

CORONAVIRUS

Dynamics-informed deconvolutional neural networks for super-resolution identification of regime changes in epidemiological time series

Jose M. G. Vilar^{1,2*} and Leonor Saiz^{3*}

The ability to infer the timing and amplitude of perturbations in epidemiological systems from their stochastically spread low-resolution outcomes is crucial for multiple applications. However, the general problem of connecting epidemiological curves with the underlying incidence lacks the highly effective methodology present in other inverse problems, such as super-resolution and dehazing from computer vision. Here, we develop an unsupervised physics-informed convolutional neural network approach in reverse to connect death records with incidence that allows the identification of regime changes at single-day resolution. Applied to COVID-19 data with proper regularization and model-selection criteria, the approach can identify the implementation and removal of lockdowns and other nonpharmaceutical interventions (NPIs) with 0.93-day accuracy over the time span of a year.

INTRODUCTION

Inferring the underlying causative signal from the observed effector response through a system extended in space, time, or both is a fundamental problem in many fields. The crucial challenge is to obtain the input signal as a function of the system output. This type of inverse problem is widespread in multiple disciplines and, in many of them, such as computer vision, has reached highly sophisticated levels (1, 2). Prominent examples include recovering the scene radiance as if there were no perturbations along the line of sight from the natural environment, as in dehazing (3), and deconvolving the spread of an image, as in super-resolution single-molecule localization microscopy (4) and deblurring (5). Compounded with the spreading and perturbations of the signal, the unavoidable presence of random noise further deteriorates the output and brings about the need to use regularization methods (6).

Inverse problems in epidemiological systems involve inferring the daily incidence, defined as new infections per day, from the number of infection-caused deaths or positive testing individuals. New daily infections over time provide the input signal. As the disease progresses in an infected individual, it has a probability of being detected and a probability of leading to death at every time point after infection. Each of these probabilities depends on the characteristics of the disease and is population specific, including age structure, clinical resources, and testing capabilities (7). Daily positive testing results and daily deaths provide the measured outputs. This problem in epidemiology compounds multiple problems such as those present in computer vision. As in blurring, the output of daily new infections is spread, in this case over time instead of over space; as in hazing, the output provides distorted information because not all infections are detected or lead to death; and, as in super-resolution, it needs to provide information at a

much finer level than the spread of the output. In addition, the output is also inherently noisy due to, among others, the fundamental stochasticity of the testing and death processes. Because of all these effects compounded together and the traditional widespread lack of reliable field data availability, the methodology in epidemiological systems has not been as developed as in computer vision and other fields.

The state of the art in the field is epitomized by the deconvolution of daily cases or mortality time series to infer infection incidence. For instance, death records have been used to infer the influenza incidence through an iterative method with early stopping as a regularization approach, which prevents the amplification of noise originating from the inherent randomness of the death process (8). Deconvolution of infection cases has also been used to reconstruct the infection curves for severe acute respiratory syndrome (SARS) epidemics (9) through the expectation maximization smoothing (EMS) method (10). In the case of coronavirus disease 2019 (COVID-19), there have also been recent refinements to account for model misspecification and censored observations (11) and to estimate new infections in real time (12). The prototypical state-of-the-art approach provides much more reliable information than simply considering incidence as the number of newly detected positive cases (13) but is not sufficient to provide a daily incidence from which to infer discontinuities in transmission, as determined by jumps in the instantaneous reproduction number, which will change immediately following the implementation of a nonpharmaceutical intervention (NPI) (14). The primary reason is that these inverse approaches consider regularization methods that promote a smooth continuous incidence.

Here, we develop a convolutional neural network approach in reverse, termed dynamics-informed deconvolutional neural network (DIDNN), to connect death records with incidence that recovers the underlying dynamics with such a precision as to allow the identification of regime changes at single-day resolution. We address major challenges that in computer vision are not as exacerbated, including the presence of a high dynamic range (three orders of magnitude, ranging from zero to over a thousand deaths

Copyright © 2023 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).

¹Biofisika Institute (CSIC, UPV/EHU), University of the Basque Country (UPV/EHU), P.O. Box 644, 48080 Bilbao, Spain. ²IKERBASQUE, Basque Foundation for Science, 48011 Bilbao, Spain. ³Department of Biomedical Engineering, University of California, 451 E. Health Sciences Drive, Davis, CA 95616, USA.

*Corresponding author. Email: j.vilar@ikerbasque.org (J.M.G.V.); lsaiz@ucdavis.edu (L.S.)

per day), increasing and decreasing exponential growth regimes (doubling times and half-lives of ~3 days), kernels widely spread over time [mean of ~20 days with standard deviation (SD) of ~10 days], and the presence of large amounts of noise (Poisson processes with mean below one). To this end, we derive an objective function to be optimized through the network that considers a suitable loss for the underlying Poisson process; is quasi-scale invariant; and includes, following a physics-informed neural network (PINN) approach, the determinants of the epidemiological dynamics. The PINN contribution to the objective function promotes sparsity in the changes in disease transmission as defined by the underlying renewal equation for the incidence. Methodologically, the approach faces the challenges of deblurring with components of dehazing, denoising, and super-resolution.

To benchmark the approach with field data, we focus on the effects of major NPIs on the global spread of the COVID-19 outbreak (15). NPIs, aimed at controlling infection transmission without medication or vaccines, can include measures as diverse as social distancing, quarantine, travel restrictions, school closures, and wearing face masks, among others. Being able to identify, with single-day resolution, the implementation time of major NPIs from the epidemiological data has not been achieved so far. It is important from the biomedical point of view to test for potential differences in the clinical evolution of the disease over time and over different countries (16). Existing analysis of restrictions, imposed quarantines, lockdowns, and other NPIs have estimated a substantially spread delay between the implementation of lockdowns and their effects on the dynamics of the outbreak as measured by the basic reproduction number from case counts (17). However, this delay and spread should not be present in the infection transmission. Other approaches based on compartmental models (18) or Bayesian analyses (19) have relied on the inclusion of temporal-specific information, namely, the date of the NPI, to estimate the amplitude of the effects (20).

The focus of our approach is to enable the acquisition of precise knowledge about the NPIs that had an extensive impact on the epidemiological dynamics, such as the earlier measures that transformed the trajectory of the pandemics from growth to decline, and how NPI effects depend on the context and the presence of other NPIs. We have focused on the early stage of the COVID-19 outbreak because of its inherent interest and because it provides stringent validation conditions for the approach. These conditions allowed us to favorably contrast our results against well-defined NPIs and well-characterized infection-to-death delays and generation times.

From a fundamental point of view, the COVID-19 outbreak provides a unique opportunity to contrast the methodology with diverse types of field data. It has been among the most, if not the most, well-studied and well-characterized pandemics in human history. There have been rapid and extensive responses from the global scientific community, as well as an unprecedented level of global cooperation and coordination in responding to the outbreak. The detailed characterization over small, controlled populations has provided detailed information from the infection to death or recovery and the chain of infections. This characterization involved testing for the causative virus SARS coronavirus 2, identifying contacts and the infection time, and following up the clinical evolution (7, 21–23). Therefore, it is highly relevant to verify to what extent this characterization holds at a large scale (country-wide level)

over diverse heterogeneous populations (different countries). A detailed understanding of the outbreak, including how clinical and epidemiological aspects are related to each other, is expected to provide a substantial background to prepare for and respond to similar situations in the future.

