

DEFSI: Deep Learning Based Epidemic Forecasting with Synthetic Information

Lijing Wang,¹ Jiangzhuo Chen,² Madhav Marathe^{2,3}

¹Computer Science, Virginia Tech

²Biocomplexity Institute & Initiative, University of Virginia

³Computer Science, University of Virginia

lijingw9@vt.edu, {chenj,marathe}@virginia.edu

Abstract

Influenza-like illness (ILI) is among the most common diseases worldwide. Producing timely, well-informed, and reliable forecasts for ILI is crucial for preparedness and optimal interventions. In this work, we focus on short-term but high-resolution forecasting and propose DEFSI (*Deep Learning Based Epidemic Forecasting with Synthetic Information*), an epidemic forecasting framework that integrates the strengths of artificial neural networks and causal methods. In DEFSI, we build a two-branch neural network structure to take both within-season observations and between-season observations as features. The model is trained on geographically high-resolution synthetic data. It enables detailed forecasting when high-resolution surveillance data is not available. Furthermore, the model is provided with better generalizability and physical consistency. Our method achieves comparable/better performance than state-of-the-art methods for short-term ILI forecasting at the state level. For high-resolution forecasting at the county level, DEFSI significantly outperforms the other methods.

Introduction

Influenza-like illness (ILI) poses a serious threat to global public health. Worldwide seasonal influenza causes three to five million cases of severe illness and 290,000 to 650,000 deaths annually. Traditionally, ILI surveillance data from the Centers for Disease Control and Prevention (CDC) has been used as reference data to predict future ILI incidence. The surveillance data is provided with a coarse resolution and is usually updated regularly. For example, in the USA it has been provided previously at HHS region level and recently at state level. Considering the heterogeneity between different subregion locations and populations, accurate predictions with a finer resolution, e.g. at county level in the USA, is crucial for local public health decision making, optimal mitigation resource allocation among subregions, as well as household or individual level preventive actions informed by neighboring prevalence. We focus on the problem of **high-resolution ILI incidence forecasting** based on ILI surveillance data of coarse resolution.

In this paper we use *flat-resolution* forecasting to denote the forecasting of ILI incidence with the same resolu-

tion as the surveillance data; and *high-resolution* forecasting to denote the forecasting with a higher geographical resolution than provided in surveillance data. To be concrete, in this paper flat-resolution means state level while high-resolution means county level, since the highest resolution CDC surveillance data is at state level.

A challenge in high-resolution disease forecasting is the lack of surveillance data at a finer spatial scale. The data driven methods suffer from this limitation, including statistical methods such as ARIMA (Benjamin, Rigby, and Stasinopoulos 2003), ARGO (Yang, Santillana, and Kou 2015; Yang et al. 2017) and artificial neural networks (ANN) methods such as Long Short Term Memory (LSTM) (Volkova et al. 2017; Venna et al. 2017; Wu et al. 2018). Even for a few states where county level surveillance data is available, training ANN methods for them is difficult due to the small size of the data. Causal methods have recently been introduced to enable high-resolution forecasting (Yang, Karspeck, and Shaman 2014; Nsoesie et al. 2013; Zhao et al. 2015). They estimate the parameters of the underlying disease model from the surveillance data. Then ILI incidence prediction is made from the output of simulations using the identified disease model. Calibrating these models is challenging due to the high dimensionality of the parameter space.

Another challenge to the epidemic forecasting domain is the generalization and physical consistency of data driven models due to the paucity of representative samples in nature. Meanwhile it is difficult for causal models to represent complex processes that are not conceptually well understood using known scientific knowledge. This motivates our work to address the shortcomings of data driven and causal models by using synthetic data generated by epidemiological theories in conjunction with observed data to improve epidemic forecasts. In a recent independent effort, Karpatne et al. (Karpatne et al. 2017) formally conceptualized the paradigm of theory-guided data science (TGDS) that seeks to exploit the promise of data science without ignoring the treasure of knowledge accumulated in scientific principles. Our work starts from a similar perspective but is different in that multi-agent models enable high-resolution forecasting and better generalizability of our models. This is important due to behavioral uncertainty of a social technical system.

