

Chickenpox Cases in Hungary: a Benchmark Dataset for Spatiotemporal Signal Processing with Graph Neural Networks

Benedek Rozemberczki
The University of Edinburgh
United Kingdom
benedek.rozemberczki@ed.ac.uk

Paul Scherer
University of Cambridge
United Kingdom
pms69@cam.ac.uk

Olivér Kiss
Central European University
Hungary
kiss_oliver@phd.ceu.edu

Rik Sarkar
The University of Edinburgh
United Kingdom
rsarkar@inf.ed.ac.uk

Tamás Ferenci
Óbuda University; Corvinus
University of Budapest
Hungary
ferenci.tamas@nik.uni-obuda.hu

ABSTRACT

Recurrent graph convolutional neural networks are highly effective machine learning techniques for spatiotemporal signal processing. Newly proposed graph neural network architectures are repetitively evaluated on standard tasks such as traffic or weather forecasting. In this paper, we propose the *Chickenpox Cases in Hungary* dataset as a new dataset for comparing graph neural network architectures. Our time series analysis and forecasting experiments demonstrate that the *Chickenpox Cases in Hungary* dataset is adequate for comparing the predictive performance and forecasting capabilities of novel recurrent graph neural network architectures.

ACM Reference Format:

Benedek Rozemberczki, Paul Scherer, Olivér Kiss, Rik Sarkar, and Tamás Ferenci. 2021. Chickenpox Cases in Hungary: a Benchmark Dataset for Spatiotemporal Signal Processing with Graph Neural Networks. In *WWW'21: Graph Learning Benchmarks Workshop, 2021, Ljubljana, SL*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Forecasting future edge and vertex attributes using spatial structure and historical values of the node and link attributes is a fundamental research problem for spatiotemporal machine learning. Recurrent graph neural networks can elegantly solve such spatiotemporal signal tasks with high predictive performance by training a graph convolutional neural network which is integrated or stacked with a recurrent neural network layer [15, 19]. These machine learning techniques have favorable practical characteristics such as online training and models that are transferable across graphs [2, 5, 10]. Hence, finding real world problems on which the forecasting performance of these architectures can be tested is crucial for fostering temporal graph representation learning research. However, recurrent graph neural networks are often iteratively evaluated and

compared using the same overutilized datasets from a restricted number of application domains such as urban traffic and weather forecasting [12, 24]. These challenges related to the public availability of suitable and relevant spatiotemporal benchmark datasets are the main motivations of the present work.

Present work. In the pursuit of advancing temporal graph neural network research we publicly release the *Chickenpox Cases in Hungary* dataset: a multivariate time series of weekly reported chickenpox cases in Hungarian counties. By utilizing this novel epidemiological dataset, the forecasting capabilities of newly proposed graph neural network models can be quantified. The intrinsic statistical characteristics of the dataset such as seasonality, spatial and temporal autocorrelation, zero inflation, heteroskedasticity and structural changes make the forecasting a challenging machine learning task. Our contribution also opens up venues to assess the predictive performance of existing spatiotemporal models.

Main contributions. The major results and contributions presented in our work can be summed up as follows:

- (1) We release *Chickenpox Cases in Hungary*, a novel spatiotemporal dataset which can be used to benchmark the forecasting performance of graph neural network architectures.
- (2) We conduct a descriptive analysis of the time series and discuss the particular spatiotemporal modeling challenges that the dataset poses.
- (3) We assess the performance of existing recurrent graph neural network architectures on multiple forecasting horizons.

The remainder of this paper has the following structure. In Section 2 we overview the related literature about chickenpox and parametric spatiotemporal statistical models. In Section 3 we carry out a descriptive analysis of the dataset, while Section 4 discusses the potential modeling challenges. We present predictive performance benchmarks on the dataset in Section 5 and the paper concludes with Section 6. The spatial adjacency matrix and the county level time series are available at https://github.com/benedekrozemberczki/spatiotemporal_datasets.