RESULTS

Convolutional approach

We have derived a physics-informed convolutional neural network able to identify and quantify the distinct growth regimes of the epidemiological dynamics as described below. Mathematically, new infections (incidence) and deaths on day t are described as Poisson processes with intensities i_t for the number of daily new infections and λ_t for the number of daily deaths. The intensities of these two processes are related to each other through the convolution

$$\lambda_t = \sum_{\tau=0}^{t-1} f_{\tau} r_{t-\tau} i_{t-\tau} = \sum_{\tau=0}^{t-1} f_{\tau} j_{t-\tau} \tag{1}$$

where f_{τ} denotes the time from infection-to-death probability mass function, which gives the probability for a death to have occurred on day τ after infection, and r_t refers to the infection fatality ratio (IFR), which gives the probability for an individual infected on day t to die according to f_{τ} . We express the convolution in terms of the scaled incidence $j_t = r_t i_t$, which represents the number of new infections on day t that will eventually result in death. Throughout the derivations below, we consider the time interval $t \in [1, T]$.

Scale-invariant probabilistic loss function

The number of daily deaths n_t is assumed to follow a Poisson distribution with parameter λ_t

$$P_{n_t} = \frac{\lambda_t^{n_t}}{n_t!} e^{-\lambda_t} \tag{2}$$

which implies a log-likelihood function for λ_t given by

$$l(n_t, \lambda_t) = n_t \ln \lambda_t - \lambda_t - \ln(n_t!) \tag{3}$$

Therefore, the intensity of the Poisson process λ_t corresponds to the expected number of daily deaths.

There are multiple loss functions that can potentially be used. The focus is on growth rate changes, namely, relative changes, which are invariant upon scaling of the population by a constant factor. To effectively approximate this invariance in the loss function, which using the Stirling approximation can be rewritten as $l(n_t, \lambda_t) \simeq n_t \ln \frac{\lambda_t}{n_t} + n_t - \lambda_t$, we weight the negative log-likelihood function at each time point by $(1 + \lambda_t)^{-1}$. We use $(1 + \lambda_t)^{-1}$ as a weight, instead of λ_t^{-1} , to avoid undetermined results for sustained time intervals with zero deaths. It leads to

$$\mathcal{L}_{\Lambda} = -\frac{1}{Z} \sum_{t=1}^T \frac{1}{1 + \lambda_t} l(n_t, \lambda_t), \text{ with } Z = \sum_{t=1}^T \frac{1}{1 + \lambda_t} \tag{4}$$

as a loss function for the overall time interval. This loss function has a minimum value for $\lambda_t = n_t, \forall t \in [1, T]$.

Dynamics and regularization

The epidemiological dynamics is characterized by the renewal equation $i(t) = \int_0^\infty k_I(t, \tau) i(t - \tau) d\tau$, with $k_I(t, \tau)$ being the rate of secondary transmissions per single primary case (24, 25). This convolution equation provides the basis for the definitions of the instantaneous reproduction number $R_t = \int_0^\infty k_I(t, \tau) d\tau$ and the probability density of the generation time $w(\tau) = \frac{k_I(t, \tau)}{R_t}$. The instantaneous reproduction number R_t accounts for the transmission of the infectious population at time t and the generation time $w(\tau)$ describes the infectiousness profile at time τ after infection.

We consider the discrete counterpart of the dynamic equations and changes in the IFR at a slower time scale than the incidence, so that the discrete renewal equation can be rewritten in terms of the normalized incidence as

$$j_t = R_t \sum_{\tau=1}^{t-1} w_\tau j_{t-\tau} \tag{5}$$

where w_τ is the probability mass function of the generation time.

We implement a regime-change-aware regularization approach through

$$\mathcal{L}_R = \frac{1}{T-1} \sum_{t=1}^{T-1} |\ln R_{t+1} - \ln R_t| \tag{6}$$

as regularization term in the objective function. This expression implies sparseness in $|\ln R_{t+1} - \ln R_t|$, and therefore, it is suitable to identify discontinuous changes in R_t . The logarithm is used to account for expanding and decaying dynamics symmetrically.

Dynamics-informed convolutional neural network

Analogously to the established PINN approach, we initially consider the loss function of the epidemiological dynamics as

$$\mathcal{L}_D = \frac{1}{T-1} \sum_{t=2}^T \left(j_t - R_t \sum_{\tau=1}^{t-1} w_\tau j_{t-\tau} \right)^2 \tag{7}$$

This equation has a minimum when the dynamics of j_t is given by Eq. 5. Strictly enforcing the epidemiological dynamics, $\mathcal{L}_D = 0$, leads to $R_t = j_t / \sum_{\tau=1}^{t-1} w_\tau j_{t-\tau}$, which allows the combination of \mathcal{L}_D and \mathcal{L}_R into a unified loss

$$\mathcal{L}_{RD} = \frac{1}{T-2} \sum_{t=2}^{T-1} \left| \ln \frac{j_{t+1}}{j_t} - \ln \frac{\sum_{\tau=1}^{t-1} w_\tau j_{t-\tau}}{\sum_{\tau=1}^{t-1} w_\tau j_{t+1-\tau}} \right| \tag{8}$$

which we use to include both the dynamics and regularization in the objective function to be optimized by the neural network.

Explicitly, we infer the estimated scaled incidence $\hat{\mathbf{j}} = (\hat{j}_1, \dots, \hat{j}_T)$ as

$$\hat{\mathbf{j}} = \underset{j_i: \forall i \in [1, T]}{\operatorname{argmin}} (\mathcal{L}_\Lambda + \gamma \mathcal{L}_{RD}) \tag{9}$$

where γ sets the strength of the regularization contribution. In this way, regularization prevents overfitting of the neural network by encouraging the model to follow the epidemiological dynamics

characterized by a renewal equation with discontinuous changes in the instantaneous reproduction number.

The structure of the neural network is shown in Fig. 1. It is based on a dual convolution neural network in reverse. Namely, the direct approach infers the convolution kernels from input and output signals. Instead, we infer the unknown signal (scaled incidence) from the known kernels (time from infection to death and generation time probability mass functions) and output signal (daily deaths). Therefore, we use as input layer the kernel with values f_τ , which is fed into a transposed convolution layer with weights j_D , which is both connected to the output layer with values λ_t and regularized according to the dynamics through a convolution with the kernel w_τ .

Model selection

Through model selection, we infer the optimal strength of the regularization term γ . We follow explicitly two approaches. Prospectively, γ is selected according to the Akaike information criterion (AIC) (26) as the value of γ that leads to the minimum of

$$L_{AIC} = \sum_{\text{locations}} \left(2 \sum_{t=2}^{T-1} \Theta(|\ln R_{t+1} - \ln R_t|) + 2T \mathcal{L}_\Lambda \right) \tag{10}$$

where $\Theta(\cdot)$ is the Heaviside step function. The term $\sum_{t=2}^T \Theta(|\ln R_{t+1} - \ln R_t|)$ represents the number of parameters in the model, as shown generally for the number of nonzero parameters in lasso problems (27), and $T \mathcal{L}_\Lambda$ is the minimum value of the negative loglikelihood function for the model, expressed as the daily average negative loglikelihood, used as a loss function, times the number of days. Retrospectively, γ is selected as the value that minimizes the overall mean squared error between inferred and actual NPI implementation times for all the locations.

Application of the approach reveals a discontinuous incidence

We applied the approach to the identification of regime changes in the COVID-19 infectious dynamics in European locations and to the quantification of how these changes correlate with the timing of NPIs. Explicitly, the expected number of daily deaths λ_t and the scaled incidence j_t were obtained by training the network over the daily deaths with the generation time w_τ and infection-to-death kernels f_τ as inputs (Fig. 1). The focus is on locations with the exact date of death on record to target the precision of the approach at the single-day level, which is substantially smaller than the observed >3-day average delay in reporting deaths (28, 29), and on the first year of the outbreak, the period that concentrated major NPIs before the widespread evolution of viral variants, substantial vaccination, and large levels of acquired immunity.

