

Neural Relational Autoregression for High-Resolution COVID-19 Forecasting

Matthew Le, Mark Ibrahim, Levent Sagun, Timothee Lacroix, Maximilian Nickel
Facebook AI Research

Corresponding author: Maximilian Nickel
maxn@fb.com

Forecasting COVID-19 poses unique challenges due to the novelty of the disease, its unknown characteristics, and substantial but varying interventions to reduce its spread. To improve the quality and robustness of forecasts, we propose a new method which aims to disentangle region-specific factors – such as demographics, enacted policies, and mobility – from disease-inherent factors that influence its spread. For this purpose, we combine recurrent neural networks with a vector autoregressive model and train the joint model with a specific regularization scheme that increases the coupling between regions. This approach is akin to using Granger causality as a relational inductive bias and allows us to train high-resolution models by borrowing statistical strength across regions. In our experiments, we observe that our method achieves strong performance in predicting the spread of COVID-19 when compared to state-of-the-art forecasts.

Introduction

Modeling the spread of COVID-19 at a high spatial and temporal resolution (i.e., confirmed cases at county or admin-3 level) has become an important task in the public health response to the disease. For instance, accurate county-level forecasts are not only central to monitor the state of the pandemic but are also important to efficiently allocate scarce resources such as ventilators, personal protective equipment, and ICU beds; and to make progress towards efficient early detection systems.

However, forecasting COVID-19 poses unique challenges – in particular when considering confirmed cases at high spatial resolution. Although there has been considerable progress towards understanding the spread of the disease, there still exists only limited data and knowledge about important factors that influence its spread. This is only exacerbated by the naturally larger noise-levels in county-level data as compared to more highly aggregated state-level data. Due to the global nature of COVID-19, the available data is also distributed among regions with very different properties, many of which may affect its spread. This includes, for instance, demographics and population densities, enacted policies, adherence to those policies, mobility patterns, and geographic features such as temperature. In addition, testing and reporting can vary considerably across regions and time. All these factors lead to considerable variability (see also Figure 2) and uncertainty in the data and makes reliable forecasting at high spatial resolution difficult.

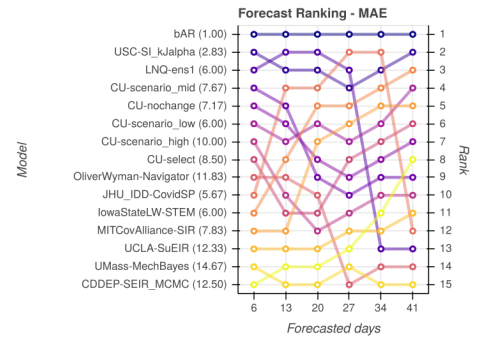


Figure 1: Ranking of county-level forecasts by average MAE over various forecast horizons. The proposed neural relational autoregressive model (β -AR) shows strong performance over all horizons when compared to state-of-the-art forecasts. Mean rank over all horizons in parentheses.

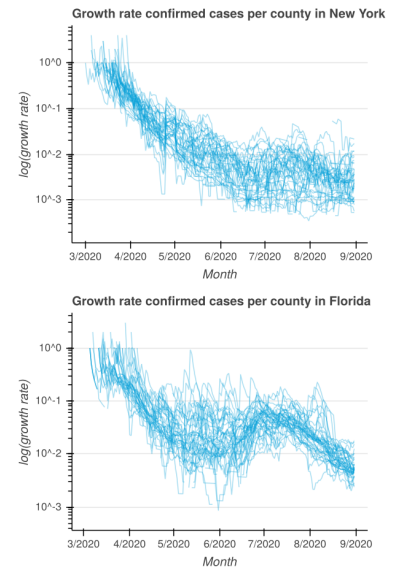


Figure 2: Variability of growth in confirmed cases per region over time. Each line represents one county in the state. For y_t denoting the number of cases at time t , growth rate is computed as $(y_{t+1} - y_t) / y_t$.

To alleviate these issues, we propose a new method for predicting the spread of COVID-19 by combining recurrent neural networks with a vector autoregressive model and a specific regularization relational scheme. Our approach is motivated by two main aspects: First, we seek to develop an end-to-end differentiable model, as this allows us to make efficient use of the limited available data while also enabling us to estimate parameters of powerful models that can capture the large variability of cases across locations and time. However, while such flexible models are needed to account for possible influencing factors there is little data to estimate them reliably and without overfitting. For this reason, we seek, second, to disentangle region- and time-specific factors from disease-inherent factors that influence its spread. This allows us to borrow statistical strength between regions by coupling their predictions – based on the assumption that once a model has correctly accounted for region-specific dynamics, information about the spread of the disease in region j can also help to improve predictions for a related region i . This approach is akin to using *Granger causality* as an inductive bias to improve forecast quality and robustness.