2 RELATED WORK

In this section we give a brief overview of related work about the epidemiology and characteristics of chickenpox and the design of recurrent graph neural network architectures.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW'21, April 19–23, 2021, Ljubljana, SL

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

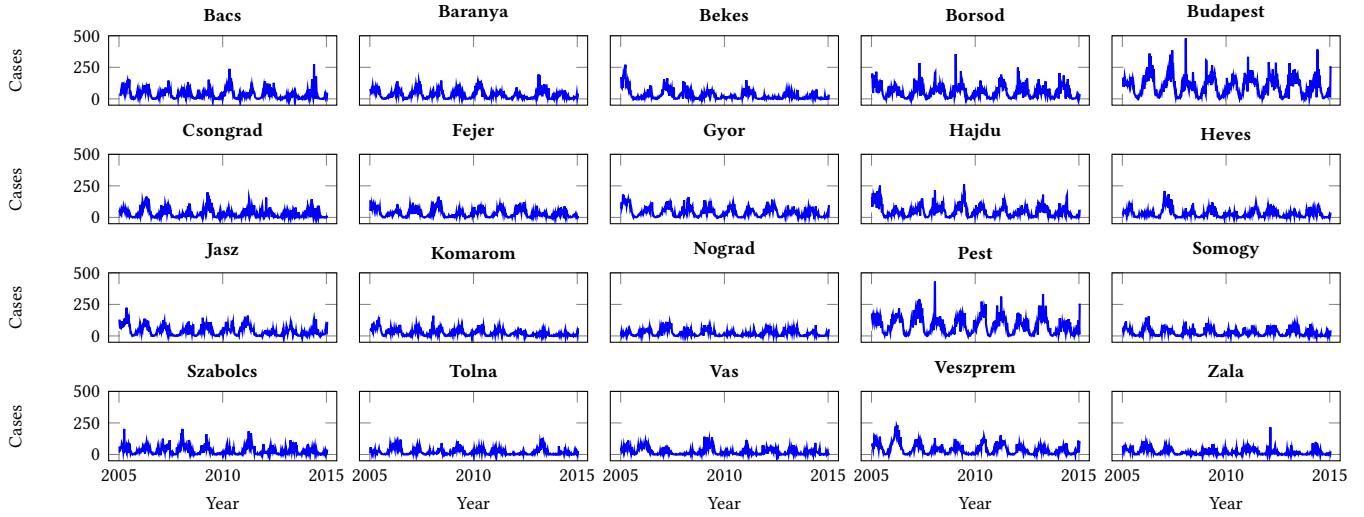


Figure 1: The weekly number of chickenpox cases in Hungarian counties and the capital between 2005 and 2015.




Figure 2: One year running mean and standard deviation of weekly chickenpox cases in selected Hungarian counties.

2.1 Chickenpox

Chickenpox or varicella is a highly contagious airborne disease caused by the varicella zoster virus (VZV) [1]. By the age of 20, more than 90 percent of the population is exposed to the VZV in developed countries [4]. While chickenpox might produce common early symptoms such as headache or nausea, its onset is characterized by the rapid appearance of an easily distinguishable skin rash [13]. Although vaccines against VZV infection are available [20], only a handful of countries have national immunization programme [6]. In Hungary there is no mandatory vaccination against chickenpox but vaccines are available and are routinely recommended to parents. Physicians have to report each case to the local centre of epidemiology which are then aggregated and publicly presented

weekly for each of the 20 counties of Hungary, resulting in an ideal data collection environment from the modeling perspective [8].

