinese	2011-01-01	2013-12-31	monthly	[105]
	Japan	Japanese	2010-10-09	2013-10-18	weekly	[106]
Influenza	Japan	Japanese	2010-06-26	2013-07-05	weekly	[106]
	Poland	Polish	2010-10-17	2013-10-23	weekly	[107]
	Thailand	Thai	2011-01-23	2014-02-01	weekly	[101]
	United States	English	2011-01-01	2014-01-10	weekly	[108]
Plague	United States	English	2011-01-22	2014-01-31	weekly	[109]
Tuberculosis	China (PRC)	Chinese	2010-12-01	2013-12-31	monthly	[105]
	Norway	Norwegian	2010-12-01	2013-12-31	monthly	[110]
	Thailand	Thai	2010-12-01	2013-12-31	monthly	[101]

This table lists the 7 diseases in 9 locations analyzed, for a total of 14 disease-location contexts. For each context, we list the language used as a location proxy, the inclusive start and end dates of analysis, the resolution of the disease incidence data, and one or more citations for those data.
doi:10.1371/journal.pcbi.1003892.t001

If we encounter a redirect during the above procedure, we use the target article. The complete selection of articles is available in the supplementary data S1.

Building and Evaluating Each Model

Our goal was to understand how well traffic for a simple selection of articles can nowcast and forecast disease incidence. Accordingly, we implemented the following procedure in Python to build and evaluate a model for each disease-location context.

1. Align the hourly article access counts with the daily, weekly, or monthly disease incidence data by summing the hourly counts for each day, week, or month in the incidence time series. This yields article and disease time series with the same frequency, making them comparable. (We ignore time zone in this procedure. Because Wikipedia data are in UTC and incidence data are in unspecified, likely local time zones, this leads to a temporal offset error of up to 23 hours, a relatively small error at the scale of our analysis. Therefore, we ignore this issue for simplicity.)
2. For each candidate article in the target language, compute Pearson's correlation r against the disease incidence time series for the target country.
3. Order the candidates by decreasing $|r|$ and select the best 10 articles.
4. Use ordinary least squares to build a linear multiple regression model mapping accesses to these 10 articles to the disease time series. No other variables were incorporated into the model. Below, we report r^2 for the multi-article models as well as a qualitative evaluation of success or failure. We also report r for individual articles in the supplementary data S1.

In order to test forecasting potential, we repeat the above with the article time series time-shifted from 28 days forward to 28 days backward in 1-day increments. For example, to build a 4-day forecasting model — that is, a model that estimates disease incidence 4 days in the future — we would shift the article time series *later* by 4 days so that article request counts for a given day are matched against disease incidence 4 days in the future. The choice of ± 28 days for lag analysis is based upon our *a priori* hypothesis that these statistical models are likely effective for a few weeks of forecasting.

We refer to models that estimate current (i.e., same-day) disease incidence as *nowcasting* models and those that estimate past disease incidence as *anti-forecasting* models; for example, a model that estimates disease incidence 7 days ago is a 7-day anti-forecasting model. (While useless at first glance, effective anti-forecasting models that give results sooner than official data can still reduce the lead time for action. Also, it is valuable for understanding the mechanism of internet-based models to know the temporal location of predictive information.) We report r^2 for each time-shifted multi-article model.

Finally, to evaluate whether translating models from one location to another is feasible, we compute a metric r_t for each pair of locations tested on the same disease. This meta-correlation is simply the Pearson's r computed between the correlation scores r of each article found in both languages; the intent is to give a gross notion of similarity between models computed for the same disease in two different languages. A value of 1 means that the two models are identical, 0 means they have no relationship, and -1 means they are opposite. We ignore articles found in only one language because the goal is to obtain a sense of feasibility: given favorable conditions, could one train a model in one location and apply it to another? Table 2 illustrates an example.

Results

Among the 14 disease-location contexts we analyzed, we found three broad classes of results. In 8 cases, the model succeeded, i.e., there was a usefully close match between the model's estimate and the official data. In 3 cases, the model failed, apparently because patterns in the official data were too subtle to capture, and in a further 3, the model failed apparently because the signal-to-noise ratio (SNR) in the Wikipedia data was too subtle to capture. Recall that this success/failure classification is based on subjective judgement; that is, in our exploration, we discovered that r^2 is insufficient to completely evaluate a model's goodness of fit, and a complementary qualitative evaluation was necessary.

Below, we discuss the successful and failed nowcasting models, followed by a summary and evaluation of transferability. (No models failed at nowcasting but succeeded at forecasting, so we omit a detailed forecasting discussion for brevity.)

Successful Nowcasting

Model and official data time series for selected successful contexts are illustrated in Figure 1. The method's good performance on dengue and influenza is consistent with voluminous prior work on these diseases; this offers evidence for the feasibility of Wikipedia access as a data source.

Success in the United States is somewhat surprising. Given the widespread use of English across the globe, we expected that language would be a poor location proxy for the United States. We speculate that the good influenza model performance is due to the high levels of internet use in United States, perhaps coupled with similar flu seasons in other Northern Hemisphere countries. Similarly, in addition to Brazil, Portuguese is spoken in Portugal and several other former colonies, yet problems again failed to arise. In this case, we suspect a different explanation: the simple absence of dengue from other Portuguese-speaking countries.

The case of dengue in Brazil is further interesting because it highlights the noise inherent in this social data source, a property shared by many other internet data sources. That is, noise in the input articles is carried forward into the model's estimate. We speculate that this problem could be mitigated by building a model on a larger, more carefully selected set of articles rather than just 10.

Finally, we highlight tuberculosis in China as an example of a marginally successful model. Despite the apparently low r^2 of 0.66, we judged this model successful because it captured the high baseline disease level excellently and the three modest peaks well. However, it is not clear that the model provides useful information at the time scale analyzed. This result suggests that additional quantitative evaluation metrics may be needed, such as root mean squared error (RMSE) or a more complex analysis considering peaks, valleys, slope changes, and related properties.

Forecasting and anti-forecasting performance of the four selected contexts is illustrated in Figure 2. In the case of dengue and influenza, the models contain significant forecast value through the limit of our 28-day analysis, often with the maximally effective lag comprising a forecast. We offer three possible reasons for this. First, both diseases are seasonal, so readers may simply be interested in the syndrome for this reason; however, the fact that models were able to correctly estimate seasons of widely varying severity provides counterevidence for this theory. Second, readers may be interested due to indirect reasons such as news coverage. Prior work disagrees on the impact of such influences; for example, Dukic et al. found that adding news coverage to their methicillin-resistant *Staphylococcus aureus* (MRSA) model had a limited effect [58], but recent Google Flu Trends failures appear to be caused in

Table 2. Transferability r_t example.

Article	Japanese	Thai
Fever	0.23	0.21
Chills	0.59	
Headache	-0.10	0.15
Influenza	0.85	0.77

This table shows simplified models for influenza in two locations: Japan, where Japanese is spoken, and Thailand, where Thai is spoken. The Japanese model yielded correlations for Japanese versions of the articles “Fever”, “Chills”, “Headache”, and “Influenza” of 0.23, 0.59, -0.10, and 0.85, respectively. The Thai model yielded correlations of 0.21, 0.15, and 0.77 for “Fever”, “Headache”, and “Influenza”, respectively. Note that the article “Chills” is not currently present in the Thai Wikipedia. Therefore, the correlation vectors are $\{0.23, -0.10, 0.85\}$ and $\{0.21, 0.15, 0.77\}$ for the two languages. The meta-correlation, r_t , for these two vectors, which provides a gross estimate of how similar the models are, is 0.97. Extending this computation to the full models yields $r_t = 0.81$, as noted below in Table 4.
doi:10.1371/journal.pcbi.1003892.t002

part by media activity [94]. Finally, both diseases have a relatively short incubation period (influenza at 1–4 days and dengue at 3–14); soon-to-be-ill readers may be observing the illness of their infectors or those who are a small number of degrees removed. It is the third hypothesis that is most interesting for forecasting purposes, and evidence to distinguish among them might be obtained from studies using simulated media and internet data, as suggested by Culotta [82].