To address the above challenges, we propose a novel epi-

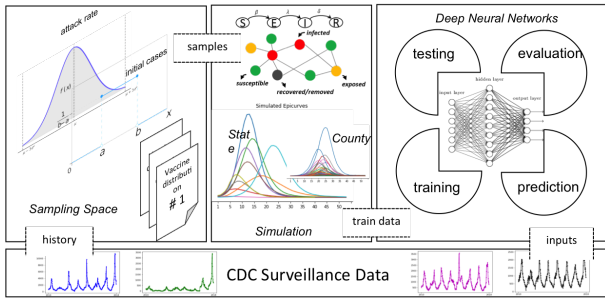


Figure 1: DEFSI framework. In this framework, a region-specific disease parameter space is constructed based on surveillance data. Then the synthetic training data consisting of both state level and county level weekly ILI incidence curves is generated by simulations parameterized by samples from the parameter space. A two-branch deep neural network model is trained on the synthetic data. The trained model takes surveillance data as the input during forecasting.

demographic forecasting framework, called **Deep Learning Based Epidemic Forecasting with Synthetic Information (DEFSI)**. It combines **multi-agent system** and **deep neural network** techniques from artificial intelligence (AI). The idea is to model the non-linear relationship between the past higher level (state) ILI incidences and the future higher (state) and lower level (county) ILI incidences with a deep neural network. The novelty of our approach is to generate the training data using a multi-agent simulation that is based on a synthetic population and contact network, where agent heterogeneities and unstructured interactions among agents are modeled. With a multi-agent model, individual or household level behavior can also be modeled as well as the public health intervention measures changing the disease dynamics. The purpose is to use simulations to create training data as similar as possible to the surveillance data observed in the real world.

To the best of our knowledge, DEFSI is the first to combine a realistic multi-agent model with deep learning for epidemic forecasting. **Our major contributions are as follows:** (1) DEFSI enables accurate high-resolution forecasting with flat-resolution observations as inputs. (2) DEFSI proposes a two-branch neural network model for ILI forecasting. It combines within-season observations (observed data points of the current season that characterize the ongoing epidemic) and between-season historical observations (observed data points from similar weeks of the past seasons that characterize general trends around the current week). (3) DEFSI constructs region-specific training datasets at multiple spatially fine-grained scale with low costs. We initialize region-specific simulations with realistic parameter settings learned from the corresponding surveillance data. (4) Extensive experiments on ILI incidence forecasting for two states of the USA show that DEFSI achieves comparable/better performance than the state-of-the-art methods at state level. For high-resolution forecasting at county level, DEFSI significantly outperforms the comparison methods.

Problem Setup

Given an observed time-series of weekly ILI incidence for a specific region, we focus on predicting ILI incidence for both the region and its subregions in short-term. Without loss of generality, in this paper we consider making predictions for a state of the USA and all counties of the state, using CDC state level ILI incidence data (CDC 2018a). In this setting, state level forecasting is flat-resolution, while county level forecasting is high-resolution. Let $\mathbf{y} = \langle y_1, y_2, \dots, y_T, \dots \rangle$ denote the sequence of weekly state level ILI incidence, where T is the last week of which the ILI incidence is given. Similarly, $\mathbf{y}^C = \langle y_1^C, y_2^C, \dots, y_T^C, \dots \rangle$ denotes the sequence of weekly ILI incidence for a particular county C within the state. Assume that there are K counties $\mathcal{D} = \{C_1, C_2, \dots, C_K\}$ in the state. Let $\mathbf{y}_t^{\mathcal{D}} = \{y_t^C | C \in \mathcal{D}\}$ denote ILI incidence of all counties in the state at week t . The consistency constraint on county level incidence is $y_t = \sum_{C \in \mathcal{D}} y_t^C$. The problem is defined as predicting both state level and county level incidence at week t , where $t = T + 1$, denoted as $\mathbf{z}_t = (y_t, \mathbf{y}_t^{\mathcal{D}})$, given historical state level incidence.

DEFSI

Framework DEFSI framework consists of three major components (shown in Fig. 1): (1) *Disease model parameter space construction*: Given an existing disease simulator, we estimate a marginal distribution for each model parameter based on the state-specific surveillance data; (2) *Synthetic training data generation*: We generate a synthetic training dataset at both flat-resolution and high-resolution scales for that state by running simulations parameterized from the parameter space; (3) *Deep neural network training and forecasting*: We design a two-branch deep neural network model trained on the synthetic training dataset and using surveillance data as inputs for predictions. We will elaborate the details in the following subsections.

