Optimizing the spatio-temporal allocation of COVID-19 vaccines: Italy as a case study

Joseph C. Lemaitre[1*], Damiano Pasetto[2], Mario Zanon[3], Enrico Bertuzzo[2], Lorenzo Mari[4], Stefano Miccoli[5], Renato Casagrandi[4], Marino Gatto[4], Andrea Rinaldo[1,6]

1 Laboratory of Ecohydrology, École Polytechnique Fédérale de Lausanne, 1015 Lausanne (CH)
2 Dipartimento di Scienze Ambientali, Informatica e Statistica, Università Ca' Foscari Venezia, 30172 Venezia-Mestre (IT)
3 IMT Scuola Alti Studi, 55100 Lucca (IT)
4 Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133 Milan (IT)
5 Dipartimento di Meccanica, Politecnico di Milano, 20156 Milan (IT)
6 Dipartimento ICEA, Università di Padova, 35131 Padova (IT)

* corresponding author: joseph.lemaitre@epfl.ch

Abstract

While SARS-CoV-2 vaccine distribution campaigns are underway across the world, communities face the challenge of a fair and effective distribution of limited supplies. We wonder whether suitable spatial allocation strategies might significantly improve a campaign's efficacy in averting damaging outcomes. To that end, we address the problem of optimal control of COVID-19 vaccinations in a country-wide geographic and epidemiological context characterized by strong spatial heterogeneities in transmission rate and disease history. We seek the vaccine allocation strategies in space and time that minimize the number of infections in a prescribed time horizon. We examine scenarios of unfolding disease transmission across the 107 provinces of Italy, from January to April 2021, generated by a spatially explicit compartmental COVID-19 model tailored to the Italian geographic and epidemiological context. We develop a novel optimal control framework to derive optimal vaccination strategies given the epidemiological projections and constraints on vaccine supply and distribution logistic. Optimal schemes significantly outperform simple alternative allocation strategies based on incidence, population distribution, or prevalence of susceptibles in each province. Our results suggest that the complex interplay between the mobility network and the spatial heterogeneities imply highly non-trivial prioritization of local vaccination campaigns. The extent of the overall improvements in the objectives grants further inquiry aimed at refining other possibly relevant factors so far neglected. Our work thus provides a proof-of-concept of the potential of optimal control for complex and heterogeneous epidemiological contexts at country, and possibly global, scales.

Author summary

The development of vaccines has sparked high hopes towards the control of SARS-CoV-2 transmission without resorting to extensive community-wide restrictions. A fundamental unanswered question concerns the best possible allocation of a limited vaccine stock in

space and time given a specific goal. We address this through an optimal control framework based on a reliable spatially explicit COVID-19 epidemiological model, where vaccine distribution is optimized under supply and deployment capacity constraints. This tool provides strategies for optimal allocations in different scenarios, yielding important improvements over considered alternatives. By accounting for spatial heterogeneities and human mobility networks, the presented approach complements currently used allocation methods based on criteria such as age or risk.

Introduction

Supply- or deployment-limited SARS-CoV-2 vaccines [1] pose the urgent question of a fair distribution of the available doses [2]. Current prioritization approaches typically target groups at higher risk of severe outcomes [3, 4], or their indirect protection by vaccinating those with higher disease transmission [3, 5, 6]. Our main hypothesis is that taking into account spatial heterogeneities in disease transmission when designing prioritization strategies will significantly improve the effectiveness of vaccination campaigns. The distribution of doses inside each country is limited by the logistic capabilities of the healthcare network and the rate at which the vaccine stock is replenished. Decisions concerning the best allocation strategies are to be taken under these constraints. Moreover, both the complex coupling between regions due to human mobility and the spatial heterogeneities in disease history and control interventions make the discovery of such optimal allocation strategies an arduous task.

We propose an optimal control framework to explore COVID-19 vaccine distribution in space and time. We study the SARS-CoV-2 epidemic in Italy, where strong spatial effects arise from the geography of the disease, heterogeneous lockdown exit strategies, and post-lockdown control measures [7]. The optimal control framework is applied to a spatial model that has proved its reliability for Italy [8, 9], updated through data assimilation of a year-long epidemiological record. This allows us to unravel the best possible vaccination strategy and probe the impact of vaccine allocations over the 107 Italian provinces.

The problem of vaccine allocation is of primary importance for public-health officials, epidemiologists, and economists [10, 11]. Roll-out strategies are conventionally based on the prioritization of individuals at risk, such as health workers and elderly people [12–15]. However, the heterogeneous ways in which different regions may be affected by each successive wave raise questions about spatial prioritization strategies. What is the best feasible spatial allocation, given supply and logistic constraints? Would that differ significantly from current non geographically-optimized plans? Should vaccines be distributed on the basis of demography or would it be better to prioritize areas currently subject to an outbreak? How relevant are the susceptibility profile and modelled future transmission in each region?

Epidemiological modeling has long been used to answer questions about the impact of vaccination campaigns, often by comparing outcomes under different scenarios [16, 17]. Optimization, i.e, the search for the best possible course of action that maximizes or minimizes an objective metric, has been carried out theoretically since the seventies [18–20]. Recent dramatic improvements of both algorithms [21] and computational power prompted applied studies using different methods to rigorously find optimal mitigation strategies: most of the time trough iterative parameter search [22, 23], but also using genetic algorithms [24] or solving the Hamilton-Jacobi-Bellman equations [25, 26].

Interesting developments have recently arose during the ongoing SARS-CoV-2 pandemic [13, 27, 28]. The urgency of effective vaccination campaigns led to the development of modeling frameworks for the optimization of vaccine allocation, based

on age or risk [3, 4, 12, 13], dose timing [29, 30], and the deployment of testing resources, using optimal control [31] or Bayesian experimental design [32], along with prioritization based on social contact networks [33].

To our knowledge, optimal spatial allocation of COVID-19 vaccines at a country scale has never been performed yet. This question is distinct from, and complementary to, risk-based prioritization. Spatial heterogeneities in disease transmission are complex, as seen during the initial outbreaks [8, 9, 34], supporting the significance of the posed problem towards an effective control of the epidemic. However, the connectivity network underlying spatial epidemiological models may generate complex large-scale control problems whose solution requires tailored formulations and efficient algorithms.

This work aims to find optimal strategies for this problem through modern optimization methods based on distributed direct multiple shooting, automatic-differentiation, and large-scale nonlinear programming [35–38]. This allows us to solve the large-scale optimization problems arising from epidemiological models, even when considering hundreds of spatial nodes.

Materials and methods

The formulation of the optimal control problem has three main components: 1) an objective function to be minimized, here the total incidence in Italy from January 11, 2021 to April 11, 2021; 2) the spatial epidemiological model [8, 9] governing the transmission dynamics with the daily vaccination rates in each province as control variables; and 3) the set of constraints that the control must satisfy, in our case the limitations on vaccine administration rate in each province and the total vaccine stock in Italy.

1) Objective function. Optimizing calls for a metric, whose selection is critical in determining the optimal solution and its outcome. The choice of an objective function relates health, economy, and ethics. Possible candidates are the minimization of e.g. DALYs (the Disability-Adjusted Life Years), the number of deaths, and economic loss [39]. All these objectives are linked and may be combined together. As the model considered for this work does not have risk-classes, without loss of generality we optimize for the minimization of the incidence on Italy from January 11, 2021 to April 11, 2021. Minimization of the deaths would yield the same results under the assumptions the model used.

2) Epidemiological model. Incidence and deaths are projected using the spatially distributed epidemiological model devised by Gatto et al. [8] and further improved by Bertuzzo et al. [9]. The model subdivides the Italian population into its 107 provinces represented as a network of connected nodes. Each province has local dynamics describing the number of individuals present in each of the model compartments: susceptible S, exposed E, pre-symptomatic P (incubating infectious), symptomatic infectious I, asymptomatic infectious A, hospitalized H, quarantined Q, recovered R, and dead D. A tenth compartment, vaccinated individuals V, is added to the original nine, as shown in Figure 1.A.

Apart from hospitalized H, quarantined Q, dead D, and symptomatic individuals I, a fraction of the other individuals commutes between provinces along the mobility network, thus we introduce node-to-node disease transmission along the network shown in Figure 1.B.

Compartments P, A, and I have different degrees of infectiousness and contribute to the force of infection (Equation SI (S3) and SI (S4)), which represents the rate at which susceptibles S become infected and, thus, enter the exposed compartment E. The force

of infection in each province is constituted of a local and a mobility component. The local component describes transmission among the local individuals. The mobility component considers that local susceptibles may enter in contact with infected individuals that are traveling, and oppositely, susceptible commuters may become infected through contact with local infected. Connected provinces contribute to this process depending on the strength of the mobility fluxes from and to the node of interest.

The estimated vaccine efficacy and immunity duration depends on the vaccine type. Because we focus on spatial patterns and differences among vaccination strategies, we model the vaccination process by assuming a one-dose vaccine with instantaneous 100% efficacy. Moreover, we assume that vaccines protection persists during the three months of projection considered.

The epidemiological model, previously calibrated during the first wave of COVID-19 in Italy [8, 9], is updated up to January 11, 2021 using an iterative particle filtering. This data assimilation scheme allows us to capture the second wave of infections that hit Italy in the Fall of 2020, a necessary requirement to generate model projections that take into account the whole epidemic history, as shown in Figure 2. In our approach, model projections are described by an ensemble of a thousand trajectories associated with different parameters, whose distributions quantify the model uncertainty.

We consider two projection scenarios characterized by two possible rates of epidemic transmission, see Figure 2. The "Optimistic" scenario assumes a constant lowering of transmission from January 11, 2021 to April 11, 2021; the "Pessimistic" scenario considers a gradual increase in transmission until mid February 2021, which results in a third wave.

