

**Document Version**

Final published version

**Licence**

CC BY

**Citation (APA)**

Schmidt, A., Zunker, H., Heinlein, A., & Kühn, M. J. (2026). Graph neural network surrogates to leverage mechanistic expert knowledge towards reliable and immediate pandemic response. *Scientific Reports*, 16(1), Article 6361. <https://doi.org/10.1038/s41598-026-39431-5>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



## OPEN Graph neural network surrogates to leverage mechanistic expert knowledge towards reliable and immediate pandemic response

Agatha Schmidt<sup>1,4</sup>, Henrik Zunker<sup>1,4</sup>, Alexander Heinlein<sup>2</sup> & Martin J. Kühn<sup>1,3</sup>✉

During the COVID-19 crisis, mechanistic models have guided evidence-based decision making. However, time-critical decisions in a dynamical environment limit the time available to gather supporting evidence. We address this bottleneck by developing a graph neural network (GNN) surrogate of an age-structured and spatially resolved mechanistic metapopulation simulation model. This combined approach complements classical modeling approaches which are mostly mechanistic and purely data-driven machine learning approaches which are often black box. Our design of experiments spans outbreak and persistent-threat regimes, up to three contact change points, and age-structured contact matrices on a spatial graph with 400 nodes representing German counties. We benchmark multiple GNN layers and identify an ARMAConv-based architecture that offers a strong accuracy–runtime trade-off. Across horizons of 30–90 day simulation and prediction, allowing up to three contact change points, the surrogate model attains 10–27% mean absolute percentage error (MAPE) while delivering (near) constant runtime with respect to the forecast horizon. Our approach accelerates evaluation by up to 28 670 times compared with the mechanistic model, allowing responsive decision support in time-critical scenarios and straightforward web integration. These results show how GNN surrogates can translate complex metapopulation models into immediate, reliable tools for pandemic response.

**Keywords** Surrogate modeling, Graph neural networks, Metapopulation, Epidemic simulation, High-performance computing, Spatio-temporal forecasting

Mathematical models can help public health experts and decision makers to explore potential future outcomes of ongoing disease dynamics. In particular during the recent COVID-19 crisis, mathematical modeling has been one of the principal forms to provide evidence on the effectiveness of public health and social interventions<sup>1</sup>. With an estimated number of more than 600 000 undiscovered viruses in mammal and avian hosts that are capable of infecting humans<sup>2</sup>, pandemic preparedness is a necessity.

Due to local outbreaks and human contact patterns, infectious disease dynamics are often heterogeneous on a spatial or demographic scale. For efficient mitigation of pandemics and accurate predictions, local transmission dynamics should be considered and corresponding data should be integrated in a mathematical model. Over the last years, a large number of authors have made contributions to predict the development of SARS-CoV-2. Simple ordinary differential equation (ODE)-based SIR (Susceptible-Infected-Removed)-type models<sup>3,4</sup> have been used for their efficiency in time-critical moments. For using more realistic, nonexponential disease stage transitions, linear chain trick (see, e.g.,<sup>5,6</sup>), generalized linear chain trick<sup>7</sup>, or flexible integro-differential (see, e.g.,<sup>8,9</sup>) equation-based models are available.

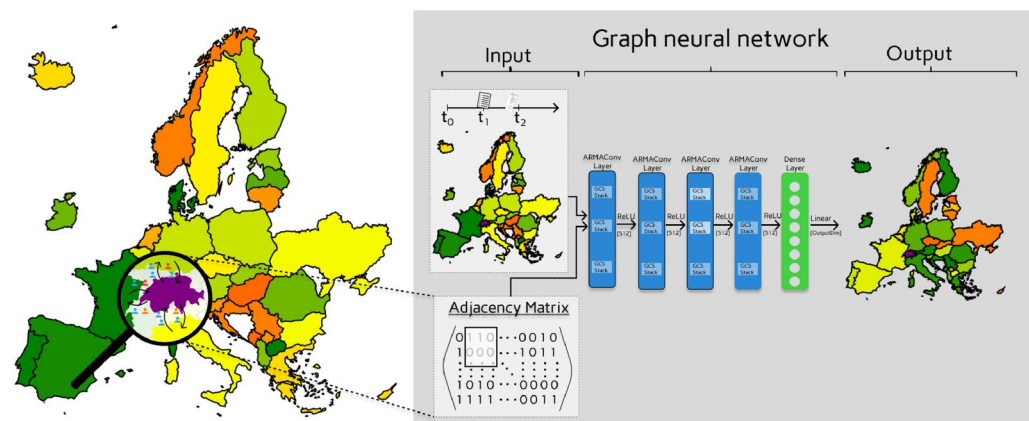
Through metapopulation modeling, various authors addressed the spatially heterogeneous spread of SARS-CoV-2 and modeled entire countries with high efficiency and small execution times<sup>10–15</sup>. For instance, Pei et al.<sup>10</sup> considered effects of intervention timing on COVID-19 spread in the United States (US), Chen et al.<sup>11</sup> considered compliance and containment of COVID-19 across towns in the US, Levin et al.<sup>12</sup> considered effects of short-term travel on COVID-19 spread in Minnesota, while Liu et al.<sup>13</sup> considered COVID-19 cases in

<sup>1</sup>Department of High-Performance Computing, Institute of Software Technology, German Aerospace Center, Cologne, Germany. <sup>2</sup>Delft University of Technology, Delft Institute of Applied Mathematics, Delft, The Netherlands. <sup>3</sup>Life and Medical Sciences Institute and Bonn Center for Mathematical Life Sciences, University of Bonn, Bonn, Germany. <sup>4</sup>Agatha Schmidt and Henrik Zunker contributed equally to this work. ✉email: martin.kuehn@dlr.de

Singapore and Zunker et al.<sup>14,15</sup> considered travel-time aware and behavior-driven metapopulation models with the 400 German counties. To include individual transmission and superspreading events in an intuitive way, different contact network-, agent-, or individual-based models have been developed; see, e.g.,<sup>16–19</sup>. For instance, Bershteyn et al. provided a very general pre-COVID-19 agent-based model (ABM) with different transmission pathways (vector, environmental, or airborne), Kerr et al.<sup>17</sup> developed a powerful ABM for COVID-19, Müller et al.<sup>18</sup> extended the MATSim ABM for the COVID-19 use case, and Kerkmann et al.<sup>19</sup> developed an ABM for respiratory diseases and direct or airborne transmission pathways. However, while more complex models allow more detailed research questions to be answered, they generally come at higher computational costs. Recent works with very similar approaches<sup>20–22</sup> combine agent- and equation-based metapopulation approaches to substantially reduce the prediction time of complex ABMs, yet these models remain the most computationally demanding among the involved approaches. Furthermore, many models using artificial intelligence (AI) have been proposed to fight the SARS-CoV-2 pandemic; see, e.g.,<sup>23</sup>. While these models often deliver accurate results in short execution times, a review on more than 650 AI-based models identified most of them to be black box models<sup>23</sup>. Unfortunately, black box models often lack interpretability or explainability and do not contribute to a better understanding of underlying processes—important when deciding and communicating interventions to the public. Nevertheless, time is a critical factor in pandemics, and evaluations of mitigation and reaction strategies have to be conducted in narrow time windows. In addition, interventions ranging from distancing recommendations to face masks (in particular locations), capacity-restrictions, and closures of venues or even vaccinations (if available) yield a tremendously large combinatorial space from which decision makers need to select. While very simplified models are executed within the range of (milli)seconds, complex models, even if highly optimized, might need seconds to minutes to be run. With automated pipelines preventing manual delay and high-performance computing<sup>24</sup>, thousands or millions of outcomes can be computed on a supercomputing facility to inform decision makers on short term. However, in a dynamically changing environment, decision makers might need expected outcomes for individual combinations of nonpharmaceutical interventions (NPIs) on-the-fly. Training a GNN on the HPC simulations eventually allows the user-driven exploration of any strictness of NPIs, based on the interpolation capabilities of the GNN. Through a low-barrier access for public health experts, e.g., a web application, this approach also drastically increases the impact of the developed (mechanistic) models. Furthermore, novel NPIs, public recommendations, or observed behavioral changes will be represented in the mechanistic models with new contact patterns and learned regularly with new simulation outputs by novel versions of the GNN.

In this paper, we combine an already validated age-structured and spatially resolved expert mechanistic model<sup>25,26</sup> with AI-based approaches to provide an on-the-fly execution of spatially resolved infectious disease models. A schematic presentation of the approach is shown in Fig. 1. In<sup>27</sup>, Angione et al. compared advantages and disadvantages of several established machine learning approaches such as random forests, decision trees, and neural networks as surrogates for agent-based models. Recently, Robertson et al.<sup>28</sup> employed random forests to calibrate a city-scaled ABM. Here, however, we build upon an already validated age-structured and spatially resolved metapopulation model<sup>25</sup>, whose simulations are used to train a machine learning model with spatial resolution.

The model proposed by Kühn et al.<sup>25</sup> interprets the metapopulation model as a graph, with nodes representing subpopulations and edges capturing the exchange between them. With the expert model executed millions of times on a supercomputing facility overnight, we obtain a diverse training set from which a neural network surrogate can learn and generalize. As the output of the model from<sup>25</sup> is already in the form of a graph, a graph-



**Fig. 1.** Visual representation of the prediction phase for a GNN surrogate model. The adjacency matrix results from the links between the considered regions, here depicted high-level as European countries. The magnifying glass highlights mobility connections between Switzerland and other countries. Colors on the national level geography represent disease outcomes such as symptomatic cases incidence, hospitalization, or deaths. Visualization of a GNN with four ARMAConv layers with three Graph Convolutional Skip (GCS) stacks and one dense output layer. Visualization based on Eurostat shape files<sup>29</sup> © EuroGeographics for the administrative boundaries.

based machine learning model is a natural choice and we therefore propose to employ graph neural networks (GNNs). GNNs are powerful tools for spatially distributed data, which have been introduced in Gori et al.<sup>30</sup> and further popularized in Scarselli et al.<sup>31</sup>. While we propose GNNs as surrogates for expert mechanistic models of infectious disease dynamics, GNNs have already excelled on data in different applications such as weather predictions<sup>32</sup>, influenza-like diseases<sup>33</sup>, or even SARS-CoV-2<sup>34–37</sup>. Fritz et al.<sup>37</sup> used GNNs in a purely data-driven fashion for spatio-temporal disease models, Gao et al.<sup>35</sup> used patients' claims data from different counties in a graph-based approach with a transmission dynamics loss, and Panagopoulos et al.<sup>36</sup> used transfer learning together with reported case data-driven GNNs.