SEIR-based Epidemic Simulation The SEIR disease model is widely used for ILI diseases (Kuznetsov and Piccardi 1994). Each person is in one of the following four health states at any time: susceptible (S), exposed (E), infectious (I), recovered or removed (R). A person v is in the susceptible state until he becomes exposed. If v becomes exposed, he remains so for $p_E(v)$ days, which is called the incubation period, during which he is not infectious. Then he becomes infectious and remains so for $p_I(v)$ days, which is called the infectious period. Finally he becomes removed (or recovered) and remains so permanently. While the SEIR model characterizes within-host disease progression, between-host disease propagation is modeled by transmission from person to person with a probability parameter τ , through either complete mixing or heterogeneous connections among people.

In this work, we adopt the agent-based simulator Epi-Fast (Bisset et al. 2009). The outputs are individual level infected cases with the infected days of a simulated season. They can be aggregated to any temporal and spatial scale, such as daily (weekly) state (county) level ILI incidence.

Vaccine intervention I_V (i.e. quantity of vaccines applied to the population and the timing of the application) is applied in our simulations. A distribution on the parameter space $\mathcal{P}(p_E, p_I, \tau, N_I, I_V)$ is estimated from CDC historical data, where N_I denotes the number of infections at the beginning of a flu season.

Disease Model Parameter Space Of the parameters in \mathcal{P} , (p_E, p_I) are known from literature (Marathe et al. 2011). We assume (τ, N_I, I_V) follows distributions that can be estimated from historical data. For clarity, we define an epidemiological week in a calendar year as **ew**, and a seasonal week in a flu season as **sw**. The historical time-series of CDC surveillance data used to construct parameter space is segmented into seasons by cutting at *ew40* of each year (i.e. *ew40* of a calendar year corresponds to *sw1* of a flu season).

Firstly, we collect observations of each parameter value by the following: (1) **Transmissibility** (τ): We compute season attack rate *ar* (i.e. fraction of population getting infected in the season) of each historical season for the target state and its neighbor states (i.e. geographically contiguous states). We calibrate a transmissibility value for each of *ar* as the solution to $\min_{\tau} |AR(EpiFast(\tau)) - ar|$, where $AR(\cdot)$ computes attack rate from the output of $EpiFast(\cdot)$. (2) **Initial Case Number** (N_I): We collect the ILI incidence of the first week of each season for the target state and its neighbors. (3) **Vaccine Intervention** (I_V): We collect 6 vaccination schedules of the past six influenza seasons in USA (CDC 2018b). Each schedule consists of timing and percent coverage of vaccine application throughout the season. We assume that the state level vaccine schedule is the same as the nationwide schedule.

Secondly, for $\tau(N_I)$, we fit the collected samples to several distributions including normal, exponential, gamma, and uniform. Then we run KS-test to choose a distribution with highest significance (refer to the supplementary file (Wang, Chen, and Marathe 2019) for more details). For I_V , we assume the 6 vaccination schedules follow a discrete uniform distribution. In this way, a region-specific parameter space \mathcal{P} is constructed.

Training Dataset from Simulations Let ℓ denote the length (number of weeks) of each flu season simulated by EpiFast. For each run of a simulation, a specific parameter setting is sampled from \mathcal{P} , and the simulator is called to generate state and county level weekly incidence, called *synthetic epicurve*. Week 1 in the synthetic epicurve corresponds to *sw1* of a flu season. Large volumes of high-resolution synthetic data are generated by repeating the sampling and simulating process.

DEFSI Neural Network Model In traditional time-series modeling problem, ILI incidences of the few previous weeks are used as the observations for the prediction of the current week. In DEFSI, we use two kinds of observations: (1) **Within-season observations**, denoted as $\mathbf{x1} = \langle y_{t-a}, \dots, y_{t-1} \rangle$, are ILI incidence from previous a weeks of the current season. They are used as the main observations to follow the weekly trend. (2) **Between-season observations**, denoted as $\mathbf{x2} = \langle y_{t-\ell*b}, \dots, y_{t-\ell*1} \rangle$, are ILI




Figure 2: DEFSI neural network architecture. This architecture consists of two branches. The left branch consists of stacked LSTM layers that encode within-season observations $\mathbf{x1}$, and the right branch is designed to be LSTM based layers that encodes between-season observations $\mathbf{x1}$. A merge layer is added to combine two branches.

incidences of the same *sw* from the past b seasons. They are used as the surrogate information to improve forecasting performance.

