
Modeling COVID-19 Dynamics in German States Using Physics-Informed Neural Networks

Phillip Rothenbeck

Sai Karthikeya Vemuri

Niklas Penzel

Joachim Denzler

Computer Vision Group, Friedrich Schiller University Jena, Germany,
firstname.lastname@uni-jena.de

Abstract

We present a three-year, state-level analysis of COVID-19 in Germany using a simple SIR model coupled with Physics-Informed Neural Networks (PINNs), trained on data from the Robert Koch Institute. For each of the 16 German federal states, we use the PINN framework to estimate transmission and recovery rates (β, α) and the time-dependent reproduction number $\mathcal{R}_t$. Our results showcase significant regional heterogeneity and an inverse relationship between vaccination uptake and both β and peak $\mathcal{R}_t$ numbers. Furthermore, we observe that the inferred progression of $\mathcal{R}_t$ aligns with the major phases of the pandemic, including the Omicron peak, followed by stabilization at or below the epidemic threshold of 1.0 by mid-2022. These findings demonstrate the utility of PINNs for localized, long-term epidemiological modeling and evaluating regional policy impacts.

1 Introduction

The COVID-19 pandemic exhibited significant regional heterogeneity in its progression and management, as observed globally and within federal systems, such as Germany [1, 2, 3]. Understanding these state-level dynamics is critical for evaluating public health responses and informing future strategies. Such granular analyses are essential for tailoring public health policies and allocating resources effectively, moving beyond a one-size-fits-all approach [4]. Epidemiological analysis often relies on compartmental models, such as the Susceptible-Infectious-Recovered (SIR) framework [5], which have been widely adapted for COVID-19 [6, 7, 8]. However, classical methods for parameter estimation struggle with noisy observational data and capturing time-varying dynamics.

Physics-Informed Neural Networks (PINNs) [9] provide a robust framework for addressing these challenges by directly embedding differential equations into the neural network’s loss function, thereby enabling parameter estimation from sparse or noisy data, e.g., [10, 11]. Previous work has applied PINNs to epidemiology, often focusing on more complex compartmental models [12, 13, 14] or different geographical regions and timeframes [15, 16]. For instance, studies on Germany have analyzed the nation as a whole or focused on the early stages of the pandemic [13, 14].

In contrast, we perform a high-resolution spatio-temporal analysis of the pandemic across all 16 German federal states over a 1,200-day period (March 2020 to June 2023). Using public data from the Robert Koch Institute [17, 3], we employ a PINN-based approach to solve the inverse problem for an SIR model. First, we estimate the state-specific transmission rate, β , and recovery rate, α . Subsequently, we compute the time-dependent reproduction number, $\mathcal{R}_t$, for each state, providing granular insights into the evolving transmission dynamics. Our analysis reveals substantial regional disparities. We find that variations in state-level vaccination rates correlate with both the estimated transmission rates β and the observed peak reproduction numbers (see Apx. A.4.2), demonstrating that local factors and interventions had a measurable impact on the pandemic’s trajectory.

2 Methodology

We employ Physics-Informed Neural Networks (PINNs) [9] to solve the inverse problem of estimating key parameters of compartmental models from observational data of the 16 German federal states. Using this framework, we train neural networks to simultaneously fit observed infections and adhere to the physical constraints imposed by the governing ordinary differential equations (ODEs) [12]. By embedding the ODEs into the loss function, the unknown model parameters, such as transmission and recovery rates, become trainable variables that are inferred directly during the optimization. We include a formal introduction in Apx. A.1. We use the following epidemiological models:

First, standard SIR (Susceptible - Infected - Recovered) [5, 18, 19] for a population of N :

$$\frac{dS}{dt} = -\beta \frac{SI}{N}, \quad \frac{dI}{dt} = \beta \frac{SI}{N} - \alpha I, \quad \frac{dR}{dt} = \alpha I, \quad (1)$$

where β is the *transmission rate* and α is the *recovery rate*.

Second, a time-rescaled form [16] with constant α , $t_s = (t-t_0)/(t_f-t_0)$, and $I(t) = c \cdot I_s(t_s)$:

$$\frac{dI_s}{dt_s} = \alpha(t_f - t_0)(\mathcal{R}_t - 1)I_s(t_s), \quad (2)$$

where $\mathcal{R}_t$ is the *time-dependent reproduction number*. This number is crucial because an outbreak is expanding if $\mathcal{R}_t > 1$ and declining if $\mathcal{R}_t < 1$ [1, 16]. Next, we detail the training objectives.