Compared to existing state-of-the-art forecasting models, our method takes a highly data-driven approach with fewer modeling assumptions as, for instance, in very detailed compartmental models. As such, we see our approach as complementary to existing models which provides strong forecasting performance at the cost of reduced interpretability.

Neural Relational Autoregression

We consider the forecasting of m time series that are different realizations of the same underlying disease process. Let $\mathcal{Y} = \{(y_i^1, \dots, y_i^T)\}_{i=1}^m$ denote the observed case counts where i indexes locations and where T denotes the maximum observation time. Furthermore, let $\mathcal{Y}(\tau) = \{(y_i^t : t \leq \tau)\}_{i=1}^m$ denote the set of all observed case counts up to time $\tau \leq T$. We then model the case counts as random variables

$$Y_i^{t+1} | \mathcal{Y}(t) \sim f(\lambda_i^t)$$

where λ_i^t denotes the *force of infection*¹ at time t in location i and where $f(x)$ denotes a probability distribution with parameter x (e.g., a Poisson or Negative Binomial distribution).

¹ Given y_i^t infected individuals, the force of infection (or hazard) models the probability that a susceptible individual at time t will become infected by time $t + 1$

Due to the different interventions during the course of the epidemic, we regard \mathcal{Y} as a time-varying process that is influenced by external factors such as policies, mobility, etc. For this reason, we decompose λ_i^t into a time-specific component β_i^t and a time-independent component λ_i such that

$$\lambda_i^t = \beta_i^t \lambda_i \quad \text{where} \quad \beta_i^t \in [0, 1], \lambda_i > 0$$

Hence, β_i^t can be understood as a dampening factor of the underlying force of infection which models the effect of interventions and depends on time and location. While some influencing factors for the evolution of β_i^t might be known (e.g., mobility, population density, etc.), we assume that the full set of influencing factors is unknown and will regard β_i^t as a latent variable.

Using this decomposition, we then model the time-independent force of infection as a autoregressive model of order p ,² i.e.,

$$\text{AR}(p) : \lambda_i = \sum_{\ell=0}^{p-1} w^\ell y_i^{t-\ell} \quad (1)$$

where $\{w^\ell > 0\}_{\ell=0}^{p-1}$ are the parameters of the model which are shared across locations i . For the time-dependent dampening β_i^t we employ recurrent neural networks (RNNs; Elman, 1990; Hochreiter and Schmidhuber, 1997; Cho et al., 2014) such that

$$\text{RNN} : \beta_i^t = f_\theta(\{x_i^k\}_{k=0}^t) \quad (2)$$

where θ are the parameters of the network which are again shared across locations and where $\{x_i^k\}_{k=0}^t$ denote observed input features to the RNN (e.g., mobility in location i at time k). Although an RNN as in Equation (2) has enough capacity to model the evolution of β_i^t , the limited data about the spread of COVID-19 makes it challenging to estimate its parameters without overfitting. We seek therefore an inductive bias which allows us to estimate β_i^t from few observations.

Relational Inductive Bias

Since all regions are affected by the same underlying process, we assume that we can borrow statistical strength between regions and use information about the spread in region i to help predicting the spread in region j – once we have accounted for time- and location-dependent dynamics. A good model of β_i^t should therefore help to improve the predictions of y_i^{t+1}/β_i^t from cases in other regions y_j^t . We interpret this as an inductive bias akin to Granger causality (Granger, 1969)³ and extend Equation (1) to a *vector autoregressive* model where it is known that Granger causality is directly linked to its coefficients. In particular, let

$$\text{VAR}(p) : \lambda_i = \sum_{\ell=0}^{p-1} \sum_{j=1}^m w_{ij}^\ell y_j^{t-\ell} \quad (3)$$

be a vector autoregressive model of order p . A *time series* y_j is then *Granger-causing* y_i if and only if $w_{ij} \neq 0$ (Seth, 2007). For causal discovery, coefficients w_{ij} are therefore often ℓ_1 -regularized. Here, we take the opposite approach and seek solutions in which many time-series can be considered

² AR models where

$$Y_i^{t+1} | \mathcal{Y}(t) \sim \text{Poisson}(\lambda_i^t)$$

can be interpreted as approximations of Reed-Frost chain binomial SIR models (Abbey, 1952). For a detailed discussion see (Bauer and Wakefield, 2018).

³ Granger causality is defined as follows: Let $X^t = \{X_i\}_{i=1}^t$, $Y^t = \{Y_i\}_{i=1}^t$, $Z^t = \{Z_i\}_{i=1}^t$ denote stochastic processes and let L denote a loss function. Furthermore, let

$$\mathcal{R}(Y^{t+1}|Y^t, Z^t) = \mathbb{E}(L(Y_{t+1}, f(Y^t, Z^t)))$$

denote the expected loss (risk) of a predictor f . We then say X *Granger-causes* Y if its inclusion in the predictor significantly improves the forecast, i.e., if

$$\mathcal{R}(Y^{t+1}|Y^t, X^t, Z^t) \ll \mathcal{R}(Y^{t+1}|Y^t, Z^t)$$

Granger-causal related. However, we do not force all time series to be related since this is likely an unrealistic constraint. Instead, we assume $\forall i \neq j : w_{ij}$ are drawn from a logit-normal distribution (Atchison and Shen, 1980), what allows us to specify a prior on the proportion of related and unrelated time series.