The Long Short Term Memory (LSTM) network (Hochreiter and Schmidhuber 1997) is adopted in our neural network architecture to capture the dynamic temporal behavior of the observations. An LSTM layer consists of a sequence of $a(b)$ cells, of which the current cell takes one ILI incidence from $\mathbf{x1}(\mathbf{x2})$ as well as the output and the cell state of its previous cell as inputs. We design a two-branch deep neural network model to combine within-season observations and between-season observations. As shown in Fig. 2, the left branch consists of stacked LSTM layers that encode within-season observations $\mathbf{x1} = \langle y_{t-a}, \dots, y_{t-1} \rangle$. The right branch is a single LSTM layer that encodes between-season observations $\mathbf{x2} = \langle y_{t-\ell*b}, \dots, y_{t-\ell*1} \rangle$. A merge layer is added before the dense layer to combine the outputs from two branches. In DEFSI model, the left and right LSTM-based branches take $\mathbf{x1}$ and $\mathbf{x2}$ as inputs respectively. The merge layer requires the outputs of two branches must be of the same dimension. The final output dimension is the same as \mathbf{z}_t .

We are interested in a predictor f , which predicts the current week's state level and county level incidences \mathbf{z}_t based on the previous a weeks of within-season state level ILI incidences $\mathbf{x1}$ and the previous b seasons of between-season state level ILI incidences $\mathbf{x2}$:

$$\hat{\mathbf{z}}_t = f([\mathbf{x1}, \mathbf{x2}]_t, \theta) \quad (1)$$

where θ denotes parameters of the predictor, $\hat{\mathbf{z}}_t$ denotes the prediction of \mathbf{z}_t . The loss function \mathcal{L} is defined as mean-square-error, with consistency constraint on outputs:

$$\min_{\theta} \mathcal{L}(\theta) = \sum_t \|\mathbf{z}_t - f([\mathbf{x1}, \mathbf{x2}]_t, \theta)\|^2, \quad (2)$$

with consistency constraint $\hat{y}_t = \sum_{C \in \mathcal{D}} \hat{y}_t^C$

Adam optimization algorithm is used to learn θ . An activity regularizer is added to $\hat{\mathbf{z}}_t$ for consistency constraint.

Variants of DEFSI Model The two-branch neural network architecture is flexible for multiple variants: (1) *DEFSI*: Two-branch neural network as shown in Fig. 2. (2) *DEFSI-L*: Only the left branch is used to take within-season observations. (3) *DEFSI-RDENSE*: Changing the LSTM layer of the right branch with Dense layers, which means that the model does not care about the temporal relationship between between-season data points. We will discuss the results of different variants in Experiments.

Multi-step Forecasting In practical situations, we are interested in making predictions for several weeks ahead. In DEFSI, the left branch of the model appends the most recent state level prediction to the input for predicting the target at the upcoming week, and the right branch uses the state level ILI incidences from the past seasons with sw equals to the upcoming week number.

Experiments

Datasets **CDC ILI incidence** : The CDC surveillance of weekly ILI incidence at state level from $ew40$, 2010 to $ew18$, 2018. **ILI Lab tested flu positive counts of New Jersey** : The state level and county level ILI Lab tested flu positive counts of season 2017-2018 in NJ. The data is available from $ew40$ to the next year’s $ew20$. **Google data** and **Weather data** are collected for comparison methods (see the supplementary file (Wang, Chen, and Marathe 2019) for details).

Comparison Methods Our method is compared with 5 state-of-the-art methods. They are *LSTM* (single layer LSTM) (Hochreiter and Schmidhuber 1997) and *AdapLSTM* (CDC + Weather data) (Venna et al. 2017) from artificial neural network methods; *ARIMA* (classic ARIMA) (Benjamin, Rigby, and Stasinopoulos 2003) and *ARGO* (CDC + Google data) (Yang, Santillana, and Kou 2015) from statistical methods; and *EpiFast* (Beckman et al. 2014) from agent-based causal models. *AdapLSTM*, *LSTM*, *ARGO* and *ARIMA* are used for state level forecasting. *EpiFast* is applied for both state level and county level forecasting.