In summary, we present a structured design of experiments covering outbreak and persistent-threat regimes using age-structured contacts and multiple contact-change points. We introduce a compact spatio-temporal input encoding for GNN surrogates and benchmark layers and architectures to identify an ARMAConv model with a strong accuracy–runtime trade-off that, eventually, demonstrates large speed-ups which enable interactive scenario exploration.

## Material and methods

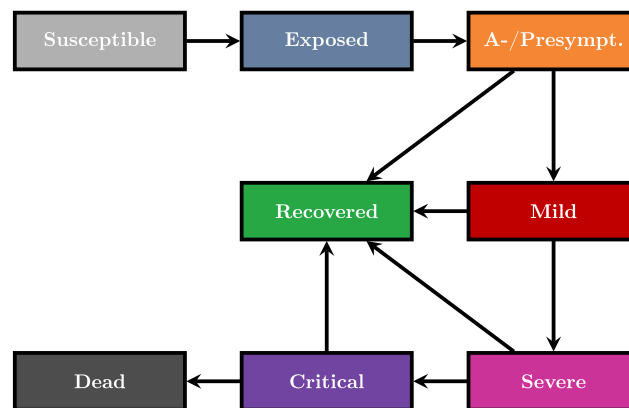
In this section, we detail the mechanistic simulator that provides the training targets, introduce the evaluation metric, describe the generation of input data, and document the neural network architectures used in the surrogate models.

### Mechanistic model

A requirement for obtaining a surrogate model that is reliable in practice is that the model, which is replaced by the surrogate, has been validated beforehand. For infectious disease dynamics, the underlying mechanistic expert model that generates the reference training has to integrate the most important disease states and properties of disease dynamics for the surrogate model to learn from. While deaths are usually monitored, several indicators such as the number of reported cases, the number of hospitalizations or intensive care unit (ICU) admissions have been considered in Germany to implement or lift nonpharmaceutical interventions against SARS-CoV-2<sup>38</sup>. Consequently, we have chosen an epidemic model based on ordinary differential equations (ODEs) that has already been validated in<sup>25,26</sup> and which categorizes subpopulations based on the infection states *susceptible*, *exposed (not infectious)*, *a- or presymptomatic and infectious*, *mildly symptomatic and infectious*, *severely infected (hospitalized)*, *critically infected (ICU)*, *dead*, and *recovered*; see Fig. 2. Aside the mentioned indicators, the model particularly considers pre- and asymptomatic transmission as studies on SARS-CoV-2 observed that viral load is peaking close to symptom onset as observed by Walsh et al. and Puhach et al.<sup>39,40</sup>. The model is efficiently implemented in C++ in the EMilio<sup>41</sup> framework, which means that it yields very fast model evaluations on a single core. We furthermore divided the population into  $G = 6$  age groups, that is, from 0 to 4, 5 to 14, 15 to 34, 35 to 59, 60 to 79, and 80+ years. While also simpler models of *SIR*-type structure and without age-structure could be learnt by a surrogate model, the resulting surrogate model could then not be evaluated for critical numbers on hospitalizations, intensive care, or deaths and without learning about most contributing age groups. In addition, the *SIR*-type structure would miss the well observed disease evolution of COVID-19 with exposed, non-infectious and nonsymptomatic, infectious states.

Transmission is modeled via the susceptible population  $S_i(t)$  becoming exposed ( $E_i(t)$ ), through infection of either infectious nonsymptomatic ( $I_{NS,i}(t)$ ) or symptomatic individuals ( $I_{Sy,i}(t)$ ),  $i \in \{1, \dots, 6\}$ . This can be described by the ODE

$$\frac{dS_i(t)}{dt} = -S_i(t)\rho_i \sum_{j=1}^6 \phi_{i,j}(t) \frac{\xi_{I_{NS,j}} I_{NS,j}(t) + \xi_{I_{Sy,j}} I_{Sy,j}(t)}{N_j - D_j(t)}, \quad (1)$$



**Fig. 2.** Disease state transition model. The disease state transition model comprises eight infection states with possibility for recovery or aggravation after each disease state following exposure state (except for death and recovery itself).

where  $\xi_{INS}$  and  $\xi_{ISy}$  denote the proportion of nonisolated a- or presymptomatic and symptomatic infectious individuals, respectively. Furthermore,  $\rho_i$  denotes the age-resolved probability of becoming infected from a contact with an infectious person, and  $\phi_{i,j}(t)$  represents the mean daily number of contacts between individuals of age group  $i$  and age group  $j$ . While, theoretically,  $\phi_{i,j}(t)$  could vary arbitrarily, in our realization, it only varies in the context of a contact change point and is constant otherwise; cf. Eq. (3). Additionally, we denote the initial population in age group  $j$  by  $N_j$  and the number of disease-related deaths in age group  $j$  up to time  $t$  by  $D_j(t)$ ,  $j = 1, \dots, 6$ .

Disease progression follows a chain of compartments, where an individual either recovers or moves to an aggravated state. We define  $z_{j,i}$  as the  $j$ -th disease state for age group  $i$  and denote the probability of transition from state  $z_{j,i}$  to aggravated state  $z_{j+1,i}$  by  $\mu_{z_{j,i}}^{z_{j+1,i}}$ , while consequently  $1 - \mu_{z_{j,i}}^{z_{j+1,i}}$  is the probability of recovery. The average time an individual spends in a given state  $z_{j,i}$  is represented by  $T_{z_{j,i}}$ . For disease states where progression or recovery is possible, i.e., excluding *susceptible*, *exposed (not infectious)*, *recovered*, and *dead*, transition follows the form

$$\frac{dz_{j,i}(t)}{dt} = \mu_{z_{j-1,i}}^{z_{j,i}} \frac{z_{j-1,i}(t)}{T_{z_{j-1,i}}} - \frac{z_{j,i}(t)}{T_{z_{j,i}}}, \tag{2}$$

where  $z_{j-1,i}$  is the former, less severe state. For the full set of equations, we refer to<sup>25</sup>.

In this study, we use parameters that were obtained for the first waves of wild-type SARS-CoV-2 in 2020. Identical to Plötzke et al.<sup>6</sup>, the parameters are either directly obtained or calculated based on<sup>25</sup>, Table 2. For the average values over all age groups, the age-specific values are weighted with respect to share of the corresponding age group in the total population provided for Germany by<sup>42</sup>. The overview on the particular values can be found in Tables 1 and 2.

Contact pattern changes are modeled within the contact rate matrix  $\phi(t) = (\phi_{i,j}(t))_{i,j}$ . Without loss of generality, let us assume that the first contact change takes place at  $c_1 \in \mathbb{R}$  with  $c_1 > c_0$ , where  $c_0$  is the initial time. For a homogeneous reduction factor  $0 \leq r \leq 1$ , we define

$$\phi_{i,j}(t) := \begin{cases} \phi_{i,j,0}, & t \leq c_1 \\ \widehat{\phi}_{i,j}(t), & t \in (c_1, c_1 + \delta) \\ (1-r)\phi_{i,j,0}, & t \geq c_1 + \delta \end{cases}, \quad 0 < \delta < 1, \tag{3}$$

where  $\phi_{i,j,0}$  is the initial contact rate,  $\delta$  defines a transition interval, and  $\widehat{\phi}_{i,j}$  represents a transition contact rate to ensure smooth transition, i.e.,  $\phi_{i,j}(t) \in C^1((c_1 - \varepsilon, c_1 + \delta + \varepsilon))$  for  $i, j \in \{1, \dots, 6\}$  with  $0 < \varepsilon \ll 1$ . As suggested in<sup>25</sup>,  $\widehat{\phi}$  is realized through a cosine function.

In order to realize the spatial resolution, we employed a generalized graph as described in<sup>25</sup>. Proceeding with this approach, we obtained a graph with one node for each spatial region, which is then equipped with a local disease dynamics model based on ODEs. Furthermore, we obtained one edge for each pair of age group and disease state, which then implements the corresponding mobility exchange. The  $400 \times 400$  mobility matrix for all of Germany's counties was extrapolated from the German Federal Employment Agency<sup>43</sup>. The adjacency matrix  $A$  needed for the GNN was obtained from this mobility data matrix by setting all nonzero connections to one.

The resulting graph, representing 400 German counties, is directed and has a density of 30.4% with 48,467 edges. Its topology is characterized by an uneven degree distribution with highly connected hubs (e.g., Frankfurt) connecting to smaller surrounding areas, indicated by a slightly negative assortativity coefficient ( $-0.096$ ). A

Parameter	Description
$T_E$	Duration (in days) in <i>exposed (not infectious)</i> state
$T_{INS}$	Duration (in days) in <i>a- or presymptomatic and infectious</i> state
$T_{Sy}$	Duration (in days) in <i>mildly symptomatic and infectious</i> state
$T_{Sev}$	Duration (in days) in <i>severely infected (hospitalized)</i> state
$T_{Cr}$	Duration (in days) in <i>critically infected (ICU)</i> state
$\rho$	Transmission risk on contact
$\xi_{INS}$	Proportion of not isolated a- or presymptomatic infectious cases
$\xi_{ISy}$	Proportion of not isolated symptomatic cases
$\mu_{INS}^{ISy}$	Proportion of symptomatic cases per all a- and presymptomatic infectious cases
$\mu_{ISy}^{ISev}$	Proportion of severe cases per symptomatic case
$\mu_{ISev}^{ICr}$	Proportion of critical cases per severe case
$\mu_{ICr}^D$	Proportion of deaths per critical case

**Table 1.** Explanation of the model parameters..

	0–4 y	5–14 y	15–34 y	35–59 y	60–79 y	80+ y	Weighted Average
$T_{E,i}$	3.335	3.335	3.335	3.335	3.335	3.335	3.335
$T_{I_{NS},i}$	2.74	2.74	2.565	2.565	2.565	2.565	2.58916
$T_{I_{Sy},i}$	7.02625	7.02625	7.0665	6.9385	6.835	6.775	6.94547
$T_{Sev}$	5	5	5.925	7.55	8.5	11	7.28196
$T_{Cr,i}$	6.95	6.95	6.86	17.36	17.1	11.6	13.066
$\rho_i(t)$	0.03	0.06	0.06	0.06	0.09	0.175	0.07333
$\mu_{I_{NS}}^{I_{Sy}}$	0.75	0.75	0.8	0.8	0.8	0.8	0.79310
$\mu_{I_{Sy}}^{I_{Sev}}$	0.0075	0.0075	0.019	0.0615	0.165	0.225	0.07864
$\mu_{I_{Sev}}^{I_{Cr}}$	0.075	0.075	0.075	0.15	0.3	0.4	0.17318
$\mu_{I_{Cr}}^D$	0.05	0.05	0.14	0.14	0.4	0.6	0.21718

**Table 2.** Age-resolved parameters for wild-type SARS-CoV-2 based on<sup>25</sup>.

high clustering coefficient (0.56) further reflects strong local community structures. From a physical perspective, the GNN learns to approximate the epidemic diffusion process on this specific topology. Since the local rules of this diffusion are consistent, the model is capable of generalizing to unseen temporal scenarios, given the local connectivity patterns remain similar.