2.2 Spatiotemporal neural models

Spatiotemporal neural models are a family of parametric statistical models which can handle data that has distinct time and spatial dimensions e.g. traffic measurements, regional epidemiological reporting or weather. The specific *recurrent graph neural network models* compared in our work fuse ideas from the design of graph convolutional neural network layers [2, 10, 11, 17, 23] and recurrent neural networks [3, 7]. These models operate on temporal sequences of spatial data; at each time step a graph neural network layer convolves the input features or hidden states of the recurrent unit. Recurrent and graph convolutional layers are trained jointly on a downstream task and the design of these architectures requires the choice of a graph neural network and a recurrent unit. Popular choices for graph neural networks are spectral graph convolutions [10] and graph attention networks [23] while the most frequently augmented recurrent neural networks include long short-term memory cells [7] and gated recurrent units [3].

3 CHARACTERISTICS OF THE DATASET

Our main contribution is the release of the *Chickenpox Cases in Hungary* dataset which consists of county level time series and a spatial graph which describes the spatial connectivity of the counties. The county level time series describe the weekly number of chickenpox cases reported by general practitioners in Hungary.

We manually collected the time series by collating the reported case counts from the digital version of the *Hungarian Epidemiological Info*¹, a weekly bulletin of morbidity and mortality of infectious diseases in Hungary. Our data collection covered the weeks between the January of 2005 and January of 2015 and the resulting time series has more than 500 entries for all of the counties without

¹<http://www.oek.hu/>

any missingness. The underlying spatial graph has 20 vertices (19 counties and the capital Budapest) and there are 61 edges between the nodes. We plotted the county level reported case count time series on Figure 1.

Main characteristics of the time series. Looking at the time series on Figure 1 we can make multiple important observations. These are the following:

- The number of reported cases is population dependent; spatial units with more inhabitants such as the capital Budapest report more cases on average.
- The time series all exhibit strong seasonality which can be a result of weather conditions or the periodicity of the school year.
- A large number of counties report no new cases in the summer months – these county level time series are zero inflated.

3.1 Structural changes and trends

The county level time series are noisy and exhibit a strong yearly seasonality, due to this, we cannot theorize whether there are structural changes without correcting the seasonality. We calculated the 52 weeks running average and standard deviation of the weekly case counts for selected counties and the capital and plotted the resulting times series on Figure 2.




Figure 3: The truncated random walk weighted spatial autocorrelation of the on level and first-order differenced chickenpox case count time series.

Main findings. We observe multiple statistical phenomena which pose modelling challenges. These can be summarized as:

- The yearly rolling average of the county level mean case count is not constant. This can be an artifact of population shift or the result of intrinsic epidemiological dynamics.
- Standard deviations of the time series change with time which means that the models have to take into consideration the time dependent variation in the outcome variable.

3.2 Measuring spatial autocorrelation

We quantify the spatial autocorrelation of chickenpox cases by using a truncated random walk weighted variant of Moran’s I index. Given an unweighted and undirected graph $G = (V, E)$ let us denote the adjacency matrix by \mathbf{A} and the diagonal degree matrix as \mathbf{D} . The row normalized adjacency matrix is defined as $\hat{\mathbf{A}} = \mathbf{D}^{-1}\mathbf{A}$ and the transition probability matrix of an r -length truncated random walk equals to $\hat{\mathbf{A}}^r$. Correspondingly the truncated random walk weighted spatial autocorrelation index [14] of the node feature vector $\mathbf{x} \in \mathbb{R}^{|V|}$ at scale r is defined by Equation (1). Here $\bar{\mathbf{x}}$ is the average of the generic vertex feature and $u, v \in V$ are vertices.

$$I = \frac{|V|}{\sum_{v \in V} \sum_{u \in V} \hat{\mathbf{A}}_{u,v}^r} \frac{\sum_{v \in V} \sum_{u \in V} \hat{\mathbf{A}}_{u,v}^r (\mathbf{x}_v - \bar{\mathbf{x}})(\mathbf{x}_u - \bar{\mathbf{x}})}{\sum_{v \in V} (\mathbf{x}_v - \bar{\mathbf{x}})^2} \quad (1)$$

We plotted on Figure 3 the spatial autocorrelation index using the first 5 proximity scales for the on level case count and first-order differenced case number time series. All of the county level time series were centralized to be 0 mean and standardized before the autocorrelation index computation.

