A Model-Based Approach to Assess Epidemic Risk

Hugo Dolan[1] · Riccardo Rastelli[1]

Abstract
We study how international flights can facilitate the spread of an epidemic to a worldwide scale. We combine an infrastructure network of flight connections with a population density dataset to derive the mobility network, and then we define an epidemic framework to model the spread of the disease. Our approach combines a compartmental SEIRS model with a graph diffusion model to capture the clusteredness of the distribution of the population. The resulting model is characterised by the dynamics of a metapopulation SEIRS, with amplification or reduction of the infection rate which is determined also by the mobility of individuals. We use simulations to characterise and study a variety of realistic scenarios that resemble the recent spread of COVID-19. Crucially, we define a formal framework that can be used to design epidemic mitigation strategies: we propose an optimisation approach based on genetic algorithms that can be used to identify an optimal airport closure strategy, and that can be employed to aid decision making for the mitigation of the epidemic, in a timely manner.

Keywords COVID-19 · SEIRS compartmental model · Genetic algorithm · Network analysis · Human mobility

1 Introduction

In recent years, the extensive development of the transportation infrastructure has radically changed how connected our world is. International flights allow individuals to travel around the globe in just a few hours or days. This has important negative implications on the spread of diseases, whereby epidemics can reach a worldwide scale before effective responses are set in place. The recent COVID-19 outbreak [1]

✉ Riccardo Rastelli
riccardo.rastelli@ucd.ie

Hugo Dolan
hugo.dolan@ucdconnect.ie

[1] University College Dublin, Dublin, Ireland

has clearly raised and emphasised this problem. As a response to the emergency, many countries have taken drastic measures to contain and slow down the spread of the virus by imposing lockdowns and airport closures [2]. While these measures have been successful in confining the epidemic, the immediate and chaotic response has blurred the actual role played by the topology of the infrastructure network on the spread of the virus.

The main goal of this paper is to create a model-based framework that can inform decision making regarding airport closures as a means to slowing down a potential epidemic without causing excessive economic damage. In particular, we introduce a new framework to study networks of international flights as potential vehicles for the spread of diseases.

In this paper, we first propose an in-depth analysis of the Open Flights network dataset [3], which describes a large number of flight connections between more than 3000 airports. We calculate a number of descriptive statistics from the data, in order to study the underlying topology of this infrastructure network and essentially to understand how individuals can move between distant locations. We use various centrality measures to identify key airports, and we test the resilience of the network when these key airports are removed.

Then, we fit a Stochastic BlockModel (SBM) to partition the airports into homogeneous groups. The SBM, originally introduced and studied by [4], is a fundamental model and tool for statistical network analysis, since it can highlight groups of nodes that exhibit similar connectivity patterns. Inference for the SBM can be performed using both classical and Bayesian approaches [5]. One fundamental aspect of this model is that it can be interpreted as a finite mixture model for networks [6], and thus, it borrows many concepts and tools from this related research area. In this context, a useful by-product of the SBM framework is that it allows us to compare the connectivity and clustering of the airports with their actual geographical location.

After the exploratory data analysis, we use the infrastructure network to create a model for the simulation of epidemics. An essential aspect of this task is the development of a statistical network model that can combine these flight routes data with the geographical distribution of the population. Our aim is to give a model-based quantification of the epidemic risk which is amplified by the travelling of individuals, and to possibly identify effective interventions that can mitigate this risk. [2] propose an approach similar to ours, in that they combine the infrastructure network with the gridded population data to study the effects of the airport closure interventions that were actioned at the beginning of this year. [2] use a tool called GLEAM [7, 8] which can combine data from different sources to predict the behaviour of the epidemic using an individual-based compartmental SEIRS model. While they focus on the effects following the actual airport closures, in this paper, we aim at defining a framework to take new decisions that can lead to optimal airport closures, or potential future airport re-openings.

Our approach relies on epidemic compartmental models [9] and in particular on a SEIRS model. This framework postulates that the population is divided in 4 ordered compartments (susceptible, exposed, infectious, recovered) and that different rates determine the flow of individuals from one compartment into the next and

eventually back into the first compartment. This family of models has been largely employed in various research fields to model the evolution of epidemics, and it has been also successfully used within the context of COVID-19 [2, 10–13].

One fundamental aspect of our epidemic model is that, similarly to [2, 14, 15], we consider a metapopulation where each subpopulation is centred at an airport location, and whereby the local epidemic is determined by a SEIRS model. We use a graph diffusion process to describe the flows between the various subpopulations, which in turn affect the local dynamics. Not only this allows us to observe the epidemic both locally (for each subpopulation) and globally, but it also allows us to appreciate the spatio-temporal progression of the virus.

We calibrate the parameters of our model so that our predictions are similar to the recent spread of COVID-19. Our method is flexible and it can be used for a diverse range of epidemic parameters, hence encompassing other relevant diseases, or COVID-19 variants of interest. We do not claim that our results are specific to the COVID-19 epidemic nor that they should be used within this context; rather, we provide a general methodology that could be employed in any epidemic setting, and, for example purposes, we recreate realistic situations that resemble the recent COVID-19 epidemic. In fact, we test the sensitivity of our model by running a number of simulations that encompass a variety of possible real epidemic scenarios.