While also this metapopulation model was implemented efficiently and did undergo several optimization steps, for large experiments we still require relevant computing resources; cf. Fig. 9C.

### Performance evaluation and objective function

Simulating the previously described mechanistic metapopulation model, we obtain time-dependent (and spatially resolved) outputs. The output vectors from all local models of the form Eqs. (1) and (2) are interpolated to full simulation days so that corresponding age-structured and spatially resolved data sets have the same dimension. To evaluate the prediction performance of the neural network-based surrogate models as well as the dummy estimator (see next subsection), we use the mean absolute percentage error (MAPE) as evaluation metric, which is defined as

$$\text{MAPE} := \frac{100}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{|y_i|}, \quad (4)$$

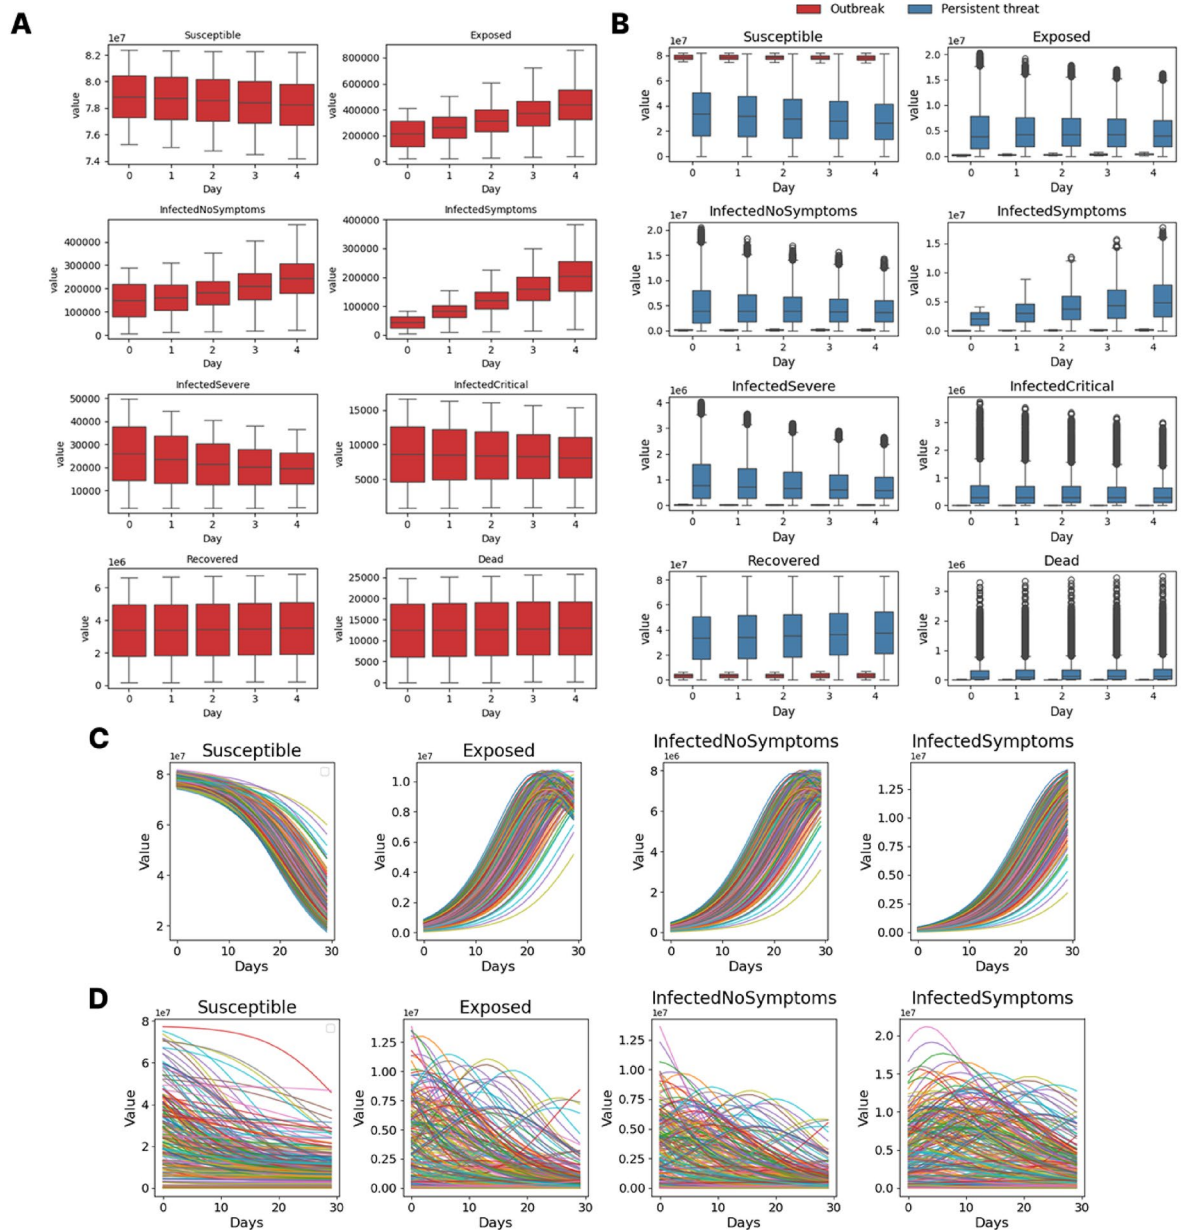
as objective function. Here,  $n$  is the total number of data points which is the product of the number of runs, the number of days, and the number of compartments,  $y_i$  and  $\hat{y}_i$  denote individual scalar values representing the interpolated “ground truth” and predicted daily values, respectively.

### Generation of input data

To develop a generalizable surrogate model for a large variety of observed disease dynamics, we require heterogeneous data sets with a large variance in the outcomes, which transfer to the input features of our surrogate models and which eventually cover the full range of all relevant scenarios. The heterogeneity of the generated data sets were assessed qualitatively by visualization and quantitatively by investigating the predictive power of a simple *dummy estimator*. We compute this dummy estimator by the averaging over 80 % of the data and evaluating its predictive performance on the trajectories of the remaining 20 % of the data using the MAPE given in Eq. (4) as the performance metric. We conclude that the higher the MAPE of the dummy estimator, the higher the variance of the input data.

We consider two epidemic regimes that define the initial five input days per sample (values per 100 000 population and stratified by six age groups unless stated otherwise). In the *outbreak* regime, currently symptomatic individuals are sampled uniformly from [7, 100] and recently exposed from [25, 500]; the remaining compartments are filled consistently to satisfy population balance. In the *persistent-threat* regime, the share of currently symptomatic individuals is drawn uniformly from [0.01%, 5%]; in the worst case, up to 50% of the population may be recently exposed, a-/presymptomatic, or symptomatic. To reflect heterogeneous historical dynamics and immunity, the initial fraction recovered per region is drawn uniformly within the feasible interval implied by the ongoing dynamics so that all compartment shares remain nonnegative and sum to one. The variability of the outbreak and persistent-threat ranges are provided in Panels A and B of Fig. 3. In Fig. 3C,D, we plot the 200 first trajectories for the outbreak scenario (Panel C) and the persistent-threat (Panel D) settings. Here, we see a much more heterogeneous development in Panel D).

Interventions are represented by contact change points. Each sample includes  $M \in \{0, 1, 2, 3\}$  changes within the first 30 days;  $M$  is drawn uniformly from this set. Change days ( $d_1, \dots, d_M$ ) are drawn without replacement from  $\{1, \dots, 30\}$  and sorted increasingly, and the homogeneous reduction factors  $r_m$  are sampled from  $\mathcal{U}[0, 1)$ . Smooth transitions over a short interval  $\delta \in (0, 1)$  are applied in the mechanistic simulator using

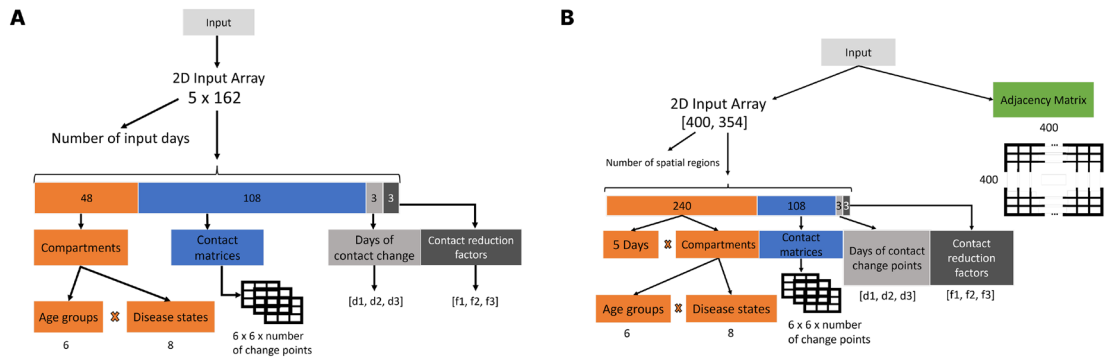


**Fig. 3.** Initial conditions over five input days and trajectories over the prediction horizon. **(A)** Initial conditions for outbreak range scenarios. **(B)** Comparison between initial conditions for outbreak range and persistent-threat range scenarios. **(C)** The two hundred first trajectories for the outbreak scenarios. **(D)** The two hundred first trajectories for the persistent-threat scenarios.

a cosine ramp (cf. Eq. (3)). Age-structured contact matrices are  $6 \times 6$  per change point; if  $M < 3$ , the unused slots are masked in the learning input.