2.1 Time-Independent Parameter Identification (α, β)

We identify global α and β rates by training a PINN to fit observational data and satisfy Eq. (1) [12]. The data fidelity loss (MSE over T days) is

$$\mathcal{L}_{\text{data}} = \frac{1}{T} \sum_{t=1}^T \left(\left\| \hat{S}^{(t)} - S^{(t)} \right\|^2 + \left\| \hat{I}^{(t)} - I^{(t)} \right\|^2 + \left\| \hat{R}^{(t)} - R^{(t)} \right\|^2 \right). \quad (3)$$

Physics consistency is enforced by the residuals of Eq. (1):

$$\mathcal{L}_{\text{physics}} = \left\| \frac{d\hat{S}}{dt} + \tilde{\beta} \frac{\hat{S}\hat{I}}{N} \right\|^2 + \left\| \frac{d\hat{I}}{dt} - \tilde{\beta} \frac{\hat{S}\hat{I}}{N} + \tilde{\alpha} \hat{I} \right\|^2 + \left\| \frac{d\hat{R}}{dt} - \tilde{\alpha} \hat{I} \right\|^2. \quad (4)$$

In Eq. (4), we reparameterize the rates via $\tilde{\beta} = \tanh(\beta)$ and $\tilde{\alpha} = \tanh(\alpha)$ to stabilize the inverse problem [20]. The overall training objective is $\mathcal{L}_{\text{SIR}} = \mathcal{L}_{\text{data}} + \mathcal{L}_{\text{physics}}$.

2.2 Time-Dependent $\mathcal{R}_t$ Estimation

Following [16], we estimate $\mathcal{R}_t$ using the reduced model in Eq. (2). The PINN takes t (days) as an input and predicts $(\hat{I}^{(t)}, \mathcal{R}_t)$, with data loss

$$\mathcal{L}_{\text{data}} = \frac{1}{T} \sum_{t=1}^T \left\| \hat{I}^{(t)} - I^{(t)} \right\|^2, \quad (5)$$

and physics loss using the residuals of Eq. (2):

$$\mathcal{L}_{\text{physics}} = \left\| \frac{d\hat{I}}{dt} - \alpha(t_f - t_0)(\mathcal{R}_t - 1)\hat{I} \right\|^2. \quad (6)$$

Training is split into two stages: first we minimize Eq. (5); then we minimize $\mathcal{L}_{\text{tSIR}} = w_0 \mathcal{L}_{\text{data}} + w_1 \mathcal{L}_{\text{physics}}$ (with loss weights w_0 and w_1) to balance data fit and model conformity [16].

3 Data Collection and Experimental Setup

Our analysis utilizes public infection data from the Robert Koch Institute (RKI) [3, 17], spanning 1200 days from March 2020 to June 2023 across all 16 German federal states. Since explicit recovery data is unavailable, we model a recovery queue to estimate the removed population. To validate our




Figure 1: **State-level results over 1,200 days (Mar 9, 2020–Jun 22, 2023).** (a) Estimated recovery rates α and transmission rates β per state from the SIR–PINN. We visualize mean and standard deviation per state, alongside national means. (b) Time-varying reproduction number $\mathcal{R}_t$ for Germany and selected states, with vaccine rollout and Alpha/Delta/Omicron peaks marked [24]. Full $\mathcal{R}_t$ plots and correlation to vaccination percentages for all states are provided in the Apxs. A.4.1 to A.4.3.

PINN-based framework, we first reproduce results from a classical ODE solver on early pandemic data [8], confirming the reliability of our approach before applying it to the full dataset (see Apx. A.2).

In our first analysis, we estimate the global transmission rate β and recovery rate α for each state. We solve the inverse problem for the standard SIR model using a PINN composed of seven hidden layers with 20 neurons each and hyperbolic tangent activations [21]. Following established practices [12], the network is trained for 10,000 iterations with an Adam optimizer [22], treating β and α as trainable parameters.

In our second analysis, we infer the time-dependent reproduction number $\mathcal{R}_t$ to capture evolving transmission dynamics. Following [16], we perform a two-stage training process totaling 50,000 iterations with balanced loss weights to ensure convergence. For this task, the network architecture uses ReLU activations [23], and we conduct the experiment using both a standard, fixed recovery rate ($\alpha = 1/14$) and the state-specific rates inferred in our first analysis.