Tuberculosis in China is another story. In this case, the model’s effectiveness is poorer as the forecast interval increases; we speculate that this is because seasonality is absent and the incubation period of 2–12 weeks is longer, diluting the effect of the above two mechanisms.

Failed Nowcasting

Figure 3 illustrates the three contexts where the model was not effective because, we suspect, it was not able to discern meaningful patterns in the official data. These suggest a few patterns that models might have difficulty with:

1. Noise. True patterns in data may be obscured by noise. For example, in the case of HIV/AIDS in China, the official data vary by a factor of 2 or more throughout the graph, and the model captures this fairly well, but the pattern seems epidemiologically strange and thus we suspect it may be merely noise. The other two contexts appear to also contain significant noise.

(Note that we distinguish noisy official data from an unfavorable signal-to-noise ratio, which is discussed below.)

2. Too slow. Disease incidence may be changing too slowly to be evident in the chosen analysis period. In all three contexts shown in Figure 3, the trend of the official data is essentially flat, with HIV/AIDS in Japan especially so. The models have captured this flat trend fairly well, but even doing so excellently provides little actionable value over traditional surveillance.

Both HIV/AIDS and tuberculosis infections progress quite slowly. A period of analysis longer than three years might reveal meaningful patterns that could be captured by this class of models. However, the social internet is young and turbulent; for example, even 3 years consumes most of the active life of some languages of Wikipedia. This complicates longitudinal analyses.

3. Too fast. Finally, incidence may be changing too quickly for the model to capture. We did not identify this case in the contexts we tested; however, it is clearly plausible. For example,

quarterly influenza data would be hard to model meaningfully using these techniques.

In all three patterns, improvements such as non-linear models or better regression techniques could lead to better results, suggesting that this is a useful direction for future work. In particular, noise suppression techniques as well as models tuned for the expected variation in a particular disease may prove fruitful.

Figure 4 illustrates the three contexts where we suspect the model failed due to a signal-to-noise ratio (SNR) in the Wikipedia data that was too low. That is, the number of Wikipedia accesses due to actual observations of infection is drowned out by accesses due to other causes.

In the case of Ebola, there are relatively few direct observations (a major outbreak has tens of cases), and the path to these observations becoming internet traces is hampered by poor connectivity in the sub-Saharan countries where the disease is active. On the other hand, the disease is one of general ongoing interest; in fact, one can observe on the graph a pattern of weekly variation (higher in midweek, lower on the weekend), which is common in online activity. In combination, these yield a completely useless model.

The United States has good internet connectivity, but plague has even lower incidence (the peak on the graph is three cases) and this disease is also popularly interesting, resulting in essentially the same effect. The cholera outbreak in Haiti differs in that the number of cases is quite large (the peak of the graph is 4,200 cases in one day). However, internet connectivity in Haiti was already poor even before the earthquake, and the outbreak was a major world news story, increasing noise, so the signal was again lost.

Performance Summary

Table 3 summarizes the performance of our models in the 14 disease-location contexts tested. Of these, we classified 8 as successful, producing useful estimates for both nowcasting and forecasting, and 6 as unsuccessful. Performance roughly broke down along disease lines: all influenza and dengue models were successful, while two of the three tuberculosis models were, and cholera, ebola, HIV/AIDS, and plague proved unsuccessful. Given the relatively simple model building technique used, this suggests that our Wikipedia-based approach is sufficiently promising to explore in more detail. (Another hypothesis is that model performance is related to popularity of the corresponding Wikipedia language edition. However, we found no relationship between r^2 and either a language’s total number of articles or total traffic.)