In particular, let $\phi(\cdot)$ denote the logistic function, let $\forall i \neq j : w_{ij} = \phi(\alpha_{ij})$, and let $\mathcal{N}(\mu, \sigma^2)$ denote the Normal distribution with mean μ and variance σ^2 . Putting everything together, we then model the full *time-varying* force of infection as

$$\beta\text{-AR}(p) : \quad \lambda_i^{t+1} = \beta_i^t \sum_{\ell=0}^{p-1} \sum_{j=1}^m w_{ij}^\ell y_j^{t-\ell} \quad (4)$$

$$\alpha_{ij} \sim \mathcal{N}(\mu, \sigma^2) \quad \forall i \neq j$$

Hence, the β -AR model consists of a standard AR component ($w_{ii} > 0$) and a relational component ($w_{ij} \in [0, 1]$) which aims to couple the different regions. The number of non-zero entries in the ‘‘adjacency matrix’’ w_{ij} can then be controlled through the logit-normal prior.

Accounting for Overdispersion

Count data such as confirmed cases is naturally modeled using Poisson distributions. However, COVID-19 case counts exhibit substantial overdispersion, i.e., the variance of the observed counts can significantly exceed their mean (e.g., see fig. 4). For this reason, we will model case counts with Negative Binomial distributions what allows us to account for varying degrees of overdispersion (Lloyd-Smith, 2007). Specifically, we set

$$y_i^{t+1} \sim \text{NB}(\lambda_i^t, \nu_i)$$

where λ_i^t and ν_i are mean and dispersion parameter of the distribution and λ_i^t is modeled using the β -AR model of eq. (4). The likelihood function in eq. (5) is then of the form

$$p_\theta(y) = \frac{\Gamma(y + \nu)}{y! \Gamma(\nu)} \left(\frac{\mu}{\mu + \nu} \right)^y \left(1 + \frac{\mu}{\nu} \right)^{-\nu} \quad \mu > 0, \nu > 0$$

Parameter Estimation and Implementation Details

To estimate the parameters of the model, we regularize the model log-likelihood such that w_{ij} is drawn from a logit-normal distribution with location μ and scale σ . Let θ denote the model parameters (i.e., α_{ij} as well as parameters of the RNN). and let $p_\theta(y)$ denote the likelihood function of the

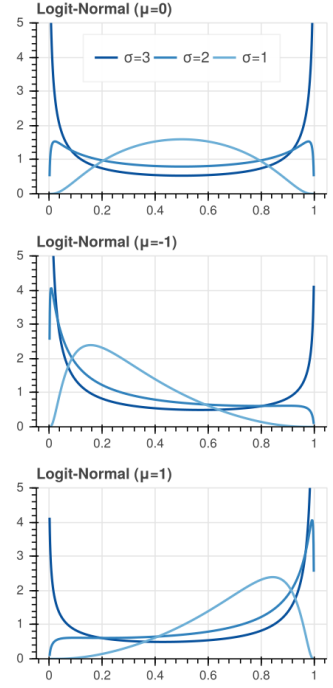


Figure 3: The Logit-Normal distribution is a probability distribution of a random variable whose logit has a normal distribution, i.e., $\phi(\mathcal{N}(\mu, \sigma))$.

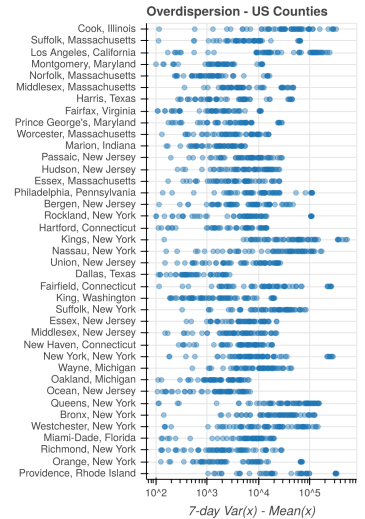


Figure 4: Overdispersion of daily case counts in US states and counties with most number of cases.