The results (Fig. 2), after model selection with the strength of the regularization term $\gamma = 2.51$, show the daily deaths fluctuating around their inferred expected values λ_t without any substantial bias, with higher relative fluctuations at low expected values. The overall average value of the relative fluctuations, quantified as $\langle (n_t - \lambda_t)^2 / \lambda_t \rangle$, is 1.22 deaths/day. The fact that this value is essentially 1 confirms the ability of the network to consistently track the nature of the underlying Poisson process over a wide dynamic

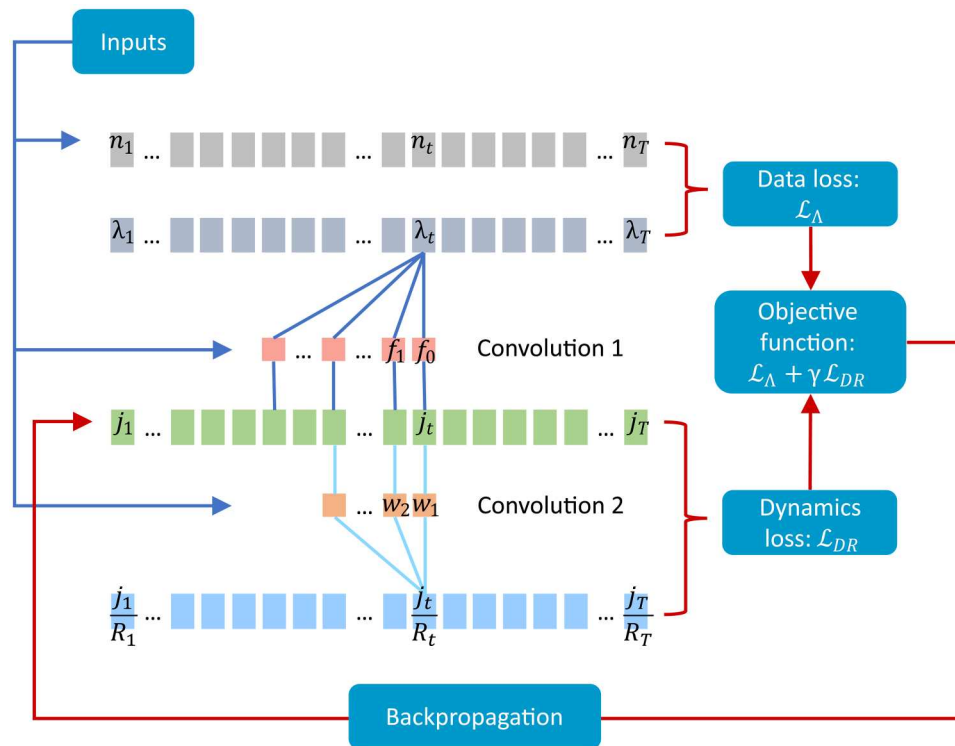


Fig. 1. Structure of the DIDNN for super-resolution identification of epidemiological regime changes. A CNN is used in reverse. The inputs consist of the filters of the two convolutions, w_t and f_t , and the daily deaths, n_t . The daily scaled incidence j_t is obtained after minimization of the objective function through backpropagation, which includes both data, \mathcal{L}_Λ , and dynamics-informed, \mathcal{L}_{DR} , losses. The forward convolutions lead to the estimated expected daily deaths, λ_t , (convolution 1), and to the estimated instantaneous reproduction number, R_t (convolution 2).

range, represented on a logarithmic scale, from values below 1 up to above 10^3 , with both expanding and decaying dynamics.

In stark contrast to the smoothness of the expected number of deaths λ_t , the network leads to a rugged, discontinuous behavior for the scaled incidence j_t . These discontinuities, as implied by the infection transmission described by Eq. 5, originate from a discontinuous behavior in the instantaneous reproduction number. The key question is if the incidence recovered from the smooth fit to daily deaths accurately describes the epidemiological dynamics. In general, the inferred changes in the instantaneous reproduction number can correspond to the regime changes or to optimal segmentations of a continuously changing dynamics. The ability of the approach to capture the effects of NPIs, which are expected to lead to regime changes, provides a stringent validation avenue.

Validation identifies regime changes with major NPIs

Explicitly, the validation of the approach is conducted by predicting the precise timing of major NPIs, including the implementation and removal of lockdowns and other restrictions in social interactions (table S1). The date of the NPI is inferred as the date of the largest R_t change in a 9-day window centered at the date of a recorded NPI. We define the offset, Δ , as the inferred minus the actual time of the NPI.

For the initial stages of the outbreak, the approach accurately identifies the dates of the major lockdowns and social restrictions, with an overall offset for all locations of 0.22 (mean) ± 0.63 (SD) days (Fig. 3). These results are particularly remarkable both in

their precision—considering that incidence stochastically spreads into deaths with a delay of 19.3 days, on average, and width of ± 9.1 (SD) days—and in their uniformity, with all the parameters of the network being the same for all the locations, including the value of the regularization term. For the whole period considered, the approach retained the accuracy at the single-day level, with an overall offset of -0.07 ± 0.92 days (Fig. 4). In addition to the high concordance across countries uncovered in the initial stages of the outbreak, identification of NPI timings over a year also highlights the consistency over time, as no systematic biases are observed.

The NPIs involved comprise a wide range of measures implemented to reduce the transmission of infections (table S1). The earliest ones identified for each country in the year 2020 illustrate this diversity. They include the quarantine of 11 municipalities in Northern Italy on 22 February; the nationwide lockdown in Spain on 14 March; prohibition of gatherings of more than 10 people and closing of all shops, except grocery stores and pharmacies in Denmark on 18 February; cancellation of classes in all educational establishments, banning all events involving more than 100 people, and partial closing of borders in Switzerland on 13 March; and multiple public event cancellations, including major sporting events in England on 13 March. Many of these measures were followed by more stringent ones, such as a subsequent nationwide lockdown in England on 23 March, or more widespread ones, such as the cancellation of exhibitions and events and the closing of all schools and universities nationwide in Italy on 4 March. Measures such as masking were identified as implemented together with other