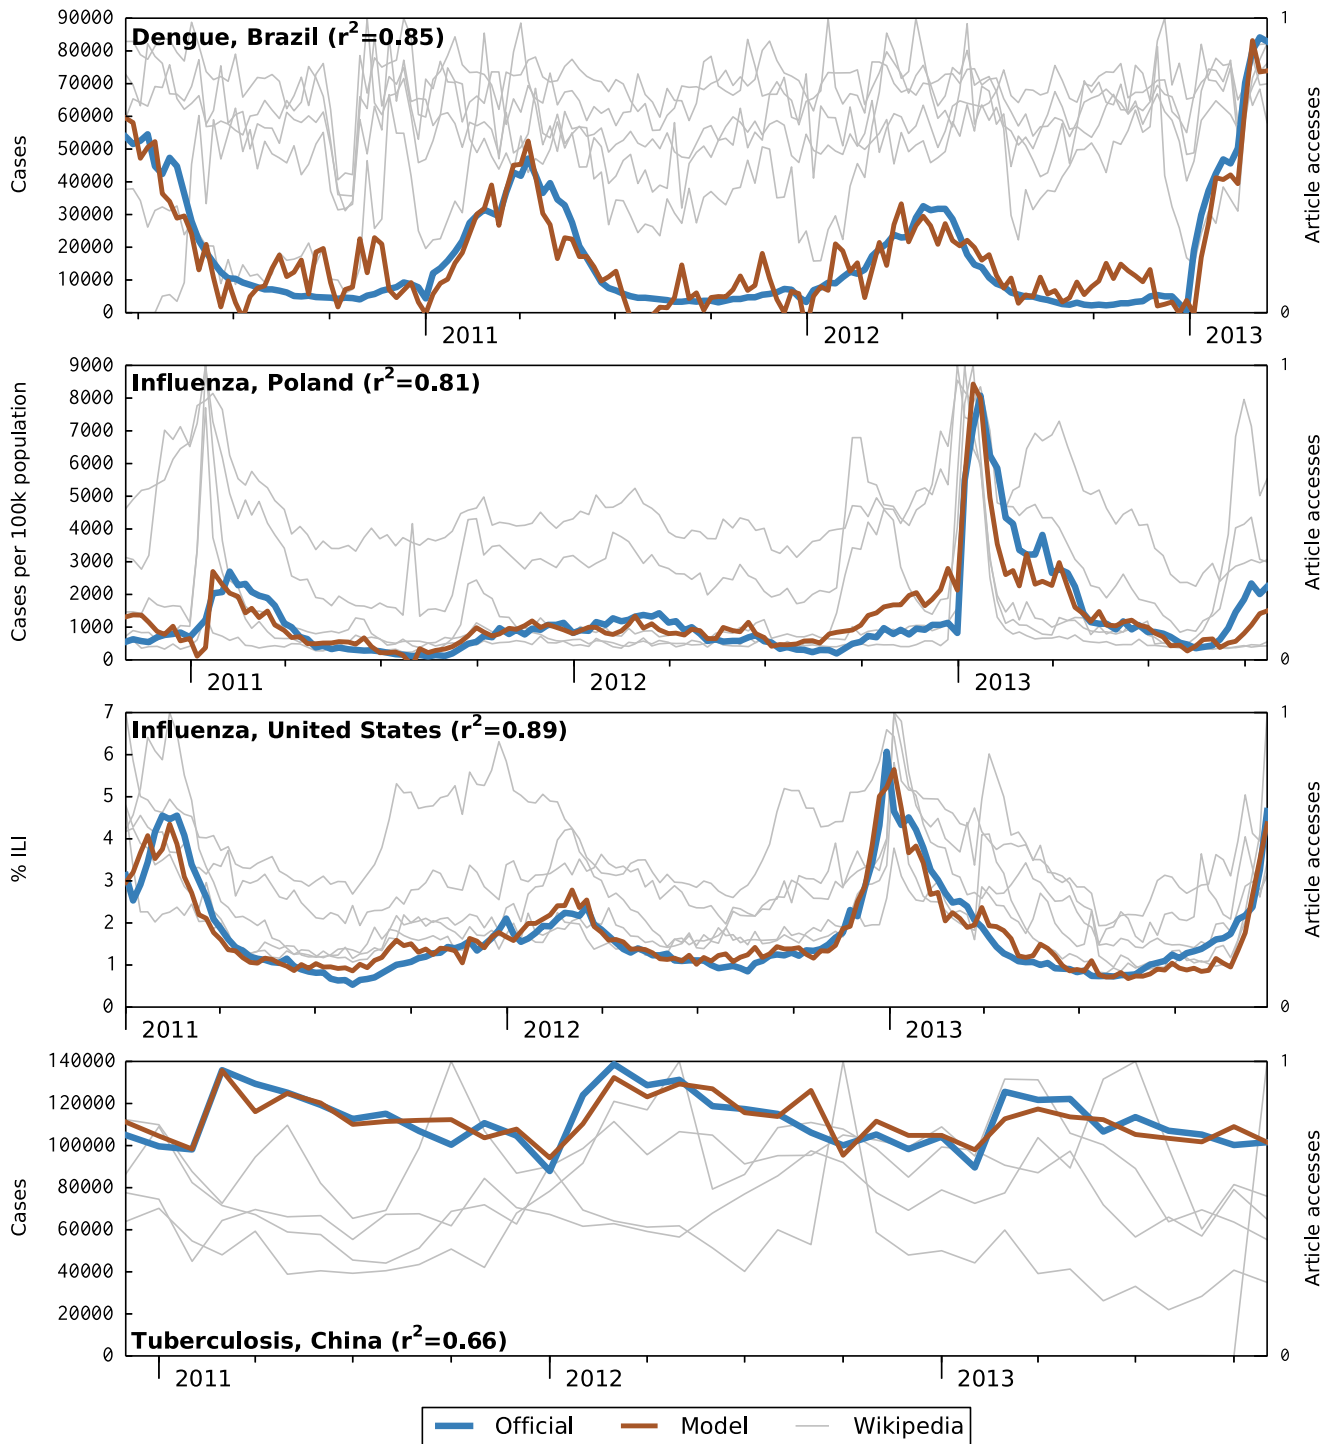


Figure 1. Selected successful model nowcasts. These graphs show official epidemiological data and nowcast model estimate (left Y axis) with traffic to the five most-correlated Wikipedia articles (right Y axis) over the 3 year study periods. The Wikipedia time series are individually self-normalized. Graphs for the four remaining successful contexts (dengue in Thailand, influenza in Japan, influenza in Thailand, and tuberculosis in Thailand) are included in the supplemental data file S1. doi:10.1371/journal.pcbi.1003892.g001

At a higher level, we posit that a successful estimation model based on Wikipedia access logs or other social internet data requires two key elements. First, it must be sensitive enough to capture the true variation in disease incidence data. Second, it must be sensitive enough to distinguish between activity traces due

to health-related observations and those due to other causes. In both cases, further research on modeling techniques is likely to yield sensitivity improvements. In particular, a broader article selection procedure — for example, using big data methods to test *all* non-trivial article time series for correlation, as Ginsberg et al.

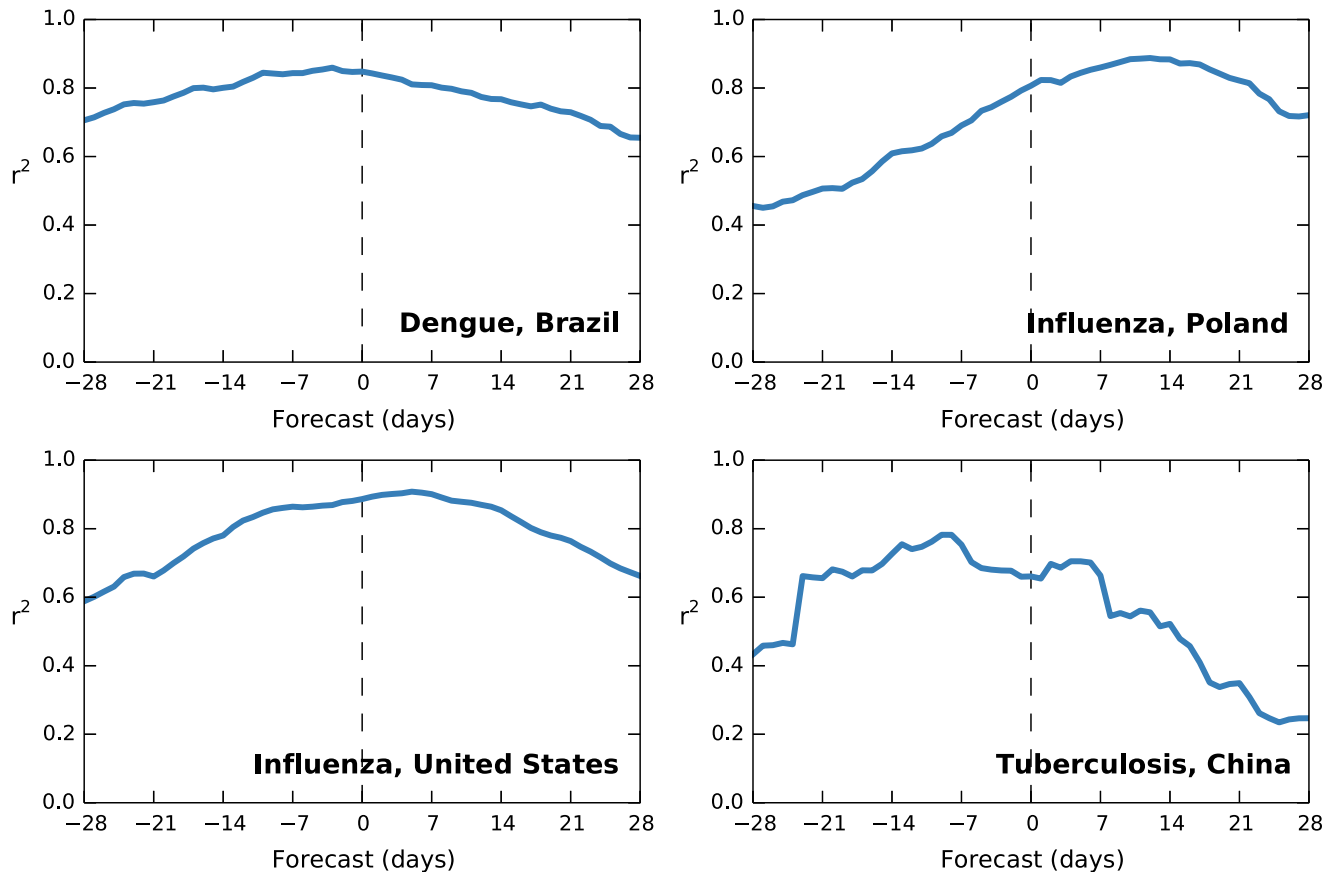


Figure 2. Forecasting effectiveness for selected successful models. This figure shows model r^2 compared to temporal offset in days: positive offsets are forecasting, zero is nowcasting (marked with a dotted line), and negative offsets are anti-forecasting. As above, figures for the four successful contexts not included here are in the supplemental data S1. doi:10.1371/journal.pcbi.1003892.g002

did for search queries [36] — is likely to prove fruitful, as might a non-linear statistical mapping.

Transferability

Table 4 lists the transferability scores r_t for each pair of countries tested on the same disease. Because this paper is concerned with establishing feasibility, we focus on the highest scores. These are encouraging: in the case of influenza, both Japan/Thailand and Thailand/United States are promising. That is, it seems plausible that careful source model selection and training techniques may yield useful models in contexts where no training data are available (e.g., official data are unavailable or unreliable). These early results suggest that research to quantitatively test methods for translating models from one disease-location context to another should be pursued.

Discussion

Human activity on the internet leaves voluminous traces that contain real and useful evidence of disease dynamics. Above, we pose four challenges currently preventing these traces from informing operational disease surveillance activities, and we argue that Wikipedia data are one of the few social internet data sources that can meet all four challenges. Specifically:

C1. Openness. Open data and algorithms are required, in order to offer reliable science as well as a flexible and robust

operational capability. Wikipedia access logs are freely available to anyone.

C2. Breadth. Thousands of disease-location contexts, not dozens, are needed to fully understand the global disease threat. We tested simple disease estimation models on 14 contexts around the world; in 8 of these, the models were successful with r^2 up to 0.92, suggesting that Wikipedia data are useful in this regard.

C3. Transferability. The greatest promise of novel disease surveillance methods is the possibility of use in contexts where traditional surveillance is poor or nonexistent. Our analysis uncovered pairs of same-disease, different-location models with similarity up to 0.81, suggesting that translation of trained models using Wikipedia's mappings of one language to another may be possible.