Group	Model	
Center for Disease Dynamics, Economics & Policy	<i>CDDP-SEIR_MCMC</i>	(Klein et al., 2020)
Columbia University	<i>CU-*</i>	(Pei and Shaman, 2020)
COVID Alliance at MIT	<i>MITCovAlliance-SIR</i>	(Baek et al., 2020)
Iowa State University Lily Wang Research Group	<i>IowaStateLW-STEM</i>	(Wang et al., 2020)
Johns Hopkins ID Dynamics COVID-19 Working Group	<i>JHU-IDD_CovidSP</i>	(Lemaitre et al., 2020)
LockNQuay	<i>LNQ-ens1</i>	(Wolfinger and Lander, 2020)
Oliver Wyman	<i>Pandemic Navigator</i>	(Koyluoglu and Milliken, 2020)
UCLA Statistical Machine Learning Lab	<i>UCLA-SuEIR</i>	(Zou et al., 2020)
University of Southern California Data Science Lab	<i>USC-SI_kJalpha</i>	(Srivastava et al., 2020)
University of Massachussets Amherst	<i>UMass-MechBayes</i>	(Sheldon et al., 2020)

Table 1: Forecasting models for confirmed cases on county-level.

β -AR model. Furthermore, let q denote the prior normal distribution for α_{ij} . We then maximize the regularized log-likelihood

$$\max_{\theta} \sum_y \log p_{\theta}(y) + \sum_{ij} \log q(\alpha_{ij} | \mu, \sigma). \quad (5)$$

We regard $\mu, \sigma > 0$ as hyperparameters which allow us to control the ratio of related and unrelated time series.

Since Equation (5) is end-to-end differentiable we can jointly estimate the parameters of the entire model using gradient-based optimization. We compute gradients via automatic differentiation using the PyTorch framework (Paszke et al., 2019). To maximize Equation (5) we then use the stochastic optimization method AdamW (Loshchilov and Hutter, 2018) where we decouple the updates of the normally distributed parameters α_{ij} from the adaptive updates of the remaining parameters.

Results

In the following, we evaluate the forecast quality of our method compared to multiple state-of-the-art forecasts for confirmed cases on county-level. All comparison forecasts are collected from the COVID-19 Forecast Hub⁴ as submitted by the respective teams. The COVID-19 Forecast Hub features county-level forecasts from July 5th onwards and we selected those models for which at least 10 forecasts were available since then. The full list of comparison forecasts is shown in Table 1.

⁴<https://github.com/reichlab/covid19-forecast-hub>

Forecast setup and model selection To compute forecasts for the different dates in the test set, we use the following fully automated model selection scheme: For each forecast date d , we perform cross-validation by holding out additional 21 days of validation data and train the model on the remaining data. We then select the best hyperparameters as measured by RMSE on the

Dataset	Source	Resolution
Confirmed Cases	The New York Times (2020) <i>Confirmed cases based on reports from state & local health agencies</i>	County
Symptom Survey	Facebook Data for Good (2020b) <i>Prevalence of COVID-like symptoms from self-reported surveys</i>	County, State
Movement Range Maps	Facebook Data for Good (2020a) <i>Mobility metrics related to physical distancing measures (change in movement and staying put)</i>	County, State
Community Mobility	Google (2020) <i>Movement trends across different categories of places (retail and recreation, groceries and pharmacies, etc.)</i>	County, State
Doctor visits	CMU COVIDcast (Farrow et al., 2015) <i>Percentage of COVID-related doctor’s visits in a given location</i>	County, State
Testing	The COVID Tracking Project (2020) <i>Total number of COVID PCR tests per state</i>	State
Weather	NOAA GHCN (Menne et al., 2012) <i>Average, minimum, maximum temperature & rainfall per county</i>	County

Table 2: Data sources for β -AR.

validation set and retrain the whole model with those hyperparameters on the combined training and validation set to compute the final forecast. When computing the forecasts, we hold all additional input data (e.g., symptom survey, mobility, weather, etc.) constant after the last observed day d .⁵ For all training details of the model, please see the supplementary material.

Input data As input features for β -AR, we use multiple data sources as listed in Table 2. Confirmed cases enter the model only in the autoregressive part. All other covariates enter the model only as input features for the time-varying β -part. For cases and weather data, we use the preprocessed data from the Google COVID-19 Open Data repository (Wahlteiz et al., 2020). All datasets are publicly available, de-identified, and aggregated at county- or state-level.

Forecast evaluation Figure 5 shows the forecast quality as measured by MAE for multiple forecast horizons.⁶ It can be seen that the proposed β -AR models shows a consistently strong performance and is for all forecasting dates and horizons either the best model or among the best. Figure 1, which shows the ranking of all models by the average MAE for each forecast horizon, further illustrates this property. It can be seen that β -AR model is consistently ranked first over all horizons. Furthermore, other models show much larger variability in their performance.

To also evaluate the performance of our model on days prior to July 5th, we

⁵ This setting places natural limits on the duration of the forecasting horizon. We reserve the joint forecasting of cases and covariates – what could extend the horizon – for future work.