For each scenario, the optimal control problem is solved for one reference model trajectory, whose parameters and state on January 11, 2021 are obtained as the median values of the 1'000 model realizations. In this way, the reference trajectory approximately represents the ensemble median in each province. Then, we assess the effectiveness of the optimal allocation on the full ensemble of trajectories.

3) Constraints   We define two types of constraints: supply constraints, which determine the weekly delivery to the national stockpile; and logistic constraints, which limits the maximum rate of local vaccine distribution in each province.

The supply constraints ensure that the model does not distribute more vaccine than what is actually available in stock. We assume that the national supply of vaccine doses is empty on January 11, 2021 and is replenished every Monday. We consider four scenarios with weekly deliveries of 479'700 (realistic, baseline value), 1M, 1.5M and 2M vaccine doses.

From the national stockpile, doses may be allocated to any of the 107 Italian provinces, but the logistic constraints limit the rate at which it is possible to distribute the vaccines in each province. We assume that the maximum number of individuals who can be vaccinated in a province per day is proportional to the province's population, such that the national maximum distribution capacity equals 500,000 doses per day, i.e., 3.5M per week if every province vaccinates at its maximum rate (which in retrospect is close to Italy's vaccination rate as of May 1st).

The objective, the model, and the constraints may be tailored to specific applications within the proposed framework.

Using state-of-the-art linear algebra solvers and automatic differentiation, we solve each scenario (optimistic and pessimistic, with different weekly stockpile deliveries) for the optimal vaccination allocation.

Fig 1. Model setup. A. Diagram representing the compartments of the epidemiological model and the possible transitions in a single province. We control for the vaccination rate (teal arrows), aiming at minimizing incident infections (pink arrow). Individuals in compartments outside of the yellow block are able to move along the mobility network shown in **B.**, hence the force of infection in a province is coupled to the dynamics of other connected provinces. To reduce the problem to a tractable size, we only consider the most important connections (red edges) when optimizing, but we use the full network (red and grey edges) to assess our strategies. A discussion on the effect of this simplification is provided in SI. Nodes size and color display each province's population, and edges width shows the straight of the coupling between each pair of province.

Optimal control problem formulation

We provide a brief methodological description of the optimal control framework. The full equations are derived in the SI, along with implementation details and the source code.

We denote n the number of spatial nodes ($n = 107$ provinces in Italy) and m the number of epidemic states in our model ($m = 9$ states). We denote as $x(t) \in \mathbb{R}_+^{n \times m}$ the state of the system, i.e., $x(t)$ is a vector containing the epidemic variables $S_i(t)$, $E_i(t)$, $P_i(t)$, $I_i(t)$, $A_i(t)$, $Q_i(t)$, $H_i(t)$, $R_i(t)$, $V_i(t)$ for every province $i = 1, ..., n$. We define $v(t) = (v_1(t), ..., v_n(t)) \in \mathbb{R}_+^n$, representing the rate of vaccine rollout for every node i at time t, as our control variable. The epidemiological model can be described by the following system of ordinary differential equations coupling disease transmission among

Fig 2. Data assimilation and scenarios for optimization. Comparison between the model outputs (95% confidence interval (CI) of the ensemble, blue shaded area) and the corresponding epidemiological data (red circles, obtained from the bulletins of the Dipartimento della Protezione Civile, `https://github.com/pcm-dpc/COVID-19`) from March 2020 to January 2021. The orange and green shaded area respectively show the ensemble dynamics (95% CI) of what we called pessimistic and optimistic transmission scenarios from January to April. The optimal vaccination strategy in the optimistic (or pessimistic) scenario is computed with respect to the the continuous green (or orange) line, representing the model trajectory obtained using the median of each ensemble parameter. **A.** The data on the daily hospitalizations is estimated as described in [9]; this data at the regional level is assimilated on a moving window of 14 days to update the model parameters describing the local transmission rates (see SI). **B.** Daily number of newly exposed individuals versus the reported positive cases. Note that, the model assumes the presence of 90% asymptomatics among the exposed individuals, who are possibly not detected by the surveillance system. **C.** Daily number of deaths.

all provinces:

$$\dot{x}(t) = F(x(t), v(t)) \tag{1}$$

The national incidence, i.e., the sum of new infections in all provinces at time t, is selected as the cost function, $L(x(t), v(t))$. Given our system with states x subject to the dynamics (1) and controls v, the optimal control problem can be formalized as:

$$\min_{v(\cdot)} \int_0^T L(x(t), v(t)) \, \mathrm{d}t \tag{2a}$$

$$\text{s.t.} \quad x(0) = \hat{x}_0, \tag{2b}$$

$$\dot{x}(t) = F(x(t), v(t)), \qquad \forall\, t \in [0, T], \tag{2c}$$

$$H(x(t), v(t)) \leq 0, \qquad \forall\, t \in [0, T], \tag{2d}$$

where we aim at minimizing the cost function over the control horizon T, while enforcing the modeled SARS-CoV-2 transmission dynamics (Equations (2b) and (2c)). Moreover, the constraints imposed by vaccine availability and the maximum vaccination

rate are lumped in function H that expands to

$$v_i(t) \geq 0, \qquad\qquad i \in \mathbb{I}_1^n, \qquad\qquad (3a)$$

$$\int_{t_\mathrm{d}}^{t_\mathrm{d}+1} v_i(t) \ \mathrm{d}t \leq v_i^\mathrm{max} \propto N_i, \qquad\qquad i \in \mathbb{I}_1^n, \ t_\mathrm{d} \in \mathbb{I}_0^T, \qquad\qquad (3b)$$

$$\int_0^t \sum_{i=1}^n v_i(t) \ \mathrm{d}t \leq D(t), \qquad\qquad \forall \, t \in [0, T], \qquad\qquad (3c)$$

where time is measured in days, and $\mathbb{I}_a^b$ is the set of all integers $a \leq k \leq b$. Equation (3a) enforces that one can only distribute a non-negative amount of vaccine doses. Equation (3b) states the logistic constraints, which limit to v_i^max the amount of individuals that can be vaccinated each day in each node; here t_d is the time at which each day starts. We impose that the daily vaccination capacity of each province is proportional to its population size N_i, assuming a fair distribution of the sanitary infrastructure among provinces, as shown in SI (Figure S1). The constraint on the national stockpile is materialized by Equation (3c), which ensures that the total vaccine allocation across all nodes does not exceed the stockpile $D(t)$. The stockpile is replenished every Monday by the delivery of new vaccines, hence $D(t)$ is a staircase function. For an overview of the possible solution approaches for optimal control problems we refer the interested reader to [40, 41]. In particular, in this work, we use a variant of direct multiple shooting [35] tailored to distributed systems [36]. We solve the optimal control problem in Equation (2) by a direct method, also called *first discretize, then optimize*, which transforms the control problem into a nonlinear programming problem. We split our time window $[0, T]$ into N intervals $[t_k, t_{k+1}]$, and we denote as $x_k = x(t_k)$ the states at time t_k, and as v_k the controls in interval $[t_k, t_{k+1}]$. The continuous-time dynamics F in Equation (1) are transformed by numerical integration into the discrete-time operator f by numerical integration. This discretization requires some care, and details are provided in the SI. We thus obtain the following nonlinear programming problem:

$$\min_{x,v} \ \sum_{k=0}^{N-1} l(x_k, v_k) \qquad\qquad\qquad (4a)$$

$$\text{s.t.} \quad x_0 = \hat{x}_0, \qquad\qquad\qquad (4b)$$

$$x_{k+1} = f(x_k, v_k), \qquad\qquad k \in \mathbb{I}_0^{N-1}, \qquad\qquad (4c)$$

$$H(x_k, v_k) \leq 0, \qquad\qquad k \in \mathbb{I}_0^{N-1}. \qquad\qquad (4d)$$

Nonlinear programming problems may be solved by readily available solvers using the primal-dual interior point method. The main difficulty in solving the proposed nonlinear programming problem (4) is the large dimension of the system and the non-linearities of the model. In order to bring the problem to a tractable form, we introduce three simplifications: (a) vaccines are administered instantaneously at the beginning of each day, rather than with a constant rate over the whole day; (b) the component of the force of infection taking into consideration the mobility of individuals across provinces is evaluated at the beginning of each day and remains constant through the day; and, (c) the mobility network is simplified, by keeping only the most important connections (see Figure 1), thus increasing the sparsity of the underlying spatial connectivity matrix. These simplifications deliver a significant computational advantage, and we verified that the impact on the model accuracy is limited. Note that, even if the optimal strategy is computed using the simplified model, its impact in terms of averted cases and deaths (shown in **Results**) is evaluated using the full epidemiological model

without any of these simplifications. A more detailed discussion on this subject is provided in the SI.

The nonlinear programming problem arising from the simplified epidemiological model is non-convex, and involves approximately 10^5 variables and 10^5 constraints. We formulate the problem using CasADi [37] and solve it using Ipopt [38] with sparsity-exploiting linear algebra solvers. In practice, solving the optimal control problem takes between two to four days on a 36-cores 2.3 GHz CPU.

Results

We obtain the optimal vaccination strategies for a set of eight scenarios drawn from the spatial model from January 11, 2021 to April 11, 2021. These scenarios are a combination of two projection scenarios (pessimistic vs optimistic) and four assumptions on the weekly stockpile delivery (479'700, 1M, 1.5M or 2M doses delivered per week). In each scenario, the optimal solution is a spatially explicit vaccine roll-out policy, i.e. an indication of the number of vaccine doses to be deployed in each province each day.