Experiment Setup Our experiments are performed on two states: Virginia (VA) and New Jersey (NJ). For each state, we separate the time sequence into 8 flu seasons from 2010-2011 to 2017-2018 by cutting at $ew40$. We use the dataset of 2010-2011 to 2016-2017 seasons as training dataset; and use the dataset of 2017-2018 season for testing. At each time step in the testing season, each model makes predictions 5 weeks ahead, i.e. $horizon = \{1, 2, 3, 4, 5\}$. In DEFSI, the training dataset is used to estimate disease parameter space, while for other comparison methods, it is used for training forecasting model directly. The collected county level ILI data of NJ is only used for evaluating. More detailed settings (including estimated parameter space in DEFSI, parameter settings for comparison methods) are elaborated in the supplementary file (Wang, Chen, and Marathe 2019). The **metrics** used to evaluate the forecasting performance are: *root mean squared error (RMSE)*, *mean absolute percentage error (MAPE)*, *Pearson correlation (PCORR)*.

Performance of Flat-resolution Forecasting We forecast state level ILI incidence for VA, 2017-2018 and NJ, 2017-2018. Fig. 3 shows the forecasting performance on RMSE, MAPE, PCORR. (1) *Performance on RMSE (left column of Fig. 3)*: In VA, DEFSI, DEFSI-L, DEFSI-RDENSE, ARIMA, LSTM achieve similar performances that are better than EpiFast, AdapLSTM. In NJ, DEFSI and its variants consistently outperform others across the horizon, which on average have 15%-64% improvement on RMSE. (2) *Performance on MAPE (middle column of Fig. 3)*: In VA, ARGO performs the best among all methods except with horizon 1 where DEFSI-L performs the best. Meanwhile, DEFSI, DEFSI-L are comparable with ARGO. In NJ, DEFSI-RDENSE achieves the best performance closely followed by DEFSI. (3) *Performance on PCORR (right column of Fig. 3)*: In both VA and NJ, DEFSI-L outperforms others on Pearson correlation (i.e. around 0.96 with horizon 1) which on average have 20%-100% improvement compared with others, followed by DEFSI, DEFSI-RDENSE, and ARGO. Overall, DEFSI and its variants make comparable/better predictions than the comparison methods at state level.

Performance of High-resolution Forecasting The performance of county level forecasting is evaluated on NJ 2017-2018. The horizon is extended to 10 for better observations. In Fig. 4, we show county level ILI forecasting performance on each metric. The metric value of each node in the figure is the average value across 21 counties in NJ. Our method consistently outperforms the comparison method EpiFast on RMSE (about 53% reduction) and PCORR (about 60% increase). However, EpiFast performs better than our method on MAPE with horizon less than 4, while the error increases dramatically as the horizon increases. Overall, our method significantly outperforms the comparison method on county level forecasting.

Discussion In general, AdapLSTM and EpiFast do not perform very well in our experiment compared with other methods. For AdapLSTM, weather factors are considered for post adjustment of LSTM outputs. As stated in (Venna et al. 2017), the weather factors are estimated using time delays computed by apriori associations and selected by the largest confidence. However, in our experiment, they all show very low confidences (less than 0.3). This may cause arbitrary adjustment for predictions and consequently poor performance. For EpiFast, one possible reason is that we did not find a good estimate of the underlying disease model for a specific region and season due to the noisy CDC observations. As we discussed in the Introduction, if the observed data is too noisy then the learned underlying model tends to make predictions with large errors. Through the results, DEFSI enables high-resolution forecasting that outperforms EpiFast. Meanwhile, it achieves comparable/better performance than the comparison methods at state level forecasting. Our experiments demonstrate that DEFSI integrates the strengths of ANN methods and causal methods to improve epidemic forecasting.

The Significance of I_V in \mathcal{P} Any intervention actions taken during a flu season will obviously affect the disease