C4. Forecasting. Effective response to disease depends on knowing not only what is happening now but also what will happen in the future. Traditional mechanistic forecasting models often cannot be applied due to missing parameters, motivating the use of simpler statistical models. We show that such statistical models based on Wikipedia data have forecasting value through our maximum tested horizon of 28 days.

This preliminary study has several important limitations. These comprise an agenda for future research work:

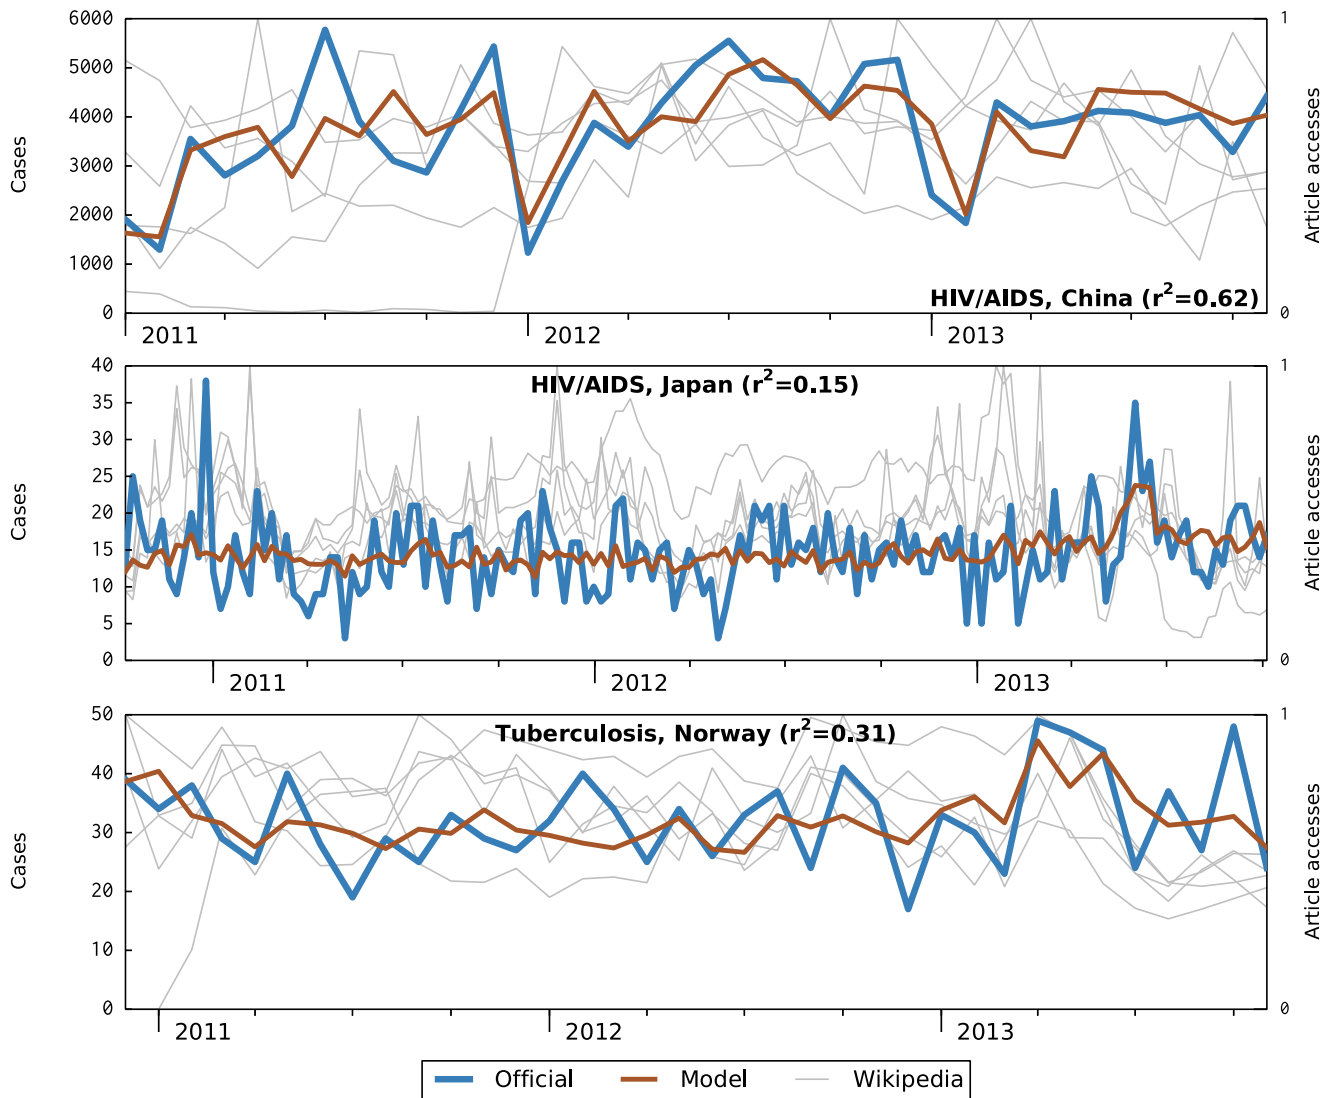


Figure 3. Nowcast attempts where the model was unable to capture a meaningful pattern in official data.
doi:10.1371/journal.pcbi.1003892.g003

1. The methods need to be tested in many more contexts in order to draw lessons about when and why this class of methods is likely to work.
 2. A better article selection procedure is needed. In the current paper, we tried a simple manual process yielding at most a few dozen candidate articles in order to establish feasibility. However, real techniques should use a comprehensive process that evaluates thousands, millions, or all plausible articles for inclusion in the model. This will also facilitate content analysis studies that evaluate which types of articles are predictive of disease incidence.
 3. Better geo-location is needed. While language as a location proxy works well in some cases, as we have demonstrated, it is inherently weak. In particular, it is implausible for use at a finer scale than country-level. What is needed is a hierarchical geographic aggregation of article traffic. The Wikimedia Foundation, operators of Wikipedia and several related projects, could do this using IP addresses to infer location before the aggregated data are released to the public. For most epidemiologically-useful granularities, this will still preserve reader privacy.
 4. Statistical estimation maps from article traffic to disease incidence should be more sophisticated. Here, we tried simple linear models mapping a single interval's Wikipedia traffic to a single interval's disease incidence. Future directions include testing non-linear and multi-interval models.
 5. Wikipedia data have a variety of instabilities that need to be understood and compensated for. For example, Wikipedia shares many of the problems of other internet data, such as highly variable interest-driven traffic caused by news reporting and other sources.
- Wikipedia has its own data peculiarities that can also cause difficulty. For example, during preliminary exploration for this paper in early July 2013, we used the inter-language link on the English article “Dengue fever” to locate the Thai version, “ไข้เลือดออกตั้งกี” (roughly, “dengue hemorrhagic fever”); article access logs indicated several hundred accesses per day for

