



Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 22/24

Fabian Heldmann, Sarah Berkhahn, Matthias Ehrhardt,
and Kathrin Klamroth

**PINN Training using Biobjective Optimization:
The Trade-off between Data Loss and Residual Loss**

(This is an updated/corrected version of Preprint 22/13)

December 11, 2022

<http://www.imacm.uni-wuppertal.de>

PINN Training using Biobjective Optimization: The Trade-off between Data Loss and Residual Loss

Fabian Heldmann^b, Sarah Berkhahn^a, Matthias Ehrhardt^a, Kathrin Klamroth^b

^a *University of Wuppertal, Chair of Applied and Computational Mathematics,
Gaußstrasse 20, 42119 Wuppertal, Germany*

^b *University of Wuppertal, Chair of Optimization, Gaußstrasse 20, 42119 Wuppertal,
Germany*

Abstract

Physics informed neural networks (PINNs) have proven to be an efficient tool to represent problems for which measured data are available and for which the dynamics in the data are expected to follow some physical laws. In this paper, we suggest a multiobjective perspective on the training of PINNs by treating the data loss and the residual loss as two individual objective functions in a truly biobjective optimization approach.

As a showcase example, we consider COVID-19 predictions in Germany and built an extended susceptibles-infected-recovered (SIR) model with additionally considered leaky-vaccinated and hospitalized populations (SVIHR model) to model the transition rates and to predict future infections. SIR-type models are expressed by systems of ordinary differential equations (ODEs). We investigate the suitability of the generated PINN for COVID-19 predictions and compare the resulting predicted curves with those obtained by applying the method of non-standard finite differences to the system of ODEs and initial data.

The approach is applicable to various systems of ODEs that define dynamical regimes. Those regimes do not need to be SIR-type models, and the corresponding underlying data sets do not have to be associated with COVID-19.

Email addresses: heldmann@uni-wuppertal.de (Fabian Heldmann), streibert@uni-wuppertal.de (Sarah Berkhahn), ehrhardt@uni-wuppertal.de (Matthias Ehrhardt), klamroth@uni-wuppertal.de (Kathrin Klamroth)

Keywords: physics-informed neural networks, compartment models, loss function, multiobjective optimization, weighting parameters, Pareto front

1. Introduction

Physics informed neural networks (PINNs) [1, 2], also called theory-inspired machine learning [3], have recently become a popular method for solving differential equations. By incorporating the residual of the differential equation into the loss function of a neural network-based surrogate model, PINNs can seamlessly combine measured data with physical constraints given by differential equations. PINNs can also be viewed as a *surrogate model* for solving differential equations by incorporating additional data or as a data-driven correction (or even discovery) of the underlying physical system.

By the end of the year 2022, we had experienced several waves of the COVID-19 pandemic with different variants of the virus prevailing at different time intervals. Various levels of interventions and protective measures were implemented to counteract the uncontrolled spreading of the disease. We focus exemplarily on the time until the fourth wave (i.e., the omicron wave) of the COVID-19 pandemic in Germany that had its peak in February and March 2022.

The B.1.617.2 (delta) variant of SARS-CoV-2, which is characterized by a higher contagiousity than the previous B.1.1.7 (alpha), B.1.351 (beta) and P.1 (gamma) variants, has been observed in Germany since March 2021 and was the predominant variant in Germany during several months in the year 2021 [4]. In the autumn of 2021, the new omicron variant was detected and classified as concerning by the World Health Organization. Three sublines (BA.1, BA.2, BA.3) and a transmission advantage with respect to the delta variant were attributed to the omicron variant at the end of November 2021 [4]. The omicron variant quickly spread worldwide. Three recombinations of the omicron and delta variant (XD, XE, XF) had already been registered as sublines [4].

Local peaks during the fourth COVID-19 wave in Germany were reached on November 28th 2021 with 693, February 14th 2022 with 2,434 and March 20th

2022 with 2,619 daily infections per 1 million people. All of these peaks were larger than the global peaks of the three previous waves experienced in spring 2020 (69 on April 2nd 2020), winter 2020/2021 (305 on December 22nd 2020) and spring 2021 (257 on April 25th 2021) [5].

The mathematical model used in this work to describe the population dynamics of COVID-19 is the *susceptible-vaccinated-infected-hospitalized-recovered* (SVIHR) model. Here, pre-symptomatic individuals are merged with symptomatic people in the infected compartment, so that we have a single infected compartment of people not hospitalized.

The contribution of our work consists of two large parts: Firstly, we establish the SVIHR model and build upon this a *Physics-Informed Neural Network* (PINN). The PINN method uses certain physics-informed constraints, expressed e.g. by differential equations, as part of the loss function of a corresponding deep neural network. Thus, the system of ODEs plays a crucial role in the training (i.e., the optimization) of the neural network.

PINNs were first introduced in the work of Raissi et al. and since then used to solve different forward and inverse problems [1]. The PINN approach trades off between the data-based and physical loss functions in the training process. This steers the search for reasonable solutions towards those that satisfy a 'physical law' to some degree, i.e., an SVIHR compartmental model in this case. The loss function of the PINN is based on reported data of recent infection events (*data loss part*), and a system of ODEs inheriting transition and transmission dynamics, from which so-called residual networks are computed using automatic differentiation (*residual loss part*). The data loss makes PINNs a data-driven technique. Here, we distinguish between training data covering the time since the outbreak of the pandemic in Germany (long-term predictions) and training data involving exclusively one peak (short-term predictions).

The PINN involves several fixed model parameters, from which transition parameters are computed. A unique feature of our approach is that we estimate the parameters that are crucial to the dynamics of the system using a *nonstandard finite difference* (NSFD) method. In other words, we augment the PINN

with a numerical approach to the ODE system to better estimate the crucial parameters inside the residual loss function. While NSFD solutions generally do not provide good approximations to the data, they can be used to better estimate model parameters like, e.g., the transmission rates before and after vaccination.

The second contribution of our work is the optimization of the parameter weighting the relation between the data and residual loss part. In situations where the physical model can only partially represent the measured data, as is the case when predicting COVID-19 infection rates, data loss and residual loss are in real conflict. While the main goal is to reproduce the measured data well, the residual loss serves more as a regularization term that helps overcome noise and outliers in the data and better predict the underlying dynamics. Choosing a reasonable weighting for these two training objectives is far from trivial. To achieve this, we interpret the training process as a biobjective optimization problem, where the residual loss and the data loss are considered as two independent objective functions. Rather than combining these two objectives with a pre-determined and fixed weighting parameter, we identify suitable weighting parameters by generating a (rough) approximation of the Pareto front. For the training process, we adopt a scalarization-based approach that transforms the biobjective problem into a series of weighted-sum scalarizations. Favorable solutions are identified by repeated training runs with adaptively selected scalarization parameters. The resulting approximation of the Pareto front provides valuable information on the trade-off between data loss and residual loss. On one hand, this information can be used to assess the suitability of the employed physical model. On the other hand, a thorough analysis of the (approximated) Pareto front supports an informed selection of a suitable compromise, focusing more on the data or more on the physical model depending on the application background and on the decision makers preferences and beliefs.

1.1. Related Research

Since the outbreak of the COVID-19 pandemic, a variety of compartmental models have been introduced as enhanced susceptible-infected-recovered (SIR) compartment models to study various aspects of the spread of SARS-CoV-2. PINNs have been applied to compartment models and studied in the context of the COVID-19 pandemic as well.

For instance, Malinzi et al. applied a PINN to a susceptible-infected-recovered-deceased (SIRD) model in order to identify the behavioural dynamics of COVID-19 in the Kingdom of Eswatini between March 2020 and September 2021. They found that their PINN outperformed all other data analysis models even when given minimal quantities of training data [6].

Kharazmi et al. [7] considered different integer-order, fractional-order and time-delay models expressed as systems of ODEs. With the aim of analyzing the past dynamics of COVID-19 in New York City, Rhode Island and Michigan states as well as Italy, they used PINNs that were reported capable of performing parameter inference and simulation of the observed and unobserved dynamics simultaneously. Their results showed that purely statistical approaches were generally not well suited for long-term predictions of epidemiological dynamics, and integer-order models seemed to be more robust than fractional-order models, that were first developed by Pang et al. [8]. Moreover, they stated that no model could accurately capture all the dynamics that play out during an extended pandemic, but models with the ability to adjust key parameters during training could lead to more useful predictions [7].

Cai, Karniadakis and Li calibrated the unknown model parameters of a *susceptible-exposed-infected-removed* (SEIR) model using the novel fractional physics-informed neural networks (fPINNs) deep learning framework in order to obtain reliable short-term predictions of the COVID-19 dynamics caused by the Omicron variant [9]. Data from the National Health Commission of the People's Republic of China covering the time from 27th February 2022 to the end of April 2022 were used. For instance, predictions were able to capture sudden changes of the tendency for the new infected cases.

On the other hand, concerning the general PINN approach, the multiobjective nature of PINN training was recognized in several recent publications. Rohrhofer et al. [10] analyze the impact of different weights in a weighted sum objective of data loss and model loss by scanning the weight interval. Also Jin et al. [11, Section 4.4.] studied the influence of weights in an experiment for turbulent channel flow, by manually tuning the weight in order to improve the results. Finally, Wang et al. [12, Algorithm 2.1] proposed an adaptive rule, called 'learning rate annealing for PINNS', for choosing the weights online during the training process. The basic idea behind this is to automatically tune the weights by using the back-propagated gradient statistics during model training to properly balance all terms in the loss function. Their numerical results on diffusion equations and Navier-Stokes equations, respectively, impressively show the impact of the weight selection on the training success. Indeed, suboptimal results are obtained for several training runs, thus leaving room for improved multiobjective training approaches. When the physical model and the data are in good correspondence (this is, for example, the case when the data is artificially generated from the model at hand), an 'ideal' solution that simultaneously minimizes data and model loss can be sought.

Maddu et al. [13] suggest a multiobjective descent method that adaptively updates the weights using an inverse Dirichlet strategy to avoid premature termination. While they do not discuss convergence guarantees, their numerical results show a good performance in comparison with recent adaptations of multiobjective descent methods [14, 15] to PINN training [16]. Stochastic multiobjective gradient descent algorithms were introduced for general NN training in [17]. We also refer to self-adaptive PINNs [18] and to PINN training in which the loss weights are regarded as hyperparameters [19].

In a more general setting, multiobjective training approaches were suggested in [20] to trade off between data loss and regularization terms in the context of image recognition. The different characteristics (slope and curvature) of the considered training goals are addressed by enhancing the stochastic multi-gradient descent approach [17] with pruning strategies, and by combining adap-

tive weighted-sum scalarizations with interval bisection. The latter supports the identification of favorable knee solutions on the Pareto front.

This paper is organized as follows: In Section 2, the compartment model for COVID-19 predictions is introduced. Firstly, the SIR model is explained in Section 2.1 to provide an insight into the basics of epidemic modelling. Then the system of ODEs of our compartment model, the SVIHR model, is defined along with the used transition rates and transmission rate in Section 2.2. All model parameters are listed in Table 1.

Section 3 is devoted to the methodological developments. The *Nonstandard Finite Difference* (NSFD) method is introduced in Section 3.1, where the concept of the scheme is explained, the so-called *denominator function* is derived and the NSFD scheme for the SVIHR model is established. Section 3.2 provides an introduction to *physics-informed neural networks* (PINNs) with a focus on the loss function and the suggested neural network structure. Section 3.3 introduces some aspects of biobjective optimization needed to examine the Pareto front that is obtained by biobjective PINN training approaches. Finally, in Section 3.4, we introduce a dichotomic search scheme aiming to quickly find near-ideal Pareto optimal solutions and supporting an informed decision on the preferable trade-off between the data loss and the residual loss.