Both experiments are repeated ten times to increase the robustness of our results. The full details on data preprocessing, specific loss weights, and hyperparameters are provided in the Apx. A.3.

4 Results and Discussion

Our main contribution is a 1,200-day, state-level temporal analysis of COVID-19 in Germany using the PINN framework. This analysis reveals persistent regional differences and long-term trends that are obscured in national aggregates. We first analyze time-independent transmission and recovery rates before examining the time-dependent reproduction number $\mathcal{R}_t$. We include extended discussions and visualizations in the Apxs. A.4.1 to A.4.3.

Time-Independent Parameter Identification Our time-independent analysis reveals significant heterogeneity in the average transmission rate β and recovery rate α across states (see Fig. 1a). States with lower vaccination rates, such as Thuringia ($\beta = 0.127$) and Saxony-Anhalt ($\beta = 0.120$), exhibit the highest transmission, while states with higher uptake, e.g., Bremen ($\beta = 0.083$), show the lowest. This relationship is quantified by a significant negative correlation between β and vaccination percentages ($r = -0.5741, p = 0.02$). These findings reveal a broader geographical pattern, with eastern states often showing higher transmission dynamics than most western and northern states. Recovery rates α largely mirror this pattern, suggesting faster turnover where transmission was more intense, while generally remaining near the 14-day recovery ($\alpha \approx 0.0714$) as noted by WHO [25].

Time-Dependent Reproduction number $\mathcal{R}_t$ To understand the temporal dynamics underlying these averages, we estimated the time-dependent reproduction number $\mathcal{R}_t$. The analysis reveals variations in outbreak severity and persistence (see Table 1 and Fig. 1b). For instance, Thuringia,

Table 1: Average number of days with $\mathcal{R}_t > 1$, and the average peak $\mathcal{R}_t$ values for all German states (MWP=Mecklenburg-Western Pomerania, NRW=North Rhine-Westphalia) and Germany, for $\alpha = 1/14$ and α_{exp} (see Table 2 in Apx. A.4.1). The vaccination percentage provided by the German Fed. Ministry for Health [26].

State name	Vaccinations [%]	days with $\mathcal{R}_t > 1$		peak $\mathcal{R}_t$	
		$\alpha = 1/14$	α_{exp}	$\alpha = 1/14$	α_{exp}
Germany	76.4	312.0	301.5	1.643	1.705
Schleswig-Holstein	79.5	352.1	355.6	1.525	1.441
Hamburg	84.5	398.3	316.0	1.689	1.577
Lower Saxony	77.6	327.9	298.4	1.637	1.682
Bremen	88.3	326.1	402.9	1.508	1.525
NRW	79.5	280.1	316.8	1.954	1.789
Hesse	75.8	344.1	308.8	1.774	1.750
Rhineland-Palatinate	75.6	341.7	335.5	1.582	1.515
Baden-Württemberg	74.5	372.2	307.0	1.617	1.608
Bavaria	75.1	342.9	321.2	1.719	1.532
Saarland	82.4	388.1	338.9	1.495	1.547
Berlin	78.1	304.7	305.7	1.686	1.485
Brandenburg	68.1	380.2	376.6	1.795	1.466
MWP	74.7	399.8	327.9	1.645	1.375
Saxony	65.1	368.1	368.9	1.696	1.523
Saxony-Anhalt	74.1	345.8	335.9	1.706	1.424
Thuringia	70.3	373.7	387.2	1.959	1.429

which had the highest estimated β , also reached the highest peak $\mathcal{R}_t$ of 1.96. Reinforcing the trend observed in our time-independent analysis, we find a negative correlation between peak $\mathcal{R}_t$ and the state vaccination rates ($r = -0.4455, p = 0.079$). Although this trend does not meet the 0.05 threshold for statistical significance, its directionality supports the conclusion that local interventions had a measurable impact. The inferred $\mathcal{R}_t$ captures key pandemic phases, culminating in a stabilization at or below the epidemic threshold of 1.0 by mid-2022 (see Apx. A.4.3).

5 Conclusions

We presented a 1200-day, state-level analysis of COVID-19 in Germany, using Physics-Informed Neural Networks (PINNs) [9] to solve the inverse problem for SIR-based models [5, 16] from public RKI data [3, 17]. Our analysis successfully inferred state-specific transmission rates β and recovery rates α , as well as the time-dependent reproduction number $\mathcal{R}_t$. The central finding is a significant negative correlation between state-level vaccination rates and transmission rates, indicating a clear and measurable link between public health interventions and pandemic dynamics. The temporal analysis of $\mathcal{R}_t$ further contextualized these findings, capturing key pandemic phases, such as the Omicron wave, before stabilizing below the epidemic threshold of 1.0. Our work demonstrates that PINNs, when paired with simple, interpretable models, are a powerful tool for robust, long-term post-hoc analysis of regional policy impacts.