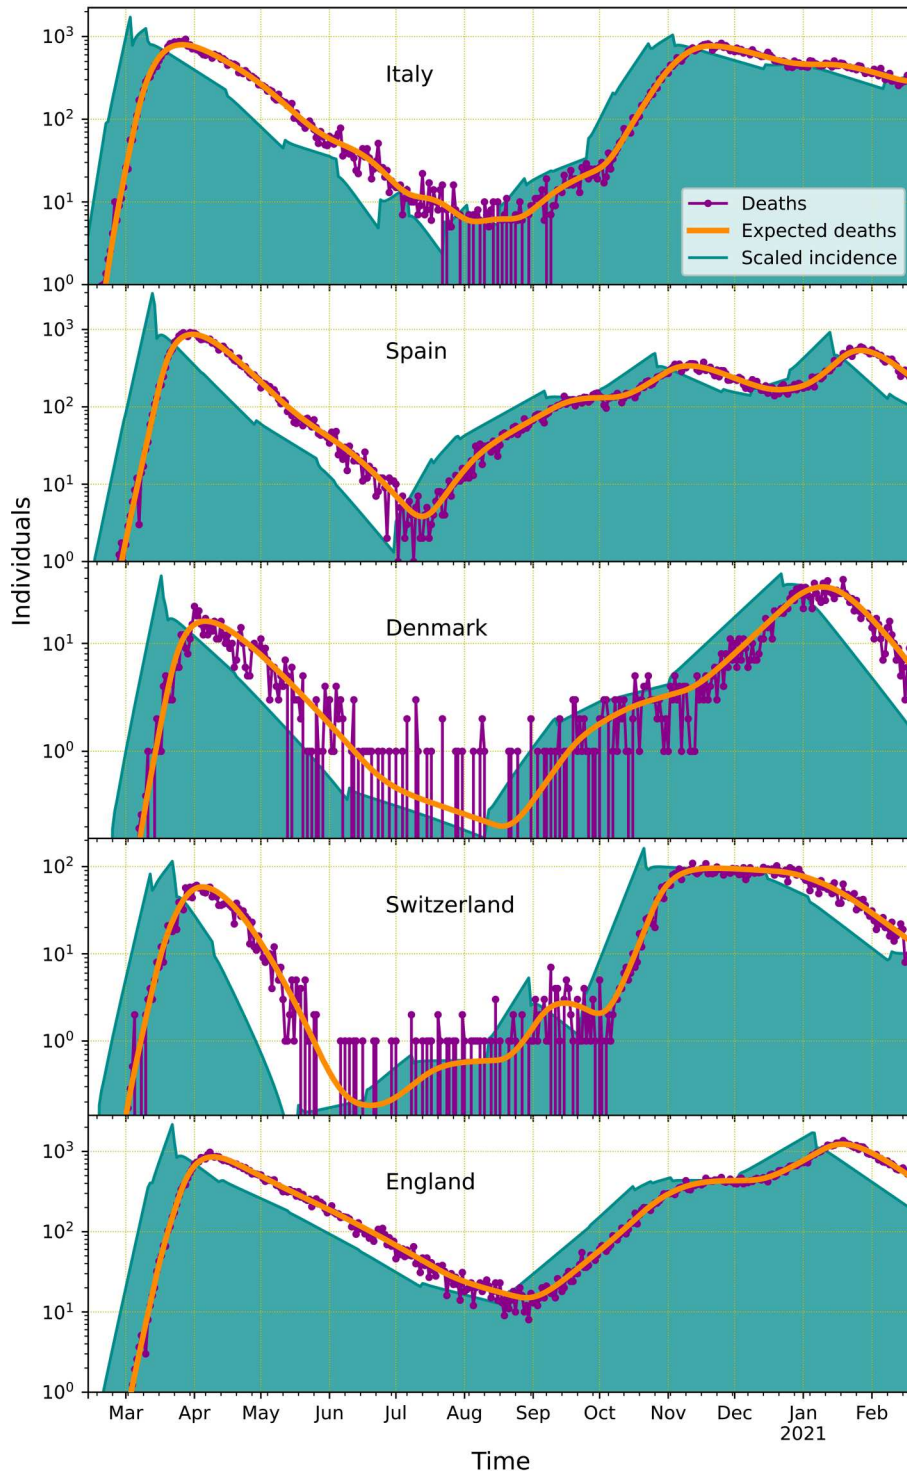


Fig. 2. The DIDNN infers an irregular, discontinuous incidence through the smoothed dynamics of the fluctuating daily deaths. The scaled incidence (filled dark cyan curve) for the different locations follows from the corresponding number of daily deaths (magenta dotted line) through the smooth time courses of their expected value (orange lines).

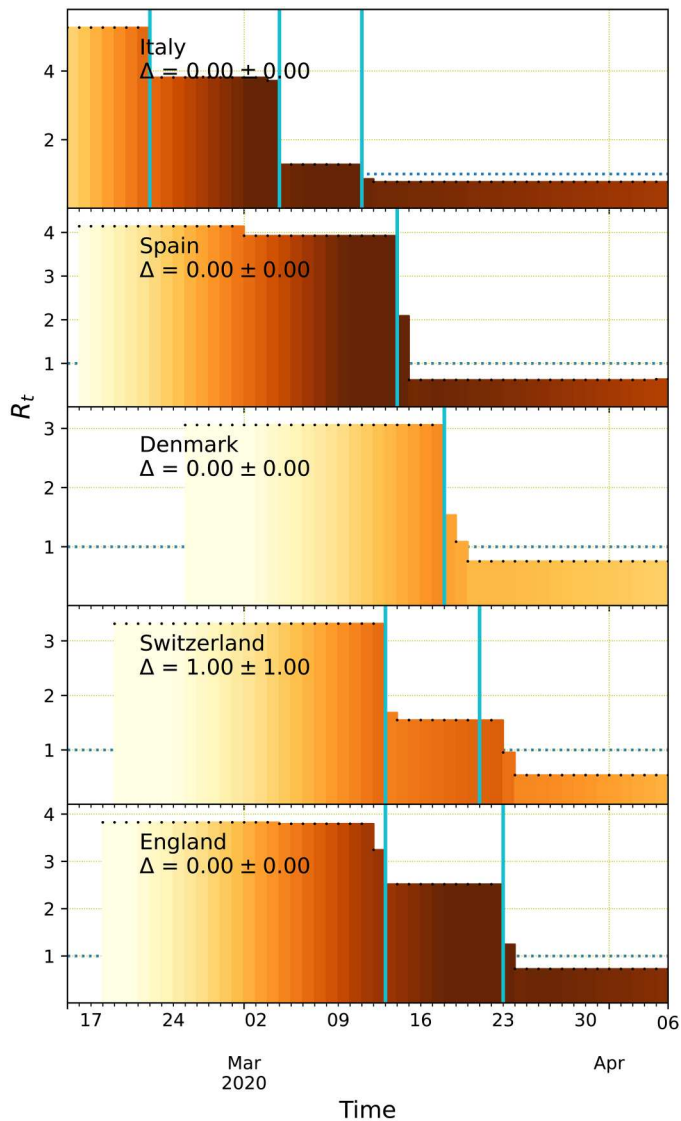


Fig. 3. Discontinuous changes in the instantaneous reproduction number are accurately associated with NPI timings. The detail of the early-stage dynamics for the scaled incidence, j_v , as in Fig. 2 and color-coded in each panel from light (1 individual/day) to dark (10^3 individuals/day) tones, shows an underlying stepwise instantaneous reproduction number, R_t (black dots). The offset, Δ , between NPI timings (cyan vertical lines) and major R_t changes is indicated for each location in days as mean \pm SD, with an overall offset for all locations of 0.22 ± 0.63 days.

measures, as in Switzerland on 23 October, or as a stand-alone measure, as in Italy on 8 October. The effects of lifting NPIs were not as marked as their implementation, partly because they were generally more progressive than their implementation. The major ones identified are the restoration of international travel in Spain on 1 July, with the concurrent restart of its large tourism industry, and the relaxation of measures allowing up to 300 people to stay in bars, restaurants, and clubs in Switzerland on 24 September.

The case of England, with an overall offset of 0.12 ± 0.93 days, is particularly relevant to illustrate the temporal invariance of the clinical and epidemiological dynamics, as all the major discontinuous changes correspond to major NPIs (Fig. 4). These NPIs included,

among others, the implementation of the first, second, and third nationwide lockdowns on 23 March, 5 November, and 6 January, respectively, and lifting the second one on 2 December. The corresponding offsets were, chronologically, 0, -2 , 1, and 1 days (table S1). Besides the clinical parameters of the disease remaining statistically invariant over time, a major reason for this excellent agreement is the number of daily deaths remaining always at high values, with the lowest value of λ_t never crossing 10 deaths/day (Fig. 2). In contrast, the approach missed the progressive lifting of NPIs from May to September in countries with extremely low deaths (Switzerland and Denmark), which reached values of λ_t below 0.2 deaths/day (Fig. 2). The quality of the data also plays a substantial role. In the case of Italy, there are major changes in June–August for R_t inferred from deaths in July–September with similar overall values to those of England but with an unusual distribution, dichotomously alternating randomly between zero and values around 10 deaths/day. Concomitantly, there was no support to potentially associate the resulting changes in R_t for Italy with obvious NPIs during this period.