⁶ MAE numbers are computed in accordance with <https://github.com/youyanggu/covid19-forecast-hub-evaluation>

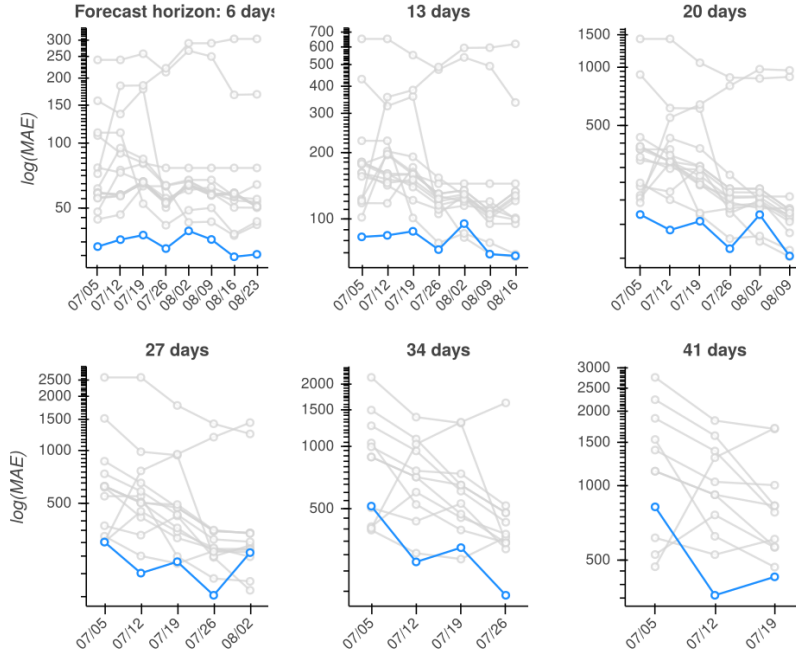


Figure 5: Comparison of β -AR model (blue) to 15 county-level models from COVID-19 forecast hub (gray). Forecast quality is measured in MAE (log-scale) where the absolute errors are averaged over all counties. For similar analysis using RMSE please see the supplementary material.

compare to forecasts of Google Cloud AI (Arik et al., 2020) and Columbia University (Pei and Shaman, 2020) which provide county-level forecasts of confirmed cases from May 11th to June 27th. Figure 6 shows the average MAE over all counties for 7 and 14 day forecasts for these models.⁷ It can be seen that the β -AR model shows again consistently strong performance on these earlier days and is typically ranked first for both 7 and 14 day forecasts.

Ablations In addition to comparisons to state-of-the-art county-level forecasts, we also evaluate the contributions of different aspects of our model. First, we test the effect of the relational autoregressive part. For this purpose, we trained additional models where we disabled the relational part (by setting $\forall i \neq j : w_{ij} = 0$) and compared their forecasts to the full model of Equation (4). To measure the relative improvement of the full model over the non-relational model, we compute then the relative error of both models, e.g.,

$$\text{Relative Mean Absolute Error} = \frac{\text{MAE}_{\text{full}}}{\text{MAE}_{\text{non-relational}}}$$

It can be seen from Figure 7 that full model offers substantial improvements over the non-relational model as the relative forecast quality grows exponentially with the forecasting horizon. While the non-relational model can offer acceptable forecast for horizons of 1-2 days, it quickly deteriorates with larger horizons. This shows the importance of the relational component for disentangling the different growth factors and learning high quality models.

⁷ For this comparison, average MAE is computed as described in (Arik et al., 2020)




Figure 6: Comparisons of β -AR model to forecasts from Google Cloud AI and Columbia for 7 and 14 day horizons and earlier forecast dates. Forecast quality is measured in MAE where the absolute errors are averaged over all counties. For a similar analysis using RMSE please see the supplementary material.

In addition to the non-relational component, we also evaluated the contributions of the logit-normal regularization method. For this purpose, we trained a model where we explicitly set the regularization parameter $\sigma = 0$. We then compare the forecast quality to the standard model where the regularization parameter has been selected via cross-validation. Figure 8 shows the results of the comparison. It can be seen that the logit-normal regularization can be very beneficial to improve forecast quality. While the differences to the standard model are much smaller than for the non-relational model, the addition of the regularization term can lead to substantial improvements, especially for horizons of 13 days and longer.

Finally, we also evaluated the contributions of the Negative Binomial distribution compared to a standard Poisson distribution for modeling confirmed cases. Similar to the logit-normalization method, we trained an additional model with Poisson likelihood and compared the forecast quality to the standard model. It can be seen from Figure 9, that Negative Binomial likelihood significantly improves the quality of the model over all forecast horizons. This is likely due to the fact that the Negative Binomial can better model the noise in the observed data, while the stricter Poisson likelihood causes the (recurrent) model to overfit to these variations.