Limitations Our choice of simpler SIR models prioritizes interpretability but does not explicitly capture complex dynamics such as age structure, waning immunity, or mobility patterns. The analysis is further constrained by the fidelity of public data, including potential under-ascertainment and the need to approximate recovery information. This methodological focus, however, makes our framework directly applicable to the most consistently reported data during the pandemic, daily infection counts [3, 17].

Outlook Future work can build upon this framework by relaxing these assumptions while retaining the state-level focus. Incorporating more complex compartmental models (e.g., SEIR, SAIRD, or SVIHR) [7, 14, 12] or integrating covariates, such as mobility data and non-pharmaceutical interventions, could provide more granular insights. Applying the PINN methodology to these more complex models promises to enhance forecasting capabilities for future public health crises.

References

- [1] Sayali Bhatkar, Mingyang Ma, Mary Zsolway, Ayush Tarafder, Sebastian Doniach, and Gyan Bhanot. Asymmetry in the peak in covid-19 daily cases and the pandemic r-parameter. *medRxiv*, 2023.
- [2] Joseph E. Gagnon, Steven B. Kamin, and John Kearns. The impact of the covid-19 pandemic on global gdp growth. *Journal of the Japanese and International Economies*, 68:101258, 2023.
- [3] Robert Koch Institute. Sars-cov-2: Virologische basisdaten sowie virusvarianten im zeitraum von 2020 - 2022. https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Virologische_Basisdaten.html?nn=13490888#doc14716546bodyText10. Accessed: 2024-09-05.
- [4] Thomas Hale, Noam Angrist, Rafael Goldszmidt, Beatriz Kira, Anna Petherick, Toby Phillips, Samuel Webster, Emily Cameron-Blake, Laura Hallas, Saptarshi Majumdar, et al. A global panel database of pandemic policies (oxford covid-19 government response tracker). *Nature human behaviour*, 5(4):529–538, 2021.
- [5] William Ogilvy Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721, August 1927.
- [6] Robert Schaback. On covid-19 modelling. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 122(3):167–205, Sep 2020.
- [7] Azizur Rahman, Md Abdul Kuddus, Ryan HL Ip, and Michael Bewong. A review of covid-19 modelling strategies in three countries to develop a research framework for regional areas. *Viruses*, 13(11):2185, 2021.
- [8] Günter Bärwolff. Modeling of covid-19 propagation with compartment models. *Mathematische Semesterberichte*, 68(2):181–219, Oct 2021.
- [9] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019.
- [10] Sai Karthikeya Vemuri, Tim Büchner, and Joachim Denzler. Estimating soil hydraulic parameters for unsaturated flow using physics-informed neural networks. In Leonardo Franco, Clélia de Mulatier, Maciej Paszynski, Valeria V. Krzhizhanovskaya, Jack J. Dongarra, and Peter M. A. Sloot, editors, *Computational Science – ICCS 2024*, pages 338–351, Cham, 2024. Springer Nature Switzerland.
- [11] Gideon Stein, Sai Karthikeya Vemuri, Yuanyuan Huang, Anne Ebeling, Nico Eisenhauer, Maha Shadaydeh, and Joachim Denzler. Investigating the effects of plant diversity on soil thermal diffusivity using physics-informed neural networks. In *ICLR Workshop on AI4DifferentialEquations In Science (ICLR-WS)*, 2024.
- [12] Sagi Shaier, Maziar Raissi, and Padmanabhan Seshaiyer. Data-driven approaches for predicting spread of infectious diseases through dinns: Disease informed neural networks. *Letters in Biomathematics*, 9(1):71–105, 2023.
- [13] Sarah Berkahn and Matthias Ehrhardt. A physics-informed neural network to model covid-19 infection and hospitalization scenarios. *Advances in Continuous and Discrete Models*, 2022(1), October 2022.
- [14] Shuai Han, Lukas Stelz, Horst Stoecker, Lingxiao Wang, and Kai Zhou. Approaching epidemiological dynamics of covid-19 with physics-informed neural networks. *Journal of the Franklin Institute*, 361(6):106671, 2024.
- [15] K. D. Olumoyin, A. Q. M. Khaliq, and K. M. Furati. Data-driven deep-learning algorithm for asymptomatic covid-19 model with varying mitigation measures and transmission rate. *Epidemiologia*, 2(4):471–489, September 2021.
- [16] Caterina Millevoi, Damiano Pasetto, and Massimiliano Ferronato. A physics-informed neural network approach for compartmental epidemiological models. *PLOS Computational Biology*, 20(9):e1012387, 2024.
- [17] Robert Koch Institute. Covid-19-strategiepapiere und nationaler pandemieplan. https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/ZS/Pandemieplan_Strategien.html. Accessed: 2024-09-06.
- [18] Herbert W. Hethcote. *Three Basic Epidemiological Models*, pages 119–144. Springer Berlin Heidelberg, Berlin, Heidelberg, 1989.