We present our numerical results in Section 4. In Section 4.1, we perform a short-term prediction of infection data, using data generated during the delta-variant wave as training data to predict the first omicron wave. In Section 4.2, we continue with the application of the dichotomic search scheme to discuss its ability to approximate a Pareto front, and identify reasonable trade-off solutions. Finally, in Section 4.3, we perform a long-term prediction using most of our available data as training data to predict the delta wave. We use the dichotomic search to find a low cost weighting parameter for both objective functions. A conclusion is drawn and an outlook to future work is given in Section 5.

2. A Compartment Model for COVID-19 Predictions

The compartment model used to compute the residual loss during PINN training in this paper is the *susceptible-vaccinated-infected-hospitalized-recovered* (SVIHR) model, which was proposed by Treibert and Ehrhardt in [21]. It is briefly derived in Section 2.2 again for the sake of completeness. Building upon the basic *susceptible-infectious-recovered* (SIR) model introduced by Kermack and McKendrick in 1927 [22], the SVIHR model enhances the SIR model to include a vaccinated and a hospitalized compartment. A general short introduction to SIR models in mathematical epidemiology is provided in Section 2.1.

In [21], a comparison between a data-driven PINN approach that takes into account a distinct training data set, and an NSFD method that approximates the SVIHR model was made with regard of the respective prediction qualities for infection and hospitalization numbers.

Treibert, Brunner and Ehrhardt [23] put the focus on the performance of the NSFD scheme for a *susceptible-vaccinated-infected-intensive care-deceased-recovered* (SVICDR) model. Here, the impact of modifications parameter bounds on the predicted prevalence was investigated, taking into account data from the pandemic in Germany and an exponentially increasing vaccination rate in the considered time window as well as trigonometric contact and quarantine rate functions. The results showed that the NSFD methods can predict a global peak solely based on the mathematical model and the defined parameters, but independently of a previously observed behavior of the infectious disease.

In this paper, we build on the SVIHR model of [21]. A novel PINN approach is presented based on updated data from [24, 25, 26] that incorporates both short-term and long-term data for the predictions. An improved network architecture is complemented by a dynamics-based parameter estimation that combines NSFD and PINN methods. Characteristic for the considered application is the often significant deviation of the measured data from the model predictions: While the SVIHR model captures the dynamics of COVID-19 infection, it cannot reflect fluctuations in the data that may have a variety of dif-

ferent causes. We address this challenge with an adaptive approach to analyze the trade-off between data loss and residual loss in the training process. This biobjective perspective on PINN training enables semi-automatic and problem-specific identification of optimized weighting parameters.

As determining the proportion of asymptomatic individuals in the total infected population is not our goal at this point, we do not incorporate a separate compartment of asymptomatic infected individuals, but assume at least very mild symptoms in infected individuals. The degree of infectivity of infected individuals can be regulated by adjusting the transmission rate in the model. Our model is adaptable to different vaccination and transmission scenarios.

2.1. The SIR Model in Mathematical Epidemiology

The basic SIR model consists of three compartments of susceptible (S), infected (I), and recovered (R) individuals. We denote with $K(t)$ the size of a compartment $K \in \{S, I, R\}$ at time t , where a time unit equals a week. Susceptible individuals have not yet become infected but may become ill. In the basic SIR model, infected individuals may infect susceptible persons, i.e. they are assumed to be infectious (without any delay) and may or may not have symptoms. Recovered individuals have overcome the disease and are assumed to be neither infectious nor ill.

The total size of the population at time t is denoted by $N(t)$. The satisfaction of the equation

$$N(t) = S(t) + I(t) + R(t) \quad \text{with } N: [0, T] \rightarrow \mathbb{N},$$

means that the number of individuals in the system is the sum of the compartment sizes at each considered time point $t \in [0, T]$. The system (2) must have initial conditions $S_0 = S(0)$, $I_0 = I(0)$, $R_0 = R(0) = 0$ to be well-defined [27, p. 11]. The population size $N(t)$ is assumed to be constant, i.e. $N(t) = N$, and the derivative of $N(t)$ is zero, which means that the system does not consider a recruitment rate Λ nor a natural death rate.

Let p be the probability that a contact with a susceptible individual results in a transmission, and let ζ be the per capita contact rate, i.e. the number of contacts made by one infectious individual. Then ζN is the number of contacts per unit of time this infectious individual makes, and $\zeta N \frac{S}{N}$ denotes the number of contacts with susceptible individuals that one infectious individual makes per unit of time. Moreover, we define a transmission rate constant β [27, p. 10] as

$$\beta = p \zeta. \quad (1)$$

For a more detailed discussion of the transmission rate and of related parameters including, among others, time dependent models, we refer to [23].

If $I(t)$ stands for the number of infected individuals at time t (*prevalence*), then $\beta S I$ denotes the number of individuals who become infected per unit of time (*incidence*). If ω_I is the recovery rate, we obtain the following system of ODEs, that describes the SIR model [27, p. 11]:

$$\begin{aligned} \frac{dS(t)}{dt} &= -\beta I(t) S(t), \\ \frac{dI(t)}{dt} &= \beta I(t) S(t) - \omega_I I(t), \\ \frac{dR(t)}{dt} &= \omega_I I(t). \end{aligned} \quad (2)$$

For the model in Equation (2), the maximum number of infected individuals that can be reached in the regarded epidemic is bounded by

$$I_{\max} = -\frac{\omega_I}{\beta} + \frac{\omega_I}{\beta} \ln\left(\frac{\omega_I}{\beta}\right) + S_0 + I_0 - \frac{\omega_I}{\beta} \ln(S_0). \quad (3)$$

Let

$$F(t) = 1 - e^{-\omega_I t}, \quad t \geq 0 \quad (4)$$

be the probability of recovering/leaving the infectious compartment in the time interval $[0, t)$ [27, p. 11]. The function $F(t)$, with $F(t) = 0$ for $t < 0$, is a probability distribution. Then $f(t) = \frac{dF(t)}{dt}$ is the respective probability density function:

$$f(t) = \omega_I e^{-\omega_I t} \quad \text{for } t > 0, \text{ and } f(t) = 0 \quad \text{for } t \leq 0. \quad (5)$$

If X denotes the average time spent in the infectious compartment, then the mean time spent in the infectious compartment can be computed as the first moment

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} t f(t) dt = \int_0^{\infty} t \omega_I e^{-\omega_I t} dt = \frac{1}{\omega_I}. \quad (6)$$

For SARS-CoV-2, the mean time of infectiousness is not clearly defined. With a mild or moderate course of the disease, contagiousness clearly declines within the ten days after symptom occurrence. Contagiousness has to be distinguished from positive test results, that can occur several weeks after catching the infection, although the infectiousness is usually on a very low level then [28].

In this basic form of the SIR model, the population is assumed to be closed so that no individual enters or leaves a compartment from the outside, and recovered individuals are completely immune so that they can never be reinfected [27, p. 13].

2.2. The SVIHR Model

The SIR model was enhanced by a compartment of hospitalized individuals H and a compartment of fully, i.e. at least twice, vaccinated individuals V in [21]. Data was obtained from the Robert Koch-Institute (RKI) [24, 25] and the German COVID-19 Vaccination Dashboard [26], based on which parameter values and compartment sizes, referred to as *reported compartment sizes* or *reported data* in the sequel, were computed.

Infected individuals remain infected for T_I days until they recover, where a proportion ξ of all individuals transiting from the infected individuals are hospitalized. The exact daily or weekly number of infectious individuals among the infected individuals is not known. The number of infections registered by the Robert Koch-Institute (RKI) is used to compute the reported size of the compartment I for all considered calendar weeks in this paper. This number is based on the number of infected individuals who are infectious enough so that the virus is usually verifiable via a rapid antigen test. Infectious and not infectious infected people are merged within the compartment I . The general

degree of infectiousness of the individuals in I depends on the transmissibility of the virus and is included in the transmission rate.

According to the RKI, the concrete time period of contagiousity is not clearly defined, but infectiousness is highest right before and after the presence of first symptoms and drastically declines after at most 10 days after the very first symptoms occur (assuming a mild or moderate course of disease) [28]. We set $T_I = 1.2$ weeks, i.e. 8.4 days, to adopt for a small time span of 1-2 days between the first showing of symptoms and getting tested. The parameter ω_1 is the rate at which individuals per unit time (week) pass from compartment I to R . It is defined as

$$\omega_1 = \frac{1 - \xi}{T_I}. \quad (7)$$

The rate η at which individuals reach the compartment H per unit of time is defined as

$$\eta = \frac{\xi}{T_I}. \quad (8)$$

Thus, we assume that $\omega_1 I(t)$ people recover and $\eta I(t)$ individuals are hospitalized within week t . It is assumed here that hospitalized individuals cannot infect susceptible individuals because they are isolated. They are assumed to remain infected T_H days from the time of their hospitalization. In a German academic survey with 1,426 COVID-19 patients with an acute respiratory disease, an average duration of hospital stay of 10 days was observed [28]. Accordingly, we set $T_H = 1.5$ weeks, i.e. 10.5 days. ω_2 is the transition rate at which persons per unit time pass from compartment H to R .

The vaccinated compartment V contains all susceptible individuals who have received a COVID-19 vaccination. It is reached from the compartment S at a rate \mathcal{V} . If vaccination does not guarantee complete immunity to infection, we speak of a *leaky vaccination*. Due to the assumed leakiness, all vaccinated individuals have a lower probability of contracting the infection than susceptible individuals in compartment S . If an all-or-nothing vaccine was assumed, vaccinated people would be completely protected from the infection to a specific portion of the susceptible class per unit time t , whereas the other susceptible

individuals did not gain any protection. Let κ denote the residual probability of infection after vaccination. The rate at which vaccinated individuals reach the infected compartment I is then $\kappa\beta I(t)$.

Furthermore, we incorporate a constant system inflow, the so-called recruitment rate Λ (e.g. birth of new individuals that can get infected), and the natural mortality rate μ . The recruitment and natural death rate are set to zero as they are regarded as equal in both [21] and in this paper, but are still included in the system of ODEs for the purpose of properly deriving the denominator function in the NSFD scheme, see Section 3.1. The total population size is kept constant like this. The corresponding system of ODEs has the following form:

$$\begin{aligned}
\frac{dS(t)}{dt} &= \Lambda - \beta I(t)S(t) - (\mathcal{V} + \mu)S(t), \\
\frac{dV(t)}{dt} &= \mathcal{V}S(t) - \kappa\beta I(t)S(t) - \mu V(t), \\
\frac{dI(t)}{dt} &= (1 + \kappa)\beta I(t)S(t) - (\eta + \omega_1 + \mu)I(t), \\
\frac{dH(t)}{dt} &= \eta I(t) - (\omega_2 + \mu)H(t), \\
\frac{dR(t)}{dt} &= \omega_1 I(t) + \omega_2 H(t) - \mu R(t).
\end{aligned} \tag{9}$$

The system (9) extends the simple system (2) by the differential equations for $V(t)$ and $H(t)$, describing the inflow into and the outflow from the compartment V or H , respectively, as well as the recruitment rate Λ and the natural death rate μ . Note that the equation for $V(t)$ in (9) can only approximate the actual dynamics, since vaccinated individuals remain vaccinated even when they become infected. We approximate the respective rates of change by using the compartment of susceptible individuals as a reference in the term $\kappa\beta I(t)S(t)$. Preliminary numerical tests, using both the NSFD scheme and the PINN approach, show that the use of $\kappa\beta I(t)V(t)$ does not significantly affect the quality of the predictions.

One can easily show that for positive parameters and positive initial data, the solution of (9) remains positive for all times. We call this the "positivity property". If we now set $\Lambda = \mu = 0$ and add all compartments in (9), we find that the total population is a quantity conserved over time. We will return to

this "conservation property" later when we discuss the NSFD scheme. This setting is an acceptable simplification because the time scale of human births and deaths is much longer than that of a COVID-19 epidemic wave.