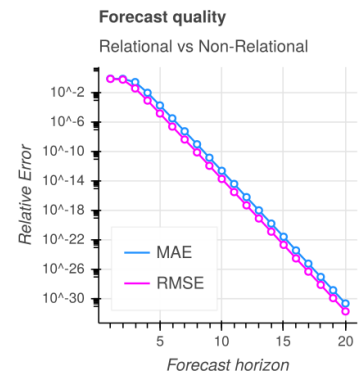


Figure 7: Relative Error (MAE and RMSE) of the fully relational β -AR model compared to a non-relational variant.

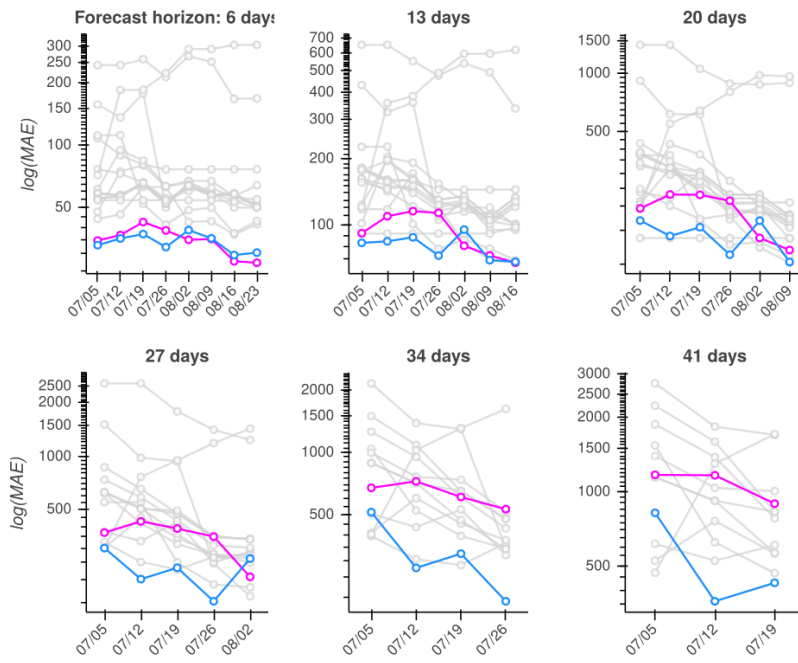


Figure 8: Comparison of β -AR model with (blue) and without (magenta) Granger regularization. Forecast quality is measured in MAE.

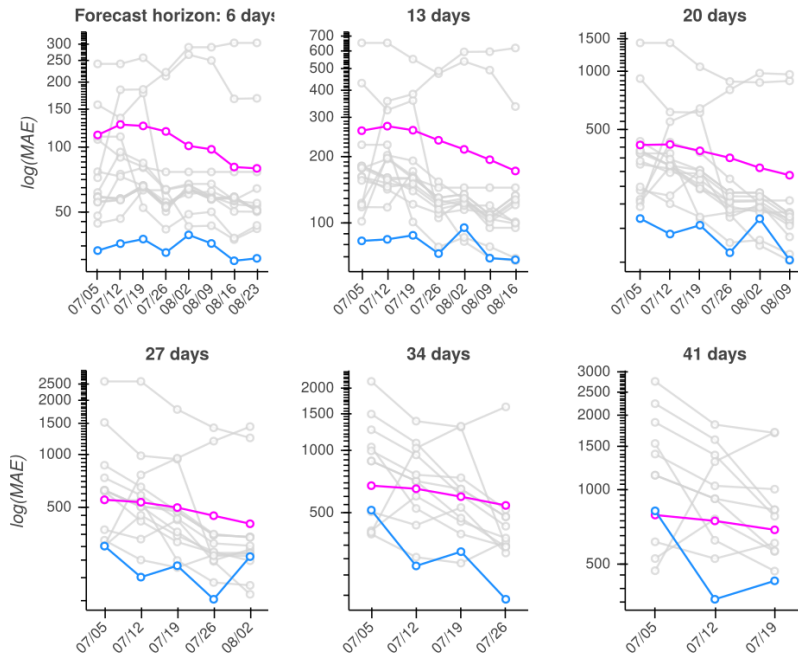


Figure 9: Comparison of β -AR model with Negative Binomial (blue) and Poisson (magenta) likelihood. Forecast quality is measured in MAE.

Related Work

We build on prior work that has proposed to use autoregressive models for spatially and temporally aggregated disease surveillance data of endemic-epidemic processes (Held et al., 2005; Meyer and Held, 2014; Meyer and Held, 2016). Such autoregressive models are, for instance, used to monitor infectious diseases by public health agencies like the Robert Koch Institute (Salmon et al., 2016).

Moreover, the negative binomial distribution has become a popular way to model infectious diseases, largely to its ability to model count data with varying degrees of overdispersion (Lloyd-Smith, 2007). Autoregressive models in combination with negative binomial distributions have, for instance, been used by Bauer and Wakefield (2018), Wakefield et al. (2019), and Held et al. (2005) to model infectious disease count data.