- [19] S. Mwalili, M. Kimathi, V. Ojiambo, D. Gathungu, and R. Mbogo. Seir model for covid-19 dynamics incorporating the environment and social distancing. *BMC Res Notes*, 13(1):352, 2020.
- [20] Shiv Ram Dubey, Satish Kumar Singh, and Bidyut Baran Chaudhuri. Activation functions in deep learning: A comprehensive survey and benchmark. *Neurocomputing*, 503:92–108, 2022.
- [21] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [22] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014.
- [23] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *International Conference on Machine Learning*, 2010.
- [24] German Federal Ministry of Health. Coronavirus-pandemie: Was geschah wann? <https://www.bundesgesundheitsministerium.de/coronavirus/chronik-coronavirus.html>. Accessed: 2024-09-05.
- [25] World Health Organization. Coronavirus disease (covid-19). https://www.who.int/health-topics/coronavirus#tab=tab_1. Accessed: 2024-09-06.
- [26] German Federal Ministry of Health. Übersicht zum impfstatus - covid-19-impfung in deutschland bis zum 8. april 2023. <https://impfdashboard.de/>. Accessed: 2024-09-08.
- [27] Victor Isakov. *Inverse Problems for Partial Differential Equations*. Springer Publishing Company, Incorporated, 3rd edition, 2018.
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [29] Robert Koch Institute. Github sars-cov-2 infektionen in deutschland. https://github.com/robert-koch-institut/SARS-CoV-2-Infektionen_in_Deutschland. Accessed: 2024-09-05.
- [30] Robert Koch Institute. Github covid-19-todesfälle in deutschland. https://github.com/robert-koch-institut/COVID-19-Todesfaelle_in_Deutschland. Accessed: 2024-09-05.
- [31] Eurostat. Population on 1 january by age, sex and nuts 2 region. <https://ec.europa.eu/eurostat/databrowser/bookmark/6b9a13dc-185a-4b23-b52a-5afae4a4e28c?lang=en>. Accessed: 2025-02-26.

A Appendix

A.1 Theoretical Foundations — PINNs for Inverse Problems

In inverse problems, the goal is to infer unknown system parameters (e.g., coefficients, boundary/initial conditions in ODEs/PDEs) from partial or noisy observations of the system’s behavior [27]. Physics-Informed Neural Networks (PINNs) [9] address these by embedding the governing differential equations directly into the neural network’s training process.

A.1.1 Mathematical Formalization:

Consider a system described by differential equations over a spatial domain Ω and time $t \in [0, T]$:

$$\mathcal{F}(u(\mathbf{x}, t), \nabla u(\mathbf{x}, t), \nabla^2 u(\mathbf{x}, t), \dots | \lambda) = 0, \quad \text{in } \Omega \times [0, T], \quad (7)$$

where $u(\mathbf{x}, t)$ is the variable of interest (e.g., number of infected people), $\mathcal{F}$ is a differential operator, and λ are unknown parameters to be inferred. Given observations $\{(\mathbf{x}_i, t_i, u_i)\}$ of the system, the inverse problem is to estimate λ such that $u(\mathbf{x}, t)$ satisfies both the physical laws (Eq. (7)) and the observed data.

PINNs approximate u with a neural network $u_\theta(\mathbf{x}, t)$ (where θ are the trainable parameters) and learn λ by minimizing a composite loss function [9, 10].