We model Germany as a graph with 400 nodes (counties). Each node carries the age-stratified compartment trajectories. The adjacency matrix  $A \in \{0, 1\}^{400 \times 400}$  is derived from the aforementioned mobility matrix by mapping nonzero flows to one; this yields a sparse graph with roughly 25% nonzero entries. The same graph is used for training and testing.

Per sample we provide the first five days as input and predict subsequent horizons of 30, 60, or 90 days. For non-spatial baselines, we use a 2D array of dimension  $5 \times 162$  array with the five input days defining the rows each with 48 age-resolved compartment values (six age groups times eight compartments) plus up to three  $6 \times 6$  contact matrices with their three contact change days and reduction factors. For the GNNs, each sample consists of node features  $X \in \mathbb{R}^{400 \times 354}$  containing the five input days of compartment values (48 per node) and the intervention descriptors (contact matrices, change days, reduction factors) broadcast per node, paired with the adjacency matrix  $A$ . This schema mirrors the input layout illustrated in Fig. 4. To balance scale differences across compartments (e.g., *Susceptible* vs. ICU), we apply a  $\log(1 + x)$  transform to all count-valued series before training; evaluation is performed on the original scale after inverse transforming.



**Fig. 4.** Non-spatial vs. spatial input encoding. **(A)** Non-spatial 2D input (5 × 162): five input days with 48 (age groups × compartments) values plus up to three 6 × 6 contact matrices, their change days, and reduction factors. **(B)** Spatial GNN input: for each node (400 counties) the same features as in Panel A are broadcast and paired with the adjacency matrix.

We generate 10,000 samples for the non-spatial baselines and 1,000 samples for the spatial GNNs. Unless stated otherwise, we use an 80–10–10 split into train/validation/test on the sample level and perform five-fold cross-validation on the training split during architecture search to avoid split artifacts; the reported test MAPE refers to the held-out test set. The dummy estimator is computed on the training split and evaluated on the held-out portion within each regime as a single-number proxy for input heterogeneity; detailed values are reported in the Results.

### Neural network surrogate model

We conduct a multi-stage model building and optimization process to build up a neural network with spatially resolved infection predictions. Before deriving suitable input data structures for learning a spatially resolved graph-based model with a GNN, we consider the temporal evolution of disease dynamics. For this, we train an MLP, an LSTM and a CNN model, each implemented using the Keras<sup>44</sup> library with TensorFlow<sup>45</sup> as backend.

With respect to the temporal evaluation, we take several substeps in our tests. In particular, we first consider a population without stratification into age groups. Then, we stratify the population into age groups, for which we consider age-dependent contact patterns and disease parameters. For each model architecture, we perform a grid search on zero to four hidden layers with 32 to 1024 neurons per layer, ranging up to seven layers and 2048 neurons for the final GNN layer type. After selecting a suitable neural network architecture, we fine-tune additional architecture-specific hyperparameters.

After our preliminary tests with simple neural network models which do not consider spatial effects, we turn to GNNs. We use the Spektral library from Grattarola and Alippi<sup>46</sup> for defining the graph data and building the GNNs and Keras to train and evaluate the model. Spektral is an open-source Python library based on the Keras API and TensorFlow<sup>45</sup> implementing GNN models. Other existing libraries, which are not considered here, are PyTorch Geometric<sup>47</sup> and DeepGraphLibrary<sup>48</sup>.

GNNs can be categorized into spectral and spatial approaches, where spectral approaches have a mathematical foundation in graph signal processing while spatial approaches directly aggregate information from the nodes’ spatial neighborhoods (see Wu et al.<sup>49</sup>), which is given by the graph. We compare the performance of four different types of GNN layers available in Spektral to build our model. As a first approach, we consider GCNConv layers, which are based on the graph convolutional network (GCN) as proposed by Kipf et al.<sup>50</sup>. In this spectral approach, a modified adjacency matrix with self-loops is used. The second type are ARMAConv layers, which are based on Bianchi et al.<sup>51</sup>. Like the GCNConv layer, this is a spectral approach, and it applies convolutional operations in combination with an autoregressive moving average (ARMA) approach for improved training performance; cf. the discussion below. The third approach uses approximative personalized propagation of neural predictions convolutional (APPNConv) layers based on Gasteiger et al.<sup>52</sup>, which can also be categorized as a spectral approach with convolutional operations. Lastly, we consider the graph attention convolutional (GATConv) layer from Velickovic et al.<sup>53</sup>, which is different from the other considered approaches, in the sense that it uses a spatial approach that implements an attention mechanism.

Since the ARMAConv layer performed best (see the Results section), we briefly provide some details on this specific type of GNN layer as well as our final architecture. In<sup>51</sup>, Bianchi et al. substituted polynomial filters in GNNs by filters which are based on the ARMA approach. In particular, in our optimized architecture, the final model consists of seven ARMAConv layers with 512 channels per layer and consisting of one stack of GCS layers, which serves as a filter on the node feature input data. Generally, we denote a ARMAConv layer with  $k$  GCS stacks as  $ARMA_k$ . The  $i$ th GCS layer in the stack is parametrized by the trainable weight matrix  $W_i$  as well as the trainable skip connection  $V_i$ , which arises from the ARMA scheme applied to the GCS stack. In the case of the  $ARMA_1$  layer with one GCS stack, the graph convolution layer corresponds to

$$\bar{X}_1^{(t+1)} = \sigma(\bar{A}\bar{X}_1^{(t)}W_1^{(t)} + \bar{X}_1^{(0)}V_1^{(t)}),$$

where  $\sigma$  is an activation function,  $\bar{X}^{(0)}$  are the initial node features,  $\bar{X}^{(t)}$  are filtered node features of the previous iteration, and  $\tilde{A}$  is the normalized adjacency matrix  $\tilde{A} = D^{-1/2}AD^{-1/2}$ ; it is normalized using the degree matrix  $D$ , which counts for each node the number of adjacent edges in the graph. After applying all GCS layers, a pooling operation is employed on the GCS stack, for instance, average pooling. With only one GCS layer, the averaging pool does not have any effect and simply applies the identity. In all configurations, the activation function was set to ReLU and no dropout was selected.

## Results

In this section, we present the results of our computational experiments. First, we compute the baseline performance to be beaten by any meaningful neural network. Then, we discuss how we optimize the architecture and hyperparameters of our models, starting with simple models accounting only for the temporal evolution and moving on to models also taking into account the spatial distribution. Finally, we show results on the computational efficiency of the models.

### Baseline performance

The design of experiments in Section 2 produces substantially different trajectories across, both, the *outbreak* and *persistent-threat* scenarios, with differently large variances in the data set. The representative ranges and trajectories are shown in Fig. 3. As mentioned in the previous section, we use the dummy estimator to compute the baseline performance to be beaten by any meaningful surrogate model. As a compact quantitative proxy for dataset variability, the mean dummy estimator (trained on 80 % of the data and evaluated on the remaining 20 %) attains MAPE values of 34.16 % for the outbreak regime and 1197.76 % for the persistent-threat regime (original scale). The latter highlights a markedly harder prediction setting and thus a useful stress test for model selection.

### Neural network performance without spatial resolution

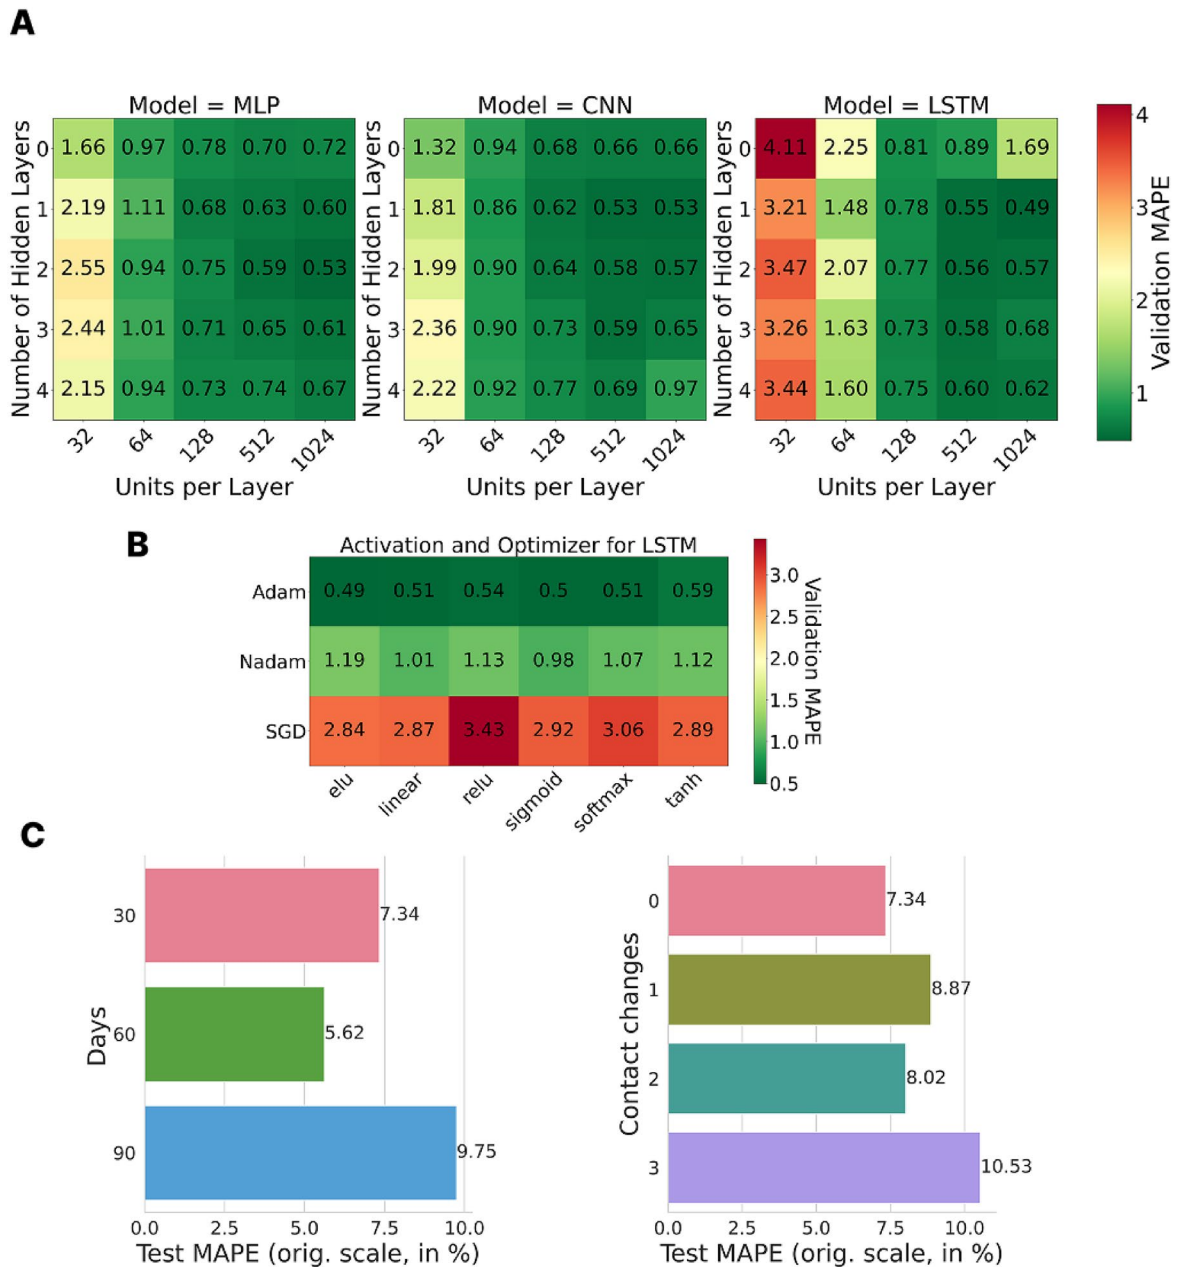
In this section, we consider only temporal evolution and neglect spatial effects to get a better idea of the potential performances on the complex datasets generated and compare the model failures against the dummy estimator outcomes of the previous section.