Valdés-Sosa et al. (2005) proposed a combination of VAR(1) models and ℓ_1 regularization to for the discovery of Granger-causal relations to understand brain connectivity. Haufe et al. (2010) proposed an improved estimator which can be applied for VAR models of order $p > 1$.

Conclusion

To improve the quality and robustness of forecasts, we propose a new method which aims to disentangle region-specific factors – such as demographics, enacted policies, and mobility – from disease-inherent factors that influence its spread. For this purpose, we combine recurrent neural networks with a vector autoregressive model and train the joint model with a specific regularization scheme that increases the coupling between regions. In our experiments, we observe that our method achieves strong performance in predicting the spread of COVID-19 when compared to state-of-the-art forecasts. Through ablations of the model, we show that the relational approach in general, the added logit-normal regularization, and the negative binomial likelihood are all important factors that contribute to the forecast quality. Compared to existing state-of-the-art forecasting models, our method takes a highly data-driven approach with fewer modeling assumptions as, for instance, in very detailed and mechanistic compartmental models. As such, we see our approach as complementary to existing models with focus on strong forecasting performance at the cost of reduced interpretability.

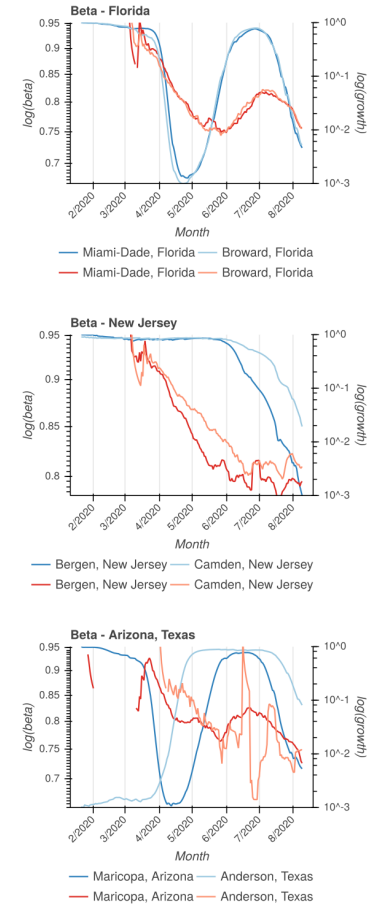


Figure 10: Evolution of β over time

References

- Abbey, H. (1952). “An examination of the Reed-Frost theory of epidemics.” In: *Human biology* 24 3, pp. 201–33.
- Arik, Sercan O. et al. (2020). *Interpretable Sequence Learning for COVID-19 Forecasting*. arXiv: 2008.00646 [cs.LG].
- Atchison, J and Sheng M Shen (1980). “Logistic-normal distributions: Some properties and uses”. In: *Biometrika* 67.2, pp. 261–272.
- Baek, Jackie et al. (2020). *The Limits to Learning an SIR Process: Granular Forecasting for Covid-19*. arXiv: 2006.06373 [stat.ME].
- Bauer, Cici and Jon Wakefield (2018). “Stratified space–time infectious disease modelling, with an application to hand, foot and mouth disease in China”. In: *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 67.5, pp. 1379–1398.
- Cho, Kyunghyun et al. (2014). “On the properties of neural machine translation: Encoder-decoder approaches”. In: *Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8), 2014*.
- Elman, Jeffrey L (1990). “Finding structure in time”. In: *Cognitive science* 14.2, pp. 179–211.
- Facebook Data for Good (2020a). *Movement Range Maps*. URL: <https://dataforgood.fb.com/tools/movement-range-maps/> (visited on 09/05/2020).
- (2020b). *Symptom Survey*. URL: <https://dataforgood.fb.com/tools/symptommap/> (visited on 09/05/2020).
- Farrow, David C. et al. (2015). *Delphi Epidata API*. URL: <https://github.com/cmu-delphi/delphi-epidata> (visited on 09/05/2020).
- Google (2020). *Community Mobility Reports*. URL: <https://www.google.com/covid19/mobility/> (visited on 09/03/2020).
- Granger, Clive WJ (1969). “Investigating causal relations by econometric models and cross-spectral methods”. In: *Econometrica: journal of the Econometric Society*, pp. 424–438.
- Haufe, Stefan et al. (2010). “Sparse causal discovery in multivariate time series”. In: *Causality: Objectives and Assessment*, pp. 97–106.
- Held, Leonhard, Michael Höhle, and Mathias Hofmann (2005). “A statistical framework for the analysis of multivariate infectious disease surveillance counts”. In: *Statistical modelling* 5.3, pp. 187–199.
- Hochreiter, Sepp and Jürgen Schmidhuber (1997). “Long short-term memory”. In: *Neural computation* 9.8, pp. 1735–1780.
- Klein, Eili, Gary Lin, and Yupeng Yang (2020). *CDDEP SEIR Markov Chain Monte Carlo*. URL: https://github.com/reichlab/covid19-forecast-hub/blob/master/data-processed/CDDEP-SEIR_MCMC (visited on 09/05/2020).