Figure 1 shows the dynamical system described by (9). Blue arrows from one compartment to another indicate a transition. Table 1 lists the model parameter

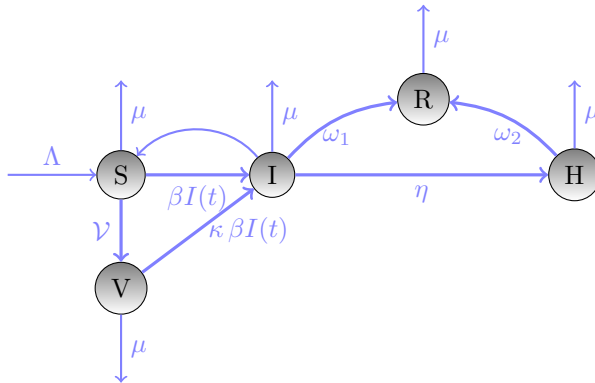


Figure 1: Illustration of the compartments and their interrelation in the SVIHR model

definitions and used values. We note that the parameter values for the SVIHR model used in this paper stated in Table 1 differ from the ones used in [21].

3. Finding Optimized Weights in PINN Approaches

We implement a physics informed neural network (PINN) that is trained both w.r.t. German COVID-19 data and w.r.t. the SVHIR model introduced in Section 2.2 above. This technique is validated using error computations with regard to reported data. The method and structure of PINNs is explained in Section 3.2. Moreover, scenarios generated using the PINN are compared to those produced using the technique of nonstandard finite difference (NSFD) schemes. A short introduction to NSFD schemes and the application on the SVIHR model are outlined in Section 3.1.

Since the PINN is trained w.r.t. two loss terms, we take a biobjective approach to investigate the influence of weighting parameters on each loss term.

Parameter	Definition	Parameter Value used in 4.3	Parameter Value used in 4.1 and 4.2
N	total population size	83,100,000 [29]	83,100,000 [29]
β	transmission risk parameter	<i>determined via NSFD</i>	<i>determined via NSFD</i>
κ	residual infection prob. after vaccination	<i>determined via NSFD</i>	<i>determined via NSFD</i>
T_I	average length of infection period	1.2 [28]	1.2 [28]
T_H	average length of hospitalization period	1.5 [28]	1.5 [28]
ν	vaccination coefficient	0.0159 [26]	0.0231 [26]
ξ	hospitalization coefficient	0.0862 [25]	0.0735 [25]
\mathcal{M}	mortality coefficient	0.0232 [24]	0.0142 [24]
ω_1	weekly recovery rate for infected people	0.7615	0.7721
ω_2	weekly recovery rate for hospitalized people	0.6512	0.6572
η	weekly hospitalization rate	0.0719	0.6125

Table 1: Parameters of the SVIHR model and their values in the implementation of the PINN.

We consider the two losses as independent objective functions and want to find weighting parameters to achieve an approximation of the Pareto front. We therefore first introduce certain aspects of bicriteria optimization in Section 3.3 and then introduce a dichotomic search scheme to efficiently approximate the Pareto front in Section 3.4.

3.1. Nonstandard Finite Difference Schemes

NSFD schemes go back to a paper by Mickens published in 1989 [30]. Their structural properties originate from investigations of special groups of differential equations for which exact finite difference schemes are not available. For the sake of completeness, the NSFD scheme for the SVIHR model is described again here. It has already been derived in a similar way in [21].

In NSFD schemes, derivatives have to be modelled by proper discrete analogues, i.e. nonstandard difference quotients of the form, cf. [31]

$$\frac{du(t)}{dt} \rightarrow \frac{u_{n+1} - \psi(h)u_n}{\phi(h)}, \quad (10)$$

where $t_n = nh$, u_n is the approximation of $u(t_n)$, and $\psi(h) = 1 + \mathcal{O}(h)$. $\phi(h)$ is a denominator function, which is explained in more detail below. Using this rather general time discretization (10) in NSFD schemes our aim is to model the asymptotic long-time behaviour of the solution. A numerical scheme for a system of first-order differential equations is called NSFD scheme if at least one of the following conditions hold [31]:

- Discrete representations for derivatives must, in general, have nontrivial denominator functions. Here, the first-order derivatives in the system are approximated by the generalized forward difference method $\frac{du_n}{dt} \approx \frac{u_{n+1} - u_n}{\phi(h)}$, where $u_n \approx u(t_n)$ and $\phi \equiv \phi(h) > 0$ is the so-called *denominator function* such that $\phi(h) = h + \mathcal{O}(h^2)$, with h the step size.
- The consistency orders of the finite difference quotients should be equal to the orders of the corresponding derivatives appearing in the differential equations.

- The nonlinear terms are approximated by non-local discrete representations, for instance by a suitable function of several points of a mesh, like $u^2(t_n) \approx u_n u_{n+1}$ or $u^3(t_n) \approx u_n^2 u_{n+1}$.
- Special conditions that hold for either the ODE and/or its solutions should also hold for the difference equation model and/or its solution, e.g. the equilibrium points of the underlying ODE system.

In contrast to conventional difference methods, NSFD schemes focus not only on stability and convergence order, but also on qualitative properties, i.e., how well the discrete model (the NSFD scheme) reproduces the most important properties of the underlying continuous model. NSFD schemes preserve the positivity property and satisfy the conservation law for $\Lambda = \mu = 0$ yielding the stability of the scheme. The equilibrium points of the ODE model (9) appear in the proposed NSFD-scheme as well.

In the sequel, we will derive an appropriate *denominator function* $\phi(h)$ for the NSFD discretization of the system (9). This function is chosen such that the numerical solution exhibits the same asymptotic behaviour as the analytic solution. To do so, we consider the total population $N = S + V + I + H + R$ of the ODE system (9). Adding the equations of (9), a differential equation describing the dynamics of the total population N is obtained as

$$\frac{dN(t)}{dt} = \Lambda - \mu N(t), \quad (11)$$

which is solved by

$$N(t) = \frac{\Lambda}{\mu} + \left(N(0) - \frac{\Lambda}{\mu}\right) e^{-\mu t} = N(0) + \left(N(0) - \frac{\Lambda}{\mu}\right) (e^{-\mu t} - 1), \quad (12)$$

where $N(0) = S(0) + V(0) + I(0) + H(0) + R(0)$. Similar to (11), adding the equations of an NSFD discretization of (9) yields the equation

$$\frac{N^{n+1} - N^n}{\phi(h)} = \Lambda - \mu N^{n+1}, \quad (13)$$

i.e.

$$\begin{aligned} N^{n+1} &= \frac{N^n + \phi(h)\Lambda}{1 + \phi(h)\mu} = N^n - \left(N^n - \frac{\Lambda}{\mu}\right) \frac{\phi(h)\mu}{1 + \phi(h)\mu} \\ &= N^n + \left(N^n - \frac{\Lambda}{\mu}\right) \left(\frac{1}{1 + \phi(h)\mu} - 1\right). \end{aligned} \quad (14)$$

The denominator function can be derived by comparing Equation (13) with the discretized version of Equation (12), that is

$$N^{n+1} = N^n + \left(N^n - \frac{\Lambda}{\mu}\right) (e^{-\mu h} - 1), \quad h = \Delta t, \quad (15)$$

such that the (positive) denominator function is defined by

$$\frac{1}{1 + \phi(h)\mu} = e^{-\mu t}, \quad (16)$$

i.e.

$$\phi(h) = \frac{e^{\mu h} - 1}{\mu} = \frac{1 + \mu h + \frac{1}{2}\mu^2 h^2 + \dots - 1}{\mu} = h + \frac{\mu h^2}{2} + \dots = h + \mathcal{O}(h^2). \quad (17)$$

Now, making use of the denominator function in (10), the NSFD discretization can be established. The implicit form of this discretization is provided in Equation (18). Here, $\phi(h)$ is given by (17)

$$\begin{aligned} \frac{S^{n+1} - S^n}{\phi(h)} &= \Lambda - \beta I^n S^{n+1} - (\mathcal{V} + \mu) S^{n+1}, \\ \frac{V^{n+1} - V^n}{\phi(h)} &= \mathcal{V} S^{n+1} - \beta \kappa I^n S^{n+1} - \mu V^{n+1}, \\ \frac{I^{n+1} - I^n}{\phi(h)} &= \beta (1 + \kappa) I^{n+1} S^{n+1} - (\eta + \omega_1 + \mu) I^{n+1}, \\ \frac{H^{n+1} - H^n}{\phi(h)} &= \eta I^{n+1} - (\omega_2 + \mu) H^{n+1}, \\ \frac{R^{n+1} - R^n}{\phi(h)} &= \omega_1 I^{n+1} + \omega_2 H^{n+1} - \mu R^{n+1}. \end{aligned} \quad (18)$$

We can rewrite the scheme in order to obtain an explicit variant of it, as to be found in (19). From the explicit representation we can deduce that this scheme

preserves the positivity.

$$\begin{aligned}
S^{n+1} &= \frac{S^n + \phi(h)\Lambda}{1 + \phi(h)(\beta I^n + \mathcal{V} + \mu)}, \\
V^{n+1} &= \frac{V^n + \phi(h)S^{n+1}(\mathcal{V} - \beta\kappa I^n)}{1 + \phi(h)\mu}, \\
I^{n+1} &= \frac{I^n}{1 + \phi(\eta + \omega_1 + \mu - \beta(1 + \kappa)S^{n+1})}, \\
H^{n+1} &= \frac{\phi(h)\eta I^{n+1} + H^n}{1 + \phi(\omega_2 + \mu)}, \\
R^{n+1} &= \frac{R^n + \phi(h)(\omega_1 I^{n+1} + \omega_2 H^{n+1})}{1 + \phi(h)\mu}.
\end{aligned} \tag{19}$$

The calculation must be implemented in exactly this order. All parameters appearing in these type of epidemic models are always non-negative.

Remark 1. Looking at the third equation of (18), one might wonder why one does not use the time discretization $I^n S^{n+1}$ instead of $I^{n+1} S^{n+1}$. Here two arguments meet in this discussion, and there is no clear right or wrong. With this "new" version, one has the exact conservation property on the discrete level. With the "old" version $I^{n+1} S^{n+1}$ (implicitly only in I^{n+1} , S^{n+1} is already calculated in the first line), on the other hand, one gets positivity conservation and thus also the stability property, which is more important. So we have chosen the first variant and have to cope with a small perturbation of the conservation of the total population.

For an actual application, where a standard solver gives false negative solution and this problem can be solved by a NSFD scheme we refer to [32].

3.2. Physics-informed Neural Networks for Compartment Models

Physics-informed neural networks (PINN) are neural networks that include the laws of dynamical systems into a deep learning framework. Machine learning has emerged as an alternative to numerical discretization in high-dimensional problems governed by partial differential equations. Nonetheless, a sufficient amount of data as required for training deep neural networks is not necessarily available. In such cases, missing data can be substituted by incorporating

additional information obtained from enforcing the physical laws of dynamical systems [33]. Such laws can be described by partial or ordinary differential equations. One example where dynamical systems can be used are populations undergoing transitions between different infected or uninfected states during an epidemic, as considered in this paper.

PINNs can approximate the solutions of differential equations by training a loss function incorporating the initial and boundary conditions and the residual at so-called collocation points [34]. Instead of approximating solutions of differential equations, PINNs can use a system of differential equations describing a certain real-world process along with time-series data sets describing the past course of such a process for the purpose of generating predictions for future progressions.

The loss function of a corresponding neural network includes not solely the so-called *data error* related to the difference between the output of the network and the reported data used, but also the so-called *residual error* related to the ODEs or PDEs.

Olumoyin et al. [35] refer to a type of feedforward neural network including epidemiological dynamics such as lockdown into their loss function by using the term *Epidemiology-Informed Neural Network (EINN)*. EINNs extend PINNs for epidemiology models and are able to capture the dynamics of the spread of the disease and the influence of the mitigation measure. The loss function is enhanced to include time-varying rates using epidemiology facts about the infectious disease [35].