Main findings. The most important empirical regularities that we can observe are the following:

- Both the county level case count and differenced case count time series exhibit spatial autocorrelation through the years.
- The spatial autocorrelation is present at multiple scales but it decreases with increasing the distance being considered.
- The strength of spatial autocorrelation is not constant – there are visible trends in the spatial autocorrelation time series.

4 THE MODELING CHALLENGES

The *Chickenpox Cases in Hungary* dataset poses a number of machine learning modeling challenges for researchers. Based on our descriptive analysis of the time series these challenges can be briefly summarized as:

- **Temporal autocorrelation.** The weekly number of new chickenpox cases is correlated with the case numbers from earlier weeks.
- **Spatial autocorrelation.** The number of newly infected children and the difference in the number of new cases are correlated across neighboring spatial units.
- **Heteroskedasticity.** The standard deviation of the county level time series is not constant over time.
- **Seasonality.** The county level count of chickenpox cases exhibit strong yearly seasonality. This can be an artifact of weather conditions or caused by the periodicity of the school year.
- **Multiple scales.** The Hungarian county system consist of spatial units which have a heterogeneous size. Budapest, the largest one, has nearly 10 times more inhabitants than Nograd which is the least populated one.
- **Count data.** The target variables describe the weekly count of chickenpox cases. The design of the graph neural network has to take this fact into account: particularly when it comes to the choice of loss function and the activation functions in the output layer.

- **Zero inflation.** Certain smaller counties report no cases during the weeks of the summer in a randomly dispersed manner, which causes challenges for traditional count data modeling.
- **Structural changes and random events.** The one and half decade long time horizon of the dataset gives plenty of space for population shift and years when the winter surge in chickenpox cases did not happen in certain counties.

Designing novel graph neural network architectures that can appropriately model a dataset with these statistical characteristics is a challenging task.

5 NEURAL BENCHMARKS

We tested the predictive performance of recurrent graph neural networks on county level chickenpox time series forecasting. Using the *PyTorch Geometric Temporal* [5, 16, 18] implementation of the models we trained on the standardized chickenpox time series and predicted it for a fixed number of weeks ahead. The input graph describes the undirected direct adjacency relations of the counties.

All recurrent models used 8 temporal lags as input features and had 32 dimensional graph convolutional filters. The output of the convolutional layer was fed to a feedforward output layer. Each model was trained for 200 epochs with the Adam optimizer [9] and we used a learning rate of 10^{-2} . In Table 4 we present average mean squared error values for various forecasting horizons calculated from 10 experimental runs.

Figure 4: The average test mean squared error with standard deviations obtained over 10/20/40 weeks long forecasting horizons calculated from a 10 experimental runs. Bold numbers denote the best performing models.

	10 weeks	20 weeks	40 weeks
GConvLSTM [19]	0.741 ± 0.005	0.403 ± 0.003	1.221 ± 0.010
GConvGRU [19]	0.757 ± 0.001	0.407 ± 0.001	1.117 ± 0.002
Evolve GCN-O [15]	0.775 ± 0.007	0.419 ± 0.004	1.120 ± 0.003
Evolve GCN-H [15]	0.766 ± 0.009	0.413 ± 0.009	1.115 ± 0.013
DynGRAE [21, 22]	0.706 ± 0.004	0.382 ± 0.002	1.112 ± 0.010
STGCN [24]	0.763 ± 0.008	0.405 ± 0.007	1.118 ± 0.005
DCRNN [12]	0.753 ± 0.003	0.395 ± 0.001	1.119 ± 0.002

Main findings. Based on the forecasting performance of graph neural networks we can make the following observations:

- The DynGRAE [21, 22] architecture works best at forecasting horizon and the advantage is significant at $\alpha = 5\%$.
- There are considerable performance differences between models; see for example Evolve GCN-O and DynGRAE.
- When the forecasting horizon is increased the predictive performance of some models is worse than random.