Super-resolution identification of implementing and lifting NPIs

Our results also show that there are no major mechanistic differences between implementing and lifting NPIs. As a clear example, lifting of the second England nationwide lockdown in December 2020 and implementation of the third one in January 2021 were both inferred to happen within 1 day of their actual date (Fig. 4 and table S1). In this case, lifting and implementing an NPI happen at a similarly high incidence, which is not generally the case. Systematic statistical differences may arise, in general, because low incidence parallels high fluctuations and lifting NPIs sides with low incidence, which makes it more challenging to infer properties of the dynamics (30). This difference is also important when concurrent measures are applied because lockdowns affect mostly regions with high incidence and would be more noticeable than releasing restrictions in low-incidence regions, which might be hidden by the high-incidence regions.

Together, the distribution of the offset of NPI inferred times is much narrower, about an order of magnitude narrower, than the distribution of infection-to-death delay used in the inference process (Fig. 5). Such a precision constates that the approach presented here, by incorporating the epidemiological dynamics, provides a super-resolution identification of NPI timings comparable to the general state of the art in traditional nondynamical scenarios, such as super-resolution localization microscopy (4), a technique used to enhance spatial resolution in imaging, and biomolecular imaging in cells (31), a method used to visualize cellular processes at a molecular level. Here, the date of the NPI is identified as the actual date, the day before, or the day after in 86% of the cases (Fig. 5). This result is remarkable also because the fundamental problem of inferring the timing of an NPI requires locating relative changes in the number of events (infections) rather than the events themselves (emission of photons).

The epidemiological dynamics plays a fundamental role through the strength of the regularization term. The value of $\gamma = 2.51$ corresponds to the retrospective model selection approach and provides the optimal constraint. This value can be estimated prospectively through the AIC. In this case, it is essential to have the proper probabilistic model for the data. Excluding Italy, which strongly deviates

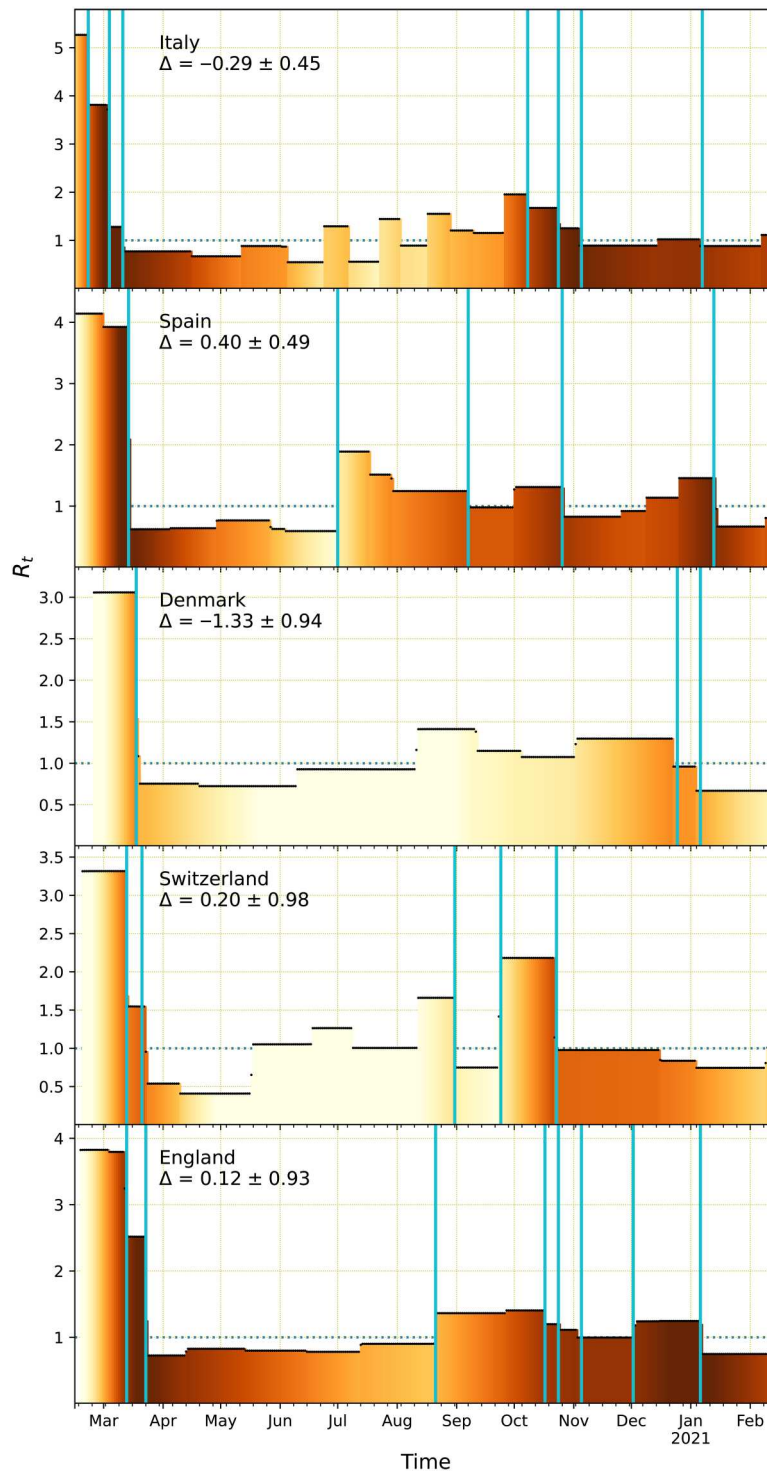


Fig. 4. The infection-to-death dynamics remains homogenous across locations and over time. The time courses of R_t and j_t accurately track the major NPIs over a year with the same distribution of infection-to-death delay (average of 19.3 days and SD of 9.1 days). R_t , j_t , and NPI timings are represented as in Fig. 3. The offset, Δ , between NPI timings and major R_t changes are indicated for each location in days as mean \pm SD, with an overall offset for all locations of -0.07 ± 0.92 days.

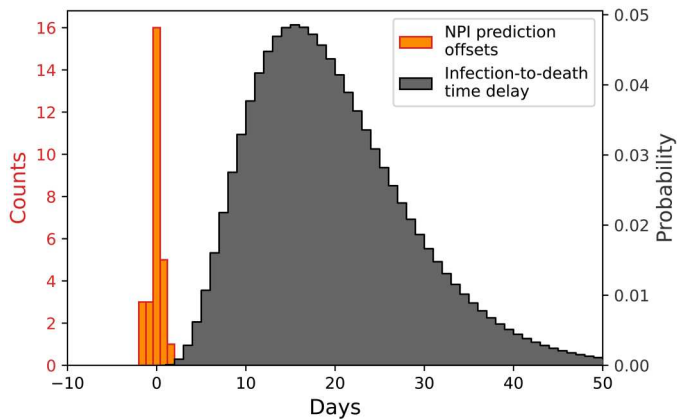


Fig. 5. Super-resolution identification of NPI timings. The distribution of the offset of NPI inferred times (orange bars, left axis), with 0.92 SD, is much narrower than the distribution of infection-to-death delay used in the inference process (gray curve, right axis), with 9.1 SD.

from the Poisson death statistics in July–September, model selection led to a similar value $\gamma = 2.00$, which resulted in an overall offset of 0.00 ± 1.10 days. This concordance indicates that the unsupervised approach can generally identify regime changes at the single-day resolution.

DISCUSSION

We have derived a DIDNN to infer the underlying epidemiological dynamics at single-day resolution from death records. The main computational challenge has been to transfer the information in highly fluctuating daily deaths backward in time to new infections and forward again to infection transmissibility changes through two kernels up to the identification of regime changes at the single-day level. The kernels have a spread as large as a 9.1-day SD with a 19.3-day average delay. In this regard, the increase in resolution of the prediction is a factor 10 \times over the information available at the daily death level. A key aspect of the success of the approach has been the incorporation of the epidemiological dynamics with proper regularization into the network through a PINN-based methodology (32–34). This development allowed the approach to infer discontinuities in the instantaneous reproduction number to account for NPIs. At a fundamental level, our results show that an apparently smooth death dynamics does not result from a similarly continuous evolution of new infections but from a discontinuous, non-differentiable incidence dynamics. Additional aspects were the consideration of the proper fluctuation model (Poisson process) and scale invariance (dynamics governed by relative changes, not absolute changes) in the objective function.