Shaier, Raissi and Seshaiyer [36] describe a type of PINN-based neural network that can be applied to increasingly complex systems of differential equations describing various known infectious diseases with the term *Disease-Informed Neural Networks (DINN)*. DINNs can be systematically applied to increasingly complex governing systems of differential equations describing infectious diseases. They are able to effectively learn the dynamics of the disease and forecast its progression a month into the future from real-life data [36].

The neural network established and applied in this paper can be described

as a special type of EINN or DINN based on a special kind of epidemic compartment model called SVIHR model regarding the susceptible, vaccinated, infected, hospitalized and recovered part of the population. We focus on the German population in this paper, using data provided by the Robert Koch-Institute [24, 25] and the German COVID-19 Vaccination Dashboard [26]. The model parameters included in (9) can be comprised in a vector ϑ :

$$\vartheta = [\beta, \mathcal{V}, \kappa, \xi, T_I, T_H, \mathcal{M}, T_H]^\top. \quad (20)$$

This vector is partitioned into fixed parameters p_f , which can be estimated directly from the observed data, and learnable parameters p , which are crucial for the dynamics of the compartment variables:

$$\begin{aligned} p_f &:= [\mathcal{V}, \xi, T_I, T_H, \mathcal{M}, T_H]^\top, \\ p &:= [\beta, \kappa]^\top. \end{aligned} \quad (21)$$

We use NSFD predictions to improve the estimates for the parameters that govern the dynamics, i.e., the transmission risk β and the residual transmission probability after vaccination κ . Their values are estimated using NSFD predictions fitted to the observed data so that the predicted and observed peak values have the same magnitude. This approach aims to combine the benefits from both worlds: The numerical predictions that follow a physical model, and the PINN predictions that rely heavily on the observed data. Since the NSFD predictions can be computed very efficiently, the parameter values can be easily adapted to different virus variants and to different pandemic waves.

We describe our PINN approach for a general compartmental model with n compartments. In the SVIHR model (9), we have $n = 5$. Let

$$\hat{\mathcal{K}}(t) = [\hat{\mathcal{K}}^1(t), \dots, \hat{\mathcal{K}}^n(t)]^\top$$

be the normalized data vector of reported compartment sizes for l time points, i.e. $t \in \{t_1, \dots, t_l\}$. In the application considered here, $t_{i+1} - t_i$ is constant (and corresponds to one week) for all $i = 1, \dots, l - 1$. Let T denote an upper bound for the time window under consideration, and let $\mathcal{K}_p : [0, T] \rightarrow \mathbb{R}^n$ denote the

vector valued function reflecting the compartment sizes $\mathcal{K}_p(t)$ over time. Then we express the vector of right-hand-sides values of (9) by

$$F(\mathcal{K}_p) = [F^1(\mathcal{K}_p), \dots, F^n(\mathcal{K}_p)]^\top.$$

The subscript p stands for the learnable model parameters that the system of ODEs depends on, cf. (21), and thus the solution will depend on, too. The system of ODEs in (9) can then be discretized as

$$\frac{d\mathcal{K}_p(t)}{dt} - F(\mathcal{K}_p) = 0, \quad t \in \{t_1, \dots, t_l\}. \quad (22)$$

Our PINN

$$\mathcal{PINN}_p^W : [0, T] \rightarrow \mathbb{R}^n$$

is used to approximate the solution $\mathcal{K}_p : [0, 1] \rightarrow \mathbb{R}^n$ of the system of ODEs (9) by performing error minimization during training [37]. The superscript W represents the weights used during the forward and backward propagation in the neural network. At time instance t , $t \in \{t_1, \dots, t_l\}$, the solution is expressed as

$$\mathcal{K}_p(t) = [\mathcal{K}_p^1(t), \dots, \mathcal{K}_p^n(t)]^\top,$$

where $\mathcal{K}_p^j(t)$ is the output of the PINN for the j^{th} compartment at time t and depends on the parameter p given by equation (21). The parameters W are optimized during the backpropagation process of the neural network such that \mathcal{PINN}_p^W fits the reported data $\hat{\mathcal{K}}$ in a least-squares sense [37]. In the i^{th} time step t_i , $i \in \{1, \dots, l\}$, with PINN output $\mathcal{PINN}_p^W(t_i)$, we compute the usual *data error* defined as

$$\text{MSE}_{\mathcal{U}} = \text{MSE}_{\mathcal{U}}(W) := \frac{1}{l} \sum_{i=1}^l \|\mathcal{PINN}_p^W(t_i) - \hat{\mathcal{K}}_p(t_i)\|^2, \quad (23)$$

where we employed in (23) the Euclidian norm in \mathbb{R}^n .

Next, let us extend the loss function of the PINN by the additional term

$$\mathcal{F}_p(\mathcal{PINN}_p^W, t_i) := \left. \frac{d\mathcal{PINN}_p^W(t)}{dt} \right|_{t=t_i} - F_p(\mathcal{PINN}_p^W(t_i)), \quad (24)$$

where

$$\mathcal{F}_p(\mathcal{PINN}_p^W, t_i) = 0 \quad \text{for all } t \in \{t_1, \dots, t_l\} \quad (25)$$

means that the PINN solves the given ODE system more reliably by enforcing additional constraints. The computation of the time derivative of the neural network output $\left. \frac{d \mathcal{PINN}_p^W(t)}{dt} \right|_{t=t_i}$ can be performed using automatic differentiation [38]. Then the physics-informed part of the loss function, the *residual error*, is given by

$$\text{MSE}_{\mathcal{F}} = \text{MSE}_{\mathcal{F}}(W) := \frac{1}{l} \sum_{i=1}^l \|\mathcal{F}_p(\mathcal{PINN}_p^W, t_i)\|^2. \quad (26)$$

We introduce a hyperparameter and weighting factor $\alpha \in [0, 1]$ weighting the data loss and residual loss in the loss function. We define the overall loss function as

$$\mathcal{L}_{\alpha} = \mathcal{L}_{\alpha}(W) := \alpha \text{MSE}_{\mathcal{U}} + (1 - \alpha) \text{MSE}_{\mathcal{F}} \quad (27)$$

and the minimization problem of the neural network as

$$\min_W (\mathcal{L}_{\alpha}). \quad (28)$$

We want to note briefly that small values of the residual loss $\text{MSE}_{\mathcal{F}}$ do not necessarily mean that the obtained solution is close to the exact solution of (9). For this, one additionally needs the well-posedness of the problem (9) and the stability of the scheme used.

For all implementations, the programming language *Python* and the deep-learning framework of *PyTorch* are used. We used a standard network architecture with three fully connected hidden layers, each with 30 neurons, see Figure 2 for illustration. For each time point $t \in [0, T]$, the network returns values for all five compartment sizes in the vector $\mathcal{K}_p(t) \in \mathbb{R}^5$. The network architecture was inspired by Ben Moseley’s harmonic oscillator PINN, see [39], and represents a reasonable compromise in terms of network complexity.

We used hyperbolic tangent activation functions and the Adam optimizer for the training. We refer to [40] for a survey of descent methods in machine learning. We also used a learning rate schedule that significantly reduces the learning rate after about 50 % of training iterations. More specifically, the learning rate

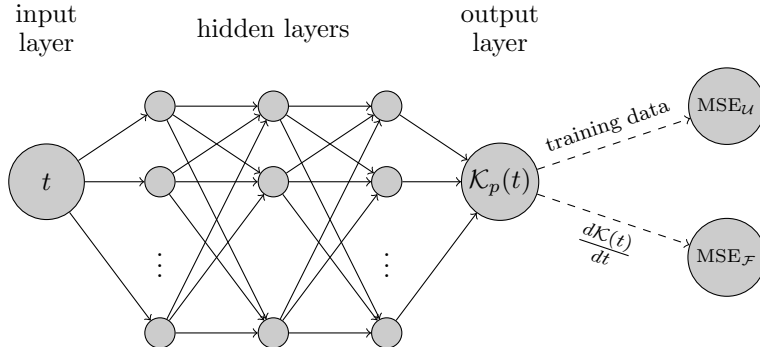


Figure 2: Network architecture of the proposed PINN. Each hidden layer consists of 30 neurons.

$t(\kappa)$ in iteration $\kappa \in \{1, \dots, \kappa_{\max}\}$ is given by

$$t(\kappa) = -(t_{\text{start}} - t_{\text{end}}) \frac{\exp\left(\frac{\kappa - 0.5 \cdot \kappa_{\max}}{0.08 \cdot \kappa_{\max}}\right)}{\exp\left(\frac{\kappa - 0.5 \cdot \kappa_{\max}}{0.08 \cdot \kappa_{\max}}\right) + 1} + t_{\text{start}}, \quad (29)$$

where t_{start} and t_{end} denote the initial and final learning rates, respectively. In our experiments, we used $t_{\text{start}} = 0.003$ and $t_{\text{end}} = 0.00015$. Figure 3 shows the learning rate schedule for $\kappa_{\max} = 1000$ iterations.

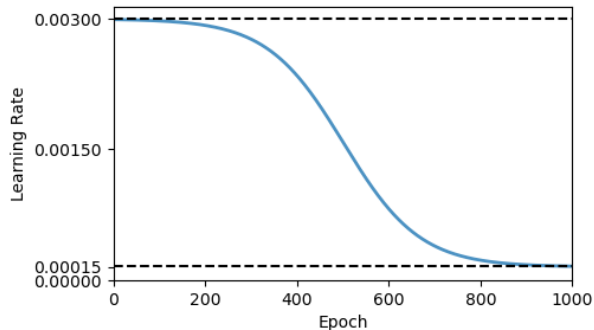


Figure 3: Our employed learning rate schedule reducing the learning rate over 1000 epochs.

The data loss was evaluated with data sets obtained from the RKI [24, 25] and the German Vaccination Dashboard [26]. These data were preprocessed and rescaled to unit intervals before calculating the data loss $\text{MSE}_{\mathcal{U}}$. Since this leads to a corresponding rescaling of the PINN predictions \mathcal{K}_p , the corresponding

scaling factors were also included in the ODE system (9) when evaluating $\text{MSE}_{\mathcal{F}}$.

The data refers to the calendar weeks 10 in 2020 to 14 in 2022. Weekly case-hospitalization, case-fatality and vaccination rates were computed based on the given data sets. The RKI registers deceased individuals, in whom the SARS-CoV-2 pathogen was detected, as people who died from COVID-19.

The weighted loss function $\mathcal{L}_\alpha = \alpha \text{MSE}_{\mathcal{U}} + (1 - \alpha) \text{MSE}_{\mathcal{F}}$ consists of the data loss and the residual loss term. As stated in [33], training using the data loss (i.e., measurements, physics-uninformed) is considered supervised learning, while training with the residual loss using the governing differential equation (physics-informed) is considered unsupervised learning.

3.3. Conflicting Training Goals: Pareto Front and Trade-Off Analysis

In this paper, we take a biobjective perspective on the optimization problem described in (27). This is the basis for a thorough trade-off analysis regarding the two loss terms $\text{MSE}_{\mathcal{U}}$ and $\text{MSE}_{\mathcal{F}}$. Rather than considering a weighted sum of these two training goals with a fixed weighting parameter $\alpha \in [0, 1]$, we consider both optimization goals independently and comprise them in a vector-valued objective function \mathcal{L} that maps every feasible solution vector (W, p) to a two-dimensional outcome vector $y = \mathcal{L}(W, p) \in \mathbb{R}^2$:

$$\min_W \mathcal{L}_\alpha(W) = \min_W (\text{MSE}_{\mathcal{U}}(W), \text{MSE}_{\mathcal{F}}(W)). \quad (30)$$

As before, W denotes the neural network weights. To analyze the biobjective optimization problem (30), we first review some basic concepts from the field of multiobjective optimization. For a more detailed introduction into this field see, e.g., [41].