Comparison of vaccination strategies

In order to measure the potential impact of the optimal allocation strategy, we compare it against three non-optimized, yet reasonable alternative approaches: vaccinate proportionally to the future incidence as projected by the epidemiological model; vaccinate proportionally to the number of susceptible individuals in each province; vaccinate proportionally to the province's population. Comparisons with additional alternative strategies are presented in SI. Spatial prioritization based on epidemiological criteria, such as past [16] or future [17] incidence, has been thoroughly used in real campaigns and prospective studies.

For each of the eight scenarios considered, we compute the number of averted infections with respect to a zero-vaccination baseline, and the number of averted infections per vaccination dose (see Table 1). In the optimistic transmission scenario, characterized by a recess of the epidemic, the vaccination campaign has a lower impact on the averted infections per dose as only a small percentage of the vaccinated individuals would have been at risk of transmission. Obviously, the impact of the vaccination campaign is more evident in the pessimistic scenario where, for all strategies, the averted infections are larger than the vaccines deployed due to indirect protection effects. By virtue of the law of diminishing returns, the number of averted infections per dose decreases when increasing the stockpile.

The optimal solution always outperforms all alternative strategies in terms of the number of averted infections and in terms of averted infections per dose (see again Table 2). Incidence-based allocation comes second, while population and susceptibility-based allocation are distant third and fourth respectively. The improvement between optimal and incidence-based allocation is significant, ranging from 8.8% (pessimistic, 2M doses/week) to 16.2% (optimistic, 479'000 doses/week). In Figure 3, the black diamonds represent the percentage of adverted infections obtained using each strategy for the reference trajectory, with respect to the averted infections resulting from the optimal strategy. We observe that, in both the optimistic and pessimistic scenarios, the optimal strategy has the largest relative benefits for the smallest stockpile.

In the pessimistic scenario (see Figure 3.A), when 479'700 doses are made available each week, the averted infections associated with the optimal strategy in the reference projection are 1.37 per dose: 25 % more compared to the strategies based on population or susceptible distributions (1.08 averted infections per dose), and more than 10 % higher compared to the strategy based on the projected infections (1.21 averted

Weekly stockpile delivery		Averted Infections		Averted Infections per doses	
	Method	Optimistic	Pessimistic	Optimistic	Pessimistic
2M	**Optimal**	**6.98M**	**30.6M**	**0.268**	**1.18**
	Incidence	6.32M	28.1M	0.243	1.08
	Population	6.02M	26.8M	0.231	1.03
	Susceptibility	5.97M	26.7M	0.229	1.02
1.5M	**Optimal**	**5.52M**	**24.1M**	**0.283**	**1.24**
	Incidence	4.89M	21.7M	0.25	1.11
	Population	4.57M	20.4M	0.234	1.05
	Susceptibility	4.51M	20.3M	0.231	1.04
1M	**Optimal**	**3.9M**	**16.9M**	**0.3**	**1.3**
	Incidence	3.41M	15.1M	0.262	1.16
	Population	3.08M	13.8M	0.237	1.06
	Susceptibility	3.02M	13.7M	0.232	1.05
479'700	**Optimal**	**2.01M**	**8.54M**	**0.322**	**1.37**
	Incidence	1.73M	7.57M	0.278	1.21
	Population	1.49M	6.74M	0.239	1.08
	Susceptibility	1.43M	6.54M	0.229	1.05

Table 1. Averted infections. Absolute number of averted infections and averted infections per dose during the first three months of 2021 as evaluated for the reference trajectory (see Figure 2, and Figure 3 for an assessment on a set of trajectories drawn for the posterior distribution). The first column represents the considered scenarios of weekly stockpile replenishment, i.e. the number of doses delivered to Italy every week, ranging from 479'700 to 2M.

infections per dose). These differences are smaller but still significant when increasing the weekly stockpile deliveries up to 2M doses; similar results are obtained also for the optimistic transmission scenario (Figure 3.B).

We recall that the optimal control strategy considered here is computed for a reference model trajectory, which is the median of an ensemble of 1000 realizations. To further investigate the effectiveness of the optimal solution, we apply it to a subset of trajectories randomly sampled from the ensemble. The box plots in Figure 3 display the main quantiles of the averted infections computed for the ensemble of trajectories. We observe that the optimal strategy still yields better results than the ensemble of projections related to the other strategies, thus suggesting that the computed solution is robust even under the presence of perturbations in the forecasts of the epidemic dynamics. More importantly, for each realization of the ensemble and for each projection scenario, the optimal strategy systematically averts more infections than any of the other control strategies.

Our results suggest that it is possible to considerably increase the impact of vaccination campaigns by optimizing the vaccine allocation in space and time. For this task, optimal control provides the best possible strategy and sets a benchmark for the theoretical potential of a vaccination campaign.

Analysis of the optimal vaccine allocation

The optimal vaccine allocation obtained as the solution of the optimal control framework is complex to analyze, and we ought to do that by unraveling the mechanism behind its performance. The strategy must obey the imposed logistic and supply constraints: 1) The vaccine stockpile is replenished every Monday by a fixed amount of doses (e.g., 479'700 doses in the baseline scenario), and 2) the maximum possible distribution capacity per province is limited, proportionally to the node population, so

Fig 3. Comparison between different vaccine allocation strategies. Percentage of averted infections per vaccine doses from January 11, 2021 to April 11, 2021 resulting from province-scale vaccine allocation strategies for both the pessimistic (panel A) and the optimistic (panel B) scenarios based on: the optimal solution, population, proportion of susceptible individuals, and projected incidence (see color codes in the legend). We optimize the vaccine allocation for a (median) reference trajectory (diamonds), and assess the performance of the computed vaccination strategy over the whole posterior (box plots). For each projection scenario, results are normalized by the number of averted infections in the reference solution (see Table 1 for absolute values). Results for alternative scenarios and vaccination strategies are shown in SI Figure S5.

that the number of doses distributed across the country can be of 0.5M per day at maximum (more details in SI Figure S1).

We display the optimal vaccination course in time for the pessimistic, 479'700 doses/week scenario in Figure 4. We observe that the optimal allocation respects the two constraints on distribution (Figure 4, A) and supply (Figure 4, B). We observe that no province is vaccinated at the maximum possible rate during the whole campaign, and that provinces display a variety of vaccination patterns. We also note that the vaccines received every Monday are always distributed during the following week, but that the rate of delivery on a national level increases with time (Figure 4, B). Surprisingly, the optimal solution favors precise targeting over speed of delivery, in order to allocate more

Fig 4. Optimal vaccine allocation for the baseline, pessimistic scenario. **A.** Cumulative proportion of vaccine doses administered in the 107 provinces, (some of which are highlighted) . The local distribution rate is limited by a rate that is proportional to population. This logistic constraint is visualized here as the maximum slope, equal for every province. **B.** Stacked cumulative absolute number of vaccines in the 107 provinces of Italy. The national stockpile is shown in black, and is replenished every week (say on Mondays) with 479'700 doses. We display the name of the provinces with a final allocation of more than 150'000 doses.

doses to those provinces where the impact of vaccines on the whole system is projected to be higher. Hence, in order to control infections, precise targeting may trump delivery speed.

Fig 5. Spatial distribution patterns for the optimal allocation (left) and alternative strategies based on population, incidence and susceptibility (additional alternative strategies are presented in SI). We show, for each province and strategy, the proportion of vaccinated individuals (top), the number of averted infections per inhabitant with respect to a no vaccination baseline (middle), and the proportion of individuals who are still susceptible at the end of the control horizon (bottom). Visual inspection suffices to reveal that the optimal allocation strategy is more alike the incidence-based allocation, but with vaccine doses spread-out to more provinces.

Furthermore, we observe in the optimal solution that every time a province is vaccinated, the rate of vaccination is equal to the maximum rate allowed by the local logistic constraint. This property is common to the alternative vaccination strategies, hence the difference in performance is due to the spatial allocation patterns. In Figure 5, one can already see by visual inspection that the optimal allocation distributes most of the available doses on a few provinces with high incidence. These provinces are neither the most connected nor the most populous in Italy. The optimal strategy makes then use of the information on the network connectivity to fine-tune the allocation, and deploys the vaccination on more provinces than the incidence-based strategy.

To further investigate these patterns, in Figure 6.A we display the number of administered doses vs the incidence projected without vaccines (the proxy variables leading to the second-best control performance), both normalized according to the resident population in each province. We observe an allocation pattern whereby provinces with a higher incidence receive more vaccines. However, the allocation is non-linear with respect to the projected incidence, suggesting that to better control the epidemic, the optimal allocation strategy takes into account other factors such as the importance of each province within the mobility network, as well as the proportion of susceptibles. When the weekly stockpile delivery is increased, as shown in Figure 6.B, this pattern shifts to the right while remaining qualitatively consistent. Hence, the

optimal allocation strategy is robust with respect to the overall vaccine availability constraint, and the same nodes are prioritized. We provide scatter plots with other covariates in SI (Figures S6–S6).

Fig 6. **Analysis of the optimal solution A.** Vaccinated population according to the optimal strategy against the projected incidence without vaccination, both normalized by provincial population size and considering the scenario with a weekly stockpile delivery of 479'700 doses. **B.** Same as in A, but considering all four scenarios of weekly stockpile replenishment. Each dot represents a province, with dot size proportional to its population.

Discussion and Conclusion

Without any constraint on supply, each country would vaccinate its population as fast as possible according to the available infrastructure. However limitations in vaccine supply and rate of delivery are a reality for every country, hence the available doses should be deployed in space and time following a fair and effective strategy.