A.1.2 Loss Function:

The total loss function for an inverse problem combines data fidelity and physics consistency and is defined as $\mathcal{L}(\theta, \lambda) = \mathcal{L}_{\text{data}}(\theta) + \mathcal{L}_{\text{physics}}(\theta, \lambda)$. The *data loss* $\mathcal{L}_{\text{data}}$ ensures that u_θ matches observations:

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{N} \sum_{i=1}^N |u_\theta(\mathbf{x}_i, t_i) - u_i|^2. \quad (8)$$

The *physics loss* $\mathcal{L}_{\text{physics}}$ penalizes deviations from the governing differential equations (Eq. (7)):

$$\mathcal{L}_{\text{physics}}(\theta, \lambda) = \frac{1}{M} \sum_{j=1}^M |\mathcal{F}(u_\theta(\mathbf{x}, t), \nabla u_\theta(\mathbf{x}, t), \nabla^2 u_\theta(\mathbf{x}, t), \dots | \lambda)|^2. \quad (9)$$

The unknown PDE parameters λ are treated as trainable variables alongside the neural network weights θ , optimized by minimizing $\mathcal{L}(\theta, \lambda)$. Additional balancing of the loss terms using scalar hyperparameters is possible.

A.1.3 Compartmental Models for Epidemiology and Inverse Problem

Compartmental models are the foundation of mathematical epidemiology [18, 5, 19] to study the spread of diseases. The SIR model [5] partitions a population of size N into three distinct compartments:

- **Susceptible (S)** individuals at risk of infection.
- **Infected (I)** individuals capable of transmitting the disease.
- **Removed (R)** individuals recovered with immunity or deceased.

The evolution of these compartments over time is governed by the following system of ordinary differential equations (ODEs):

$$\frac{dS}{dt} = -\beta \frac{SI}{N}, \quad \frac{dI}{dt} = \beta \frac{SI}{N} - \alpha I, \quad \frac{dR}{dt} = \alpha I, \quad (10)$$

where β is the **transmission rate** and α is the **recovery rate**.

Real-world infectious diseases exhibit dynamic behavior due to changing intervention strategies, population immunity, and viral mutations. A key metric to measure this is the effective reproduction number $\mathcal{R}_t$ [7], quantifying the average secondary infections from one infected individual at time t :

$$\mathcal{R}_t = \frac{\beta(t)}{\alpha(t)} \cdot \frac{S(t)}{N}. \quad (11)$$

Specifically, an outbreak is expanding if $\mathcal{R}_t > 1$ and declining if $\mathcal{R}_t < 1$ [1, 16].

To model time-varying dynamics, Millevoi et al. [16] reformulated SIR using a rescaled time-dependent formulation, assuming a constant recovery rate α . Let t_s be a normalized time variable $t_s = (t - t_0)/(t_f - t_0)$, for an interval $t \in [t_0, t_f]$. Then the infected compartment $I(t)$ is scaled by a constant c as $I(t) = c \cdot I_s(t_s)$, where the dynamics of the scaled infected compartment I_s are given by:

$$\frac{dI_s}{dt_s} = \alpha(t_f - t_0)(\mathcal{R}_t - 1)I_s(t_s). \quad (12)$$

A.2 Reproducing Pandemic Parameters using PINNs

In order to validate our method, we reproduced the results of a traditional method on real-world data. Bärwolff et al. [8] employ the damped Gauss-Newton method to find β from a time series of data points. Furthermore, they provide the time points, which they used to derive β for the time span between February 13, 2020, and March 19, 2020, together with the corresponding $\beta_{\text{true}} = 0.22658$.

Just like the original study, we set $\alpha = 0.07$ for the experiment. Our model consists of 12 hidden layers with 64 neurons each and hyperbolic tangent [21] activation layers. We trained using a




Figure 2: Visualization of the real numbers of infectious individuals in Germany and the prediction of the training with $\alpha = 0.07$ for the time span between February 13, 2020, and March 19, 2020.

polynomial scheduler from PyTorch [28], the Adam optimizer [22], an initial learning rate of $1e^{-3}$, and 15K iterations. The data loss is weighted by $1e^1$ in the total loss.

As the provided data consists only of infectious data, we generated the corresponding data for the susceptible and removed compartments by utilizing Eq. (1). The population size is $N = 70M$ individuals, and the provided initial amount of infectious individuals is $I_0 = 15$.

In Fig. 2, we visualize our results from the training, demonstrating that the model successfully fits the predictions to the observed data. We repeated the experiment ten times, which resulted in a mean of $\beta_{\text{PINN}} = 0.22822$ and a standard deviation of $\sigma_{\text{PINN}}^2 = 1.03367 \times 10^{-5}$. A comparison with the provided value of $\beta_{\text{true}} = 0.22658$ indicates that our method has a satisfying accuracy. Hence, we will employ PINNs to investigate longer time frames and regional variations in our main study.