We consider simple neural networks, such as multilayer perceptrons (*MLPs*), convolutional neural networks (*CNNs*) with one-dimensional convolutions in the time dimension, and long short-term memory (*LSTM*) networks. We conduct grid searches on the number of hidden layers (0 to 4) and neurons (neurons, channels, or units) per layer (32 to 1024) for a setting with six age groups and simulated over 30 days without contact change points. We used an 80-10-10 split of the dataset into training, validation, and test sets. All grid searches are conducted by using a five-fold cross validation (CV) to avoid artifacts from the data split. The logarithmically scaled MAPE values of the cross-validation are shown in Fig. 5A. The MAPE on the validation sets will be denoted as *validation MAPE*, the MAPE on a corresponding test set will be denoted as *test MAPE*. A further analysis revealed Adam as the best optimizer and ReLU as the best activation function; see Fig. 5B. While we evaluated the models on both datasets, outbreak and persistent-threat, we only provide the latter case, which is the substantially more complex one. The best architecture in our comparison was a simple LSTM with one hidden layer and 1024 units which achieved a MAPE of 0.49 % across the grid search and, thus, 3-4 magnitudes better than our dummy baseline estimator. We further test it on longer time horizons with up to three contact change points; see Fig. 5C. While the architecture is optimized on a much simpler model, we obtain test MAPEs between only 5-10 % for the input features depicted in Fig. 4A. In the supplementary material, Supplementary Figures 1 and 2, we provide predictions for the first two trajectories with and without stratification into age groups to visually assess the meaning of the corresponding MAPE values.

### Optimized GNN architecture

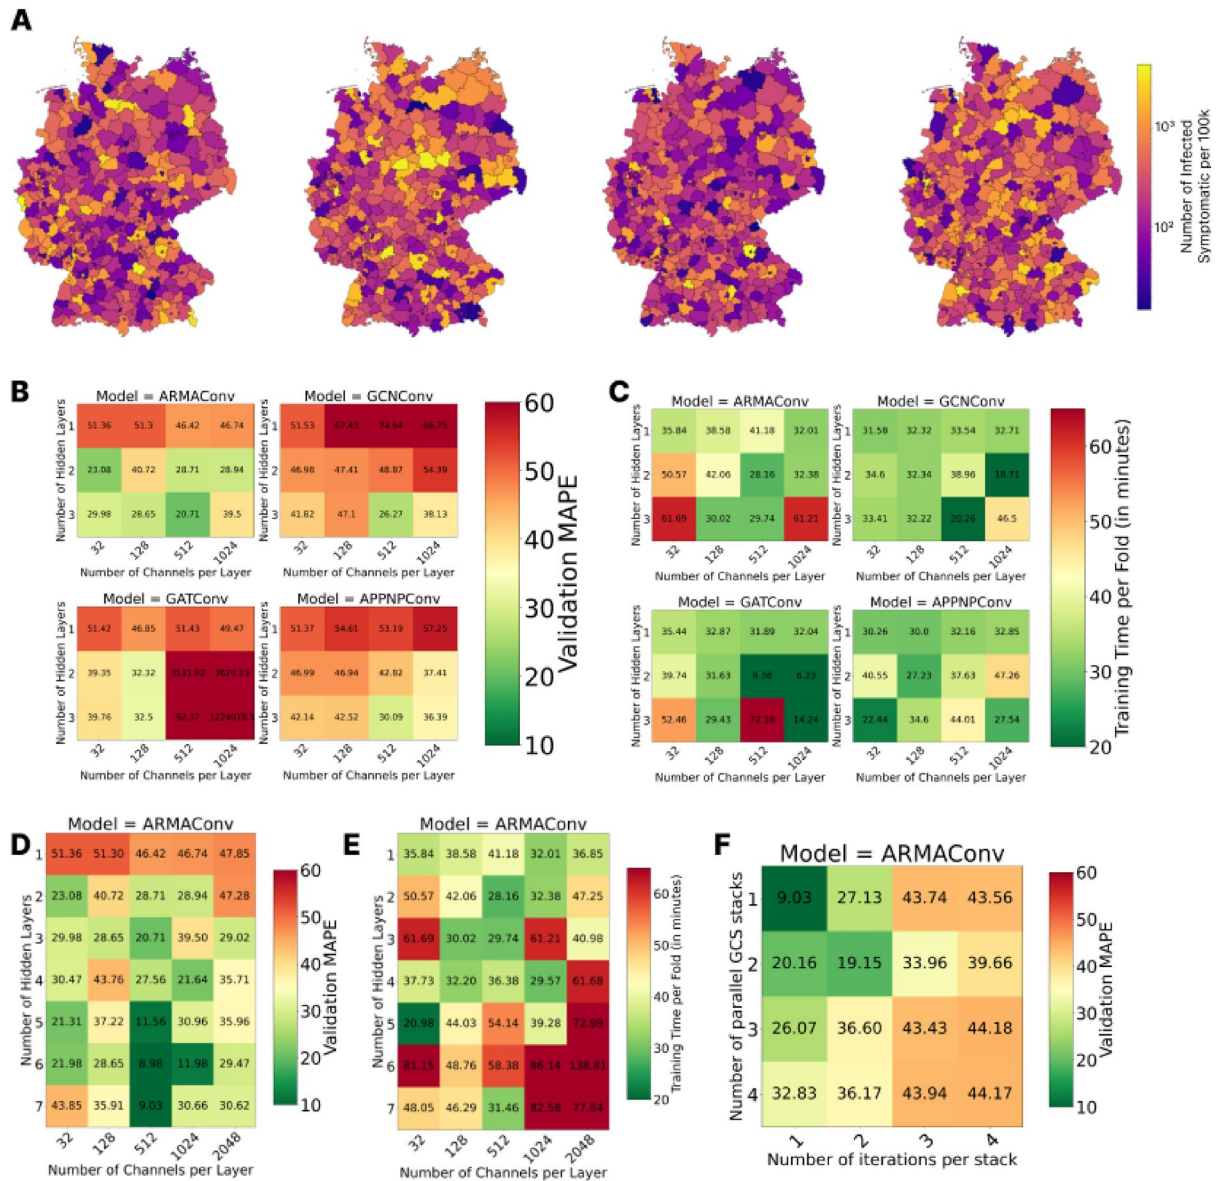
Combining the results from our explorational experiments of the simple MLP and LSTM neural networks with some minor adaptations, i.e., eliminating the redundant storage of the contact matrices of the 2D dimensional  $5 \times 162$  and transforming it to a one dimensional vector of length 354, we use the input data set structure for the GNNs as presented in the previous section and shown in Fig. 4B. For a GNN to be trained or executed, we provide an (additional) adjacency matrix  $A$  for the considered regions. The adjacency matrix is naturally obtained from the graph-based expert modeling approach. The dimension of the adjacency matrix is  $400 \times 400$ , which comes from the particular application on the 400 German counties. Similarly, the first array dimension of the remaining input is 400. Eventually, to adequately integrate heterogeneous disease dynamics over an entire country, we randomize historic and ongoing disease dynamics over all 400 counties. The ongoing disease dynamics are initialized per county with the persistent-threat scenario, thus allowing between very small and very high infectious disease dynamics. For the first four initializations, see Fig. 6A. To additionally randomize for the historic dynamics and immunities, we draw the number of recoveries uniformly from the interval between zero and one minus the maximum ongoing dynamics (infected and dead) that are possible in the persistent-threat range.

Having derived a suitable set of input features for a GNN surrogate model, we selected an appropriate GNN model and optimized its architecture and hyperparameters. We conducted a grid search to identify the model architecture with the best performance, also taking into account the respective training time, which can be considerable. To avoid unnecessarily long training times, we used an early stopping criterion after 200 iterations without improvement (i.e., patience of 200 iterations). Again, averaged validation MAPEs are presented in log-scale.



**Fig. 5.** MAPE results of the grid search for the simple neural networks and simple input data structure. **(A)** Validation MAPE results (in log-scale) for a grid search for simple models with six age groups and no contact change on a horizon over 30 days. **(B)** Validation MAPE results (in log-scale) for the grid search for different optimizers and activation functions. **(C)** Test MAPE results (in original scale) for the final LSTM model on different prediction horizons (30, 60, and 90 days) and for different numbers of contact changes (0, 1, 2, 3 changes in 30 days).