In stockpile-limited settings, like most current vaccination campaigns worldwide, careful allocation may significantly increase the number of averted infections and deaths. The goal is to distribute the vaccines where they have the strongest beneficial impact on the dynamics of the epidemic. However, deriving an algorithm capable of computing spatially optimal allocation strategies in real, heterogeneous settings is far from trivial and our approach is, to the best of our knowledge, the first attempt in this direction.

We developed an novel optimal control framework that delivers the best vaccination strategy under constraints on supply and logistics. This allows us to compute the allocation strategy that maximizes the number of averted infections during a projection of the COVID-19 epidemic in Italy from January 11, 2021 to April 11, 2021. Our results show that the optimal strategy has a complex structure that mainly reflects the

projected incidence of each province, but also takes into account the spatial connectivity provided by the mobility network and the landscape of acquired population immunity. Although the reason why this strategy is optimal is not immediately intuitive, our simulations clearly outline its better overall performances over other, more straightforward strategies. This comparison suggests that the simplicity underlying intuitive vaccination strategies may undermine their effectiveness, and calls for complementing these simple approaches with rigorous and objective mathematical tools, like optimal control, that allow a full account of the complexity of the problem.

With the present work, we show that it is possible to solve optimal control problems for spatially explicit dynamical models of infectious diseases at a national scale, thus overcoming the computational limitations that, up to now, precluded this kind of applications. The proposed framework can account for any compartmental epidemic model, with up to hundreds of connected spatial nodes. Supply and logistic constraints can be adapted to the actual landscape of decisions faced by the stakeholders, such as no/reduced vaccine delivery on weekends, or the need for fairness in vaccine distribution, e.g. by ensuring that each region receives at least a fixed fraction of the available vaccines. This is especially important as in our optimal allocation scenarios, some provinces receive no vaccine at all. Moreover, the optimization can be carried for single-dose vaccines, as done here, or for two-dose vaccines, where one could potentially optimize the time between the first and second dose (and if a second dose should be administered at all), clearly also considering the intervals recommended by the health authorities.

Our method is obviously not devoid of limitations. The main one is that the optimal vaccination strategy strongly depends on the projection of the underlying epidemiological model. These projections are subject to several sources of uncertainty, especially for long term horizons, for example due to model design and calibration [42], the generation of baseline transmission scenarios, and unforeseen events that may change the course of the epidemic (such as the importation of cases, the emergence of new virus variants, changes in disease awareness or social distancing policies). The optimal vaccination strategy is thus reliable only if the projections given by the underlying model dynamics are sufficiently accurate. A successful approach developed by the automatic control community to tackle that issue, named Model Predictive Control [43], consists in compensating the performance losses expected over long horizons by constantly adapting the optimal strategy. In this context, Model Predictive Control might be implemented using the following steps: (a) at the beginning of each week, the state of the system is estimated by using newly acquired epidemiological data; (b) the optimization problem is solved over a fixed prediction horizon using the estimated state as initial condition; (c) the optimal strategy for the first week is applied and, as soon as the next weeks starts, these steps are repeated starting from (a). This method corrects the model inaccuracies by constantly resetting the initial state to the estimated one. Moreover, constraints may be updated to account for unexpected deliveries or new orders. Future work will aim at further evaluating the benefits of implementing this scheme for the design of optimal vaccination strategies.

Moreover, the epidemiological model underlying our control optimization has known validity and limitations [8,9]. An additional limitation of the model for the specific scopes of this work is that it does not account explicitly for risk-based classes, and thus does not account for the heterogeneities that may result from the demography of the population, as well as from the age-related transmission and clinical characteristics associated with COVID-19. While surely limiting for operational use of the tools, we note that the scope of this paper is to provide a proof-of-concept of the relevance of spatial effects, which have not been addressed so far in the literature. To that end, we are confident that our results support the relevance of the research question posed. Our

framework can anyway be extended to optimize across both spatial and risk   396
heterogeneities.   397

A counter-factual assumption in this work is that we consider a one-dose vaccine   398
with full and instantaneous efficacy. At the time of development, the details about   399
COVID-19 vaccines were not released, and this hypothesis allowed us to demonstrate   400
our framework in a simple setting. Our framework can be further extended to account   401
also for the simultaneous deployment of different vaccine types, some of which may   402
require the administration of two doses. This extension too is subject of ongoing   403
research, in particular to extend the modeling tools described here to accommodate the   404
peculiarities of each authorized vaccine candidate while designing effective   405
spatiotemporal deployment strategies.   406

In conclusion, in this work we have optimized vaccine allocation at country scale on   407
different scenarios of epidemic transmission and vaccine availability. Using a data   408
assimilation scheme, we updated a spatially explicit compartmental model that had   409
already been successfully used to describe the COVID-19 pandemic in Italy. To this   410
aim, we have discretized, transformed and simplified the model and constructed a   411
pipeline to perform large-scale nonlinear optimization on vaccine allocation, subject to   412
stockpile and logistic constraints. Solving this problem yielded a complex solution that   413
outperforms other strategies by a significant margin and proves robust across posterior   414
realizations of the underlying model. As such, beside inherent limitations, it provides a   415
benchmark against which other, possibly simpler vaccine rollout strategies can be   416
usefully compared.   417

Supporting information   418

Optimal control problem   419

The proposed framework is constituted of two disease transmission models, one "true"   420
model and a simplified one used for control: 421

- The full model is a COVID-19 model, as designed in [8,9]. This model is   422
  ODE-based, includes full connectivity based on mobility data, and is implemented   423
  in MATLAB using the adaptive step ode45 integration scheme. Using data   424
  assimilation, we obtain the joint-posterior distribution for all parameters of this   425
  model.   426

- The simplified model used for optimal control is an approximation of the above   427
  model, integrated using an explicit Runge-Kutta 4 method with fixed stepsize. We   428
  simplified the problem by limiting the connectivity to the largest mobility fluxes   429
  (see Figure 1 B) and optimizing only one realization of the posterior. This model   430
  is implemented in Python with the CasADi library.   431

This division is necessary in order to solve the optimal control problem (OCP) in a   432
reasonable time. To adapt our framework to another model/country, one would need to   433
update the "true" model to a suitable candidate (which could be a stochastic model, a   434
Hidden Markov model, or any other kind) and design a tractable approximation of this   435
new model to be solved by optimal control.   436

In order to evaluate the effectiveness of our approach, we first compute the optimal   437
vaccination course that minimizes the objective based on the simplified model. Then,   438
we assess this strategy and the alternative ones on the full model, for different posterior   439
realisations. If the simplified model is sufficiently accurate, the performance loss is small 440
and the proposed strategy outperforms simpler strategies, as shown in our simulation   441
results.   442

It the subsections below, we first detail the full COVID-19 model, then we describe the optimal control framework and the simplifications we have introduced to bring the problem to a tractable form.

COVID-19 model

The optimal control framework may be used with any compartmental SARS-CoV-2 transmission model that can be approximated by ordinary differential equations. To demonstrate its usefulness, we use a complex model based on previous work that was aimed to describe the first wave of COVID-19 infections in Italy [8,9]. We consider the 107 Italian provinces and the spatial connections induced by human mobility fluxes. In each province, the human population is subdivided according to infection status into the epidemiological compartments of susceptible S, exposed E, pre-symptomatic P (incubating infectious), symptomatic infectious I, asymptomatic infectious A, quarantined Q, hospitalized H, recovered R, dead D, and vaccinated V. The possible transitions between these compartments are shown in Figure 1.A. Individuals in compartments P, A and I are infectious and contribute differently to the force of infection, driving susceptible S into incubating individuals E.

The COVID-19 transmission dynamics are described by the following set of ordinary differential equations in each node i:

$$
\begin{aligned}
\dot{S}_i &= -\lambda_i(t)S_i - r_i^v(t)S_i \\
\dot{E}_i &= \lambda_i(t)S_i - (\delta^E + r_i^v(t))E_i \\
\dot{P}_i &= \delta^E E_i - (\delta^P + r_i^v(t))P_i \\
\dot{I}_i &= \sigma\delta^P P_i - (\gamma^I + \eta)I_i \\
\dot{A}_i &= (1-\sigma)\delta^P P_i - (\gamma^A + r_i^v(t))A_i \\
\dot{Q}_i &= \zeta\eta I_i - \gamma^Q Q \\
\dot{H}_i &= (1-\zeta)\eta I_i - (\gamma^H + \alpha^H)H \\
\dot{R}_i &= \gamma^I I_i + \gamma^A A_i + \gamma^H H_i + \gamma^Q Q_i - r_i^v(t)R_i \\
\dot{V}_i &= r_i^v(t) \cdot (S_i + E_i + P_i + A_i + R_i).
\end{aligned}
\tag{S1}
$$

We define N_i the population of province i. Susceptible individuals get exposed to the pathogen at rate $\lambda_i(t)$, corresponding to the force of infection for community i, thus becoming latently infected (but not infectious yet). Exposed individuals transition to the post-latent, infectious stage at rate δ^E. Post-latent individuals progress to the next infectious classes at rate δ^P, developing an infection that can be either symptomatic—with probability σ—or asymptomatic—with probability $1-\sigma$. Symptomatic infectious individuals recover from infection at rate γ^I and may seek treatment at rate η. Asymptomatic individuals recover at rate γ^A. Infected individuals who sought treatment are either hospitalized (rate $1-\zeta$) or quarantined (rate ζ) at home and are considered to be effectively removed from the community, thus not contributing to disease transmission. Individuals who recover from the infection are assumed to have long-lasting immunity to reinfection at the timescale studied, but possible loss of immunity can be easily included in the model. Hospitalized individuals die at rate α_H and recover at rate γ^H.

