ARTICLE IN PRESS

Epidemics xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Epidemics

journal homepage: www.elsevier.com/locate/epidemics

Results from the second year of a collaborative effort to forecast influenza seasons in the United States

Matthew Biggerstaff[a,*], Michael Johansson[b], David Alper[c], Logan C. Brooks[d], Prithwish Chakraborty[e], David C. Farrow[f], Sangwon Hyun[g], Sasikiran Kandula[h], Craig McGowan[a], Naren Ramakrishnan[e], Roni Rosenfeld[i], Jeffrey Shaman[h], Rob Tibshirani[j], Ryan J. Tibshirani[k], Alessandro Vespignani[l], Wan Yang[h], Qian Zhang[l], Carrie Reed[a]

[a] Epidemiology and Prevention Branch, Influenza Division, Centers for Disease Control and Prevention, Atlanta, GA, USA
[b] Dengue Branch, Division of Vector-Borne Diseases, Centers for Disease Control and Prevention, Atlanta, GA, USA
[c] Everyday Health, New York City, NY, USA
[d] Department of Computer Science, Carnegie Mellon University, Pittsburg, PA, USA
[e] Discovery Analytics Center, Department of Computer Science, Virginia Tech, Blacksburg, VA, USA
[f] Department of Computational Biology, Carnegie Mellon University, Pittsburg, PA, USA
[g] Deptartment of Statistics, Carnegie Mellon University, Pittsburg, PA, USA
[h] Department of Environmental Health Sciences, Mailman School of Public Health, Columbia University, New York, NY, USA
[i] Deptartment of Machine Learning, Department of Language Technologies, Department of Computational Biology, Department of Computer Science, Carnegie Mellon University, Pittsburg, PA, USA
[j] Department of Health Research and Policy, Department of Statistics, Stanford University, Stanford, CA, USA
[k] Deptartment of Statistics, Department of Machine Learning, Carnegie Mellon University, Pittsburg, PA, USA
[l] Northeastern University, Boston, MA, USA

ARTICLE INFO

Keywords:
Influenza
Epidemics
Forecasting
Prediction
Modeling

ABSTRACT

Accurate forecasts could enable more informed public health decisions. Since 2013, CDC has worked with external researchers to improve influenza forecasts by coordinating seasonal challenges for the United States and the 10 Health and Human Service Regions. Forecasted targets for the 2014–15 challenge were the onset week, peak week, and peak intensity of the season and the weekly percent of outpatient visits due to influenza-like illness (ILI) 1–4 weeks in advance. We used a logarithmic scoring rule to score the weekly forecasts, averaged the scores over an evaluation period, and then exponentiated the resulting logarithmic score. Poor forecasts had a score near 0, and perfect forecasts a score of 1.

Five teams submitted forecasts from seven different models. At the national level, the team scores for onset week ranged from < 0.01 to 0.41, peak week ranged from 0.08 to 0.49, and peak intensity ranged from < 0.01 to 0.17. The scores for predictions of ILI 1–4 weeks in advance ranged from 0.02–0.38 and was highest 1 week ahead. Forecast skill varied by HHS region.

Forecasts can predict epidemic characteristics that inform public health actions. CDC, state and local health officials, and researchers are working together to improve forecasts.

1. Introduction

Preparing for and responding to influenza epidemics and pandemics are critical functions of public health agencies. The Centers for Disease Control and Prevention (CDC) currently tracks influenza activity through a nationwide influenza surveillance system (Centers for Disease Control and Prevention, 2014a). Together with information on historic experiences, these data are used for situational awareness and assessing needs for the near future. However, these data lag behind real-time flu activity and give no direct insight on what might happen next. Accurate, timely, and reliable influenza forecasts could enable more informed public health and emergency response decisions during both influenza seasons and pandemics, including the development and use of pharmaceutical (e.g., vaccine and influenza antivirals) and non-

Abbreviations: CDC, centers for disease control and prevention; HHS, health and human service; ILI, influenza-like illness; MMWR, morbidity and mortality weekly report; ILINet, U.S. outpatient influenza-like illness surveillance network

* Corresponding author at: Centers for Disease Control and Prevention, 1600 Clifton Road NE MS A-32, Atlanta, GA 30333, USA.
E-mail addresses: zmo2@cdc.gov, mbiggerstaff@cdc.gov (M. Biggerstaff).

https://doi.org/10.1016/j.epidem.2018.02.003
Received 1 May 2017; Received in revised form 6 February 2018; Accepted 20 February 2018
1755-4365/ Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/BY-NC-ND/4.0/).