6 CONCLUSIONS

We introduced *Chickenpox Cases in Hungary*, a longitudinal dataset for benchmarking the predictive performance of spatiotemporal graph neural network architectures. Our exploratory analysis highlighted the unique statistical characteristics of the dataset which make predicting the weekly number of cases a challenging task.

We evaluated the forecasting capabilities of the state-of-the-art recurrent graph neural networks. Our findings demonstrate that the current design of graph neural networks is moderately well suited for solving this task.

REFERENCES

- [1] A M Arvin. 1996. Varicella-Zoster Virus. *Clinical Microbiology Reviews* 9, 3 (1996), 361–381. <https://doi.org/10.1128/CMR.9.3.361>
- [2] Aleksandar Bojchevski, Johannes Klicpera, Bryan Perozzi, Amol Kapoor, Martin Blais, Benedek Rózemberczki, Michal Lukasik, and Stephan Günnemann. 2020. Scaling Graph Neural Networks with Approximate PageRank. In *KDD*.
- [3] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations Using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1724–1734.
- [4] Peter W. Choo, James G. Donahue, JoAnn E. Manson, and Richard Platt. 1995. The Epidemiology of Varicella and Its Complications. *The Journal of Infectious Diseases* 172, 3 (09 1995), 706–712.
- [5] Mathias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [6] Andrew Platt and Judy Breuer. 2012. Varicella Vaccines. *British Medical Bulletin* 103, 1 (08 2012), 115–127.
- [7] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [8] J. Karsai, Rita Csoma-Kovács, Ágnes Dánielisz, Z. Molnár, János Dudás, Teodóra Borsos, and G. Röst. 2020. Modeling the Transmission Dynamics of Varicella in Hungary. *Journal of Mathematics in Industry* 10 (2020), 1–16.
- [9] Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations*.
- [10] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.
- [11] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. 2019. Combining Neural Networks with Personalized PageRank for Classification on Graphs. In *ICLR*.
- [12] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations*.
- [13] Monica L. McCrary, Jessica Severson, and Stephen K. Tyring. 1999. Varicella Zoster Virus. *Journal of the American Academy of Dermatology* 41, 1 (1999), 1–16.
- [14] Patrick AP Moran. 1950. Notes on Continuous Stochastic Phenomena. *Biometrika* 37, 1/2 (1950), 17–23.
- [15] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao B Schardl, and Charles E Leiserson. 2020. EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs. In *AAAI*. 5363–5370.
- [16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*. 8024–8035.
- [17] Benedek Rozemberczki, Peter Englert, Amol Kapoor, Martin Blais, and Bryan Perozzi. 2021. Pathfinder Discovery Networks for Neural Message Passing. In *WWW*.
- [18] Benedek Rozemberczki and Paul Scherer. 2020. PyTorch Geometric Temporal. https://github.com/benedekrozemberczki/pytorch_geometric_temporal.
- [19] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. 2018. Structured Sequence Modeling with Graph Convolutional Recurrent Networks. In *International Conference on Neural Information Processing*. Springer, 362–373.
- [20] Jane Seward. 2002. Varicella Disease After Introduction of Varicella Vaccine in the United States, 1995–2000. *JAMA* 287 (02 2002), 606.
- [21] Aynaz Taheri and Tanya Berger-Wolf. 2019. Predictive Temporal Embedding of Dynamic Graphs. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 57–64.
- [22] Aynaz Taheri, Kevin Gimpel, and Tanya Berger-Wolf. 2019. Learning to Represent the Evolution of Dynamic Graphs with Recurrent Models. In *Companion Proceedings of The 2019 World Wide Web Conference (WWW '19)*. 301–307.
- [23] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.
- [24] Bing Yu, Haoqiang Yin, and Zhanxing Zhu. 2018. Spatio-Temporal Graph Convolutional Networks: a Deep Learning Framework for Traffic Forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 3634–3640.