Individuals in compartments S, E, P, A, R might receive vaccine doses. If the chosen strategy allocates $v_i(t)$ doses in node i at time t, the vaccination rate is

$$
r_i^v(t) = \frac{v_i(t)}{S_i(t) + E_i(t) + P_i(t) + A_i(t) + R_i(t)}
\tag{S2}
$$

Vaccinated individuals are moved at rate $r_i^v(t)$ from their original compartments to compartment V, where they do not contribute to the infection anymore.

The force of infection of the full model is specified as in Gatto et al. [8], Bertuzzo et al. [9]. In addition to province's local dynamics, it also considers that local susceptibles may enter in contact with infected individuals that are traveling, and oppositely, susceptible commuters may become infected through contact with local infected. The force of infection of the OCP model slightly simplified, and detailed thereafter.

We split the force of infection $\lambda_i(t)$ as the sum of the local force of infection $\lambda_i^L(t)$, from infected in node i and a mobility-driven force of infection from the network $\lambda_i^N(t)$, hence $\lambda_i(t) = \lambda_i^L(t) + \lambda_i^M(t)$. We observe while running our model that $\lambda_i^M(t) \ll \lambda_i^L(t)$. Hence this artificial separation will be exploited when simplifying our model. As described below, we update $\lambda_i^M(t)$ every day whereas $\lambda_i^L(t)$ is updated at each integration step.

As for the formulations of the force of infection, we recall here the formulas designed by [8,9]. The local force of infection reads:

$$\lambda_i^L(t) = C_{i,i}\beta_0\beta_i(t) \cdot \frac{C_{i,i}(P_i + \epsilon_A A_i) + \epsilon_I I_i}{C_{i,i} \cdot (S_i + E_i + P_i + R_i + A_i + V_i) + I_i}, \tag{S3}$$

and the influence of other provinces on province i is written as:

$$\lambda_i^M(t) = \sum_{m,m \neq i} \left(C_{i,m} \cdot \frac{\sum_{n,n \neq m}[C_{n,m} \cdot \beta_0\beta_n(t)(P_n + \epsilon_A A_n)] + \epsilon_I \beta_0 \beta_m(t) I_m}{\sum_{l,l \neq m}[C_{l,m} \cdot (S_l + E_l + P_l + R_l + A_l + V_l)] + I_m} \right), \tag{S4}$$

where β_0 is the baseline transmission rate, while $\beta_i(t)$ is a spatially distributed and time-varying parameter describing site- and time-specific variations in transmissibility due to non-pharmaceutical interventions or other exogenous factors like variants. The parameters ϵ_A and ϵ_I represent the reduction of transmission respectively for asymptomatic and symptomatic individuals with respect to pre-symptomatic individual transmissions. Matrix C accounts for mobility: each element $C_{i,j}$ of the matrix $(i \neq j)$ represents the proportion of individuals moving from i to j, while the diagonal elements $C_{i,i}$ are the proportions of individual who do not move in each node i.

The objective for our model is to minimize the total incidence of infections, i.e., $\int_{t_i}^{t_f} \sum_i \lambda_i(t) S_i$. Note that for the present model, this is equivalent to optimizing the total deaths or hospital admissions, as without risk-classes the sizes of these two compartments are proportional to each other.

Optimal control

We lump the epidemiological compartments of each node i in variable $x_i(t) = (S_i(t), E_i(t), P_i(t), I_i(t), A_i(t), Q_i(t), H_i(t), R_i(t), V_i(t))$ and we define $v_i(t)$ as our control variable, representing the number of vaccines administered in node i at time t. We express the dynamics of the epidemiological model (Equation (S1)) as an ordinary differential equation in each province i:

$$\dot{x}_i(t) = F_i(x_i(t), v_i(t), m_i(t), t), \tag{S5}$$

where $m_i(t)$ carries the contribution of other provinces to the force of infection of node i. For simplicity, we drop the time dependence in the equations below, and we define the state and control variables for the full system as

$$x = (x_1, \ldots, x_n), \qquad\qquad v = (v_1, \ldots, v_n),$$

Fig S1. Local maximum vaccination rate $v_i^{\max}$ for each province. This logistic constraint bounds the maximum number of vaccines to 0.5M of doses per day, with a local rate that is proportional to the node population. Here we show the maximum vaccination rate for each province (the constraint the solution has to comply with), in red, and the maximum rate prescribed by the optimal solution while simulating the pessimistic scenario with a stockpile delivery of 479'700 doses, in black. The optimal solution uses the maximal capacity of the logistic network, while respecting the constraint defined.

where n is the number of spatial node considered ($n = 107$). The global dynamics for all provinces are denoted:

$$F(x, v) = (F_1(x_1, v_1, m_1), \ldots, F_n(x_n, v_n, m_n)).$$

The coupled force of infection in node i is denoted λ_i. We define the cost function as the sum of total incidence of infections (transitions $S_i \longrightarrow E_i$) for every node i, i.e.,

$$L(x, v) = \sum_{i=1}^{n} \lambda_i S_i.$$

For the sake of generality, we introduce the terminal cost M, which can be used to ensure that we leave the system in a proper state instead of optimizing for short-term gain. Since properly designing the terminal cost could require a long analysis, for simplicity we do not use it in this work, hence $M(\cdot) = 0$.

Given our dynamical system with states x, controls v, and dynamics F, the OCP is:

$$\min_{v(\cdot)} \quad \int_0^T L(x(t), v(t)) \, \mathrm{d}t + M(x(T)) \tag{S6a}$$

$$\text{s.t.} \quad x(0) = \hat{x}_0, \tag{S6b}$$

$$\dot{x}(t) = F(x(t), v(t)), \qquad\qquad \forall\, t \in [0, T], \tag{S6c}$$

$$H(x(t), v(t)) \leq 0, \qquad\qquad \forall\, t \in [0, T], \tag{S6d}$$

where we aim at minimizing the cost function over the prediction horizon T, while enforcing the modeled SARS-CoV-2 transmission dynamics (Equations (S6b) and (S6c)). Moreover, constraints on vaccine availability and maximum vaccination rate are lumped in function H, which reads:

$$v_i(t) \geq 0, \qquad\qquad i \in \mathbb{I}_1^n, \tag{S7a}$$

$$\int_{t_{\mathrm{d}}}^{t_{\mathrm{d}}+1} v_i(t) \, \mathrm{d}t \leq v_i^{\max} \propto N_i, \qquad\qquad i \in \mathbb{I}_1^n, \ t_{\mathrm{d}} \in \mathbb{I}_0^T, \tag{S7b}$$

$$\int_0^t \sum_{i=1}^{n} v_i(t) \, \mathrm{d}t \leq D(t), \qquad\qquad \forall\, t \in [0, T], \tag{S7c}$$

where time is measured in days, and $\mathbb{I}_a^b$ is the set of all integers $a \leq k \leq b$. Equation (S7a) enforces that one can only distribute a non-negative amount of vaccine doses.

Equation (S7b) states the logistic constraints, which limit the amount of individuals that can be vaccinated each day in each node to $v_i^{\max}$; here t_{d} is the time at which each day starts. We assume that the daily capacity of each province is proportional to the population size of each node N_i, because we assume a fair distribution of the sanitary infrastructure among provinces with regard to population, as shown in SI Figure S1. The stockpile is materialized by Equation (S7c), which ensures that the total vaccine allocation across every node does not exceed the total availability $D(t)$. The stockpile is replenished every Monday by the delivery of new vaccines, hence $D(t)$ is a staircase function.

We convert our problem formulation (S6) to a nonlinear programming problem using direct multiple shooting. Standard multiple shooting splits the time horizon $[0, T]$ using a time grid $t_0, \ldots, t_N$, with $N + 1$ points and $t_0 = 0$, $t_N = T$. The control function is parameterized using basis functions with local support. Common choices are a uniform time grid, i.e., $t_{k+1} = t_k + \delta_t$ and a piecewise constant control function, i.e., $v(t) = v_k$, $t \in [t_k, t_{k+1}]$. The system dynamics are then discretized to obtain a discrete-time system

$$x_{k+1} = f(x_k, v_k),$$

satisfying $x_k = x(t_k)$ for all $k = 0, \ldots, N$. Moreover, the cost function is also discretized, to obtain

$$l(x_k, v_k) = \int_{t_k}^{t_k+1} L(x(t), v(t)).$$

We perform the discretization using numerical integration techniques (such as a fourth-order Runge-Kutta scheme, with 50 steps per days) to obtain a good approximation of the true trajectory and cost. Finally, the path constraints H are relaxed and imposed at a finite amount of time instants, here coinciding with the time grid $t_0, \ldots, t_N$. We ought to observe that, since in our case the constraints only involve the controls, we are not introducing any approximation by enforcing these constraints only on this uniform grid. The OCP (S6) is then approximated by the nonlinear programming problem

$$\min_{x,v} \ M(x_N) + \sum_{k=0}^{N-1} l(x_k, v_k) \tag{S8a}$$

$$\text{s.t.} \ x_0 = \hat{x}_0 \tag{S8b}$$

$$x_{k+1} = f(x_k, v_k), \qquad\qquad k \in \mathbb{I}_0^{N-1}, \tag{S8c}$$

$$H(x_k, v_k), \qquad\qquad k \in \mathbb{I}_0^{N-1}. \tag{S8d}$$

In (S8), both the states $x = (x_0, \ldots, x_N)$ and the controls $v = (v_0, \ldots, v_{N-1})$ are defined as optimization variables, which is a distinguishing trait of multiple shooting as opposed to single shooting.