We start by considering different graph layers as described in the *Materials and Methods* section and vary the number of layers between 1 and 3 and the number of channels per layer from 32 to 1024. We observe that the ARMAConv model with three layers and 512 channels per layer performs best with respect to the validation MAPE (Fig. 6B), while only the GATConv models with the highest MAPEs need substantially less training time (Fig. 6C). We thus refine the grid search for the ARMAConv model to seven layers and observe that the models with six or seven layers and 512 channels per layer perform best. While the validation MAPE for seven layers, 9.03 %, is slightly larger than that for six layers, 8.98 % (Fig. 6D), the training time for seven layers, 31.46 minutes, is reduced by almost a factor of two compared with the model with six layers, which requires 58.38 minutes of training (Fig. 6E). This counterintuitive reduction in total training time results from the early stopping mechanism: although the computational cost per epoch increases with model depth, the 7-layer model converged in substantially fewer epochs across the cross-validation folds compared to the 6-layer model. To




**Fig. 6.** Spatially distributed disease dynamics and grid search of different GNN infrastructures. (A) Four random initializations of heterogeneous disease dynamics over the 400 counties of Germany. (B) Validation MAPEs (in log-scale) for the four different GNN layers. (C) Training times corresponding to (B). (D) Extended grid search for ARMAConv model (in log-scale). (E) Training times corresponding to (D). (F) Further hyperparameter tuning for ARMAConv model with seven layers and 512 channels per layer (in log-scale).

ensure that this large architecture remains stable, we analyze the learning dynamics. As shown in Fig. 7, training and validation errors decrease without divergence. At the best epoch (1277), the validation MAPE is comparable to the training MAPE, confirming robust generalization without overfitting.

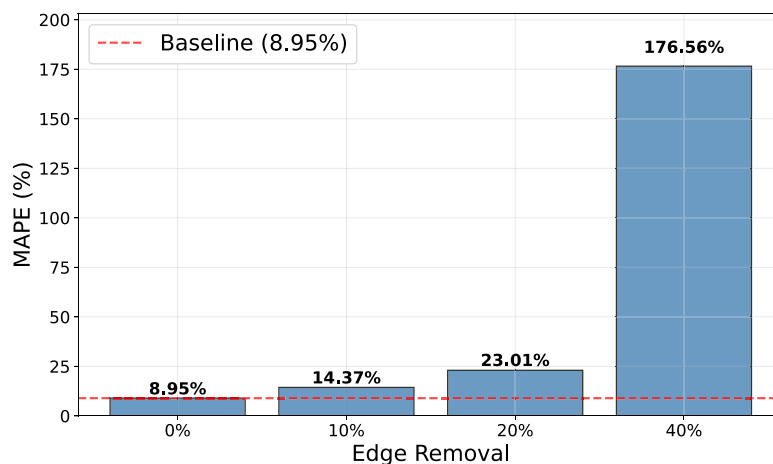
In Fig. 6F, we report an additional assessment on hyperparameter tuning, testing the parameters that are specific to the ARMAConv model: the numbers of parallel GCS stacks and iterations per stack. For details, we again refer to the *Materials and Methods* section and the original publication of the ARMAConv model<sup>51</sup>.

The final model architecture is then set to seven ARMAConv layers with 512 channels per layer containing one GCS stack whose filtered response is computed over one iteration.

Analogously to the simple neural networks employed for predictions without a spatial distribution, we tested the performance of the GNN on tasks with various prediction horizons (30, 60, 90 days) and numbers of contact change points (0, 1, 2, 3) within a time span of 30 days. Here, we used a classic 80-10-10 split of the data set into training in, validation, and test data. The performance of the model was relatively constant for different prediction horizons. As depicted in Fig. 9A, we obtain test MAPEs between 10.33 % and 17.89 % for up to two contact changes on 30 days or without contact changes on 30, 60, and 90 days horizon. For a prediction horizon of 30 days and three contact changes, we obtain the highest MAPE with 27.11 %; see Fig. 9.



**Fig. 7.** Learning dynamics of the optimized ARMAConv model (7 layers, 512 channels). The training (red) and validation (green) MAPEs show consistent convergence.



**Fig. 8.** Generalization of GNN surrogate under edge removal. The bar chart shows MAPE (original scale) for four scenarios: baseline, 10% edge removal, 20% edge removal, and 40% edge removal. Bars represent mean prediction error across 20 independent epidemic simulations per scenario using the pre-trained model without retraining. The red dashed horizontal line marks baseline performance on the original 400-county German mobility network.

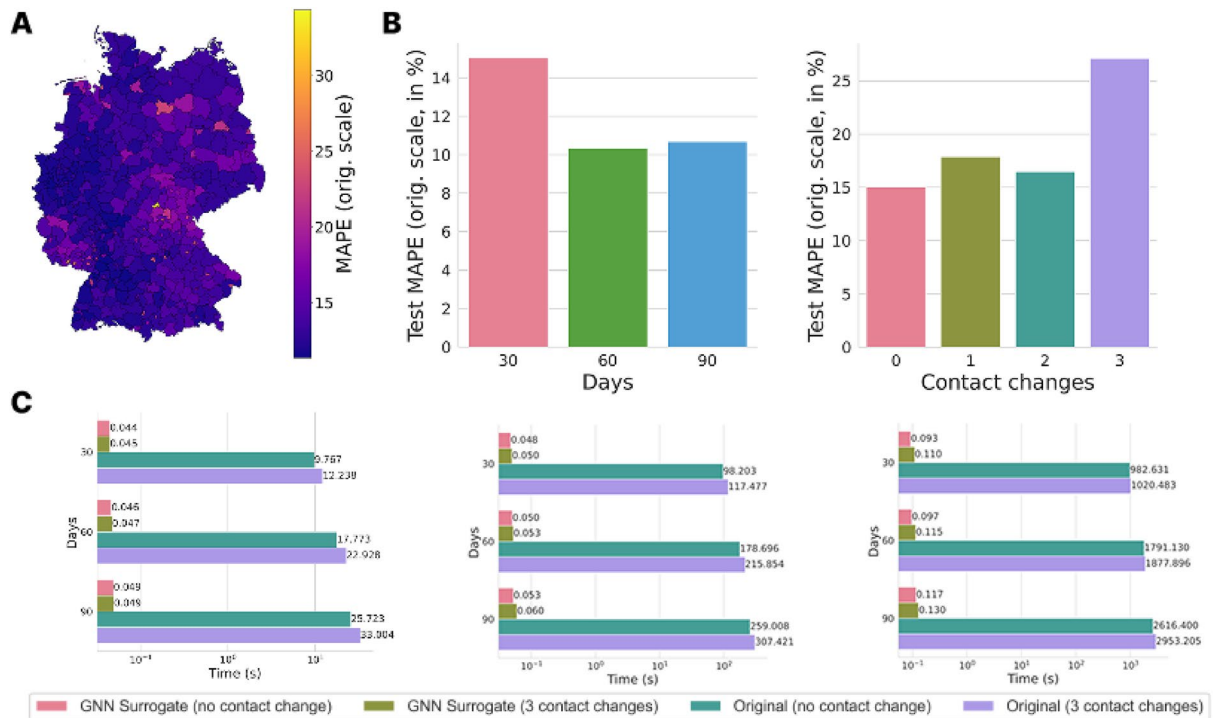
### Generalization to modified networks

In this section, we briefly report on the generalization capabilities of the trained GNN. We use the model which was trained on a fixed graph topology where all 1,000 training samples used the same mobility network.

We then created three modified mobility scenarios by randomly removing edges from the original network with (b) 10% edge removal, (c) 20% edge removal, and (d) 40% edge removal, always comparing against the baseline (0% reduction).

For each scenario, we generated 20 independent simulations using the mechanistic ODE metapopulation model with the modified mobility matrices using the persistent threat scenario. We then evaluated the pre-trained GNN model on these novel test sets and measured prediction accuracy using MAPE on the original scale. The results are presented in Fig. 8 and demonstrate generalization capabilities for graph perturbations with 10 or 20% less edges and clear degradation as topological changes become substantial. With 10% edge removal, MAPE increases from 8.95% to 14.37%. The model maintains reasonable prediction accuracy, suggesting robustness to minor network disruptions (e.g., temporary mobility restrictions in specific regions). With 20% edge removal, MAPE rises to 23.01%. While noticeable, predictions may still provide useful scenario guidance with the GNN showing sensitivity to moderate topological changes. With 40% edge removal, MAPE jumps to 176.56%, indicating high prediction errors and when nearly half of the mobility network is removed.

In summary, the surrogate is reasonably robust to modest mobility changes (e.g., 10 to 20% edge reduction for simulating localized travel restrictions) but sensitive to major structural changes of the mobility network. On



**Fig. 9.** Test MAPE for final ARMAConv GNN and runtime comparison between GNN surrogate model and expert mechanistic model. **(A)** Test MAPE over the 400 counties. **(B)** Test MAPE over prediction horizons of 30, 60, and 90 days without contact change and with 0, 1, 2, and 3 contact changes over 30 days horizon. **(C)** Runtime comparison between final GNN and mechanistic expert model of<sup>25</sup> for one, ten, and 100 model executions with either zero or three contact changes over different prediction horizons.

the other side, the mechanistic model itself is robust to any topology change, as it simulates dynamics from first principles and the surrogate trades this flexibility for computational speed.

### Execution time performance

While the time and energy needed for training a surrogate model should also be taken into account, the execution time performance is the critical factor if the model is to be used with a graphical user interface (GUI) that needs to respond immediately. The best GNN surrogate model from the previous section yielded test MAPEs between approximately 10% and 27%, when predicting age-structured and spatially resolved disease dynamics over the 400 counties of Germany. In this section, we present result that indicate how the surrogate model outperforms the expert mechanistic model in terms of execution times. Timings were measured on a GPU-accelerated Intel Xeon Scalable Processor “Skylake” Silver node with four NVIDIA Tesla V100 SXM2 GPUs with 16 GB, utilizing a single GPU for the measurements.