A.3 Experimental Setup Details

We use public Robert Koch Institute (RKI) infection data [3, 17], preprocessing raw infections [29] per federal state and German death cases [30] separately. Lacking explicit recovery data, we model a recovery queue to transition infected individuals to the removed group, aligning with typical recovery periods noted by WHO [25]. We use state population sizes from 2020 [31], and the initial number of infectious individuals is taken from the original cases recorded on March 9, 2020. Our analysis spans March 9, 2020, to June 22, 2023 (1200 days), covering the most active phases of the COVID-19 pandemic [3, 17].

To estimate pandemic parameters for each German state, we employ Physics-Informed Neural Networks (PINNs) to fit SIR models to observed case data, as detailed in Apx. A.1. Crucially, before applying this framework to the extended RKI data, we validate its fundamental capabilities. Thus, we first replicate an experiment from [8], where Germany’s early pandemic dynamics were analyzed using a classical dampened Gauss-Newton method. Using this approach, they approximate the transmission rate β as 0.22658, whereas our PINN-based approach yields a consistent result of 0.22822 on their data. This replication confirms that **our PINN-based approach reproduces results obtained by established PDE solvers**, providing confidence in its application to more complex scenarios. We include the full details of this validation in the supplementary material.

With our methodology validated, the core contribution of our work is the fine-grained, spatiotemporal analysis of COVID-19 dynamics across all German federal states over a three-year period. Our focus is thus not on comparative benchmarking with other solvers, which often operate on synthetic or short-term data. Instead, we aim to demonstrate the practical utility of PINNs for extracting detailed insights from extensive, real-world epidemiological data. To ensure the robustness of our findings,

Table 2: Pandemic parameter means and standard deviations for Germany, and each German state (MWP=Mecklenburg-Western Pomerania, NRW=North Rhine-Westphalia). Furthermore, we include the vaccination percentage provided by the German Fed. Ministry for Health [26] and corresponding population sizes [31].

State	N [10^6]	α	β	Vaccinations [%]
Germany	83.16	0.080 ± 0.000	0.104 ± 0.001	76.4
Schleswig-Holstein	2.90	0.077 ± 0.000	0.097 ± 0.000	79.5
Hamburg	1.84	0.084 ± 0.000	0.108 ± 0.001	84.5
Lower Saxony	7.99	0.074 ± 0.001	0.097 ± 0.001	77.6
Bremen	0.68	0.061 ± 0.000	0.083 ± 0.000	88.3
NRW	17.94	0.078 ± 0.000	0.100 ± 0.001	79.5
Hesse	6.29	0.066 ± 0.001	0.086 ± 0.001	75.8
Rhineland-Palatinate	4.08	0.079 ± 0.001	0.102 ± 0.001	75.6
Baden-Württemberg	11.07	0.086 ± 0.000	0.113 ± 0.001	74.5
Bavaria	13.10	0.080 ± 0.001	0.109 ± 0.002	75.1
Saarland	0.99	0.072 ± 0.000	0.098 ± 0.001	82.4
Berlin	3.67	0.090 ± 0.001	0.112 ± 0.001	78.1
Brandenburg	2.52	0.086 ± 0.001	0.109 ± 0.001	68.1
MWP	1.61	0.092 ± 0.000	0.118 ± 0.000	74.7
Saxony	4.07	0.081 ± 0.001	0.109 ± 0.001	65.1
Saxony-Anhalt	2.20	0.093 ± 0.000	0.120 ± 0.000	74.1
Thuringia	2.13	0.097 ± 0.001	0.127 ± 0.001	70.3

each experiment is repeated ten times per state. Further, we investigate the pandemic’s evolution under two separate paradigms:

A.3.1 Time-Independent Parameter Identification:

First, we estimate the transmission (β) and recovery (α) rates for the entire pandemic by optimizing a PINN to fit the SIR model and the observed infection data, as detailed in Sec. 2.1. Here α and β are trainable variables initialized within the PINN training process.

We employ a PINN architecture comprising seven hidden layers (each with 20 neurons) and hyperbolic tangent activations [21]. Adhering to hyperparameter settings from [12], the model is subsequently trained for 10K iterations using a 0.001 learning rate with a polynomial scheduler [28]. The complete training process is repeated ten times for each German state to obtain state-wise estimates of α and β .