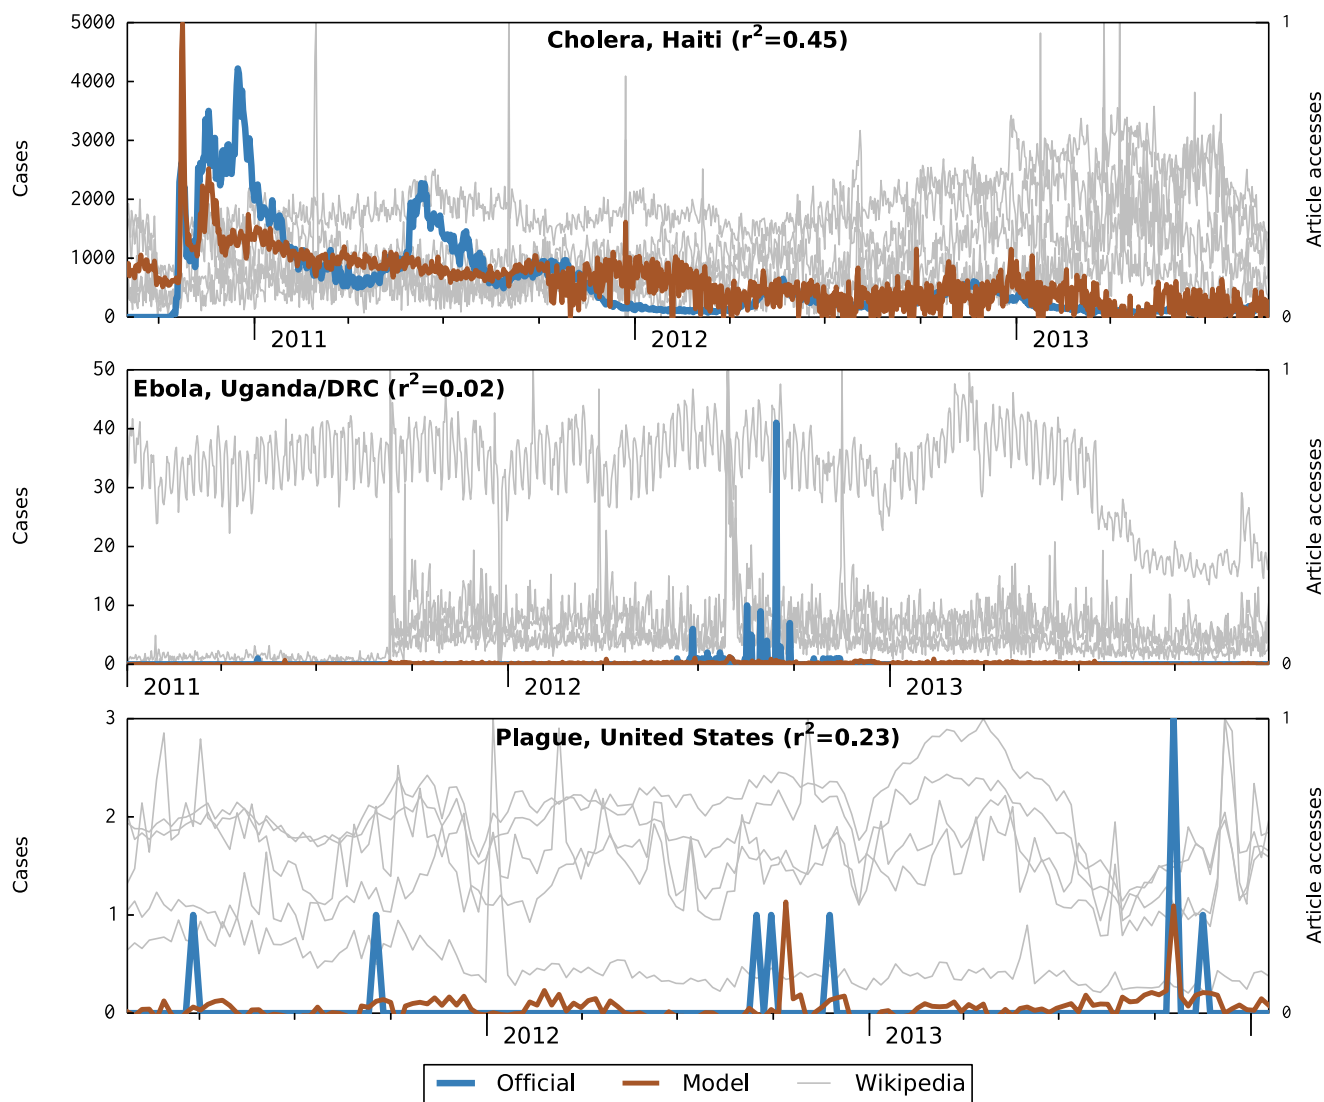


Figure 4. Nowcast attempts with poor performance due to unfavorable signal-to-noise ratio.
doi:10.1371/journal.pcbi.1003892.g004

this article in the month of June 2013. When we repeated the same process in March 2014, the inter-language link led to a page with the same content, but a different title, “ไข้เด็งกี” (roughly, “dengue fever”). As none of the authors are Thai speakers, and Google Translate renders both versions as “dengue fever”, we did not notice that the title of the Thai article had changed and were alarmed to discover that the article’s traffic in June 2013 was essentially zero.

The explanation is that before July 23, 2013, “ไข้เด็งกี” was a redirect to “ไข้เลือดออกเด็งกี”; on that day, the direction of the redirect was reversed, and almost all accesses moved over to the new redirect target over a period of a few days. That is, the article was the same all along, but the URL under which its accesses were recorded changed.

Possible techniques for compensation include article selection procedures that exclude such articles or maintaining a time-aware redirect graph so that different aliases of the same article can be merged. Indeed, when we tried the latter approach by manually summing the two URLs’ time series, it improved nowcast r^2 from 0.55 to 0.65. However, the first technique is likely to discard useful

information, and the second may not be reliable because complete history for this type of article transformation is not available [98].

In general, ongoing, time-aware re-training of models will likely be helpful, and limitations of the compensation techniques can be evaluated with simulation studies that inject data problems.

6. We have not explored the full richness of the Wikipedia data. For example, complete histories of each language edition are available, which include editing metadata (timestamps, editor identity, and comments), the text of each version, and conversations about the articles; these would facilitate analysis of edit activity as well as the articles’ changing text. Also, health-related articles are often mapped to existing ontologies such as the International Statistical Classification of Diseases and Related Health Problems (ICD-9 or ICD-10).
7. Transferability of models should be tested using more realistic techniques, such as simply building a model in one context and testing its performance in another.

Finally, it is important to recognize the biases inherent in Wikipedia and other social internet data sources. Most importantly,

Table 3. Model performance summary.

Disease	Location	Result	r^2 at forecast				Best forec.	
			0	7	14	28	Days	r^2
Cholera	Haiti	Failure (SNR)	0.45	0.39	0.41	0.48	26	0.50
Dengue	Brazil	Success	0.85	0.81	0.77	0.65	-3	0.86
	Thailand	Success	0.55	0.54	0.57	0.74	28	0.74
Ebola	Uganda/DRC	Failure (SNR)	0.02	0.01	0.02	0.02	5	0.14
HIV/AIDS	China (PRC)	Failure (Official data)	0.62	0.48	0.34	0.31	-1	0.63
	Japan	Failure (Official data)	0.15	0.19	0.15	0.05	9	0.22
Influenza	Japan	Success	0.82	0.92	0.86	0.52	8	0.92
	Poland	Success	0.81	0.86	0.88	0.72	12	0.89
	Thailand	Success	0.79	0.76	0.67	0.48	-2	0.80
	United States	Success	0.89	0.90	0.85	0.66	5	0.91
Plague	United States	Failure (SNR)	0.23	0.03	0.05	0.07	0	0.23
Tuberculosis	China (PRC)	Success	0.66	0.66	0.52	0.25	-9	0.78
	Norway	Failure (Official data)	0.31	0.41	0.40	0.42	20	0.48
	Thailand	Success	0.68	0.68	0.69	0.69	9	0.69

This table summarizes the performance of our estimation models. For each disease and location, we list the subjective success/failure classification as well as model r^2 at nowcasting (0-day forecast) and 7-, 14-, and 28-day forecasts. We also list the temporal offset in days of the best model (again, a positive offset indicates forecasting) along with that model's r^2 .
doi:10.1371/journal.pcbi.1003892.t003

Table 4. Transferability scores r_t for paired models.

Disease	Location 1	Location 2	r_t
Dengue	Brazil	Thailand	0.39
HIV/AIDS	China (PRC)	Japan	-0.06
Influenza	Japan	Poland	0.45
	Japan	Thailand	0.81
	Japan	United States	0.62
	Poland	Thailand	0.48
	Poland	United States	0.44
	Thailand	United States	0.76
Tuberculosis	China (PRC)	Norway	0.19
	China (PRC)	Thailand	-0.20
	Norway	Thailand	n/a

This table lists the transferability scores r_t for each tested pair of countries within a disease. Countries that did not share enough articles to compute a meaningful r_t are indicated with *n/a*.

doi:10.1371/journal.pcbi.1003892.t004

the data strongly over-represent people and places with good internet access and technology skills; demographic biases such as age, gender, and education also play a role. These biases are sometimes quantified (e.g., with survey results) and sometimes completely unknown. Again, simulation studies using synthetic internet data can quantify the impact and limitations of these biases.