Towards this end, we denote by Y the *outcome set of problem* (30) that includes all possible outcome vectors $y = \mathcal{L}_\alpha(W) \in \mathbb{R}^2$. A solution \hat{W} (i.e., a set of NN-weights \hat{W} resulting from the training) is called *efficient* or *Pareto optimal* if there exists no other solution W (i.e., no other set of NN-weights W) such that $\text{MSE}_{\mathcal{U}}(W) \leq \text{MSE}_{\mathcal{U}}(\hat{W})$ and $\text{MSE}_{\mathcal{F}}(W) \leq \text{MSE}_{\mathcal{F}}(\hat{W})$, where at

least one of these two inequalities is strict. If \hat{W} is Pareto optimal, then the corresponding outcome vector $y = (\text{MSE}_{\mathcal{U}}(\hat{W}), \text{MSE}_{\mathcal{F}}(\hat{W}))$ is called *nondominated point*. Hence, Pareto optimal solutions are those solutions (i.e., NN weights) that can not be improved in one loss function without deterioration in the other loss function. The set of all Pareto optimal solutions (nondominated points, respectively) is denoted by X_E (Y_N , respectively). Note that Y_N is also often referred to as *Pareto front*. The ultimate goal is the efficient, i.e., fast approximation of so-called *knee solutions* that provide near-optimal values for both data loss *and* residual loss. This is realized by a dichotomic search strategy specifically tailored to NN training.

An approximation of the Pareto front yields an approximation of the knee solution. It can be found by solving a series of parametric single-objective subproblems, so-called *scalarizations* (see again, e.g., [41]). To keep these subproblems simple, we use a *weighted sum* approach that leads to the single-objective optimization problem (28) with the objective function (27), i.e.,

$$\min_W (\mathcal{L}_\alpha(W)) = \min_W \{ \alpha \text{MSE}_{\mathcal{U}}(W) + (1 - \alpha) \text{MSE}_{\mathcal{F}}(W) \} \quad (28)$$

It is a well-known fact [41] that for all weighting parameters $\alpha \in (0, 1)$, an optimal solution of (28) is always part of Y_N . However, the converse statement is only true under convexity assumptions. Indeed, the complete set Y_N can be generated by varying the weighting parameter $\alpha \in (0, 1)$ whenever the set Y is \mathbb{R}_{\geq}^2 -convex, a property that can generally not be guaranteed in neural network training. We recall that Pareto optimal solutions $\hat{W} \in X_E$ are called *supported* if there is some $\alpha \in (0, 1)$ such that $\hat{W} \in X_E$ is an optimal solution of (28). The sets of all supported efficient solutions and supported nondominated points are denoted X_{sE} and Y_{sN} , respectively. Note that all supported nondominated points are located on the boundary of the convex hull of Y (see, e.g., [42]). Figure 4 shows an example of a set of supported nondominated points Y_{sN} within a set of outcome vectors Y in \mathbb{R}^2 as well as an unsupported nondominated point.

The choice of the weighting parameter α in (28) is both a crucial and highly

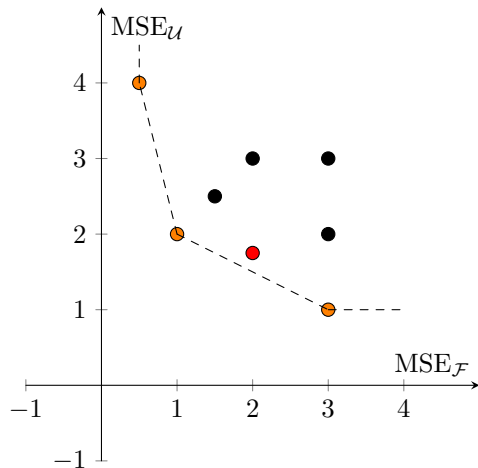


Figure 4: Schematic illustration of a potential set of outcome vectors $(\text{MSE}_{\mathcal{F}}, \text{MSE}_{\mathcal{U}})$ obtained from NN training, Points belonging to $Y_{s,N}$ are depicted in orange, the unsupported nondominated point is shown in red. The dashed line outlines the convex hull of the set of outcome vectors $Y + \mathbb{R}_{\ge}^2$.

challenging task. While a too large value of α overemphasizes data loss and therefore often leads to overfitting, a too small value of α may give too much priority to a possibly only approximate physical model. When trying to approximate the Pareto front by scanning over potential weights $\alpha \in (0, 1)$, even for biobjective and convex problems predefined weighting parameters may lead to very un-evenly distributed points on the Pareto front, see, e.g., [43]. Moreover, this approach does not scale well to higher-dimensional problems that include more than two training objectives. Indeed, when searching through all reasonable weights, the required number of weighting vectors α (and hence the individual NN training with respect to a weighted sum objective (28)) grows exponentially with the number of objective functions.

In the next section we present an efficient approach that supports an adaptive selection of weighting parameters to identify knee solutions for the case of two optimization goals (the data loss and the residual loss in our case). We emphasize that our approach can be generalized to more than two objective functions (e.g. for the PINN solution of PDEs) and thus has the potential to

significantly reduce the number of individual NN trainings required to identify near-ideal NN weights.

3.4. Dichotomic Search for Adaptive Pareto Front Approximations

Near-ideal NN-solutions can be identified by adapting the scalarization-based dichotomic search algorithm described in [42]. The idea is to compute adaptively weighting parameters α that refine the current approximation of the Pareto front in the most promising regions. This approach aims at an automatic adaptation to the curvature and scaling of the problem in order to quickly find a diverse set of solutions, and to quickly identify near-ideal knee solutions. Moreover, dichotomic search can be easily integrated in an interactive procedure that allows to zoom in into specific parts of the Pareto front that are most interesting to the decision maker, see, e.g., [44].

The following ideas can be found in [42]. In the biobjective case, the dichotomic search comes down to solving a sequence of weighted sum scalarizations (28) with $\alpha \in (0, 1)$ and makes use of the fact that in the two-dimensional case, for two nondominated points y^r and y^s it holds that $y_1^r < y_1^s$ implies $y_2^r > y_2^s$. A weighted sum scalarization (28) with $\alpha = (y_2^r - y_2^s)/c > 0$ and $1 - \alpha = (y_1^s - y_1^r)/c > 0$, where $c = y_2^r - y_2^s + y_1^s - y_1^r$, is then solved to find new supported points between y^r and y^s . The weighting parameter α hence defines a normal vector to the line segment connecting y^r and y^s since

$$\begin{aligned} (y^r - y^s)^\top \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix} &= (y_1^r - y_1^s, \quad y_2^r - y_2^s) \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix} \\ &= c \cdot \left(-(1 - \alpha), \quad \alpha \right) \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix} = 0. \end{aligned}$$

This is illustrated in the left of Figure 5. Solving the weighted sum problem with the new weighting parameter α leads to a nondominated point y^t (if the problem is solved to global optimality) for which two cases can occur:

1. If $(\alpha, 1 - \alpha)^\top y^t < (\alpha, 1 - \alpha)^\top y^r$, then y^t is a new supported nondominated point. Two new subproblems are generated, one of which is defined by y^r

and y^t while the other one is defined by y^t and y^s . This case is illustrated in the right of Figure 5.

2. If $(\alpha, 1 - \alpha)^\top y^t = (\alpha, 1 - \alpha)^\top y^r = (\alpha, 1 - \alpha)^\top y^s$, then y^t lies on the line segment connecting y^r and y^s and the search can stop in this interval.

The dichotomic search progresses in *levels*, where level 1 contains the outcome vectors of the weighted sum scalarization with the two initial weights α_1 and α_2 , $\alpha_1 < \alpha_2$, and one dichotomy step (see the left part of Figure 5 for an illustration of the associated weighting parameter). In level 2, weighted sum scalarizations are solved for all weights defined by the line segments comprising the convex hull of the current approximation of the Pareto front (see the right part of Figure 5 for an illustration). The search is repeated until a predefined number of levels has been evaluated, or until no new subproblems have been generated.

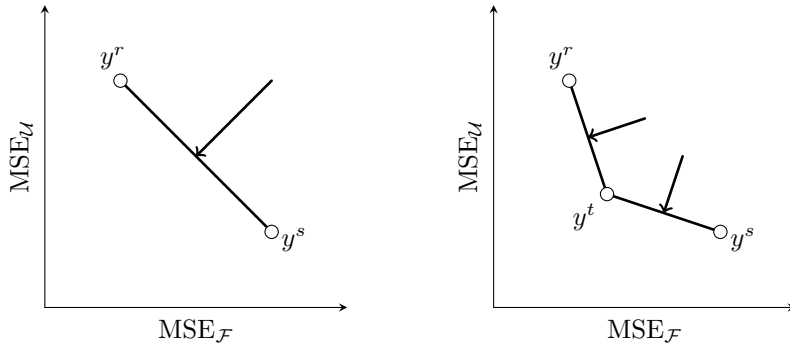


Figure 5: On the left: The connecting line segment between y^r and y^s is perpendicular to the vector $(\alpha, 1 - \alpha)^\top$ that is here multiplied by a negative scalar to show the minimization direction. On the right: Since $(\alpha, 1 - \alpha)^\top y^t < (\alpha, 1 - \alpha)^\top y^r$, the new point y^t is nondominated and two new subproblems are generated, see [42, p. 7] (own illustration).

Algorithm 1 summarizes the implementation of the bisection enhanced dichotomic search (BEDS) that enhances the dichotomic search scheme by an occasional bisection step. It is based on [20] and considers the fact that NN training with a certain weight parameter α may yield a local minimum, i.e., the training (here the Adam optimizer) may terminate in a local minimum and not in a global minimum as assumed in the general dichotomic search procedure.

Algorithm 1: Bisection Enhanced Dichotomic Search (BEDS)

Data: Training data, hyperparameter settings for Adam solver, depth of the search (levels), initial weighting parameters $\alpha_1, \alpha_2 \in (0, 1)$, $\alpha_1 < \alpha_2$, to approximate extremal solutions focusing on $\text{MSE}_{\mathcal{U}}$ and $\text{MSE}_{\mathcal{F}}$, respectively (see (27))

Result: Approximation of the Pareto front and corresponding PINN parameters

```
1  $\Lambda \leftarrow \{\alpha_1, \alpha_2\}$ ;
2 cand  $\leftarrow \emptyset$ ;
3 for  $l = 1, \dots, \text{levels}$  do
4   for  $\alpha \in \Lambda$  do
5     train with weighted sum objective  $\mathcal{L}_\alpha$ ;
6     add objective vector  $(\text{MSE}_{\mathcal{U}}, \text{MSE}_{\mathcal{F}})$  to cand;
7   delete all dominated points in cand;
8   sort cand by second objective function (in increasing order);
9   if  $l < \text{levels}$  then
10    for  $i \in \{2, \dots, |\text{cand}|\}$  do
11       $\text{diff} \leftarrow \text{cand}(i) - \text{cand}(i - 1)$ ;
12       $\alpha_{\text{new}} \leftarrow \left\{ \frac{-\text{diff}_1}{\text{diff}_2 - \text{diff}_1} \right\}$ ;
13      if  $(0.999 - \alpha_{\text{new}}) < 0.001$  or  $\alpha_{\text{new}} < 0.8$  then // failed
14        weight setting
15         $\alpha_{\text{new}} \leftarrow (\text{succ}(\alpha) + \text{pred}(\alpha))/2$ ;
16        // replace unsucc.  $\alpha$  by bisecting parent
17        interval;
18       $\Lambda \leftarrow \Lambda \cup \alpha_{\text{new}}$ ;
19    resort  $\Lambda$  increasingly
20 Return nondominated points from cand to illustrate trade-offs;
```

To overcome the numerical difficulties arising from this fact, the dichotomic scheme is enhanced by bisection steps that generate new - and promising -

weighting parameters α when needed. The bisection step occurs whenever the newly found weight falls outside the original search interval $[\alpha_1, \alpha_2]$. This can happen if the calculated slope of the dichotomic search becomes too small, leading to numerical problems. Algorithm 1 performs several full training runs, resetting the weights of the network after line 6. The variable `cand` is initialized with an empty set and stores all objective vectors $y = (\text{MSE}_{\mathcal{U}}, \text{MSE}_{\mathcal{F}})$ generated during training for different weighting parameters α of the loss function \mathcal{L}_{α} .