Previous approaches to identifying the effects of NPIs relied primarily on tracking the case-based reproduction number over time. In these instances, it has been argued that the effects of NPIs are delayed from 7 days (35, 36) up to 3 weeks (17, 37) for the implementation and an extra week (17) up to a total of 4 weeks (37) for the lifting. In part, these delays arise from considering incidence as the number of newly reported cases rather than new infections, which introduces at least a delay of a week from the average time from infection to symptom onset. This delay is compounded with the time lag from symptom onset to reporting. On the methodological side,

there are also delays introduced by averaging in real time, both by noncentered moving averages over the data and by the inherent averaging procedure used to compute the instantaneous reproduction number (38). Overall, all these processes lead to delayed blurring of the actual incidence. The advantages of using case counts are that their numbers are substantially larger than deaths, therefore reducing the noise, and that the average delay from infection to illness is shorter than to death, therefore allowing for an easier mathematical estimation without performing the needed deconvolutions. The disadvantages are the potential lack of testing capabilities, which were extreme in the early stages of the COVID-19 pandemic, and the still undetermined nature of the delays. Uncertainties in the case-based approach become even greater in a contact-tracing context, as asymptomatic cases can typically be detected up to 4 weeks after infection (39). Therefore, case-based approaches are generally not suitable for the time period we considered nor to target resolutions at the single-day level.

Our results, in contrast, show that the effects of major NPIs generally happen instantaneously in the disease transmission, as generally expected (14) and as assumed in previous approaches that relied on death counts to infer the effects of NPIs with known implementation dates (19). The instantaneous effect of NPIs is a critical component of the response to infectious diseases and has far-reaching implications beyond reducing transmission. Therefore, our approach provides an unsupervised avenue for quantifying this effect in disease control and for better understanding its correlation with the resulting downstream effects, such as the impact on health care systems. Having precise information is crucial for weighing the potential benefits of NPIs against their broad impact on other aspects of society, including economics, psychology, and public opinion. It enables informed decision-making by policymakers and economists, facilitating the assessment of NPIs and their impact on economic stability. In addition, the knowledge of how NPIs can lead to regime changes in the epidemiological dynamics can influence individual and collective behavior, fostering public trust, compliance, and psychological well-being. Furthermore, this information informs public discourse, potentially shaping opinions and fostering a deeper understanding of the complex interplay between public health measures and societal impacts.

From the epidemiological side, our results show that the same kernel for the infection-to-death time distribution can accurately account for the timing of major NPIs over multiple countries over the time span of a year. This consistency implies the same properties for the underlying clinical features of the disease over time and over similar populations (European countries that recorded death dates). In this regard, the kernel we used is the same as the one used in previous studies (19) with the consideration that deaths occur, on average, 3.6 days earlier than their reporting.

It is also possible to directly apply our approach to populations under conditions that strongly differ from those of the European countries considered here explicitly, as in countries in which limited access to health care led to faster and more deaths, compounded with substantial underreporting. The limiting step is the availability of the kernel that connects incidence with death records. If the probabilistic descriptions for the delay in reporting deaths, the underreported deaths, and the infection to death were available, then it would be possible to obtain a composite kernel, so that our approach could be applied straightforwardly. With this information, it would also be possible to estimate, as reflected

in departures from the expected regime changes in the inferred epidemiological dynamics, to what extent NPI implementation was erratic and partial in effectiveness in these countries.

In the absence of essential information to estimate the kernel that connects incidence with death records, it would still be possible to use the approach, swapping the inference goal, to find the kernel that best captures the timing of well-characterized NPIs rather than to infer the NPIs themselves. This strategy would involve, for instance, a grid-search of a parametrized kernel that minimizes the mean squared distance of the timings between reference NPIs and the corresponding regime changes. The analysis of the inferred kernel would be potentially useful to get insights into the disease progression and its reporting that were not readily available.

From the decision-making point of view, the most direct implication of our approach is to provide precise knowledge about the NPIs that had an extensive impact on the epidemiological dynamics, such as the earlier measures identified that changed the course of the pandemics from growing to decaying in the five countries analyzed explicitly. A most interesting result is the lack of regime changes after the early lockdowns up to the summer in Spain and England, which retained sustained instantaneous reproduction numbers below one. During this period, protective masks became mandatory in public spaces, and both countries progressively lifted social interaction restrictions. Therefore, the use of protective masking had minor effects under lockdown conditions, and the subsequent lifting of many restrictions under masking conditions had limited effects. These results highlight that NPI effects in the epidemiological dynamics are context dependent and can effectively saturate. We identified mask requirements as having a strong impact in an expanding phase, as observed with the mandatory use of protective masks in Italy on 8 October 2020.

We have focused on the early stage of the COVID-19 outbreak because of its inherent interest and because it provides stringent validation conditions for the approach. These conditions allowed us to favorably contrast our results against well-defined NPIs and well-characterized infection-to-death delays and generation times. As the pandemic evolved into the year 2023, vaccination, acquired immunity, and new viral variants changed the clinical progression and epidemiological dynamics. Therefore, the applicability of the approach to this new scenario would need the infection-to-death and generation-time kernels to be reparametrized. Vaccines, NPIs, and other mitigation strategies at play have also substantially reduced the number of deaths, which in the case of England, for instance, has gone from more than 1000 daily deaths in April 2020 and January 2021 to a sustained ~50 to 100 daily deaths in the last 6 months of the year 2022. Therefore, Poisson noise has become higher than earlier in the pandemic, and consequently, the difficulty of extracting signals from death records has increased. However, the reported number of cases has become more reliable and better characterized than in the early stages, which would make it possible to mitigate potential Poisson noise effects by adapting the approach to deconvolve both death and, with an appropriate infection-to-case reporting kernel, case records simultaneously.

From the computational point of view, our results have paved the way for the inclusion of fully-integral dynamic equations into PINNs, which have traditionally been restricted to systems with differential time evolution, such as PDEs, ODEs, and integrodifferential equations (40). More broadly, our approach has provided an avenue to generally implement complex model-aware

regularization protocols in deconvolution methodologies (41) and to explicitly use the dynamics of the system as a part of the regularization process.

MATERIALS AND METHODS

Field data

The data were downloaded from the “The Demography of COVID-19 Deaths” curated database (42) except for England. The United Kingdom reports by default deaths within 28 days of a positive test, which misses a substantial part of deaths. For England, we considered daily “deaths within 60 days of a positive test by death date” (43), which provides a more stringent compromise between deaths by and deaths with COVID-19.

Data for NPIs were obtained from online resource (44) and NPI dataset (20). Multiple consecutive NPIs were assigned the middle-point date rounded to days. The 9-day width of the window to infer the NPI date was chosen as the middle point of the range that does not affect the results, which extends from 5 to 13 days. Five days is the minimum value that covers the maximum offset of ± 2 days, and 13 days is the maximum value that does not include the two NPIs closest to each other, which are 7 days apart.

Kernel parametrization

The kernels are parametrized according to the available clinical and field data and are based on previous parametrizations (19). The generation time kernel, w_τ , follows a gamma distribution with a mean of 6.3 days and an SD of 4.2 days (19). This parametrization is the same as for the serial interval (45) and is consistent also with the characterization using U.K. household data, which led to similar average (5.9 days, 95% credible interval 5.2–7.0 days) and SD (4.8 days, 95% credible interval 4.0–6.3 days) (46).