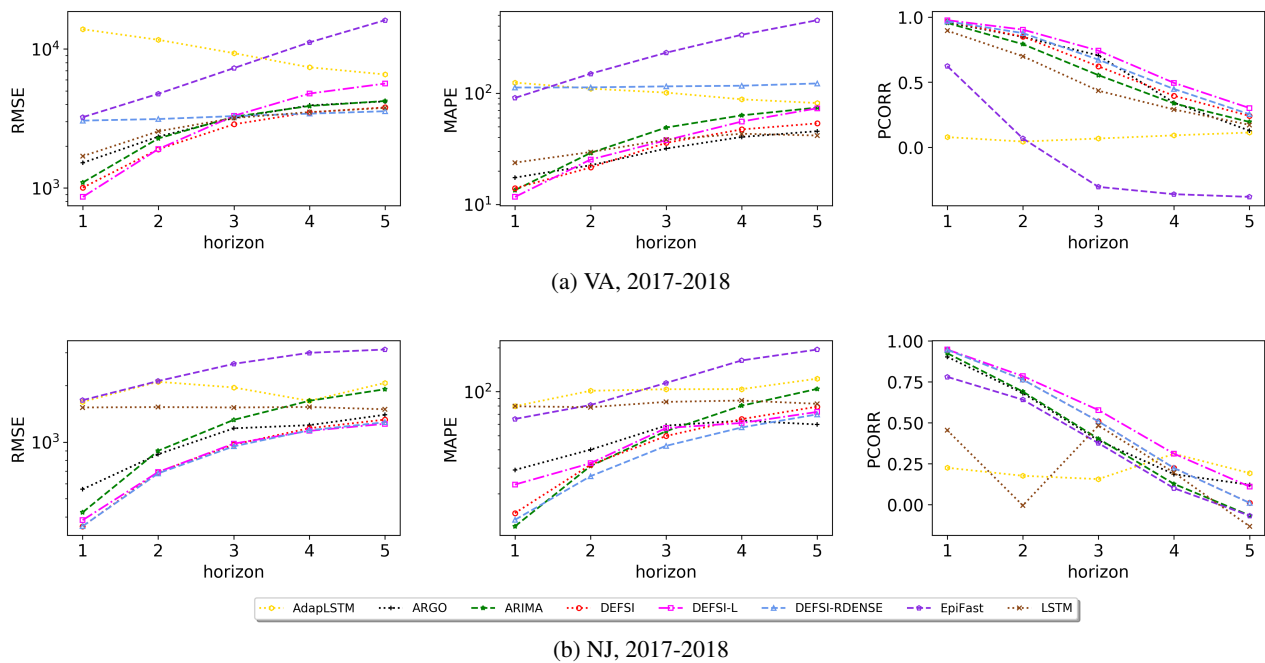


Figure 3: State level ILI incidence forecasting performance on RMSE, MAPE, PCORR. A log y-scale is used in RMSE and MAPE.




Figure 4: County level ILI incidence forecasting performance on RMSE, MAPE, PCORR for NJ, 2017-2018. A log y-scale is used in RMSE and MAPE.

spread. Our method described above constructs a parameter space $\mathcal{P}(p_E, p_I, \tau, N_I, I_V)$ that includes the vaccine intervention I_V . In this section, we investigate how I_V affects DEFSI model by generating two synthetic training datasets: (1) *vaccine-case*: simulations with I_V (the training dataset used in Experiments); (2) *base-case*: simulations that share the common settings of p_E, p_I, τ, N_I with vaccine-case except $I_V = \emptyset$.

We train DEFSI on vaccine-case and base-case respectively with the same settings described in Experiment Setup, denoted as *DEFSI-vac* and *DEFSI-base*. State level forecasting on VA 2017-2018 is evaluated. Fig. 5 shows the performance on RMSE, MAPE, and PCORR. DEFSI-vac consistently outperforms DEFSI-base on RMSE and PCORR, especially with small horizon less than 5. On MAPE, DEFSI-vac performs better than DEFSI-base with small horizon less than 5, while the situation inverses as the horizon increases.

Our experiments show the significance of realistic interventions in \mathcal{P} for good forecasting performance of our method. Our proposed framework is extensible for further available realistic interventions, such as school closure and quarantine, to further improve forecasting performance.

Conclusion

In this paper we propose DEFSI, a novel epidemic forecasting framework combining deep neural network methods with causal models. In DEFSI, a two-branch neural network model and its variants are designed to combine within-season and between-season observations. The model is trained on a region-specific synthetic dataset constructed at multiple spatially fine-grained scale. The trained model enables accurate high-resolution forecasting with flat-resolution observations as inputs. Extensive experiments on NJ and VA showed that DEFSI achieved comparable/better




Figure 5: Performance comparison between DEFSI models trained on base-case synthetic training dataset (DEFSI-base) and vaccine-case synthetic training dataset (DEFSI-vac). VA, 2017-2018.

performance than the state-of-the-art methods on state level forecasting and consistently better performance than others on county level forecasting. A direction for **future work** is to further investigate the use of synthetic data generated by social, epidemiological and behavioral theories in conjunction with observed data to improve epidemic forecasts.