- Koyluoglu, Ugur and John Milliken (2020). *Oliver Wyman Pandemic Navigator*. URL: <https://github.com/reichlab/covid19-forecast-hub/blob/master/data-processed/OliverWyman-Navigator/> (visited on 09/05/2020).
- Lemaitre, Joseph Chadi et al. (2020). “A scenario modeling pipeline for COVID-19 emergency planning”. In: *medRxiv*. URL: <https://www.medrxiv.org/content/early/2020/06/12/2020.06.11.20127894>.
- Lloyd-Smith, James O. (Feb. 2007). “Maximum Likelihood Estimation of the Negative Binomial Dispersion Parameter for Highly Overdispersed Data, with Applications to Infectious Diseases”. In: *PLOS ONE* 2.2, pp. 1–8. URL: <https://doi.org/10.1371/journal.pone.0000180>.
- Loshchilov, Ilya and Frank Hutter (2018). “Decoupled Weight Decay Regularization”. In: *International Conference on Learning Representations*.
- Menne, Matthew J et al. (2012). “An overview of the global historical climatology network-daily database”. In: *Journal of Atmospheric and Oceanic Technology* 29.7, pp. 897–910.
- Meyer, Sebastian and Leonhard Held (Sept. 2014). “Power-law models for infectious disease spread”. In: *Ann. Appl. Stat.* 8.3, pp. 1612–1639. URL: <https://doi.org/10.1214/14-AOAS743>.
- (Dec. 2016). “Incorporating social contact data in spatio-temporal models for infectious disease spread”. In: *Biostatistics* 18.2, pp. 338–351. ISSN: 1465-4644. URL: <https://doi.org/10.1093/biostatistics/kxw051>.
- Paszke, Adam et al. (2019). “Pytorch: An imperative style, high-performance deep learning library”. In: *Advances in neural information processing systems*, pp. 8026–8037.
- Pei, Sen and Jeffrey Shaman (2020). “Initial Simulation of SARS-CoV2 Spread and Intervention Effects in the Continental US”. In: *medRxiv*. URL: <https://www.medrxiv.org/content/early/2020/03/27/2020.03.21.20040303>.
- Salmon, Maëlle, Dirk Schumacher, and Michael Höhle (2016). “Monitoring Count Time Series in R: Aberration Detection in Public Health Surveillance”. In: *Journal of Statistical Software, Articles* 70.10, pp. 1–35. ISSN: 1548-7660. URL: <https://www.jstatsoft.org/v070/i10>.
- Seth, Anil (2007). “Granger causality”. In: *Scholarpedia* 2.7. revision #127333, p. 1667.
- Sheldon, Dan, Graham Gibson, and Nick Reich (2020). *UMass MechBayes*. URL: <https://github.com/reichlab/covid19-forecast-hub/blob/master/data-processed/UMass-MechBayes/> (visited on 09/05/2020).
- Srivastava, Ajitesh, Tianjian Xu, and Viktor K. Prasanna (2020). *Fast and Accurate Forecasting of COVID-19 Deaths Using the SIKJa Model*. arXiv: 2007.05180 [q-bio.PE].

- The COVID Tracking Project (2020). *State testing data (CC-BY 4.0 license)*. URL: <https://covidtracking.com/> (visited on 09/05/2020).
- The New York Times (2020). *Coronavirus (Covid-19) Data in the United States*. URL: <https://github.com/nytimes/covid-19-data> (visited on 09/05/2020).
- Valdés-Sosa, Pedro A et al. (2005). “Estimating brain functional connectivity with sparse multivariate autoregression”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 360.1457, pp. 969–981.
- Wahlteiz, Oscar et al. (2020). “COVID-19 Open-Data: curating a fine-grained, global-scale data repository for SARS-CoV-2”. In: *Work in progress*. URL: <https://github.com/GoogleCloudPlatform/covid-19-open-data>.
- Wakefield, Jon, Tracy Qi Dong, and Vladimir N Minin (2019). “Spatio-temporal analysis of surveillance data”. In: *Handbook of Infectious Disease Data Analysis*, pp. 455–476.
- Wang, Li et al. (2020). *Spatiotemporal Dynamics, Nowcasting and Forecasting of COVID-19 in the United States*. arXiv: 2004.14103 [stat.AP].
- Wolfinger, Russ and David Lander (2020). *LNQ-ens1*. URL: <https://github.com/reichlab/covid19-forecast-hub/blob/master/data-processed/LNQ-ens1/> (visited on 09/05/2020).
- Zou, Difan et al. (2020). “Epidemic Model Guided Machine Learning for COVID-19 Forecasts in the United States”. In: *medRxiv*. URL: <https://www.medrxiv.org/content/early/2020/05/25/2020.05.24.20111989>.