Despite these limitations, we have established the utility of Wikipedia access logs for global disease monitoring and forecasting, and we have outlined a plausible path to a reliable, scientifically sound, operational disease surveillance system. We look forward to collaborating with the scientific and technical community to make this vision a reality.

Supporting Information

Dataset S1 Input data, raw results, and additional figures. This archive file contains: (a) inter-language article

mappings, (b) figures for the 4 successful contexts not included above, (c) official epidemiological data used as input, (d) complete correlation scores r , (e) wiki input data, and (f) a text file explaining the archive content and file formats.

(ZIP)

Acknowledgments

Mac Brown's participation in discussions improved this work significantly.

Author Contributions

Conceived and designed the experiments: NG GF AD SYDV RP. Performed the experiments: NG GF RP. Analyzed the data: NG GF RP. Contributed reagents/materials/analysis tools: NG GF RP. Wrote the paper: NG GF RP. Proofread the paper: AD SYDV.

References

- Lopez AD, Mathers CD, Ezzati M, Jamison DT, Murray CJ (2006) Global and regional burden of disease and risk factors, 2001: Systematic analysis of population health data. *The Lancet* 367: 1747–1757.
- Thompson M, Shay D, Zhou H, Bridges C, Cheng P, et al. (2010) Estimates of deaths associated with seasonal influenza — United States, 1976–2007. *Morbidity and Mortality Weekly Report* 59.
- Molinari NAM, Ortega-Sanchez IR, Messonnier ML, Thompson WW, Wortley PM, et al. (2007) The annual impact of seasonal influenza in the US: Measuring disease burden and costs. *Vaccine* 25: 5086–5096.
- Kman NE, Bachmann DJ (2012) Biosurveillance: A review and update. *Advances in Preventive Medicine* 2012.
- Centers for Disease Control and Prevention (CDC) (2012). Overview of influenza surveillance in the United States. URL <http://www.cdc.gov/flu/pdf/weekly/overview.pdf>.
- Bravata DM, McDonald KM, Smith WM, Rydzak C, Szeto H, et al. (2004) Systematic review: Surveillance systems for early detection of bioterrorism-related diseases. *Annals of Internal Medicine* 140: 910–922.
- Borchardt SM, Ritger KA, Dworkin MS (2006) Categorization, prioritization, and surveillance of potential bioterrorism agents. *Infectious Disease Clinics of North America* 20: 213–225.
- Hutwagner L, Thompson W, Seaman GM, Treadwell T (2003) The bioterrorism preparedness and response early aberration reporting system (EARS). *Journal of Urban Health* 80: i89–i96.
- Wikipedia editors (2013). Wikipedia. URL <http://en.wikipedia.org/w/index.php?title=Wikipedia&oldid=587396222>.
- Alexa Internet, Inc (2013). Alexa top 500 global sites. URL <http://www.alexa.com/topsites>. Accessed December 23, 2013.
- Wikimedia Foundation (2013). Page views for Wikipedia, all platforms, normalized. URL <http://stats.wikimedia.org/EN/TablesPageViewsMonthlyCombined.htm>. Accessed December 23, 2013.
- Giles J (2005) Internet encyclopaedias go head to head. *Nature* 438: 900–901.
- Ahn BG, Van Durme B, Callison-Burch C (2011) WikiTopics: What is popular on Wikipedia and why. In: Proc. Workshop on Automatic Summarization for Different Genres, Media, and Languages (WASDGM). Association for Computational Linguistics, p. 33–40. URL <http://dl.acm.org/citation.cfm?id=2018987.2018992>.
- Ciglan M, Norvåg K (2010) WikiPop: Personalized event detection system based on Wikipedia page view statistics. In: Proc. Information and Knowledge Management (CIKM). ACM, p. 1931–1932. doi:10.1145/1871437.1871769.
- Holaker MR, Emanuelsen E (2013) Event Detection using Wikipedia. Master's thesis, Institutt for datateknikk og informasjonsvitenskap. URL <http://www.diva-portal.org/smash/record.jsf?pid=diva2:655606>.
- Osborne M, Petrovic S, McCreddie R, Macdonald C, Ounis I (2012) Bieber no more: First story detection using Twitter and Wikipedia. In: Proc. SIGIR Workshop on Time-aware Information Access (TAIA). ACM. URL <http://www.dcs.gla.ac.uk/~craigm/publications/osborneTAIA2012.pdf>.
- Althoff T, Borth D, Hees J, Dengel A (2013) Analysis and forecasting of trending topics in online media streams. In: Proc. Multimedia. ACM, p. 907–916. doi:10.1145/2502081.2502117.
- Priedhorsky R, Chen J, Lam SK, Panciera K, Terveen L, et al. (2007) Creating, destroying, and restoring value in Wikipedia. In: Proc. Supporting Group Work (GROUP). ACM. doi: 10.1145/1316624.1316663.
- Thij Mt, Volkovich Y, Lamiado D, Kaltenbrunner A (2012) Modeling page-view dynamics on Wikipedia. arXiv:12125943 [physics].