Acknowledgment

The authors would like to thank members of the Network Simulation Science and Advanced Computing (NSSAC) Division. This work has been partially supported by Defense Threat Reduction Agency (DTRA) Grant HDTRA1-17-D-0023, National Institutes of Health (NIH) Grant IR01GM109718, NSF BIG DATA Grant IIS-1633028, NSF DIBBS Grant ACI-1443054.

References

Beckman, R. J.; Bisset, K. R.; Chen, J.; Lewis, B. L.; Marathe, M. V.; and Stretz, P. E. 2014. ISIS: a networked-epidemiology based pervasive web app for infectious disease pandemic planning and response. In *KDD*.

Benjamin, M. A.; Rigby, R. A.; and Stasinopoulos, D. M. 2003. Generalized autoregressive moving average models. *Journal of the American Statistical Association* 98(461): 214–223.

Bisset, K. R.; Chen, J.; Feng, X.; Kumar, V. A.; and Marathe, M. V. 2009. EpiFast: A fast algorithm for large scale realistic epidemic simulations on distributed memory systems. In *Proceedings of the 23rd ICS*, 430–439. ACM.

CDC. 2018a. Fluview interactive. <https://www.cdc.gov/flu/weekly/fluviewinteractive.htm>. Accessed April 20, 2018.

CDC. 2018b. Historical seasonal influenza vaccine schedule. <https://www.cdc.gov/flu/professionals/vaccination/vaccinesupply.htm>. Accessed November 1, 2017.

Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.

Karpatne, A.; Atluri, G.; Faghmous, J. H.; Steinbach, M.; Banerjee, A.; Ganguly, A.; Shekhar, S.; Samatova, N.; and Kumar, V. 2017. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering* 29(10):2318–2331.

Kuznetsov, Y. A., and Piccardi, C. 1994. Bifurcation analysis of periodic SEIR and SIR epidemic models. *Journal of Mathematical Biology* 32(2):109–121.

Marathe, A.; Lewis, B.; Chen, J.; and Eubank, S. 2011. Sensitivity of household transmission to household contact structure and size. *PLoS ONE* 6.

Nsoesie, E. O.; Beckman, R. J.; Shashaani, S.; Nagaraj, K. S.; and Marathe, M. V. 2013. A simulation optimization approach to epidemic forecasting. *PLoS ONE* 8(6):1–10.

Venna, S. R.; Tavanaei, A.; Gottumukkala, R. N.; Raghavan, V. V.; Maida, A.; and Nichols, S. 2017. A novel data-driven model for real-time influenza forecasting. *bioRxiv*.

Volkova, S.; Ayton, E.; Porterfield, K.; and Corley, C. D. 2017. Forecasting influenza-like illness dynamics for military populations using neural networks and social media. *PLoS ONE* 12(12):1–22.

Wang, L.; Chen, J.; and Marathe, M. 2019. DEFSI: Deep Learning Based Epidemic Forecasting with Synthetic Information (Supplement). https://github.com/christa60/defsi/blob/master/supp_IAAI.DEFSI.pdf.

Wu, Y.; Yang, Y.; Nishiura, H.; and Saitoh, M. 2018. Deep learning for epidemiological predictions. In *SIGIR*.

Yang, S.; Santillana, M.; Brownstein, J. S.; Gray, J.; Richardson, S.; and Kou, S. C. 2017. Using electronic health records and internet search information for accurate influenza forecasting. *BMC Infectious Diseases* 17(1):332.

Yang, W.; Karspeck, A.; and Shaman, J. 2014. Comparison of filtering methods for the modeling and retrospective forecasting of influenza epidemics. *PLoS Computational Biology* 10(4):1–15.

Yang, S.; Santillana, M.; and Kou, S. C. 2015. Accurate estimation of influenza epidemics using Google search data via ARGO. *PNAS* 112(47):14473–14478.

Zhao, L.; Chen, J.; Chen, F.; Wang, W.; Lu, C. T.; and Ramakrishnan, N. 2015. SimNest: Social Media Nested Epidemic Simulation via Online Semi-supervised Deep Learning. *Proceedings of IEEE ICDM* 2015:639–648.