The main difficulty in solving (S8) in the context of this paper is the large dimension of the system and the nonlinearity of the model, which can pose severe issues to the numerical solvers. In the following, we will thus introduce a few simplifications, and we will verify through numerical simulations that these simplifications do not imply large errors in the solution of the OCP.

We discretize the OCP using a uniform grid with sampling time $\delta t = 1$ day. We assume that (a) vaccinations are administered instantaneously at the beginning of each day, rather than with a constant rate over the whole day; (b) the force of infection associated with mobility is constant over each day; and (c) the weakest mobility links can be pruned. Thus, each node dynamics can be made independent of the other nodes

dynamics by introducing an auxiliary control variable z that is constrained to match the force of infection due to the other nodes at the beginning of each time interval. Then, the dynamics of the decoupled system in each node can be written as:

$$\dot{x}_i(t) = F_i(x_i(t), z_{i,k}), \qquad\qquad t \in [t_k, t_{k+1}]$$
$$x_i(t_k) = x_{i,k} + g_i(v_{i,k}), \qquad\qquad z_{i,k} = e_i(x).$$

Discussion on Simplification (a). We ought to remark that, realistically, vaccinations will occur at least eight hours per day. Our assumption, while justified as a computationally convenient approximation of reality, is not a priori worse than assuming that vaccine administration takes place over the whole day. More refined approximations, while in principle possible, pose severe issues because of the nature of the system dynamics. While for most initial values the system dynamics can be easily simulated with time-continuous vaccinations, the system becomes stiff by construction once almost the entire population has been vaccinated. In this case, numerical integration errors can drive the size of some compartments to be negative, which violates the model assumptions and makes the result of the numerical integration meaningless. The main issue in this case is that the optimizer will exploit these inaccuracies in order to reduce the cost. Therefore, this issue is much more evident when solving optimal control problems than when simply simulating the system dynamics. We have investigated some simple approaches to tackle this issue, but no technique yielded satisfactory performances. It is our impression that ad-hoc integration strategies will be required in order to reliably simulate and optimize dynamics with continuous vaccination rates. While this will be the subject of future research, the results obtained with the current approximation have yielded sufficient accuracy.

Discussion on Simplification (b). This simplification has been proposed in [36] as an approach to solve distributed optimal control problems by means of multiple shooting. In the original version, the coupling variable z is not necessarily piecewise constant, but rather piecewise polynomial. We have observed in simulations that, for this problem, the piecewise constant parametrization yielded sufficient accuracy.

We discretize the dynamics of each node using an explicit Runge-Kutta integrator of order four, with 50 integration steps per day. Alternative integrators such as explicit Euler, or implicit Runge-Kutta integrators, yielded similar results. Furthermore, in order to verify the accuracy of the integrator and the impact of the introduced simplifications on the solution accuracy, we simulated the system in open-loop, i.e. we applied the optimal control trajectory to the full model starting from the initial condition provided by the data assimilation scheme.

Discussion on Simplification (c). We sparsify the mobility matrix by pruning element below a threshold (see Figure S2). This operation reduces the number of connection between nodes. Also in this case, we verified through numerical simulations that the introduced simplification had a small impact on the prediction and control accuracy.

Possible further improvements Applying optimal control in open loop, i.e., solving the optimization problem once and applying the control input over the whole time interval, may lead to poor performance due to model inaccuracy and external perturbations. A common remedy consists in closing the loop by repeatedly solving the OCP by using the most updated information on the initial states. This is the principle behind Model Predictive Control (MPC) [43]. In this context, the state would be

Fig S2. Simplification of the mobility matrix to obtain a sparse and tractable problem. After the optimization, we assess the effectiveness of the optimal control strategy on the full model.

estimated on a daily, weekly, or monthly basis so as to solve again the OCP and correct the optimal strategy.

Implementation of the OCP

We implement the optimal control framework using the automatic differentiation framework CasADi [44], the interior-point solver ipopt [38], and the HSL ma86 large sparse symmetric indefinite solver [45]. The full framework and analysis code is available here: `https://github.com/jcblemai/COVID-19_italy-vaccination-oc` (a zenodo DOI will be added after reviews).

Solving the OCP is both CPU and RAM intensive. For numerical computations, we used the Helvetios cluster the EPFL HPC facility (one problem per computing node, each equipped with 36 2.3 GHz cores and 192 GB RAM). On this cluster, it takes approximately four days to solve the large-scale OCP just presented. It should be possible to solve even larger problems with more RAM available.

Data assimilation and model parameters

The regional transmission rates are the main parameters governing the force of infection of the model and, thus, the daily exposed individuals. To better track possible changes in the transmission rates, we adopt a data assimilation strategy based on an iterative particle filter [46] used on a moving window of 14 days. The filter starts considering $N_r = 1000$ model realizations at time t_0 (February 21st, 2020), whose state variables are $x_0^{(j)}$, $j = 1, \ldots, N_r$, where the superscript (j) is the realization index and the subscript is the temporal index. Each realization is associated with a parameter combination that is randomly sampled from the posterior distribution evaluated in [9], indicated with $\theta^{(j)}$. Possible spatial heterogeneities in regional transmission on a given day t_k are obtained multiplying the transmission parameter by a coefficient $\phi_{k,i}^{(j)}$, where i is the region's index. At time t_0, the coefficients $\phi_{0,i}^{(j)}$ are sampled from a truncated normal distribution (mean $\mu_0 = 1$, standard deviation 0.4, bounds 0.8μ-$1.2\mu_0$). At time t_k, we assume to know the state variables $x_k^{(j)}$ and coefficients $\phi_{k,i}^{(j)}$, the latter having ensemble mean $\mu_{k,i}$. To update state variables and coefficients at time t_{k+1}, we consider the observations

Fig S3. Modeled daily hospitalizations (blue) the versus hospitalization data (red dots), regional detail of Figure 2.A in the main text. The optimistic and pessimistic transmission scenarios are represented in green and yellow, respectively.

Fig S4. Modeled daily incidence (blue) versus the daily reported cases (red dots), regional detail of Figure 2.B in the main text. The optimistic and pessimistic transmission scenarios are represented in green and yellow, respectively.

(daily hospitalizations) collected in a temporal window of $\tau = 14$ days, $(t_k, t_k + \tau)$. New coefficients from the truncated normal distribution (mean $\mu_k = 1$, standard deviation 0.4, bounds $0.8\mu_k$-$1.2\mu_k$) are sampled at time $\tau = t_0 + 14$ days. For each realization, we run the model during the window of 14 days, assuming that the coefficients change linearly for a week, from $\phi_{0,i}^{(j)}$ to $\tilde{\phi}_{0,i}^{(j)}$, and remain constant afterwards. The regional likelihood of each realization is then evaluated during these two weeks considering that the daily hospitalizations follow a gamma distribution (as in [9]). A resampling step (systematic resampling , see, e.g. [47]) selects and duplicates the coefficients $\tilde{\phi}_{0,i}^{(j)}$ associated with the largest likelihood values. These coefficients are then used to update the mean value μ_k. Finally, the simulation is repeated on the same temporal window by sampling new coefficients $\tilde{\phi}_{0,i}^{(j)}$ from the truncated normal distribution with the updated mean μ_k. This set of coefficients is used to compute state variables and parameters at time t_k, and then as starting condition to produce the projections used in the main text.

Model parameters (in the absence of vaccination) are taken from a paper [9] where they were inferred in a Bayesian framework for the period February 24th – May 1st, 2020, on the basis of the official epidemiological bulletins released daily by Dipartimento della Protezione Civile [48] (data available online at `https://github.com/pcm-dpc/COVID-19`) and the bulletins of Epicentro, at Istituto Superiore di Sanità [49, 50]. All the parameters estimated for the initial phase of the Italian COVID-19 epidemic, including the transmission rates, are spatially homogeneous [9]. This parameterization has been used to produce all the results presented in the main text.

Spatial set-up

The modeling tools described in the following sections are applied to the Italian COVID-19 epidemic at the scale of second-level administrative divisions, i.e. provinces and metropolitan cities (currently, as of 2021, 107 spatial units). Official data about resident population at the provincial level is produced yearly by the Italian National Institute of Statistics (Istituto Nazionale di Statistica, ISTAT; data available at `http://dati.istat.it/Index.aspx?QueryId=18460`). The latest update (January 1, 2019) has been used to inform the spatial distribution of the population.

The data to quantify nation-wide human mobility come from ISTAT (specifically, from the 2011 national census; data available online at `https://www.istat.it/it/archivio/139381`). Mobility fluxes, mostly reflecting commuting patterns related to work and study purposes, are provided at the scale of third-level administrative units (municipalities) [51, 52]. These fluxes were upscaled to the provincial level following the administrative divisions of 2019, and used to evaluate the fraction p_i of mobile people in each node i, as well as the fraction q_{ij} of mobile people who move between i and all other administrative units j (see Supplementary Material in [8]).

Alternative strategies

We designed alternative strategies to compare the optimal solutions. Each strategy uses a decision variable, $\mathcal{V}_i$, as a basis for the allocation of vaccines among provinces. The decision variable is one of:

- **modelled future incidence, absolute**: the modelled total future incidence in a no-vaccination scenario. This is equivalent to the objective of the optimal control problem with no control;

- **modelled future incidence, per population**: as above, but normalized by the resident population in each node;

- **modelled initial susceptibility, absolute**: the modelled number of susceptibles in each province at the start of the vaccination campaign;

- **modelled initial susceptibility, per population**: as above, but normalized by the resident population in each node;

- **province's population**.