Please cite this article as: Biggerstaff, M., Epidemics (2018), https://doi.org/10.1016/j.epidem.2018.02.003

pharmaceutical (e.g., school closures and social distancing, travel restrictions) countermeasures, communication, deployment of Strategic National Stockpile assets (e.g., ventilators), and hospital resource management (e.g., inventory and staff management) (Chretien et al., 2014).

CDC's Influenza Division began working in 2013 to advance influenza forecasting efforts by engaging with members of the scientific community who were developing innovative methods to predict influenza activity (Brooks et al., 2015; Shaman et al., 2009; Shaman and Karspeck, 2012; Kandula et al., 2017; Tizzoni et al., 2012; Balcan et al., 2009; Nsoesie et al., 2014). This effort launched with the "Predict the Influenza Season Challenge," a contest which encouraged participants to predict the timing, peak, and intensity of the 2013–14 influenza season using social media data (e.g., Twitter, internet search data, web surveys, etc.) along with data from CDC's routine flu surveillance systems (Centers for Disease Control and Prevention, 2013). Eleven teams participated in the original CDC competition, and team members developed their own models to predict flu activity based on a variety of data sources (Biggerstaff et al., 2016). This challenge identified a number of research gaps limiting forecasting model development, evaluation, and adoption by decision-makers, including the need to develop standardized metrics to assess forecast accuracy and standardized ways to communicate forecasts and their uncertainty.

To address these gaps, CDC and original challenge participants worked together through a collaborative challenge to forecast the 2014–15 influenza season. The objectives of this challenge were to continue to improve the accuracy of influenza forecasts, develop standardized metrics to assess and communicate forecast accuracy and uncertainty, and to identify the types of decisions best aided by forecasts. Challenge participants were asked to forecast seasonal milestones (the onset, peak, and intensity) and short-term activity during the 2014–15 influenza season for the United States as a country and for each of the 10 Health and Human Services (HHS) regions. In this report, we present the results and lessons learned from the challenge.

2. Methods

Teams that participated in CDC's 2013–14 Predict the Influenza Season Challenge were invited to continue to work with CDC to provide forecasts for the 2014–15 influenza season in the United States. This group of teams and CDC collaboratively defined a set of forecast targets and established evaluation metrics to assess accuracy prior to the challenge. Participating groups then submitted weekly forecasts for the 2014–2015 influenza season beginning October 20, 2014, and ending May 25, 2015. Forecasting targets were selected to ensure they were feasible for forecasting models and provided information for public health decision making.

All forecasting targets were based on data from the U.S. Outpatient Influenza-like Illness (ILI) Surveillance Network (ILINet). ILINet provides accurate information on the timing and impact of influenza activity each season and consists of more than 2000 outpatient healthcare providers around the country who report data to CDC weekly on the number of patients with ILI and the total number of patients seen in their practices (Centers for Disease Control and Prevention, 2014a; Brammer et al., 2011). ILINet data are based on a Morbidity and Mortality Weekly Report (MMWR) surveillance week that starts on Sunday and ends on Saturday; data are reported online through CDC's FluView surveillance report the following Friday (or Monday if federal holidays delay publication) (Centers for Disease Control and Prevention, 2014b). Further information on ILINet is available elsewhere (Centers for Disease Control and Prevention, 2014a; Brammer et al., 2011). Teams could use any other data sources available to them, including digital (e.g., Twitter data, mining internet search term data, Internet-based surveys), meteorological, and traditional surveillance.