4. Results

This section is subdivided into three parts. In Section 4.1, our PINN is validated. We first use the data generated during the delta variant (85th to 100th considered week or week 41/2021 to week 4/2022) and predict the first omicron wave (100th to 104th considered week or week 4/2022 to week 8/2022). We use the NSFD scheme to determine fitting parameter values for β and κ .

In Section 4.2, we present the results of dichotomic search to investigate the impact of the weighting parameter α on the training objective (27). We discuss the ability of dichotomic search to approximate a Pareto front, which can help in finding reasonable trade-off solutions.

Finally, in Section 4.3, we present a result using a long training dataset, starting from the beginning of the pandemic in January 2020. The differences between the results obtained with the two training datasets are discussed. Thus, we distinguish between training datasets that cover the time since the outbreak of the pandemic in Germany, which we refer to as *long-term training data*, and training datasets that contain only data from a particular wave or even a peak reached during the pandemic, which we refer to as *short-term training data*.

4.1. Validation of Scenarios Generated with Short-Term Training Data

In this part, we use short-term training data and use it to predict the infected compartment I for the following weeks. The reason we use short-term data is that our physical model is built to describe the behavior of exactly one wave

of infection with one maximum. Since there were many waves of different virus variants described by different infection parameters in the COVID-19 pandemic, using only one wave for training to predict the next wave is a more realistic approach to incorporate our model. Nevertheless, we will present results with a longer training period of training data for comparison in Section 4.3. It is important to note that data covering the increase in infection rates due to omicron spread must be included in the training data to predict the further increase. The PINN adjusts its trained parameters to the underlying data set.

We first use the NSFD scheme to determine appropriate parameter values for the infection rates β and κ . Figure 6 shows (normalized) infected data in the time frame we use to train and validate our PINN, along with the predictions we obtained using the NSFD scheme (see Eq. (19)) with the parameters in the right column of Table 1, i.e., $\beta = 0.00000001476$ and $\kappa = 0.001$.

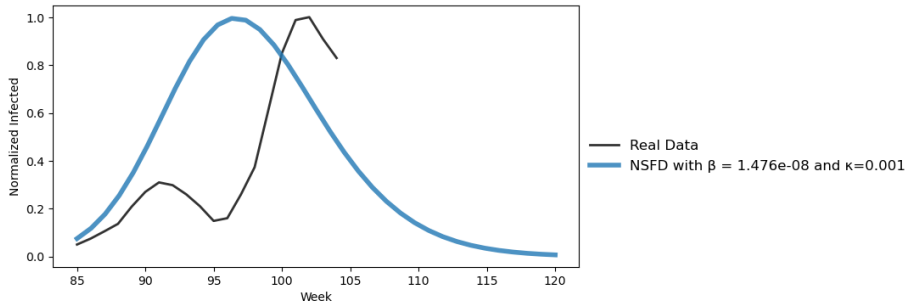


Figure 6: Normalized infection numbers from the 85th considered week to the 104th considered week and the normalized infected part of the NSFD scheme.

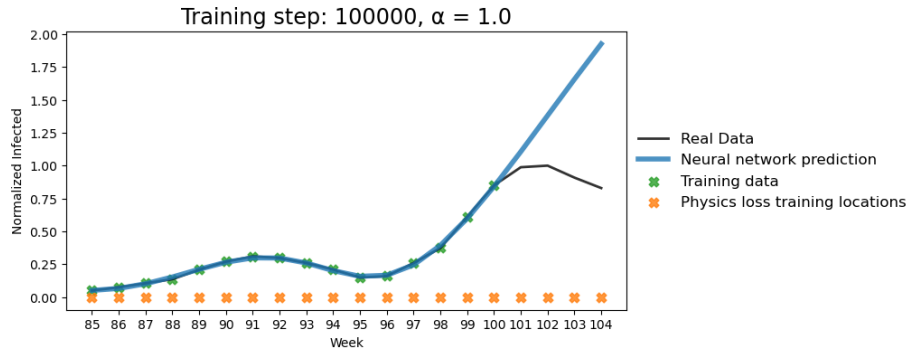
We see that this parameter choice reproduces well both the initial slope of the delta wave and the maximum of the omicron wave. The NSFD scheme with a constant transmission risk obviously predicts one maximum in the course of one wave. We note that the inclusion of a time-varying transmission rate in the NSFD approach would facilitate the prediction of multiple peaks within a wave, as shown in [21]. For simplicity, we have not included a time-varying rate in the PINN or NSFD scheme used in this work; this is left for future research.

Figure 7 shows the results of a short-term prediction in which the PINN was trained on data ranging from the 85th week under consideration to the 100th week under consideration and used to predict the following four weeks (weeks 101 – 104). We used three different values for the weighting parameter $\alpha \in \{1.0, 0.995, 0.99\}$, see Eq. (27). Each training run included 100000 iterations.

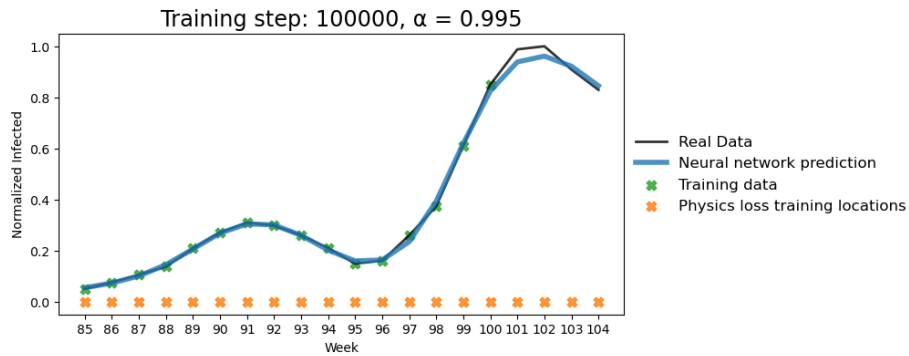
We note that the weighting parameter α plays an important role in the quality of the prediction. In Figure 7(a), $\alpha = 1$ was chosen, so the loss only considered $\text{MSE}_{\mathcal{U}}$. In this case, PINN succeeds in fitting the training data almost perfectly, but it fails to predict the decline in infection numbers in the following weeks because it has no information about the physical properties of the system, leading to overfitting. Figure 7(b) shows that using $\alpha = 0.995$ leads to a slightly less perfect fit to the training data, but a good prediction of the infection numbers in the following four weeks. Finally, Figure 7(c) shows that overweighting $\text{MSE}_{\mathcal{F}}$ again leads to suboptimal predictions, which can be explained by the fact that the PINN has learned the physical parameters of the delta wave, which are different from those of the omicron wave. It can therefore be observed that the data loss must be weighted much higher than the residual loss to obtain a good prediction in applications where the physical model only partially explains the true dynamics.

However, complete neglect of residual loss still does not lead to the best performance of our PINN. Although no mathematical model can optimally describe an infectious disease because not every epidemiological detail relevant to transmission is known, we can use residual loss to incorporate systematic knowledge about the spread and transmission dynamics of the disease into our neural network. Since $\alpha = 1$ does not yield the smallest errors in our validation runs, the inclusion of the residual error is still justified and reasonable. In general, predictions of future pandemic dynamics due to unknown mutational variants are subject to many uncertainties, changing intervention measures or compliance of the population and new vaccination strategies.

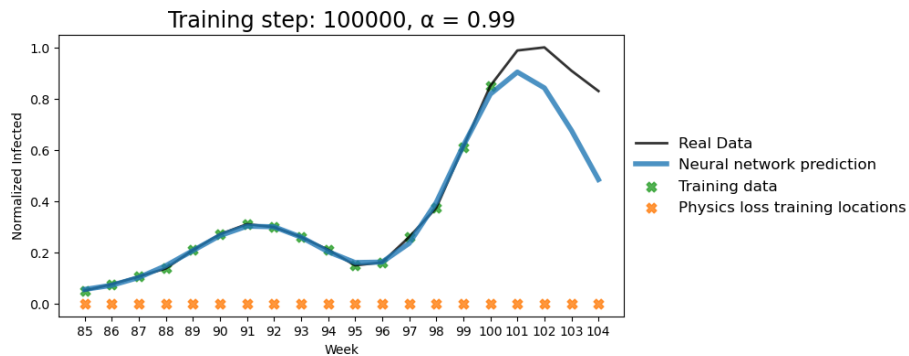
Table 2 shows the mean squared error between the infected compartment of our PINN prediction $\mathcal{K}_p^3(t_i)$ and the reported infection data $\hat{\mathcal{K}}_p^3(t_i)$ in the



(a) Perfect fit of training data, weak prediction.



(b) Good fit of training data, good prediction.



(c) The prediction gets worse if $MSE_{\mathcal{F}}$ is weighted too heavily.

Figure 7: Infection numbers obtained from the reported data (black) or the training of the PINN with the loss terms $MSE_{\mathcal{U}}$ and $MSE_{\mathcal{F}}$ (blue) from the 85th considered week to the 104th week for three different assignments of the weighting parameter α . The training data set covers the calendar weeks 41 in 2021 to 4 in 2022 (85th to 100th considered week).

Value of α	1.0	0.999	0.998	0.997	0.996	0.995	0.994
MSE _{val}	0.0961	0.0553	0.0505	0.0031	0.0012	0.0003	0.0024
Value of α	0.993	0.992	0.991	0.99	0.95	0.9	0.8
MSE _{val}	0.0016	0.0097	0.0109	0.0104	0.0772	0.0853	0.1030

Table 2: Mean squared errors between the reported infection data and the PINN results using the 85th to 104th considered week, depending on the weighting parameter α . This can be considered a validation because it includes the error for the prediction for weeks 100 to 104 in addition to the data used for training (weeks 85 to 100).

weeks 85 to 104 considered (i.e., for the entire time frame for which data were available, not just the time frame used for training), where the mean squared error is calculated as follows

$$\text{MSE}_{\text{val}} := \frac{1}{l} \sum_{i=1}^l \|\mathcal{K}_p^3(t_i) - \hat{\mathcal{K}}_p^3(t_i)\|^2. \quad (31)$$

To create Table 2, we manually modified the weighting parameter $\alpha \in (0.8, 1)$. It is remarkable that the smallest error is achieved with $\alpha = 0.995$ (among the considered parameter values.) With $\alpha = 0.996$, the error becomes four times as large and with $\alpha = 0.994$ it becomes eight times as large. Note, however, that the individual training runs do not necessarily terminate with a (globally) optimal solution; so the reported error values can only approximate the best possible error for the respective choices of weighting parameters. Nevertheless, we can observe clear reductions in the performance of the network when setting $\alpha = 0.9$ or $\alpha = 1.0$. For instance, we obtain an approximate eight-fold increase of MSE_{val} if the weighting parameter α is decreased from $\alpha = 0.99$ to $\alpha = 0.9$.

The above discussion shows that the choice of the weighting parameter α plays an important role in the prediction quality. Note that the prediction quality here is measured by comparisons with the real data. It is therefore not surprising that the prediction quality is higher when the data loss MSE_U is highly weighted in the weighted sum training objective (27). However, the results also show that the residual loss MSE_F should not be ignored. This

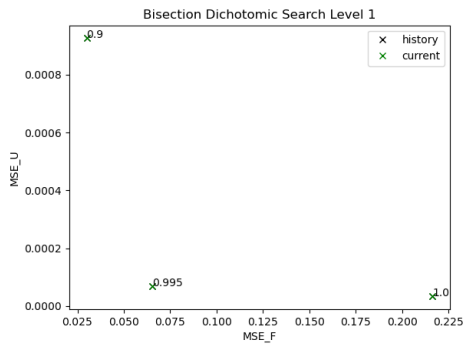
motivates a more detailed analysis of the trade-off between data loss and residual loss using dichotomic search in the following section.