To allow for just-in-time, web-based, exploration of different scenarios by manually adapting one parameter at a time, the model requires one single model prediction for every parameter change. When predefining multiple scenarios in a joint space over different parameters, for instance, the implementation of interventions to be evaluated on different days (e.g., today, in three days, and in one week) and with varying strictness (e.g., restricting gatherings to ten, 20, or 50 persons), a batch of multiple predictions are needed. With a factor space in just two or three dimensions and two or three discrete values, we then need between four and 27 model predictions. In Fig. 9, we thus provide the measured run times for one, ten, and 100 executions of the expert and the surrogate model, respectively. Firstly, we observe that the expert model runtime scaled—as expected—linearly with the number of days while the GNN surrogate model provided results in (almost) constant time independent of the horizon. Furthermore, the execution time of the expert model also scaled linearly with the number of executions (without parallelization). Again, the scaling of the GNN surrogate model is clearly superior. While the timings are similar for one and ten executions, it only doubles to tripled from ten to 100 model executions. The prediction times of both models are almost independent of the number of contact change points. Overall, the neural network speeds up the prediction process by 222 to 673 times for a single model execution and by 2046 to 5123 times for ten model executions, going up to a maximal speed-up factor of 28 670 for 100 executions and 90 days prediction horizon. With less than or approximately 0.1 seconds for the surrogate model to be executed, it can be employed on-the-fly for up to 100 model predictions. Although the absolute execution speed on less powerful hardware setups will be slower, we expect that the efficiency of the GNN surrogate model still allows for rapid and interactive exploration scenarios, even in resource-poor environments.

## Discussion

Surrogate models are a cost-efficient technique to approximate the behavior of complex models. Usually, the more complex the reference model is, the higher the requirement of computational resources. The goal of the surrogate model is to approximate the original model's behavior, achieving accuracies close to the original model, at substantially reduced computational cost and time. In time-critical epidemic or pandemic situations, the evaluation of a large set of potential reactions is recommended, and swift reaction is crucial<sup>54</sup>. Unfortunately, supercomputing resources might not be available on model execution or predictions might be needed by stakeholders which do not have access to supercomputing facilities. To overcome these barriers and provide just-in-time answers to urgent questions, (pre)trained, e.g., overnight, surrogate models can become of great aid for decision makers and substantially increase the impact of mathematical-epidemiological modeling. With a regularly calibrated mechanistic model that integrates most recent viral parameter knowledge, mechanisms, and changed contact behavior, the surrogate model complements classical, purely data-driven ML models through the learning step of the simulator.

In this article, we present a surrogate model based on neural networks that is capable of handling temporal and spatial resolution. While several recent surrogate models are based on the concept of physics-informed neural networks<sup>55–57</sup>, our surrogate model is entirely data-driven and uses a “physics”-based simulator only to create a wide range of reliable training data. On the one hand, this means that physical constraints such as population conservation are not considered explicitly. On the other hand, our approach offers a substantial advantage, as particular modifications to the original model do not require changes in the surrogate model architecture; only new training data has to be generated and the model has to be retrained, potentially using the previous model as a starting point. Since (epidemiological) parameters usually do not change on a daily basis, retraining is only necessary if these parameters change substantially. In future work, we will investigate if our surrogate model could further benefit from integrating “physical” constraints such as population conservation into the loss function.

Although we use age-resolved models and disease parameters, contact reductions in this study are applied homogeneously to all age groups. As NPIs, such as school closures or remote work policies, often target particular age groups, heterogeneous contact reductions should be studied in future research. While an implementation of heterogeneously changing contact patterns would be straightforward, we have not investigated how this would affect the performance of the surrogate model. Similarly, we have not considered spatially heterogeneous interventions. Both scenarios will be considered in future research. Moreover, when increasing the complexity of the contact restrictions in those scenarios, we expect that a larger data set and perhaps a larger network would be required to obtain a similar model performance. Another limitation of the current study is that we have not yet considered vaccination, waning immunity, or reinfection. This means that the model in its current form has not been validated for realistic, later stage epidemic situations or more input settings.

In this study, we first compare four different types of GNN layers using the binary adjacency matrix modeling the mobility between different spatial regions as input. Therefore, although already providing convincing results, none of the considered approaches takes the highly varying values for traveling and commuting activities between the different spatial regions into account. We expect that incorporating this information as edge weights in the input graph could further improve the performance. Promising directions for incorporating the edge weights could be the CrystalConv<sup>58</sup> or XENetConv<sup>59</sup> layers. However, due to hardware limitations, specifically GPU memory on a single NVIDIA Quadro RTX 6000, we were unable to train models incorporating large graphs with 40,000 edges. Addressing this limitation will be a focus of future work.

This study mainly sheds light on the performance of surrogate model using the same graph connectivity in training and evaluation processes. We only briefly explored the generalization potential of the GNN model to other graphs than the one used for training, which is one of the strengths of GNNs. We considered the mobility network of 400 German counties and removed different quantities of edges, observing that a modest topological change with 20 % removed edges still yields acceptable prediction results. In particular, research on generalization with several realistic mobility networks and edge weights will be an avenue for future research.

With certain limitations and directions for future work, we present a surrogate model using methods of artificial intelligence that naturally integrate spatial information. The presented GNN performs well on the considered data sets for an initial or persistent disease outbreak and nonpharmaceutical interventions on all 400 counties of Germany. It is able to make predictions with low MAPEs over time horizons between 30 and 90 days, integrating contact pattern changes of the population on up to three different days during the considered period.

Our surrogate model delivers an age-structured and spatially resolved disease dynamics for a prediction horizon of several months and speeds up the prediction process substantially. As these results can be computed within a fraction of a second, our model is capable of being integrated in web-based application<sup>60,61</sup>.

## Data availability

All data and code used to generate the results presented in this work are publicly available as part of the MEMilio framework. The specific implementation used in this study, including all relevant scripts and pretrained models, can be found under a persistent doi at zenodo.

Received: 6 November 2025; Accepted: 5 February 2026

Published online: 13 February 2026

## References

1. European Centre for Disease Prevention and Control. *Public Health and Social Measures For Health Emergencies and Pandemics in the EU/EEA: Recommendations for Strengthening Preparedness Planning* TechRep (2024).