The minimum set of forecasts required of all participants were national-level forecasts of the onset week, peak week, and peak intensity of the influenza season (collectively referred to in the paper as seasonal targets), and short-term forecasts of the weekly percentage of outpatient ILINet visits due to ILI one, two, three, and four weeks after the week most recently reported by ILINet in FluView (collectively referred to in the paper as short-term targets). Participants also had the option of submitting forecasts of the same targets for each of the 10 HHS regions. We defined the onset of the season as the first surveillance week in ILINet where the ILINet percentage was at or above the baseline value (which is developed by calculating the mean percentage of patient visits for ILI during non-influenza weeks for the previous three seasons and adding two standard deviations (Centers for Disease Control and Prevention, 2014a) and remained there for at least two additional weeks. We defined the peak week of the season as the surveillance week that the ILINet percentage was the highest; if more than one week achieved the highest value, all such weeks were considered peak weeks. We defined the peak value as the highest numeric value that the ILINet percentage reached (Centers for Disease Control and Prevention, 2014b).

Each forecast included a point estimate and a probability distribution within pre-defined bins for each target. For onset and peak weeks, each bin represented a single week (e.g., week 1, week 2). For start week, an additional bin was used for the probability that the onset week definition would not be met during the influenza season. For the peak percentage of outpatient visits due to ILI and the weekly percentage of ILI one to four weeks in advance, 11 bins were used; 10 bins represented semi-open 1% intervals (e.g., 3% < = ILI peak value < 4.0%) from 0% to 10% while the final bin represented all values greater than or equal to 10%. Teams were also required to submit a narrative describing the methodology of the forecasting model. The forecasting methodology could be changed during the course of the season if an updated narrative describing the changes was provided; no team indicated that they changed their methodology during the 2014–15 season.

We used the logarithmic scoring rule to measure the accuracy of the probability distribution of a forecast (Gneiting and Raftery, 2007; Rosenfeld et al., 2018). If $\mathbf{p}$ is the set of probabilities across all bins for a given forecast, and p_i is the probability assigned to the observed outcome, i, the logarithmic score is $S(\mathbf{p},i) = \ln(p_i)$. For example, a forecast that assigned a probability of 0.6 to the correct influenza season onset week would receive a score of $\ln(0.6) = -0.51$. Undefined natural logs (which occur when the probability assigned to the observed outcomes was 0), missing forecasts, and forecasts that summed to probabilities less than 0.9 or greater than 1.1 were assigned a value of -10. Logarithmic scores were averaged across different combinations of seasonal and short-term targets, geographic locations, and time periods. For the seasonal targets, the evaluation period was chosen post hoc to represent periods when the forecasts would be most useful and began with the first forecast submission on October 20, 2014, while the end of the evaluation period varied by seasonal target. The evaluation period end for the onset target was the forecast received after the week in which peak occurred in the final ILINet data, and the evaluation period end for the peak week and peak percent targets was the forecast received after the final week ILINet was above baseline (Table 1 and Supplemental Tables 1–10). For the short-term forecasts, time periods were chosen to represent forecasts that were received during the weeks that ILINet was above baseline (Table 1 and Supplemental Tables 1–10). Evaluation results for national- and regional-level targets using forecasts from the entire forecast period (October 20, 2014 to May 25, 2015) are found in Supplemental Table 11. Because ILINet data for past weeks may change as more reports are received, we used the ILINet data weighted on the basis of state population reported on week 34 of 2015 (the week ending August 29) for forecast evaluation.

To aid in interpretation, we exponentiated the mean log score to indicate forecast skill on a 0–1 scale. Perfect forecasts (i.e. forecasted probability of 1.0 for the observed outcome across all forecasts) have a log score of 0 and a forecast skill of 1. For forecasts with low

Download English Version:

https://daneshyari.com/en/article/8950977

Download Persian Version:

https://daneshyari.com/article/8950977

Daneshyari.com