4.2. Results of the Dichotomic Search

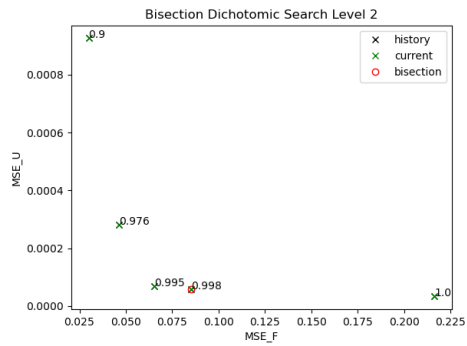
In this section, we present the results of dichotomic search using Algorithm 1. Based on the discussion in the previous sections, we focus on short-term training data from calendar week 41 in 2021 to 4 in 2022 (85th to 100th considered week). See Section 4.1 and Figure 7 for comparison. Four levels of Pareto front approximations (based on three, five, nine and 13 training runs, respectively) are shown in in Figure 8. Each training run was performed with 100000 epochs and the Adam optimizer with a learning rate schedule as defined in equation (29) with $t_{\text{start}} = 0.003$ and $t_{\text{end}} = 0.00015$. The search was initiated with $\alpha_1 = 0.9$ and $\alpha_2 = 1.0$. Other search windows can be used depending on preferences or if additional information is available. In addition, the search window can be used to zoom into a specific part of the Pareto front.

The results confirm that a pronounced Pareto front with a diverse set of outcome vectors and a clear knee solution was approximated after only a few training runs. Indeed, level 2 based on five training runs (see Figure 8(b)) already provides a rough approximation of the Pareto front. At level 3 (Figure 8(c)), after nine training runs, we already have a very good approximation of the Pareto front. The final level 4 is based on 13 training runs (Figure 8(d)). On a machine with an AMD Ryzen 7 3700X 8-Core Processor and NVIDIA GeForce RTX 3060 graphic card, each training run took about 3 minutes, so the final level needed a computation time of about 39 minutes. The computation time of each training run is of course dependent on the specific problem and network, however, the goal here was to show that one can gain insight into the complete Pareto front of the problem with rather few training runs.

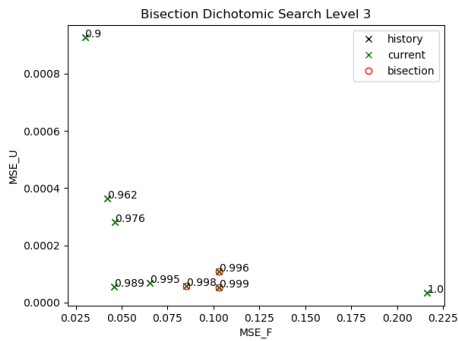
This shows that the dichotomic search scheme was successful in obtaining a diverse set of solutions with only very few training runs necessary, allowing the decision maker insight into the effect of changing the weighting parameter α that reflects the importance of each loss function.



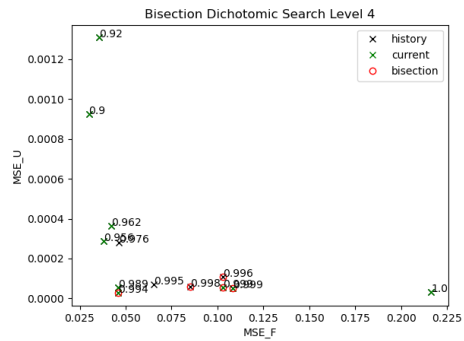
(a) Level 1: Lexicographic optimization and one weighted sum (three training runs)



(b) Level 2: Five training runs, bisection used



(c) Level 3: Nine training runs



(d) Level 4: Final Level with 13 training runs

Figure 8: Pareto front approximation using dichotomic search with a starting learning rate of 0.003 and learning rate schedule with starting weights $\alpha_1 = 0.9$ and $\alpha_2 = 0.999$. Favorable trade-offs can be identified starting from Level 3.

The results also show that it is possible to significantly improve the loss with respect to $MSE_{\mathcal{F}}$ without losing much in the $MSE_{\mathcal{U}}$ part of the loss function. While with $\alpha = 1$ we obtain $MSE_{\mathcal{F}} = 0.2176$ and $MSE_{\mathcal{U}} = 7.495 \cdot 10^{-5}$, $\alpha = 0.994$ leads to $MSE_{\mathcal{F}} = 0.0471$ and $MSE_{\mathcal{U}} = 7.599 \cdot 10^{-5}$. We can therefore use the BEDS algorithm to quickly and automatically identify favorable trade-offs in the parameter weighting of PINN loss functions.

Overall, it can be seen that the variation of α has a large impact on the relative importance of data loss and residual loss.

4.3. Validation of Scenarios Generated with Long-Term Training Data

Finally, we want to use our results to make a long-term prediction using most of our available data. We use the 1st to 89th week under consideration (week 10/2020 to week 45/2021) as our training data and let the PINN predict the following delta wave (89th to 95th week under consideration or week 45/2021 to week 51/2021). First, we reapply the NSFD scheme to reevaluate the parameter values for β and κ . Figure 9 shows the results of this procedure.

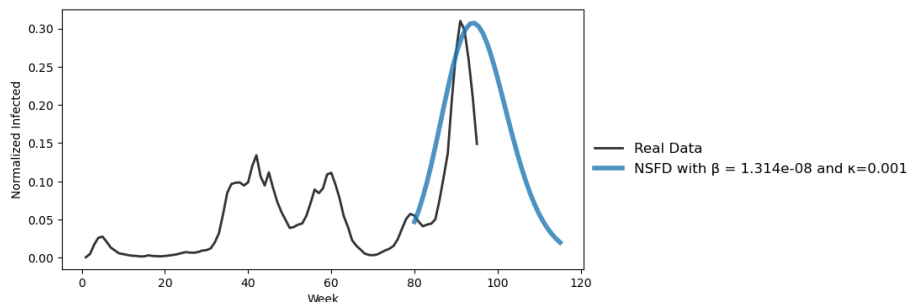
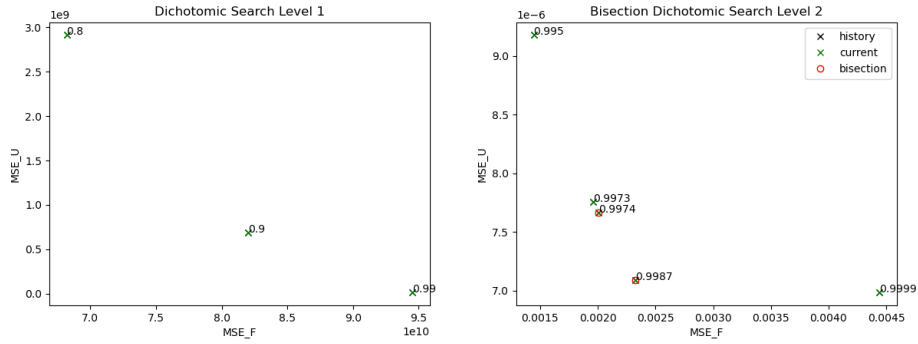


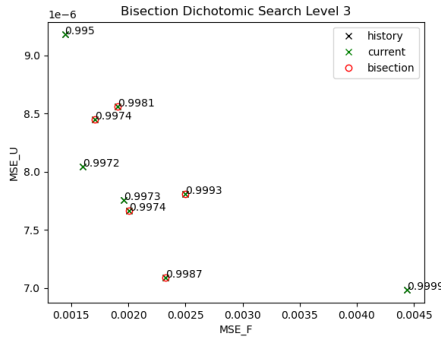
Figure 9: Normalized infection numbers from the 1st considered week to the 95th considered week and the normalized infected part of the NSFD scheme.

We choose $\beta = 0.00000001314$ and $\kappa = 0.001$. With these values, we start another dichotomic search to find a loss weighting with a good trade-off. Since the underlying physical system in this case fits the data less well than in the previous short-term prediction (since the long-term data consist of multiple waves with different infection peaks), we chose a search window closer to $\alpha = 1$

with $\alpha = [0.995, 0.9999]$ to weight the data loss even higher. The results can be seen in Figure 10.



(a) Level 1: Lexicographic optimization and one weighted sum (three training runs) (b) Level 2: Five training runs, bisection used



(c) Level 3: Nine training runs

Figure 10: Pareto front approximation using dichotomic search with a starting learning rate of 0.003 and learning rate schedule with starting weights $\alpha_1 = 0.995$ and $\alpha_2 = 0.9999$. Favorable trade-offs can be identified starting from level 3.

We identify the weighting $\alpha = 0.9987$ as a good trade-off solution that does not lose much in the $MSE_{\mathcal{U}}$ -loss and is much better than $\alpha = 0.9999$ in the $MSE_{\mathcal{F}}$ -loss. We decide to use this weighting for our long-term prediction. The results can be seen in Figure 11.

We note that the quality of the prediction is lower compared to the short-term prediction. This is not surprising given that the underlying SVIHR sys-

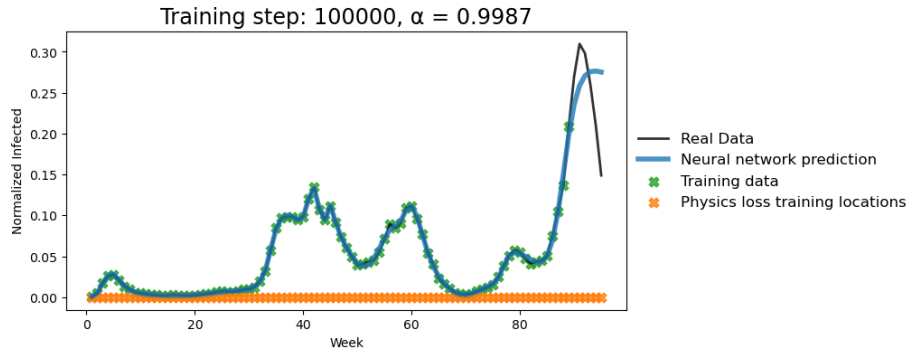


Figure 11: Infection numbers obtained from the reported data (black) or the training of the PINN with the loss terms $MSE_{\mathcal{U}}$ and $MSE_{\mathcal{F}}$ (blue) from the 1st considered week to the 95th week for $\alpha = 0.9987$. The training data set covers the calendar weeks 10 in 2020 to 45 in 2021 (1st to 89th considered week).

tem is designed to simulate one wave of infection. To improve these results, time-varying functions for vaccination and transmission rates would need to be included, which we plan to investigate in future work.

5. Conclusion and Outlook

We consider *physics-informed neural networks* (PINNs), a Deep Learning technique that combines data and physical knowledge. The predictions for COVID-19 infection rates are used as a case study. Our main contribution is a new perspective on the trade-off between data loss and residual loss in PINN training. We present an inherently biobjective method that efficiently identifies near-ideal knee solutions. This is complemented by an inverse modeling approach to determine the model parameters that govern the dynamics of the system, based on a numerical solution of the associated system of differential equations.

In our approach, the data loss is computed for data on COVID-19 infection rates in Germany. The PINN predicts the sizes of the compartments of an established *susceptible-vaccinated-infected-hospitalized-recovered* (SVIHR) model, with a focus on the compartment of infected individuals. The residual loss is

derived from a system of ODEs based on the proposed SVIHR model, which mathematically describes the dynamics of transitions between different compartments and infectious disease transmission in a population affected by the COVID-19 pandemic. We propose an NSFD scheme especially designed for the numerical solution of the SVIHR model. Its solution is used to estimate the transmission risk β and the residual transmission probability after vaccination κ that govern the dynamics.