2. Daszak, P. et al. *Workshop Report on Biodiversity and Pandemics of the Intergovernmental Platform on Biodiversity and Ecosystem Services (IPBES)*. Tech. Rep. <https://doi.org/10.5281/zenodo.7432079> (2020).
3. Bauer, S. et al. Relaxing restrictions at the pace of vaccination increases freedom and guards against further COVID-19 waves. *PLoS Comput. Biol.* **17**, e1009288 (2021) (Publisher: Public Library of Science San Francisco, CA USA.).
4. Schüler, L., Calabrese, J. M. & Attinger, S. Data driven high resolution modeling and spatial analyses of the COVID-19 pandemic in Germany. *PLoS ONE* **16**, e0254660 (2021).
5. MacDonald, N. *Periodic Solutions* 39–49 (Springer, 1978).
6. Plötzke, L., Wendler, A., Schmieding, R. & Kühn, M. J. Revisiting the linear chain trick in epidemiological models: Implications of underlying assumptions for numerical solutions. *Math. Comput. Simul.* **239**, 823–844 (2026).
7. Hurtado, P. J. & Kiro Singh, A. S. Generalizations of the 'Linear Chain Trick': Incorporating more flexible dwell time distributions into mean field ODE models. *J. Math. Biol.* **79**, 1831–1883 (2019).
8. Medlock, J. & Kot, M. Spreading disease: Integro-differential equations old and new. *Math. Biosci.* **184**, 201–222 (2003).
9. Wendler, A., Plötzke, L., Tritschschak, H. & Kühn, M. J. A nonstandard numerical scheme for a novel seir integro-differential equation-based model allowing nonexponentially distributed stay times. *Appl. Math. Comput.* **509**, 129636 (2026).
10. Pei, S., Kandula, S. & Shaman, J. Differential effects of intervention timing on COVID-19 spread in the United States. *Sci. Adv.* **6**, eabd6370 (2020).
11. Chen, X., Zhang, A., Wang, H., Gallaher, A. & Zhu, X. Compliance and containment in social distancing: Mathematical modeling of COVID-19 across townships. *Int. J. Geogr. Inf. Sci.* **35**, 446–465 (2021).
12. Levin, M. W., Shang, M. & Stern, R. Effects of short-term travel on COVID-19 spread: A novel SEIR model and case study in Minnesota. *PLoS ONE* **16**, e0245919 (2021).
13. Liu, J., Ong, G. P. & Pang, V. J. Modelling effectiveness of COVID-19 pandemic control policies using an area-based SEIR model with consideration of infection during interzonal travel. *Transport. Res. Part A*. **161**, 25–47 (2022).
14. Zunker, H. et al. Novel travel time aware metapopulation models and multi-layer waning immunity for late-phase epidemic and endemic scenarios. *PLoS Comput. Biol.* **20**, 1012650 (2024).
15. Zunker, H., Dönges, P., Lenz, P., Contreras, S. & Kühn, M. J. Risk-mediated dynamic regulation of effective contacts de-synchronizes outbreaks in metapopulation epidemic models. *Chaos Solitons Fractals* **199**, 116782 (2025).
16. Bershteyn, A. et al. Implementation and applications of EMOD, an individual-based multi-disease modeling platform. *Pathogens Dis.* **76**, 059 (2018).
17. Kerr, C. C. et al. Covasim: An agent-based model of COVID-19 dynamics and interventions. *PLoS Comput. Biol.* **17**, 1–32 (2021).
18. Müller, S. A. et al. Predicting the effects of COVID-19 related interventions in urban settings by combining activity-based modelling, agent-based simulation, and mobile phone data. *medRxiv* (2021).
19. Kerkmann, D. et al. Agent-based modeling for realistic reproduction of human mobility and contact behavior to evaluate test and isolation strategies in epidemic infectious disease spread. *Comput. Biol. Med.* **193**, 110269 (2025).
20. Bicker, J., Schmieding, R., Meyer-Hermann, M. & Kühn, M. J. Hybrid metapopulation agent-based epidemiological models for efficient insight on the individual scale: A contribution to green computing. *Infect. Dis. Model.* **10**, 571–590 (2025).
21. Kehrer, K., Weiser, M. & Conrad, T. Hybrid PDE–ODE models for efficient simulation of infection spread in epidemiology. *Proc. R. Soc. A* **481**, 20240421. <https://doi.org/10.1098/rspa.2024.0421> (2025).
22. Bostanci, I. & Conrad, T. Integrating agent-based and compartmental models for infectious disease modeling: A novel hybrid approach (2024). [arXiv:2407.20993](https://arxiv.org/abs/2407.20993).
23. Tayarani, N.M.-H. Applications of artificial intelligence in battling against COVID-19: A literature review. *Chaos Solitons Fractals* **142**, 110338 (2021).
24. Memon, S. Automated. et al. Automated Processing of Pipelines Managing Now- and Forecasting of In563 fectionous Diseases. In 2024 47th MIPRO ICT and Electronics Convention (MIPRO) 1157–564 1162 (IEEE, 2024)
25. Kühn, M. J. et al. Assessment of effective mitigation and prediction of the spread of SARS-CoV-2 in Germany using demographic information and spatial resolution. *Math. Biosci.* 108648 (2021).
26. Koslow, W. et al. Appropriate relaxation of non-pharmaceutical interventions minimizes the risk of a resurgence in SARS-CoV-2 infections in spite of the Delta variant. *PLoS Comput. Biol.* **18**, e1010054 (2022).
27. Angione, C., Silverman, E. & Yaneske, E. Using machine learning as a surrogate model for agent-based simulations. *PLoS ONE* **17**, e0263150 (2022).
28. Robertson, C., Saffa, C., Collier, N., Ozik, J. & Ray, J. Bayesian calibration of stochastic agent based model via random forest (2024). [ArXiv:2406.19524](https://arxiv.org/abs/2406.19524) [cs, stat].
29. Eurostat. Nuts 2024. <https://ec.europa.eu/eurostat/de/web/gisco/geodata/statistical-units/territorial-units-statistic>. Accessed: 2024-10-10.
30. Gori, M., Monfardini, G. & Scarselli, F. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005* vol. 2, 729–734 vol. 2 (2005). ISSN: 2161-4407.
31. Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M. & Monfardini, G. The graph neural network model. *IEEE Trans. Neural Netw.* **20**, 61–80 (2009) (Conference Name: IEEE Transactions on Neural Networks).
32. Lam, R. et al. Learning skillful medium-range global weather forecasting. *Science* **382**, 1416–1421 (2023).
33. Deng, S., Wang, S., Rangwala, H., Wang, L. & Ning, Y. Cola-gnn: Cross-location attention based graph neural networks for long-term ili prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, 245–254 (Association for Computing Machinery, 2020).
34. Kapoor, A. et al. Examining COVID-19 Forecasting using Spatio-Temporal Graph Neural Networks (2020).
35. Gao, J. et al. STAN: Spatio-temporal attention network for pandemic prediction using real-world evidence. *J. Am. Med. Inform. Assoc.* **28**, 733–743 (2021).
36. Panagopoulos, G., Nikolentzos, G. & Vazirgiannis, M. Transfer graph neural networks for pandemic forecasting. *Proc. AAAI Conf. Arti. Intell.* **35**, 4838–4845 (2021) (Number: 6).
37. Fritz, C., Dorigatti, E. & Rügamer, D. Combining graph neural networks and spatio-temporal disease models to improve the prediction of weekly covid-19 cases in germany. *Sci. Rep.* **12** (2022).
38. Robert Koch Institute (RKI). ControlCOVID: Strategie und Handreichung zur Entwicklung von Stufenkonzepten bis Frühjahr 2021. Tech. Rep., Robert Koch Institute (2021). [https://www.rki.de/DE/Content/InfAZ/N/Neuartiges\\_Coronavirus/Downloads/Stufenplan-Fruhjahr21.pdf?\\_blob=publicationFile](https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Downloads/Stufenplan-Fruhjahr21.pdf?_blob=publicationFile). Stand 19.03.2021, Änderung: Grenzwerte des Indikators Hospitalisierungsinzidenz.
39. Walsh, K. A. et al. The duration of infectiousness of individuals infected with SARS-CoV-2. *J. Infect.* **81**, 847–856 (2020).
40. Puhach, O., Meyer, B. & Eckerle, I. SARS-CoV-2 viral load and shedding kinetics. *Nat. Rev. Microbiol.* **21**, 147–161 (2023) (Publisher: Nature Publishing Group).
41. Bicker, J. et al. MEmilio - A high performance Modular EpideMIcs simuLatiOn software for multi-scale and comparative simulations of infectious disease dynamics (2026). To be submitted.
42. Regionaldatenbank Deutschland. Fortschreibung des Bevölkerungsstandes: 12411-04-02-4-b Bevölkerung nach Geschlecht und Altersjahren (79) - Stichtag 31.12. - (ab 2011) regionale Ebenen (2024). Key date used: 31.12.2020.
43. Bundesagentur für Arbeit. Pendlerverflechtungen der sozialversicherungspflichtig Beschäftigten nach Kreisen - Deutschland (Jahreszahlen) (2020). [https://statistik.arbeitsagentur.de/SiteGlobals/Forms/Suche/Einzelheftsuche\\_Formular.html?nn=24390&opic\\_f=beschaeftigung-pendler-krpend](https://statistik.arbeitsagentur.de/SiteGlobals/Forms/Suche/Einzelheftsuche_Formular.html?nn=24390&opic_f=beschaeftigung-pendler-krpend).
44. Chollet, F. et al. Keras (2015). <https://keras.io>.

45. Abadi, M. *et al.* TensorFlow: Large-scale machine learning on heterogeneous systems (2015). <https://www.tensorflow.org/>. Software available from tensorflow.org.
46. Grattarola, D. & Alippi, C. Graph Neural Networks in TensorFlow and Keras with Spektral [Application Notes]. *IEEE Comput. Intell. Mag.* **16**, 99–106 (2021).
47. Fey, M. & Lenssen, J. E. Fast Graph Representation Learning with PyTorch Geometric (2019). [https://github.com/pyg-team/pytorch\\_geometric](https://github.com/pyg-team/pytorch_geometric). Original-date: 2017-10-06T16:03:03Z.
48. Wang, M. *et al.* Deep Graph Library: A Graph-Centric, Highly-Performant Package for Graph Neural Networks (2020). [ArXiv:1909.01315](https://arxiv.org/abs/1909.01315) [cs, stat].
49. Wu, Z. *et al.* A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **32**, 4–24 (2021).
50. Kipf, T. N. & Welling, M. Semi-Supervised Classification with Graph Convolutional Networks (2017). [ArXiv:1609.02907](https://arxiv.org/abs/1609.02907) [cs, stat].
51. Bianchi, F. M., Grattarola, D., Livi, L. & Alippi, C. Graph neural networks with convolutional ARMA filters. *IEEE Trans. Pattern Anal. Mach. Intell.* **PP**, 1–1 (2021).
52. Gasteiger, J., Bojchevski, A. & Günnemann, S. Predict then Propagate: Graph Neural Networks meet Personalized PageRank (2022). [ArXiv:1810.05997](https://arxiv.org/abs/1810.05997) [cs, stat].
53. Veličković, P. *et al.* Graph Attention Networks (2018). [ArXiv:1710.10903](https://arxiv.org/abs/1710.10903) [cs, stat].
54. Kühn, M. J. *et al.* Regional opening strategies with commuter testing and containment of new SARS-CoV-2 variants in Germany. *BMC Infect. Dis.* **22**, 333 (2022).
55. Calzolari, G. & Liu, W. Deep learning to replace, improve, or aid CFD analysis in built environment applications: A review. *Build. Environ.* **206**, 108315 (2021).
56. Sun, L., Gao, H., Pan, S. & Wang, J.-X. Surrogate modeling for fluid flows based on physics-constrained deep learning without simulation data. *Comput. Methods Appl. Mech. Eng.* **361**, 112732 (2020).
57. Shi, R., Mo, Z., Huang, K., Di, X. & Du, Q. A physics-informed deep learning paradigm for traffic state and fundamental diagram estimation. *IEEE Trans. Intell. Transp. Syst.* **23**, 11688–11698 (2022).
58. Xie, T. & Grossman, J. C. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys. Rev. Lett.* **120**, 145301 (2018) (Publisher: American Physical Society).
59. Maguire, J. B., Grattarola, D., Mulligan, V. K., Klyshko, E. & Melo, H. XENet: Using a new graph convolution to accelerate the timeline for protein design on quantum computers. *PLoS Comput. Biol.* **17**, e1009037 (2021).
60. Bicker, J. *et al.* MEMilio v2.1.0 - A high performance Modular EpideMIcs simulatIOn software. *zenodo* (2025). <https://eliv.dlr.de/209739/>; Release version 2.1.0 code: <https://doi.org/10.5281/zenodo.16684608>.
61. Schmidt, A., Zunker, H., Heinlein, A. & Kühn, M. J. Supplementary code to graph neural network surrogates to leverage mechanistic expert knowledge towards reliable and immediate pandemic response (2025). <https://doi.org/10.5281/zenodo.15055792>.

## Acknowledgements

The authors thank Martin Siggel for the initial idea on developing AI surrogate models. AS thanks Simon Volland for exchange on neural network models and technical considerations to run them. This work was supported by the Initiative and Networking Fund of the Helmholtz Association (grant agreement number KA1-Co-08, Project LOKI-Pandemics), by the German Federal Ministry for Digital and Transport under grant agreement FKZ19F2211A (Project PANDEMOS) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) (grant agreement 528702961).

## Author contributions

Conceptualization: AH, MK. Data Curation: AS, HZ. Formal Analysis: AS, HZ, AH, MK. Funding Acquisition: MK. Investigation: AS, HZ, AH, MK. Methodology: AS, HZ, AH, MK. Project Administration: MK. Resources: MK. Software: AS, HZ. Supervision: AH, MK. Validation: AS, HZ, AH, MK. Visualization: AS, HZ, MK. Writing – Original Draft: AS, HZ, AH, MK. Writing – Review & Editing: AS, HZ, AH, MK.

## Funding

Open Access funding enabled and organized by Projekt DEAL. This work was supported by the Initiative and Networking Fund of the Helmholtz Association (grant agreement number KA1-Co-08, Project LOKI-Pandemics), by the German Federal Ministry for Digital and Transport (grant agreement FKZ19F2211A, Project PANDEMOS), and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) (grant agreement 528702961). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

## Declarations

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-026-39431-5>.

**Correspondence** and requests for materials should be addressed to M.J.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026