We define two strategies to distribute the doses:

- **Focused** Where every province is sorted (higher on top) according to its decision variable $\mathcal{V}_i$. We then allocate the maximum local rate v_i^{max} to every province going down through the list, until the stockpile is empty. In other words, assuming we have an amount K of vaccines in the stockpile, we find the province index i that satisfy $\max_i \mathcal{V}_i$, and we assign to province i $M_i = \min(v_i^{max}, K)$ vaccines. Then, we find the province j that satisfy $\max_{j,j \neq i} \mathcal{V}_j$ and we assign it $M_j = \min(v_j^{max}, K - M_i)$. And so on, until no vaccine remains in the stockpile. This strategy will concentrate the allocation on nodes with the highest values of the considered decision variable.

- **Proportional** In this case, assuming that on a given day there is a quantity of vaccine K in the stockpile, we assign to each province i an amount $M_i = \min(v_i^{max}, K \cdot \frac{\mathcal{V}_i}{\sum_j \mathcal{V}_j})$. This approach vaccinate each node proportionally to the value of its decision variable $\mathcal{V}_i$.

In the main text, we show the results for three alternative strategies, namely *proportional absolute incidence*, *proportional population*, and *proportional susceptibility*—named respectively Incidence, Population, and Susceptibility. These strategy are good performers across scenarios, and show how different choices for the decision variables may affect the outcomes of the OCP. In the next sections, we will show the results for all these alternative strategies.

Additional results for the spatial model

We present the results for all these strategies in Table 2, and we show them side-by-side in Figure S5. The optimal solutions outperforms all the others solution. In fact, for every given posterior realisation, the optimal control solution always outperforms all other allocation strategies. Even if we observe some scatter when sampling the posterior, the performances of optimal strategies are clearly separated from the rest of the alternatives.

To further investigate the features of the optimal solution, we present a linear scatter plot of the optimal proportion of vaccinated individuals per province (sorting variable) side by side with the province population, the projected incidence without vaccination, and the proportion of susceptible individuals at the start of the simulation. We present these results for the optimistic scenario in Figure S6 and for the pessimistic scenario in Figure S7. We find no clear visual pattern associating these covariates to the optimal proportion vaccinated, highlighting again that the optimal allocation uses the epidemiological variable in a non-straightforward way, different from every simple strategy we could come up with.

Finally, to highlight the temporal dimension of the prioritization strategy for the deployment of vaccine doses, we present an example of epidemiological dynamics in every compartment of the model for an optimal scenario (Figure S8), and a stackplot of the proportion of vaccine dose allocated in each province according to the optimal solution (Figure S9).

		Averted Infections		Averted Infections per dose	
Scenario	Method	Optimistic	Pessimistic	Optimistic	Pessimistic
2M	Optimal	6.98M	30.6M	0.268	1.18
	Incidence	6.32M	28.1M	0.243	1.08
	Proportional Incidence	6.23M	27.5M	0.239	1.06
	Focused Susceptibility	6.03M	26.9M	0.232	1.03
	Focused Proportional Susceptibility	6.03M	26.9M	0.232	1.03
	Focused Proportional Incidence	6.03M	26.9M	0.232	1.03
	Focused Population	6.03M	26.9M	0.232	1.03
	Focused Incidence	6.03M	26.9M	0.232	1.03
	Population	6.02M	26.8M	0.231	1.03
	Susceptibility	5.97M	26.7M	0.229	1.02
	Proportional Susceptibility	5.6M	25.3M	0.215	0.971
1.5M	Optimal	5.52M	24.1M	0.283	1.24
	Incidence	4.89M	21.7M	0.25	1.11
	Proportional Incidence	4.82M	21.3M	0.246	1.09
	Focused Population	4.58M	20.5M	0.235	1.05
	Focused Incidence	4.58M	20.5M	0.235	1.05
	Focused Proportional Incidence	4.58M	20.5M	0.235	1.05
	Focused Proportional Susceptibility	4.58M	20.5M	0.235	1.05
	Focused Susceptibility	4.58M	20.5M	0.235	1.05
	Population	4.57M	20.4M	0.234	1.05
	Susceptibility	4.51M	20.3M	0.231	1.04
	Proportional Susceptibility	4.18M	19.0M	0.214	0.975
1M	Optimal	3.9M	16.9M	0.3	1.3
	Incidence	3.41M	15.1M	0.262	1.16
	Proportional Incidence	3.34M	14.7M	0.257	1.13
	Focused Population	3.09M	13.9M	0.238	1.07
	Focused Susceptibility	3.09M	13.9M	0.238	1.07
	Focused Proportional Susceptibility	3.09M	13.9M	0.238	1.07
	Focused Incidence	3.09M	13.9M	0.238	1.07
	Focused Proportional Incidence	3.09M	13.9M	0.238	1.07
	Population	3.08M	13.8M	0.237	1.06
	Susceptibility	3.02M	13.7M	0.232	1.05
	Proportional Susceptibility	2.75M	12.6M	0.211	0.972
479'700	Optimal	1.96M	8.39M	0.314	1.34
	Focused Proportional Incidence	1.95M	7.74M	0.312	1.24
	Proportional Incidence	1.69M	7.32M	0.271	1.17
	Incidence	1.63M	7.21M	0.262	1.15
	Focused Incidence	1.59M	6.64M	0.254	1.06
	Focused Population	1.57M	6.85M	0.251	1.09
	Population	1.45M	6.57M	0.233	1.05
	Focused Susceptibility	1.45M	6.53M	0.232	1.04
	Susceptibility	1.41M	6.43M	0.225	1.03
	Focused Proportional Susceptibility	1.28M	6.09M	0.204	0.973
	Proportional Susceptibility	1.26M	5.89M	0.202	0.944

Table 2. Absolute number of averted infections for each scenario

Acknowledgments

JCL and AR acknowledge funding from the Swiss National Science Foundation via the project "Optimal control of intervention strategies for waterborne disease epidemics" (200021–172578). AR, EB and DP acknowledge funding from Fondazione Cassa di Risparmio di Padova e Rovigo (IT) through its grant 55722. JCL acknowledge Nicolas Richart for help on scaling up computations on the EPFL high-performance computing clusters.

References

1. Khamsi R. If a coronavirus vaccine arrives, can the world make enough? Nature. 2020;580(7805):578–580. doi:10.1038/d41586-020-01063-8.

2. National Academies of Sciences aM Engineering. Framework for Equitable Allocation of COVID-19 Vaccine. Kahn B, Brown L, Foege W, Gayle H, editors. The National Academies Collection: Reports funded by National Institutes of Health. Washington (DC): National Academies Press (US); 2020. Available from: http://www.ncbi.nlm.nih.gov/books/NBK562672/.

3. Spassiani I, Gubian L, Palù G, Sebastiani G. Vaccination Criteria Based on Factors Influencing COVID-19 Diffusion and Mortality. Vaccines. 2020;8(4):766. doi:10.3390/vaccines8040766.

Fig S5. Comparison of different allocation strategies. Percentages of averted infections per vaccine dose from January 11, 2021 to April 11, 2021 using different vaccine distribution strategies for the pessimistic (panel A) and the optimistic (panel B) scenario based on: the optimal solution, the spatial distribution of the population, the amount of susceptible individuals at the beginning of the vaccination campaign, and the projected disease incidence in the absence of control. We optimize a median realization of the modeled posterior (diamonds), and assess the performance on the whole posterior (box plots). The results are normalized by the number of averted infections in the optimized solution (see Table 2 for absolute values).

4. Matrajt L, Eaton J, Leung T, Brown ER. Vaccine optimization for COVID-19: who to vaccinate first? medRxiv. 2020; p. 2020.08.14.20175257. doi:10.1101/2020.08.14.20175257.

5. Gallagher ME, Sieben AJ, Nelson KN, Kraay ANM, Orenstein WA, Lopman B, et al. Indirect benefits are a crucial consideration when evaluating SARS-CoV-2 vaccine candidates. Nature Medicine. 2021;27(1):4–5. doi:10.1038/s41591-020-01172-x.

6. Tuite AR, Zhu L, Fisman DN, Salomon JA. Alternative dose allocation strategies to increase benefits from constrained COVID-19 vaccine supply. Annals of Internal Medicine. 2021; p. 1–18.

7. Marziano V, Guzzetta G, Rondinone BM, Boccuni F, Riccardo F, Bella A, et al. Retrospective analysis of the Italian exit strategy from COVID-19 lockdown.

Fig S6. Control and co-variates for the optimistic scenario with a stockpile delivery of 479'700 vaccine doses.

Proceedings of the National Academy of Sciences. 2021;118(4). doi:10.1073/pnas.2019617118.

8. Gatto M, Bertuzzo E, Mari L, Miccoli S, Carraro L, Casagrandi R, et al. Spread and dynamics of the COVID-19 epidemic in Italy: Effects of emergency containment measures. Proceedings of the National Academy of Sciences. 2020;doi:10.1073/pnas.2004978117.

9. Bertuzzo E, Mari L, Pasetto D, Miccoli S, Casagrandi R, Gatto M, et al. The geography of COVID-19 spread in Italy and implications for the relaxation of confinement measures. Nature Communications. 2020;11(1):4264. doi:10.1038/s41467-020-18050-2.

10. Emanuel EJ, Persad G, Kern A, Buchanan A, Fabre C, Halliday D, et al. An ethical framework for global vaccine allocation. Science. 2020;369(6509):1309–1312. doi:10.1126/science.abe2803.

11. Lipsitch M, Dean NE. Understanding COVID-19 vaccine efficacy. Science. 2020;370(6518):763–765. doi:10.1126/science.abe5938.

Fig S7. Control and co-variates for the pessimistic scenario with a stockpile delivery of 479'700 vaccine doses.