The time from infection-to-death kernel, f_τ , has been estimated to approximately follow a gamma distribution with 22.9-day mean and 9.1-day SD for reported deaths (19). Because we use actual death dates, we shortened the mean to 19.3 days to account for the reporting delay (28, 29). The values selected are consistent with U.K. field data estimates for the time from infection-to-death distributions, with average values in the range from 17.4 to 24.7 days (47).

Computational implementation

The overall approach was implemented in Keras (48) with custom losses coded in TensorFlow (49). The network was trained with the AdaMax optimization algorithm (50) and variable learning rate, decreasing from 10^4 to 10^{-4} with an exponential schedule. Logarithms of quantities susceptible of being zero were calculated with 10^{-1} added to them to avoid divergences during the training process. For each country, we considered 41 independent values of the regularization term, equally spaced on a logarithmic scale, from $\gamma = 10^{-1}$ to $\gamma = 10$. In model selection, the quantity $|\ln R_{t+1} - \ln R_t|$ was considered to be zero for values below 10^{-3} .

Because sparse regularization is nonlocal, the initial growing trend was propagated backward in time to avoid the effects of the widespread loss of the death counts (lack of diagnosis) during the very early stages of the outbreak. The cutoff time, T_c , to initiate backward propagation was set for each country as the last day the cumulative number of deaths was smaller than 1% of the

maximum daily deaths. The growing trend, $\kappa = \sum_{T_c+8}^{T_c+14} n_t / \sum_{T_c+1}^{T_c+7} n_t$, was

quantified as the total number of deaths in the second week over those in the first week after the cutoff time. The propagation was performed backward in time for the m week before the cutoff time by taking the scaled 3-day centered moving average of the first week after the cutoff as $\tilde{n}_{t-7m} = \kappa^{-m} \frac{1}{3} (n_{t-1} + n_t + n_{t+1})$ for $t \in [T_c + 1, T_c + 7]$.

Supplementary Materials

This PDF file includes:

Legend for table S1

Other Supplementary Material for this manuscript includes the following:

Table S1

REFERENCES AND NOTES

- R. Szeliski, *Computer Vision: Algorithms and Applications*. (Springer-Verlag, ed. 2, 2022).
- C. Payer, D. Stern, H. Bischof, M. Urschler, Regressing heatmaps for multiple landmark localization using CNNs, in *Medical Image Computing and Computer-Assisted Intervention* (2016), vol. 9901, pp. 230–238.
- B. Li, W. Ren, D. Fu, D. Tao, D. Feng, W. Zeng, Z. Wang, Benchmarking single-image dehazing and beyond. *IEEE Trans. Image Process.* **28**, 492–505 (2019).
- L. von Diezmann, Y. Shechtman, W. E. Moerner, Three-dimensional localization of single molecules for super-resolution imaging and single-particle tracking. *Chem. Rev.* **117**, 7244–7275 (2017).
- J. Biemond, R. L. Lagendijk, R. M. Mersereau, Iterative methods for image deblurring. *Proc. IEEE* **78**, 856–883 (1990).
- M. R. Banham, A. K. Katsaggelos, Digital image restoration. *IEEE Signal Process. Mag.* **14**, 24–41 (1997).
- R. Verity, L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G. Cuomo-Dannenburg, H. Thompson, P. G. T. Walker, H. Fu, A. Dighe, J. T. Griffin, M. Baguelin, S. Bhatia, A. Boonyasiri, A. Cori, Z. Cucunubá, R. F. John, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, D. Laydon, G. Nedjati-Gilani, S. Riley, S. van Elsland, E. Volz, H. Wang, Y. Wang, X. Xi, C. A. Donnelly, A. C. Ghani, N. M. Ferguson, Estimates of the severity of coronavirus disease 2019: A model-based analysis. *Lancet Infect. Dis.* **20**, 669–677 (2020).
- E. Goldstein, J. Dushoff, J. Ma, J. B. Plotkin, D. J. Earn, M. Lipsitch, Reconstructing influenza incidence by deconvolution of daily mortality time series. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 21825–21829 (2009).
- P. S. F. Yip, K. F. Lam, Y. Xu, P. H. Chau, J. Xu, W. Chang, Y. Peng, Z. Liu, X. Xie, H. Y. Lau, Reconstruction of the infection curve for SARS epidemic in Beijing, China using a back-projection method. *Commun. Stat. B: Simul. Comput.* **37**, 425–433 (2008).
- N. G. Becker, L. F. Watson, J. B. Carlin, A method of non-parametric back-projection and its application to aids data. *Stat. Med.* **10**, 1527–1542 (1991).
- A. C. Miller, L. A. Hannah, J. Futoma, N. J. Foti, E. B. Fox, A. D'Amour, M. Sandler, R. A. Saurous, J. A. Lewnard, Statistical deconvolution for inference of infection time series. *Epidemiology* **33**, 470–479 (2022).
- M. Jahja, A. Chin, R. J. Tibshirani, Real-time estimation of COVID-19 infections: Deconvolution and sensor fusion. *Stat. Sci.* **37**, 207–228 (2022).
- K. M. Gostic, L. M. Gough, E. B. Baskerville, S. Abbott, K. Joshi, C. Tedijanto, R. Kahn, R. Niehus, J. A. Hay, P. M. De Salazar, J. Hellewell, S. Meakin, J. D. Munday, N. I. Bosse, K. Sherratt, R. N. Thompson, L. F. White, J. S. Huisman, J. Scire, S. Bonhoeffer, T. Stadler, J. Wallinga, S. Funk, M. Lipsitch, S. Cobey, Practical considerations for measuring the effective reproductive number, *Rt*. *PLOS Comput. Biol.* **16**, e1008409 (2020).
- C. Fraser, Estimating individual and household reproduction numbers in an emerging epidemic. *PLOS ONE* **2**, e758 (2007).
- N. Perra, Non-pharmaceutical interventions during the COVID-19 pandemic: A review. *Phys. Rep.* **913**, 1–52 (2021).
- COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University. 2020 (<https://github.com/CSSEGISandData/COVID-19>).
- Y. Li, H. Campbell, D. Kulkarni, A. Harpur, M. Nundy, X. Wang, H. Nair; Centre for Global Health Research; Usher Institute; Deane of Molecular, Genetic and Population Health Sciences; Global Health Academy, The temporal association of introducing and lifting non-pharmaceutical interventions with the time-varying reproduction number (R) of SARS-CoV-2: A modelling study across 131 countries. *Lancet Infect. Dis.* **21**, 193–202 (2021).
- R. Li, S. Pei, B. Chen, Y. Song, T. Zhang, W. Yang, J. Shaman, Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science* **368**, 489–493 (2020).
- S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, M. Monod; Imperial College COVID-19 Response Team, A. C. Ghani, C. A. Donnelly, S. Riley, M. A. C. Vollmer, N. M. Ferguson, L. C. Okell, S. Bhatt, Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* **584**, 257–261 (2020).
- J. M. Brauner, S. Mindermann, M. Sharma, D. Johnston, J. Salvatier, T. Gaveniak, A. B. Stephenson, G. Leech, G. Altman, V. Mikulik, A. J. Norman, J. T. Monrad, T. Besiroglu, H. Ge, M. A. Hartwick, Y. W. Teh, L. Chindelevitch, Y. Gal, J. Kulveit, Inferring the effectiveness of government interventions against COVID-19. *Science* **371**, eabd9338 (2021).
- J. T. Wu, K. Leung, M. Bushman, N. Kishore, R. Niehus, P. M. de Salazar, B. J. Cowling, M. Lipsitch, G. M. Leung, Estimating clinical severity of COVID-19 from the transmission dynamics in Wuhan, China. *Nat. Med.* **26**, 506–510 (2020).
- C. Rothe, M. Schunk, P. Sothmann, G. Bretzel, G. Froeschl, C. Wallrauch, T. Zimmer, V. Thiel, C. Janke, W. Guggemos, M. Seilmaier, C. Drosten, P. Vollmar, K. Zwirgmaier, S. Zange, R. Wölfel, M. Hoelscher, Transmission of 2019-nCoV infection from an asymptomatic contact in Germany. *N. Engl. J. Med.* **382**, 970–971 (2020).
- M. M. Böhmer, U. Buchholz, V. M. Corman, M. Hoch, K. Katz, D. V. Marosevic, S. Böhm, T. Woudenberg, N. Ackermann, R. Konrad, U. Eberle, B. Treis, A. Dangel, K. Bengs, V. Fingerle, A. Berger, S. Hörmansdorfer, S. Ippisch, B. Wicklein, A. Grahl, K. Pörtner, N. Müller, N. Zeitlmann, T. S. Boender, W. Cai, A. Reich, M. A. der Heiden, U. Rexroth, O. Hamouda, J. Schneider, T. Veith, B. Mühlemann, R. Wölfel, M. Antwerpen, M. Walter, U. Protzer, B. Liebl, W. Haas, A. Sing, C. Drosten, A. Zapf, Investigation of a COVID-19 outbreak in Germany resulting from a single travel-associated primary case: A case series. *Lancet Infect. Dis.* **20**, 920–928 (2020).
- N. C. Grassly, C. Fraser, Mathematical models of infectious disease transmission. *Nat. Rev. Microbiol.* **6**, 477–487 (2008).
- H. W. Hethcote, The mathematics of infectious diseases. *SIAM Rev.* **42**, 599–653 (2000).
- H. Akaike, A new look at the statistical model identification. *IEEE Trans. Autom. Control* **19**, 716–723 (1974).
- R. J. Tibshirani, J. Taylor, Degrees of freedom in lasso problems. *Ann. Stat.* **40**, 1198–1232 (2012).
- J. M. G. Vilar, L. Saiz, Reliably quantifying the evolving worldwide dynamic state of the COVID-19 outbreak from death records, clinical parametrization, and demographic data. *Sci. Rep.* **11**, 19952 (2021).
- A. T. Levin, W. P. Hanage, N. Owusu-Boaitey, K. B. Cochran, S. P. Walsh, G. Meyerowitz-Katz, Assessing the age specificity of infection fatality rates for COVID-19: Systematic review, meta-analysis, and public policy implications. *Eur. J. Epidemiol.* **35**, 1123–1138 (2020).
- J. M. G. Vilar, L. Saiz, Ascertaining the initiation of epidemic resurgences: An application to the COVID-19 second surges in Europe and the Northeast United States. *R. Soc. Open Sci.* **8**, 210773 (2021).
- Y. M. Sigal, R. Zhou, X. Zhuang, Visualizing and discovering cellular structures with super-resolution microscopy. *Science* **361**, 880–887 (2018).
- M. Raissi, P. Perdikaris, G. E. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning. *Nat. Rev. Phys. Ther.* **3**, 422–440 (2021).
- E. Kharazmi, M. Cai, X. Zheng, Z. Zhang, G. Lin, G. E. Karniadakis, Identifiability and predictability of integer- and fractional-order epidemiological models using physics-informed neural networks. *Nat. Comput. Sci.* **1**, 744–753 (2021).
- Y. Liu, C. Morgenstern, J. Kelly, R. Lowe; CMMID COVID-19 Working Group, M. Jit, The impact of non-pharmaceutical interventions on SARS-CoV-2 transmission across 130 countries and territories. *BMC Med.* **19**, 40 (2021).
- A. Pan, L. Liu, C. Wang, H. Guo, X. Hao, Q. Wang, J. Huang, N. He, H. Yu, X. Lin, S. Wei, T. Wu, Association of public health interventions with the epidemiology of the COVID-19 outbreak in Wuhan, China. *JAMA* **323**, 1915–1923 (2020).
- S. Singh, M. Shaikh, K. Hauck, M. Miraldo, Impacts of introducing and lifting nonpharmaceutical interventions on COVID-19 daily growth rate and compliance in the United States. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2021359118 (2021).
- A. Cori, N. M. Ferguson, C. Fraser, S. Cauchemez, A new framework and software to estimate time-varying reproduction numbers during epidemics. *Am. J. Epidemiol.* **178**, 1505–1512 (2013).