20. Tran KN, Christen P (2013) Cross language prediction of vandalism on Wikipedia using article views and revisions. In: Pei J, Tseng VS, Cao L, Motoda H, Xu G, editors, *Advances in Knowledge Discovery and Data Mining*, Springer. pp. 268–279. URL http://link.springer.com/chapter/10.1007/978-3-642-37456-2_23.
21. Burdziej J, Gawrysiak P (2012) Using Web mining for discovering spatial patterns and hot spots for spatial generalization. In: Chen L, Felfernig A, Liu J, Raś ZW, editors, *Foundations of Intelligent Systems*, Springer. pp. 172–181. URL http://link.springer.com/chapter/10.1007/978-3-642-34624-8_21.
22. Tinati R, Tiropanis T, Carr L (2013) An approach for using Wikipedia to measure the flow of trends across countries. In: *Proc. World Wide Web (WWW) Companion*. ACM, p. 1373–1378. URL <http://dl.acm.org/citation.cfm?id=2487788.2488177>.
23. Yasserli T, Bright J (2013) Can electoral popularity be predicted using socially generated big data? arXiv:13122818 [physics].
24. Mestyań M, Yasserli T, Kertész J (2013) Early prediction of movie box office success based on Wikipedia activity big data. *PLOS ONE* 8: e71226.
25. Moat HS, Curme C, Avakian A, Kenett DY, Stanley HE, et al. (2013) Quantifying Wikipedia usage patterns before stock market moves. *Scientific Reports* 3.
26. Leithner A, Maurer-Ertl W, Glehr M, Friesenbichler J, Leithner K, et al. (2010) Wikipedia and osteosarcoma: A trustworthy patients' information? *Journal of the American Medical Informatics Association* 17: 373–374.
27. Rajagopalan MS, Khanna VK, Leiter Y, Stott M, Showalter TN, et al. (2011) Patient-oriented cancer information on the internet: A comparison of Wikipedia and a professionally maintained database. *Journal of Oncology Practice* 7: 319–323.
28. Lutsky K, Bernstein J, Beredjikian P (2013) Quality of information on the internet about carpal tunnel syndrome: An update. *Orthopedics* 36: e1038–e1041.
29. Kupferberg N, Protus BM (2011) Accuracy and completeness of drug information in Wikipedia: An assessment. *Journal of the Medical Library Association* 99: 310–313.
30. Thomas GR, Eng L, de Wolff JF, Grover SC (2013) An evaluation of Wikipedia as a resource for patient education in nephrology. *Seminars in Dialysis* 26: 159–163.
31. Tausczik Y, Faasse K, Pennebaker JW, Petrie KJ (2012) Public anxiety and information seeking following the H1N1 outbreak: Blogs, newspaper articles, and Wikipedia visits. *Health Communication* 27: 179–185.
32. Laurent MR, Vickers TJ (2009) Seeking health information online: Does Wikipedia matter? *Journal of the American Medical Informatics Association* 16: 471–479.
33. Aitken M, Altmann T, Rosen D (2014) Engaging patients through social media. *Tech report*, IMS Institute for Healthcare Informatics.
34. McIver DJ, Brownstein JS (2014) Wikipedia usage estimates prevalence of influenza-like illness in the United States in near real-time. *PLOS Computational Biology* 10: e1003581.
35. Mniszewski SM, Valle SYD, Priedhorsky R, Hyman JM, Hickman KS (2014) Understanding the impact of face mask usage through epidemic simulation of large social networks. In: Dabbaghian V, Mago VK, editors, *Theories and Simulations of Complex Social Systems*, Springer. pp. 97–115. URL http://link.springer.com/chapter/10.1007/978-3-642-39149-1_8.
36. Ginsberg J, Mohebbi MH, Patel RS, Brammer L, Smolinski MS, et al. (2008) Detecting influenza epidemics using search engine query data. *Nature* 457.
37. Rice RE (2006) Influences, usage, and outcomes of internet health information searching: Multivariate results from the Pew surveys. *International Journal of Medical Informatics* 75: 8–28.
38. Fox S (2006) *Online health search 2006*. Technical report, Pew Research Center. URL <http://www.pewinternet.org/Reports/2006/Online-Health-Search-2006.aspx>.
39. Collier N, Doan S, Kawazoe A, Goodwin RM, Conway M, et al. (2008) BioCaster: Detecting public health rumors with a Web-based text mining system. *Bioinformatics* 24: 2940–2941.
40. Zhou X, Li Q, Zhu Z, Zhao H, Tang H, et al. (2013) Monitoring epidemic alert levels by analyzing internet search volume. *IEEE Transactions on Biomedical Engineering* 60: 446–452.
41. Ritterman J, Osborne M, Klein E (2009) Using prediction markets and Twitter to predict a swine flu pandemic. In: *Proc. 1st International Workshop on Mining Social Media*. URL <http://homepages.inf.ed.ac.uk/miles/papers/swine09.pdf>.
42. Shaman J, Karspeck A, Yang W, Tamerius J, Lipsitch M (2013) Real-time influenza forecasts during the 2012–2013 season. *Nature Communications* 4: 2837.
43. Olson DR, Konty KJ, Paladini M, Viboud C, Simonsen L (2013) Reassessing Google Flu Trends data for detection of seasonal and pandemic influenza: A comparative epidemiological study at three geographic scales. *PLOS Computational Biology* 9: e1003256.
44. Freifeld CC, Mandl KD, Reis BY, Brownstein JS (2008) HealthMap: Global infectious disease monitoring through automated classification and visualization of internet media reports. *Journal of the American Medical Informatics Association* 15: 150–157.
45. Chunara R, Aman S, Smolinski M, Brownstein JS (2013) Flu Near You: An online self-reported influenza surveillance system in the USA. *Online Journal of Public Health Informatics* 5: e133.
46. Chunara R, Chhaya V, Bane S, Mekaru SR, Chan EH, et al. (2012) Online reporting for malaria surveillance using micro-monetary incentives, in urban India 2010–2011. *Malaria Journal* 11: 43.
47. Ayers JW, Althouse BM, Allem JP, Rosenquist JN, Ford DE (2013) Seasonality in seeking mental health information on Google. *American Journal of Preventive Medicine* 44: 520–525.
48. Seifter A, Schwarzwald A, Geis K, Aucott J (2010) The utility of “Google Trends” for epidemiological research: Lyme disease as an example. *Geospatial Health* 4: 135–137.
49. Lindh J, Magnusson M, Grünewald M, Hulth A (2012) Head lice surveillance on a deregulated OTC-sales market: A study using Web query data. *PLOS ONE* 7: e48666.
50. Althouse BM, Ng YY, Cummings DAT (2011) Prediction of dengue incidence using search query surveillance. *PLOS Neglected Tropical Diseases* 5: e1258.
51. Breyer BN, Sen S, Aaronson DS, Stoller ML, Erickson BA, et al. (2011) Use of Google Insights for Search to track seasonal and geographic kidney stone incidence in the United States. *Urology* 78: 267–271.
52. Carneiro HA, Mylonakis E (2009) Google Trends: A Web-based tool for real-time surveillance of disease outbreaks. *Clinical Infectious Diseases* 49: 1557–1564.
53. Chan EH, Sahai V, Conrad C, Brownstein JS (2011) Using Web search query data to monitor dengue epidemics: A new model for neglected tropical disease surveillance. *PLOS Neglected Tropical Diseases* 5: e1206.
54. Cho S, Sohn CH, Jo MW, Shin SY, Lee JH, et al. (2013) Correlation between national influenza surveillance data and Google Trends in South Korea. *PLOS ONE* 8: e81422.
55. Cooper CP, Mallon KP, Leadbetter S, Pollack LA, Peipins LA (2005) Cancer internet search activity on a major search engine, United States 2001–2003. *Journal of Medical Internet Research* 7.
56. Desai R, Hall AJ, Lopman BA, Shimsoni Y, Rennick M, et al. (2012) Norovirus disease surveillance using Google internet query share data. *Clinical Infectious Diseases* 55: e75–e78.
57. Dugas AF, Jalalpour M, Gel Y, Levin S, Torcaso F, et al. (2013) Influenza forecasting with Google Flu Trends. *PLOS ONE* 8: e56176.
58. Dukic VM, David MZ, Lauderdale DS (2011) Internet queries and methicillin-resistant *Staphylococcus aureus* surveillance. *Emerging Infectious Diseases* 17: 1068–1070.
59. Eysenbach G (2006) Infodemiology: Tracking flu-related searches on the Web for syndromic surveillance. *AMIA Annual Symposium* 2006: 244–248.
60. Hagihara A, Miyazaki S, Abe T (2012) Internet suicide searches and the incidence of suicide in young people in Japan. *European Archives of Psychiatry and Clinical Neuroscience* 262: 39–46.
61. Hulth A, Andersson Y, Hedlund KO, Andersson M (2010) Eye-opening approach to norovirus surveillance. *Emerging Infectious Diseases* 16: 1319–1321.
62. Hulth A, Rydevik G, Linde A (2009) Web queries as a source for syndromic surveillance. *PLOS ONE* 4: e4378.
63. Hulth A, Rydevik G (2011) Web query-based surveillance in Sweden during the influenza A(H1N1)2009 pandemic, April 2009 to February 2010. *Euro Surveillance* 16.
64. Jia-xing B, Ben-fu L, Geng P, Na L (2013) Gonorrhea incidence forecasting research based on Baidu search data. In: *Proc. Management Science and Engineering (ICMSE)*. IEEE, pp. 36–42. doi:10.1109/ICMSE.2013.6586259.
65. Kang M, Zhong H, He J, Rutherford S, Yang F (2013) Using Google Trends for influenza surveillance in South China. *PLOS ONE* 8: e55205.
66. Ocampo AJ, Chunara R, Brownstein JS (2013) Using search queries for malaria surveillance, Thailand. *Malaria Journal* 12: 390.
67. Pelat C, Turbelin C, Bar-Hen A, Flahault A, Valleron AJ (2009) More diseases tracked by using Google Trends. *Emerging Infectious Diseases* 15: 1327–1328.
68. Polgreen PM, Chen Y, Pennock DM, Nelson FD, Weinstein RA (2008) Using internet searches for influenza surveillance. *Clinical Infectious Diseases* 47: 1443–1448.
69. Walcott BP, Nahed BV, Kahle KT, Redjal N, Coumans JV (2011) Determination of geographic variance in stroke prevalence using internet search engine analytics. *Journal of Neurosurgery* 115: E19.
70. Wilson K, Brownstein JS (2009) Early detection of disease outbreaks using the internet. *Canadian Medical Association Journal* 180: 829–831.
71. Xu W, Han ZW, Ma J (2010) A neural network [sic] based approach to detect influenza epidemics using search engine query data. In: *Proc. Machine Learning and Cybernetics (ICMLC)*. IEEE, pp. 1408–1412. doi:10.1109/ICMLC.2010.5580851.
72. Xu D, Liu Y, Zhang M, Ma S, Cui A, et al. (2011) Predicting epidemic tendency through search behavior analysis. In: *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI, p. 2361–2366. doi:10.5591/978-1-57735-516-8/IJCAI11-393.
73. Yang AC, Tsai SJ, Huang NE, Peng CK (2011) Association of internet search trends with suicide death in Taipei City, Taiwan, 2004–2009. *Journal of Affective Disorders* 132: 179–184.
74. Yuan Q, Nsoesie EO, Lv B, Peng G, Chunara R, et al. (2013) Monitoring influenza epidemics in China with search query from Baidu. *PLOS ONE* 8: e64323.
75. Zheluk A, Quinn C, Hercz D, Gillespie JA (2013) Internet search patterns of human immunodeficiency virus and the digital divide in the Russian Federation: Infoveillance study. *Journal of Medical Internet Research* 15: e256.