12. Bubar KM, Reinholt K, Kissler SM, Lipsitch M, Cobey S, Grad YH, et al. Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. Science. 2021;doi:10.1126/science.abe6959.

13. Fitzpatrick MC, Galvani AP. Optimizing age-specific vaccination. Science. 2021;doi:10.1126/science.abg2334.

14. Baden LR, El Sahly HM, Essink B, *et al* . Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. New England Journal of Medicine. 2020;384:403–416.

15. Yang J, Zheng W, Shi H, Yan X, Dong K, You Q, et al. Who should be prioritized for COVID-19 vaccination in China? A descriptive study. BMC Medicine. 2021;19(1):45. doi:10.1186/s12916-021-01923-8.

16. Lee EC, Chao DL, Lemaitre JC, Matrajt L, Pasetto D, Perez-Saez J, et al. Achieving coordinated national immunity and cholera elimination in Haiti through vaccination: a modelling study. The Lancet Global Health. 2020;8(8):e1081–e1089. doi:10.1016/S2214-109X(20)30310-7.

Fig S8. Example of the dynamics in all compartments for every node in the pessimistic scenario with a stockpile delivery of 479'700 doses. The lower right plot shows the control variable, the number of doses per day in each province.

Fig S9. Time allocation for the pessimistic scenario with a stockpile delivery of 479'700. We see for each week, how the 479'700 doses are spread accross the provinces, as percent. This view unravel the temporal pattern in the allocation.

17. Pasetto D, Finger F, Camacho A, Grandesso F, Cohuet S, Lemaitre JC, et al. Near real-time forecasting for cholera decision making in Haiti after Hurricane Matthew. PLOS Computational Biology. 2018;14(5):e1006127. doi:10.1371/journal.pcbi.1006127.

18. Morton R, Wickwire KH. On the optimal control of a deterministic epidemic. Advances in Applied Probability. 1974;6(4):622–635. doi:10.2307/1426183.

19. Sethi SP, Staats PW. Optimal Control of Some Simple Deterministic Epidemic Models. Journal of the Operational Research Society. 1978;29(2):129–136. doi:10.1057/jors.1978.27.

20. Greenhalgh D. Some results on optimal control applied to epidemics. Mathematical Biosciences. 1988;88(2):125–158. doi:10.1016/0025-5564(88)90040-5.

21. Quirynen R, Vukov M, Diehl M. Multiple shooting in a microsecond. In: Carraro T, Geiger M, Körkel S, Rannacher R, editors. Multiple shooting and time domain decomposition methods. Cham: Springer International Publishing; 2015. p. 183–201.

22. Sah P, Medlock J, Fitzpatrick MC, Singer BH, Galvani AP. Optimizing the impact of low-efficacy influenza vaccines. Proceedings of the National Academy of Sciences. 2018;115(20):5151–5156. doi:10.1073/pnas.1802479115.

23. Medlock J, Galvani AP. Optimizing Influenza Vaccine Distribution. Science. 2009;325(5948):1705–1708. doi:10.1126/science.1175570.

24. Patel R, Longini IM, Elizabeth Halloran M. Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. Journal of Theoretical Biology. 2005;234(2):201–212. doi:10.1016/j.jtbi.2004.11.032.

25. Zakary O, Rachik M, Elmouki I. On the analysis of a multi-regions discrete SIR epidemic model: an optimal control approach. International Journal of Dynamics and Control. 2017;5(3):917–930. doi:10.1007/s40435-016-0233-2.

26. Miller Neilan R, Lenhart S. Optimal vaccine distribution in a spatiotemporal epidemic model with an application to rabies and raccoons. Journal of Mathematical Analysis and Applications. 2011;378(2):603–619. doi:10.1016/j.jmaa.2010.12.035.

27. Thul L, Powell W. Stochastic Optimization for Vaccine and Testing Kit Allocation for the COVID-19 Pandemic. arXiv:210101204 [cs]. 2021;.

28. Moore S, Hill EM, Tildesley MJ, Dyson L, Keeling MJ. Vaccination and Non-Pharmaceutical Interventions: When can the UK relax about COVID-19? medRxiv. 2021; p. 2020.12.27.20248896. doi:10.1101/2020.12.27.20248896.

29. Saad-Roy CM, Morris SE, Metcalf CJE, Mina MJ, Baker RE, Farrar J, et al. Epidemiological and evolutionary considerations of SARS-CoV-2 vaccine dosing regimes. Science. 2021;doi:10.1126/science.abg8663.

30. Kadire SR, Wachter RM, Lurie N. Delayed Second Dose versus Standard Regimen for Covid-19 Vaccination. New England Journal of Medicine. 2021;0(0):e28. doi:10.1056/NEJMclde2101987.

31. Acemoglu D, Fallah A, Giometto A, Huttenlocher D, Ozdaglar A, Parise F, et al. Optimal adaptive testing for epidemic control: combining molecular and serology tests. arXiv:210100773 [physics, q-bio]. 2021;.

32. Chatzimanolakis M, Weber P, Arampatzis G, Wälchli D, Kičić I, Karnakov P, et al. Optimal allocation of limited test resources for the quantification of COVID-19 infections. Swiss Medical Weekly. 2020;150(5153). doi:10.4414/smw.2020.20445.

33. Chen J, Hoops S, Marathe A, Mortveit H, Lewis B, Venkatramanan S, et al. Prioritizing allocation of COVID-19 vaccines based on social contacts increases vaccination effectiveness. medRxiv. 2021; p. 2021.02.04.21251012. doi:10.1101/2021.02.04.21251012.

34. Li R, Pei S, Chen B, Song Y, Zhang T, W Y, et al. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV2). Science. 2020;368(6490):489–493. doi:10.1126/science.abb3221.

35. Bock HG, Plitt KJ. A multiple shooting algorithm for direct solution of optimal control problems. In: Proceedings 9th IFAC world congress Budapest. Pergamon Press; 1984. p. 242–247.

36. Savorgnan C, Romani C, Kozma A, Diehl M. Multiple shooting for distributed systems with applications in hydro electricity production. Journal of Process Control. 2011;21:738–745. doi:DOI: 10.1016/j.jprocont.2011.01.011.

37. Andersson JAE, Gillis J, Horn G, Rawlings JB, Diehl M. CasADi – A software framework for nonlinear optimization and optimal control. Mathematical Programming Computation. 2018;.

38. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical Programming. 2006;106(1):25–57. doi:10.1007/s10107-004-0559-y.

39. Du Z, Pandey A, Bai Y, Fitzpatrick MC, Chinazzi M, Piontti APy, et al. Comparative cost-effectiveness of SARS-CoV-2 testing strategies in the USA: a modelling study. The Lancet Public Health. 2021;6(3):e184–e191. doi:10.1016/S2468-2667(21)00002-5.

40. Betts JT. Practical Methods for Optimal Control and Estimation Using Nonlinear Programming. 2nd ed. SIAM; 2010.

41. Biegler LT. Nonlinear programming. MOS-SIAM series on optimization. SIAM; 2010.

42. Cramer EY, Ray EL, Lopez VK, Bracher J, Brennen A, Rivadeneira AJC, et al. Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the US. medRxiv. 2021; p. 2021.02.03.21250974. doi:10.1101/2021.02.03.21250974.

43. Rawlings JB, Mayne DQ, Diehl M. Model Predictive Control: Theory, Computation, and Design. 2nd ed. Nob Hill Publishing; 2017.

44. Andersson J, Åkesson J, Diehl M. CasADi: A Symbolic Package for Automatic Differentiation and Optimal Control. In: Recent Advances in Algorithmic Differentiation. Lecture Notes in Computational Science and Engineering. Springer, Berlin, Heidelberg; 2012. p. 297–307. Available from: https://link.springer.com/chapter/10.1007/978-3-642-30023-3_27.

45. HSL. A collection of Fortran codes for large scale scientific computation.;. Available from: http://www.hsl.rl.ac.uk/.

46. Manoli G, Rossi M, Pasetto D, Deiana R, Ferraris S, Cassiani G, et al. An iterative particle filter approach for coupled hydro-geophysical inversion of a controlled infiltration experiment. Journal of Computational Physics. 2015;283(C):37–51. doi:10.1016/j.jcp.2014.11.035.

47. Douc R, Cappe O. Comparison of resampling schemes for particle filtering. In: ISPA 2005. Proceedings of the 4th international symposium on image and signal processing and analysis, 2005.; 2005. p. 64–69.

48. Dipartimento della Protezione Civile. Emergenza Coronavirus: la risposta nazionale;. Available from: `http://www.protezionecivile.gov.it/attivita-rischi/rischio-sanitario/emergenze/coronavirus`.

49. Istituto Superiore di Sanità. Coronavirus: ultimi aggiornamenti; 2020. Available from: `https://www.epicentro.iss.it/coronavirus/aggiornamenti`.

50. Palmieri L, Andrianou X, Barbariol P, Bella A, Bellino S, Benelli E, et al.. Characteristics of COVID-19 patients dying in Italy Report based on available data on October 4th, 2020; 2020. Available from: `https://www.epicentro.iss.it/en/coronavirus/sars-cov-2-analysis-of-deaths`.

51. Pepe E, Bajardi P, Gauvin L, Privitera F, Lake B, Cattuto C, et al. COVID-19 outbreak response, a dataset to assess mobility changes in Italy following national lockdown. Scientific Data. 2020;7(1). doi:10.1038/s41597-020-00575-2.

52. Vollmer MAS, others. Report 20: Using mobility to estimate the transmission intensity of COVID-19 in Italy: A subnational analysis with future scenarios. Imperial College London; 2020. Available from: `https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/report-20-italy/`.