39. R. Wölfel, V. M. Corman, W. Guggemos, M. Seilmaier, S. Zange, M. A. Müller, D. Niemeyer, T. C. Jones, P. Vollmar, C. Rothe, M. Hoelscher, T. Bleicker, S. Brünink, J. Schneider, R. Ehmann, K. Zwirgmaier, C. Drosten, C. Wendtner, Virological assessment of hospitalized patients with COVID-2019. *Nature* **581**, 465–469 (2020).
40. L. Lu, X. Meng, Z. Mao, G. E. Karniadakis, DeepXDE: A deep learning library for solving differential equations. *SIAM Rev.* **63**, 208–228 (2021).
41. M. D. Zeiler, D. Krishnan, G. W. Taylor, R. Fergus, in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2010), pp. 2528–2535.
42. A. Caporali, J. Garcia, É. Couppié, S. Poniakina, M. Barbieri, F. Bonnet, C. G. Camarda, E. Cambois, I. Hourani, D. Korotkova, F. Meslé, O. Penina, J.-M. Robine, M. Sauerberg, C. Torres, The demography of COVID-19 deaths database, a gateway to well-documented international data. *Sci. Data* **9**, 93 (2022).
43. The official UK government website for data and insights on coronavirus (COVID-19). 2022 (<https://coronavirus.data.gov.uk/>).
44. Source collection and A3M events covering the COVID-19 pandemic. 2022 (<https://www.global-monitoring.com/en/collection-covid-19/>).
45. Q.-X. Long, X.-J. Tang, Q.-L. Shi, Q. Li, H.-J. Deng, J. Yuan, J.-L. Hu, W. Xu, Y. Zhang, F.-J. Lv, K. Su, F. Zhang, J. Gong, B. Wu, X.-M. Liu, J.-J. Li, J.-F. Qiu, J. Chen, A.-L. Huang, Clinical and immunological assessment of asymptomatic SARS-CoV-2 infections. *Nat. Med.* **26**, 1200–1204 (2020).
46. W. S. Hart, S. Abbott, A. Endo, J. Hellewell, E. Miller, N. Andrews, P. K. Maini, S. Funk, R. N. Thompson, Inference of the SARS-CoV-2 generation time using UK household data. *eLife* **11**, e70767 (2022).
47. T. Ward, A. Johnsen, Understanding an evolving pandemic: An analysis of the clinical time delay distributions of COVID-19 in the United Kingdom. *PLOS ONE* **16**, e0257978 (2021).
48. F. Chollet, *Deep learning with Python*. (Manning Publications Co., 2018).
49. M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, TensorFlow: A system for Large-Scale machine learning, in *12th USENIX Symposium on Operating Systems Design and Implementation* (OSDI, 2016), pp. 265–283.
50. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization. arXiv:1412.6980 [cs.LG] (22 December 2014).

Acknowledgments

Funding: J.M.G.V. acknowledges support from Ministerio de Ciencia e Innovación under grants PGC2018-101282-B-I00 and PID2021-128850NB-I00 (MCI/AEI/FEDER, UE). **Author contributions:** J.M.G.V. and L.S. designed research, performed research, analyzed data, and wrote the paper. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data and code used in the analysis and needed to evaluate the conclusions in the paper are publicly available (GitHub repository: <https://github.com/DIDNN/NPIs>; Zenodo: <https://doi.org/10.5281/zenodo.7737141>). All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Submitted 27 September 2022

Accepted 14 June 2023

Published 14 July 2023

10.1126/sciadv.adf0673