76. Zhou Xc, Shen Hb (2010) Notifiable infectious disease surveillance with data collected by search engine. *Journal of Zhejiang University SCIENCE C* 11: 241–248.
77. Achrekar H, Gandhe A, Lazarus R, Yu SH, Liu B (2011) Predicting flu trends using Twitter data. In: *Proc. Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, pp. 702–707. doi:10.1109/INFOCOMW.2011.5928903.
78. Achrekar H, Gandhe A, Lazarus R, Yu SH, Liu B (2012) Twitter improves seasonal influenza prediction. In: *Proc. Health Informatics (HEALTHINF)*. p. 61–70. URL http://www.cs.uml.edu/~bliu/pub/healthinf_2012.pdf.
79. Aramaki E, Maskawa S, Morita M (2011) Twitter catches the flu: Detecting influenza epidemics using Twitter. In: *Proc. Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, p. 1568–1576. URL <http://dl.acm.org/citation.cfm?id=2145432.2145600>.
80. Broniatowski DA, Paul MJ, Dredze M (2013) National and local influenza surveillance through Twitter: An analysis of the 2012–2013 influenza epidemic. *PLOS ONE* 8: e83672.
81. Chunara R, Andrews JR, Brownstein JS (2012) Social and news media enable estimation of epidemiological patterns early in the 2010 Haitian cholera outbreak. *American Journal of Tropical Medicine and Hygiene* 86: 39–45.
82. Culotta A (2013) Lightweight methods to estimate influenza rates and alcohol sales volume from Twitter messages. *Language Resources and Evaluation* 47: 217–238.
83. Doan S, Ohno-Machado L, Collier N (2012) Enhancing Twitter data analysis with simple semantic filtering: Example in tracking influenza-like illnesses. In: *Proc. Healthcare Informatics, Imaging and Systems Biology (HISB)*. IEEE, pp. 62–71. doi:10.1109/HISB.2012.21.
84. Gomide J, Veloso A, Meira Jr W, Almeida V, Benevenuto F, et al. (2011) Dengue surveillance based on a computational model of spatio-temporal locality of Twitter. In: *Proc. Web Science Conference (WebSci)*. ACM, p. 1–8. URL http://www.websci11.org/fileadmin/websci/Papers/92_paper.pdf.
85. Hirose H, Wang L (2012) Prediction of infectious disease spread using Twitter: A case of influenza. In: *Proc. Parallel Architectures, Algorithms and Programming (PAAP)*. IEEE, pp. 100–105. doi:10.1109/PAAP.2012.23.
86. Kim EK, Seok JH, Oh JS, Lee HW, Kim KH (2013) Use of Hangeul Twitter to track and predict human influenza infection. *PLOS ONE* 8: e69305.
87. Lamb A, Paul MJ, Dredze M (2013) Separating fact from fear: Tracking flu infections on Twitter. In: *Proc. Human Language Technologies (NAACL-HLT)*. North American Chapter of the Association for Computational Linguistics, p. 789–795. URL <http://www.aclweb.org/anthology/N/N13/N13-1097.pdf>.
88. Lamos V, Cristianini N (2012) Nowcasting events from the social Web with statistical learning. *Transactions on Intelligent Systems and Technology* 3: 72: 1–72: 22.
89. Lamos V, Cristianini N (2010) Tracking the flu pandemic by monitoring the social Web. In: *Proc. Cognitive Information Processing (CIP)*. IEEE, pp. 411–416. doi:10.1109/CIP.2010.5604088.
90. Nagel AC, Tsou MH, Spitzberg BH, An L, Gawron JM, et al. (2013) The complex relationship of realspace events and messages in cyberspace: Case study of influenza and pertussis using Tweets. *Journal of Medical Internet Research* 15: e237
91. Paul MJ, Dredze M (2011) You are what you Tweet: Analyzing Twitter for public health. In: *Proc. Weblogs and Social Media (ICWSM)*. AAAI.
92. Signorini A, Segre AM, Polgreen PM (2011) The use of Twitter to track levels of disease activity and public concern in the U.S. during the influenza A H1N1 pandemic. *PLOS ONE* 6: e19467.
93. Johnson HA, Wagner MM, Hogan WR, Chapman W, Olszewski RT, et al. (2004) Analysis of Web access logs for surveillance of influenza. *Studies in Health Technology and Informatics* 107: 1202–1206.
94. Butler D (2013) When Google got flu wrong. *Nature* 494: 155–156.
95. Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. *Machine Learning Research* 3: 993–1022.
96. Zachte E (2012). *readme.txt*. URL <http://dumps.wikimedia.org/other/pagecounts-cz/projectcounts/readme.txt>. Accessed December 24, 2013.
97. Dean J, Ghemawat S (2008) MapReduce: Simplified data processing on large clusters. *Communications of the ACM (CACM)* 51: 107–113.
98. Wikipedia editors (2014). *Wikipedia:Moving a page*. URL http://en.wikipedia.org/w/index.php?title=Wikipedia:Moving_a_page&oldid=602263434. Section “Moving over a redirect”.
99. Ministère de la Santé Publique et de la Population. *Centere de documentation*. URL <http://mspp.gouv.ht/newsite/documentation.php>. Accessed January 23, 2014.
100. Ministério da Saúde. *Portal da saúde*. URL <http://portalsaude.saude.gov.br/portalsaude/index.cfm/?portal=pagina.visualizarArea&codArea=347>. Accessed September 26, 2013.
101. Bureau of Epidemiology. *Weekly epidemiological surveillance report, Thailand*. URL <http://www.boe-wesr.net/>. Accessed March 25, 2014.
102. World Health Organization (WHO) (2011). *Ebola in Uganda*. URL http://www.who.int/csr/don/2011_05_18/en/. Accessed March 25, 2014.
103. World Health Organization (WHO) (2012). *Ebola outbreak in Democratic Republic of Congo – Update*. URL http://www.who.int/csr/don/2012_10_26/en/. Accessed March 25, 2014.
104. Ministère de la Santé Publique, République Démocratique du Congo (2012). *Fièvre hémorragique à virus Ebola dans le district de haut uele*. URL <http://www.minisanterdc.cd/new/index.php/accueil/78-secgerale/92-fievre-hemorragique-a-virus-ebola-dans-le-district-de-haut-uele>. Accessed March 25, 2014.
105. Chinese Center for Disease Control and Prevention. *Notifiable infectious diseases statistic data*. URL <http://www.chinacdc.cn/tjsj/fdcrrbbg/>. Accessed September 23, 2014.
106. National Institute of Infectious Diseases, Japan. *Infectious diseases weekly report*. URL <http://www.nih.go.jp/niid/en/idwr-e.html>. Accessed January 24, 2013.
107. National Institute of Public Health. *Influenza in Poland*. URL <http://www.pzh.gov.pl/oldpage/epimeld/grypa/aindex.htm>. Accessed September 24, 2013.
108. Centers for Disease Control and Prevention (CDC). *FluView*. URL <http://gis.cdc.gov/grasp/fluview/fluportaldashboard.html>. Accessed March 25, 2014.
109. Centers for Disease Control and Prevention (CDC). *Morbidity and mortality weekly report (MMWR) tables*. URL <http://wonder.cdc.gov/mmwr/mmwr morb.asp>. Accessed March 25, 2014.
110. Meldingssystem for Smittsomme Sykdommer. *MSIS statistikk*. URL <http://www.msis.no/>. Accessed January 28, 2014.