A.3.2 Time-Dependent Reproduction Number ($\mathcal{R}_t$) Estimation:

Following [16], we estimate the time-dependent reproduction number $\mathcal{R}_t$, assuming a constant recovery rate of $\alpha = 1/14$ for normal conditions, as noted by WHO [25]. Additionally, we conduct a second experiment using the state-wise α_{exp} values determined in our first time-independent experiment.

We use the same PINN architecture, employing ReLU activations [23]. To ensure the model effectively learns the infection compartment, we first optimize only the data loss ($\mathcal{L}_{\text{data}}$) for 30K iterations. Next, we train using the combined loss $\mathcal{L}_{\text{rSIR}}$ for 20K iterations, initializing $\mathcal{R}_t$ from the time-independent experiment results. In $\mathcal{L}_{\text{rSIR}}$, we find balancing $\mathcal{L}_{\text{data}}$ and $\mathcal{L}_{\text{physics}}$ to be crucial for ensuring convergence given the different magnitudes of the loss terms. Specifically, we multiply the data loss by $w_0 = 10^2$ and scale the physics term by $w_1 = 1 \times 10^{-6}$. For the federal states, we observe improvements when increasing both weights to 10^3 and 4×10^{-6} respectively, due to the smaller population sizes. As before, each experiment is repeated ten times to ensure robustness.




Figure 3: Higher vaccination coverage coincides with lower pandemic effects. (Left): The correlation between the vaccination rate and the corresponding mean transmission rate β for each federal state. (Right): The correlation between the vaccination rate and the peak $\mathcal{R}_t$ value for each state.

A.4 Additional Results

A.4.1 Time-Independent Parameter Estimates

Table 2 presents results for our first setup, showing the estimated recovery rate α and transmission rate β for each German state, along with vaccination percentages as reported by the German Federal Ministry for Health [26]. The results highlight significant regional variations in pandemic dynamics.

A.4.2 Correlation Analysis

Our analysis reveals a negative correlation between state-level vaccination rates and key transmission metrics, as illustrated in Fig. 3. We computed a statistically significant Pearson correlation of $r = -0.5741$ ($p = 0.02$) between the estimated transmission rate β and the vaccination percentage for each state. A similar negative trend can be observed between the peak reproduction number $\mathcal{R}_t$ and vaccination rates ($r = -0.4455$, $p = 0.079$), although this result was not statistically significant at the 0.05 level. These regional differences suggest that local interventions had a measurable impact on transmission dynamics, highlighting the value of localized analyses.

A.4.3 Additional $\mathcal{R}_t$ Visualizations

In this section, we present all results of our experiments for the estimation of the time-dependent reproduction number $\mathcal{R}_t$ in Sec. 4. Specifically, we visualize the $\mathcal{R}_t$ trends in Figs. 4 to 6. We find distinct regional variations in transmission intensity and pandemic duration across the German federal states. Eastern states such as Saxony, Thuringia, and Saxony-Anhalt experienced prolonged periods where $\mathcal{R}_t > 1$, aligning with their high transmission rates and lower vaccination coverage. In contrast, northern states like Bremen, Schleswig-Holstein, and Lower Saxony exhibited lower peak values and shorter transmission durations, reflecting the effectiveness of their higher vaccination rates and public health measures. Southern states, including Bavaria and Baden-Württemberg, saw strong waves during Alpha and Delta but recovered faster post-Omicron, likely due to a combination of vaccine uptake and healthcare capacity. Western states, particularly North Rhine-Westphalia and Hesse, had moderate outbreaks but were able to manage transmission effectively, keeping $\mathcal{R}_t$ under control for longer periods. Berlin displayed higher-than-expected peak values despite strong vaccination efforts, likely influenced by its high population density and mobility patterns, whereas Brandenburg exhibited prolonged transmission, suggesting spillover effects from Berlin. These findings emphasize the importance of considering regional differences in pandemic response planning, as factors such as mobility, healthcare infrastructure, and policy measures played a significant role in shaping the trajectory of COVID-19 across Germany.




Figure 4: All visualizations of the $\mathcal{R}_t$ value from Sec. 4. (part 1)




Figure 5: All visualizations of the $\mathcal{R}_t$ value from Sec. 4. (part 2)




Figure 6: All visualizations of the $\mathcal{R}_t$ value from Sec. 4. (part 3)