Our results show that the prediction quality of the PINN is highly dependent on the considered data, and on the weighing parameter α that combines the two conflicting loss functions in a weighted sum \mathcal{L}_α . A combination of short-term data and an optimized choice of the weighting parameter α leads to a very good prediction of the first omicron wave.

The approach is based on a dichotomic search method that iteratively approximates the Pareto front and identifies near-ideal knee solutions (i.e., trained networks) with comparatively few training runs. We found that the preferred values of the weighting parameter α in the total loss function were greater than 0.99 in most cases, thus giving greater weight to data loss than to residual loss. This numerical value has to be interpreted with care due to the different scalings of the data loss and the residual loss. We emphasize that the biobjective approach to PINN training can be extended to PINNs regardless of the number of loss terms or application field. While appropriate weighting parameters α could theoretically be obtained by scanning the weights of reasonable candidate values, this would require a prohibitively large number of individual neural network trainings, growing exponentially with the number of loss functions.

Finally, we use NSFD parameter estimation and dichotomic search to perform long-term prediction for COVID-19 infection rates, using most of the available data as training data for the delta wave prediction. We find that the PINN provides reasonable results even for long-term predictions. However, since the dynamics in the long term are affected by many different aspects and measures, the physical model is even less accurate in this case, which explains the slightly worse performance compared to the short term predictions.

In future work, time-varying functions for vaccination and transmission rates will be included, i.e. $\beta = \beta(t)$, $\kappa = \kappa(t)$, to account for seasonal and variant-dependent fluctuations, cf. [45]. Time-variability in the transmission rate should especially be taken into consideration with respect to long-term predictions, in which the training data set covers multiple months and thus includes different mutations, lockdown measures, and seasons. To this end, it is also necessary to incorporate a reflux from the recovered compartment R to the susceptible compartment S . It is worth noting that we only dealt with two loss terms here and therefore used a biobjective optimization model. The concept of multiobjective PINN training can easily be extended to consider more than two objective functions. This occurs in more advanced applications such as solving PDEs by PINNs, e.g., using gradient-enhanced PINNs (gPINNs) [46], in which case the number of loss terms is greater than two and a multiobjective approach is required. For example, in [46, Section 2.2], the number of loss terms is $3 + d$, where d is the dimension of the spatial domain. Here, one has to identify similar and conflicting training objectives in order to avoid an overly large parameter set. This extension to multiobjective optimization approaches will be the content of a follow-up article.

References

- [1] M. Raissi, P. Perdikaris, G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707. doi:10.1016/j.jcp.2018.10.045.
- [2] J. Blechschmidt, O. Ernst, Three ways to solve partial differential equations with neural networks – A review, *GAMM-Mitteilungen* 44 (2) (2021) e202100006.
- [3] J. Hoffer, A. Ofner, F. Rohrhofer, et al., Theory-inspired machine learning – towards a synergy between knowledge and data, *Weld World* (2022). doi:10.1007/s40194-022-01270-z.

- [4] Robert Koch-Institute, last access: May 6, 2022 (2022). [link].
URL https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Virologische_Basisdaten.html;jsessionid=960297181B52EB351E833DD09CD96CAB.internet081?nn=13490888#doc14716546bodyText6
- [5] Ourworldindata.de, last access: May 5, 2022 (2022). [link].
URL <https://ourworldindata.org/coronavirus>
- [6] J. Malinzi, S. Gwebu, S. Motsa, Determining COVID-19 dynamics using physics informed neural networks, *Axioms* 11 (3) (2022) 121. doi:10.3390/axioms11030121.
- [7] E. Kharazmi, M. Cai, X. Zheng, G. Lin, G. Karniadakis, Identifiability and predictability of integer- and fractional-order epidemiological models using physics-informed neural networks, *Nature Comput. Sci.* 1 (11) (2021) 744–753. doi:10.1101/2021.04.05.21254919.
- [8] G. Pang, L. Lu, G. E. Karniadakis, fPINNs: fractional physics-informed neural networks, *SIAM J. Sci. Comput.* 41 (4) (2019) A2603–A2626. doi:10.1137/18M1229845.
- [9] M. Cai, G. E. Karniadakis, C. Li, Fractional SEIR model and data-driven predictions of COVID-19 dynamics of Omicron variant, *Chaos* 32 (7) (2022) 071101. doi:10.1063/5.0099450.
- [10] F. Rohrhofer, S. Posch, B. Geiger, On the Pareto front of physics-informed neural networks, arXiv preprint 2105.00862 (2021).
- [11] X. Jin, S. Cai, H. Li, G. E. Karniadakis, NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations, *J. Comput. Phys.* 426 (2021) 109951. doi:10.1016/j.jcp.2020.109951.

- [12] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, *SIAM J. Sci. Comput.* 43 (5) (2021) A3055–A3081. doi:10.1137/20M1318043.
- [13] S. Maddu, D. Sturm, C. Müller, I. Sbalzarini, Inverse Dirichlet weighting enables reliable training of physics informed neural networks, *Machine Learning: Sci. Techn.* 3 (2021) 015026. doi:10.1088/2632-2153/ac3712.
- [14] J.-A. Désidéri, Multiple-gradient descent algorithm (MGDA) for multi-objective optimization, *Compt. Rend. Math.* 350 (2012) 313–318. doi:10.1016/j.crma.2012.03.014.
- [15] J. Fliege, B. Svaiter, Steepest descent methods for multicriteria optimization, *Math. Meth. Oper. Res.* 51 (3) (2000) 479–494. doi:10.1007/s001860000043.
- [16] O. Sener, V. Koltun, Multi-task learning as multi-objective optimization, *CoRR* (2018). arXiv:1810.04650v2.
- [17] S. Liu, L. Vicente, The stochastic multi-gradient algorithm for multi-objective optimization and its application to supervised machine learning, *Annals Oper. Res.* (2021). doi:10.1007/s10479-021-04033-z.
- [18] L. McClenny, U. Braga-Neto, Self-adaptive physics-informed neural networks using a soft attention mechanism, arXiv preprint 2009.04544 (2020).
- [19] A. F. Psaros, K. Kawaguchi, G. E. Karniadakis, Meta-learning PINN loss functions, *J. Comput. Phys.* 458 (2022) 111121. doi:10.1016/j.jcp.2022.111121.
- [20] M. Reiners, K. Klamroth, F. Heldmann, M. Stiglmayr, Efficient and sparse neural networks by pruning weights in a multiobjective learning approach, *Comput. Oper. Res.* 141 (2022) 105676. doi:10.1016/j.cor.2021.105676.

- [21] S. Berkhahn, M. Ehrhardt, A physics-informed neural network to model COVID-19 infection and hospitalization scenarios, *Adv. Cont. Discr. Mod. Theo. Appl.* 2022 (61) (2022). doi:10.1186/s13662-022-03733-5.
- [22] W. Kermack, A. McKendrick, Contributions to the mathematical theory of epidemics, *Bull. Math. Bio.* 53 (1) (1991) 700–721. doi:10.1016/S0092-8240(05)80040-0.
- [23] S. Treibert, H. Brunner, M. Ehrhardt, A nonstandard finite difference scheme for the SVICDR model to predict COVID-19 dynamics, *Math. Biosci. Engrg.* 19 (2022) 1213–1238. doi:10.3934/mbe.2022056.
- [24] Robert Koch-Institute, last access: March 29, 2022 (2022). [link].
URL https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Daten/Fallzahlen_Kum_Tab.html
- [25] Robert Koch-Institute, last access: March 29, 2022 (2022). [link].
URL https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Daten/Klinische_Aspekte.html
- [26] COVID-19 Vaccination Dashboard, last access: May 6, 2022 (2022). [link].
URL <https://impfdashboard.de/daten>
- [27] M. Martcheva, *An Introduction to Mathematical Epidemiology*, 1st Edition, Springer, 2015. doi:10.1007/978-1-4899-7612-3.
- [28] Robert Koch-Institute, last access: March 29, 2022 (2022). [link].
URL https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Steckbrief.html
- [29] Federal Statistical Office, Germany, last access: January 10, 2022 (2022). [link].
URL https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Bevoelkerung/Bevoelkerungsstand/_inhalt.html

- [30] R. E. Mickens, Exact solutions to a finite-difference model of a nonlinear reaction-advection equation: Implications for numerical analysis, *J. Diff. Eqs. Appl.* 9 (11) (2003) 313–325. doi:10.1080/1023619031000146959.
- [31] R. E. Mickens, *Applications of nonstandard finite difference schemes*, World Scientific (2000).
- [32] M. Maamar, M. Ehrhardt, L. Tabharit, A nonstandard finite difference scheme for a time-fractional model of Zika virus transmission, *IMACM Preprint 22/21* (2022).
- [33] G. E. Karniadakis, Y. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nat. Rev. Phys.* 3 (2021) 422–440. doi:10.1038/s42254-021-00314-5.
- [34] S. Cuomo, V. S. di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics-informed neural networks: Where we are and what’s next (2022). doi:10.48550/ARXIV.2201.05624.
- [35] K. D. Olumoyin, A. Q. M. Khaliq, K. M. Furati, Data-driven deep-learning algorithm for asymptomatic COVID-19 model with varying mitigation measures and transmission rate, *Epidemiologia* 2 (4) (2021). doi:10.3390/epidemiologia2040033.
- [36] S. Shaier, M. Raissi, P. Seshaiyer, Data-driven approaches for predicting spread of infectious diseases through DINNs: Disease informed neural networks (2021). doi:10.48550/ARXIV.2110.05445.
- [37] V. Grimm, A. Heinlein, A. Klawonn, M. Lanser, J. Weber, Estimating the time-dependent contact rate of SIR and SEIR models in mathematical epidemiology using physics-informed neural networks, *Electr. Trans. Numer. Anal.* 56 (2022) 1–27. doi:10.1553/etna_vol156s1.
- [38] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, J. M. Siskind, Automatic differentiation in machine learning: a survey, *J. Mach. Learn. Res.* 18 (2017) 5595–5637.

- [39] Ben Moseley, last access: December 12, 2022 (2021). [link].
URL <https://github.com/benmoseley/harmonic-oscillator-pinn>
- [40] S. Ruder, An overview of gradient descent optimization algorithms (2016).
doi:10.48550/ARXIV.1609.04747.
- [41] M. Ehrgott, Multicriteria Optimization, 2nd Edition, Springer, 2005.
- [42] A. Przybylski, K. Klamroth, R. Lacour, A simple and efficient dichotomic search algorithm for multi-objective mixed integer linear programs, arXiv preprint 1911.08937 (2019).
- [43] I. Das, J. Dennis, A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems, Struct. Optim. 14 (1997) 63–69. doi:10.1007/BF01197559.
- [44] K. Klamroth, K. Miettinen, Integrating approximation and interactive decision making in multicriteria optimization, Operations Research 56 (2008) 222–234. doi:10.1287/opre.1070.0425.
- [45] M. Jagan, M. deJonge, O. Krylova, D. Earn, Fast estimation of time-varying infectious disease transmission rates, PLoS Comput. Biol. 16 (9) (2020) e1008124. doi:10.1371/journal.pcbi.1008124.
- [46] J. Yu, L. Lu, X. Meng, G. Karniadakis, Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems, Comput. Meth. Appl. Mech. Engrg. 393 (2022) 114823. doi:10.1016/j.cma.2022.114823.
- [47] T. Nagel, M. F. Huber, Kalman-Bucy-informed neural network for system identification, arXiv preprint arXiv:2210.03424 (2022).
- [48] L. Yang, X. Meng, G. E. Karniadakis, B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data, J. Comput. Phys. 425 (2021) 109913. doi:10.1016/j.jcp.2020.109913.