The fundamental contribution of this work regards the study of optimal epidemic mitigation strategies. Once we possess a model which is calibrated to a realistic setting, we explore an optimisation approach to identify what could be the optimal airport closure strategies that should be implemented. We use the predictions from our epidemic model to construct an objective function that takes into account measures for the spread of the disease as well as economic losses. We perform the optimisation using Genetic Algorithms (GAs) [16]. GAs are heuristic stochastic optimisation algorithms that explore new candidate solutions by selecting and transforming a set of current solutions using some basic principles of evolution and natural selection. In our context, GAs are especially convenient since the problem that we address is a combinatorial one, where we want to find the optimal subset of airports that should be closed to minimise our objective function.

The software that we have used in this paper is maintained by the authors and is available from the GitHub repository: [17].

2 Network Topology

In this section, we propose an exploratory data analysis and basic statistical modelling of the Open Flights dataset [3]. The Open Flights dataset contains information on 3425 airports globally, including a database of 37,594 commercial routes between these airports collected in 2014. The dataset is transformed into an adjacency matrix with nodes representing airports and directed edges representing whether there exists a direct route between any two airports. In Fig. 1, we present the network visually, and, on initial inspection, it is clear that the network exhibits extremely high degree of connectivity, with the plot of degree distribution indicating that over 20% exhibit a degree greater than 10.
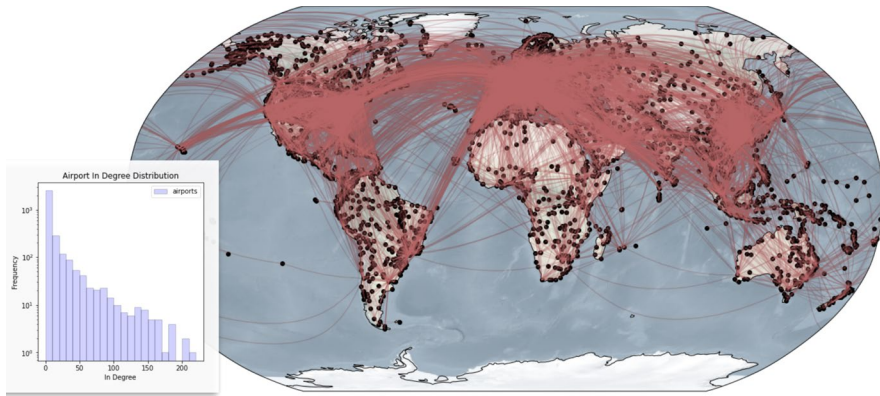
Fig. 1 Open Flights Network Visualisation and its in-degree distribution

Identifying airports which are important to the overall connectivity of the network is crucial in gaining a better understanding of the network's topology. We consider several metrics to measure the importance of nodes. These include the Page-Rank centrality, betweenness centrality, coreness ranks, as well as the in-degrees and out-degrees of nodes [18]. We present a table of the 20 most important airports according to Page-Rank in Table 1. Airports with high Page-Rank are clearly major international destinations, and they form an extremely well-connected subnetwork with a coreness of over 60, meaning that every airport in this subnetwork has 60 or more connections, or, equivalently, since the graph is approximately symmetric (the majority of air routes run return flights), we can say that every airport in this subnetwork has an out-degree close to, or exceeding, 30. It is interesting to note that airports with high betweenness centrality (Charles De Gaulle, Dubai, Beijing, Amsterdam, Los Angeles, Toronto, Frankfurt) are also major connecting flight hubs, exhibiting out-degrees of over 200.

We identify homogeneous subgroups of airports within the network by employing a SBM framework: we use a Python implementation of an efficient Markov chain Monte Carlo method, which is suitable for inferring SBMs in large networks, as described by [19]. The optimal SBM partition and the corresponding block matrix are shown in Fig. 2a and b, respectively. It is clear that the communities found are very strongly associated with the geographical location of the airports and with their region or province. This is quite surprising as this information is not encoded explicitly in the data provided to the algorithm. This would strongly suggest a high degree of connectivity of airports not only globally but also within regions or geographical areas. We also note from the block matrix in Fig. 2b that the majority of connections are not only within relatively large communities representing the geographic clustering observed in Fig. 2a, but also towards the lower right corner there is significant disassortative behaviour, likely these nodes are large international hubs such as the small community of London, Frankfurt, Amsterdam, Charles De Gaulle which share connections to many cities across the world.

Table 1 Summary Statistics for the top 20 airports, sorted according to Page-Rank centrality values

| Airport	Page-Rank	Betw.	Coreness	**deg**(+)	**deg**(−)	Country
ATL	0.0047	0.0294	60	216	**217**	United States
IST	0.0044	**0.0412**	62	230	**227**	Turkey
ORD	0.0043	**0.0474**	60	203	**206**	United States
DEN	0.0043	0.0262	60	168	169	United States
DFW	0.0042	0.0251	60	185	187	United States
DME	0.0041	0.0294	62	189	189	Russia
CDG	0.0039	**0.0617**	62	233	**237**	France
FRA	0.0038	**0.051**	62	238	**239**	Germany
PEK	0.0038	**0.0492**	60	206	**206**	China
DXB	0.0036	**0.0594**	62	182	188	United Arab Emirates
AMS	0.0036	**0.0427**	62	231	**232**	Netherlands
IAH	0.0035	0.023	60	168	169	United States
LAX	0.0033	**0.0662**	60	148	149	United States
SYD	0.0032	0.0326	45	83	85	Australia
YYZ	0.003	**0.0425**	60	146	147	Canada
JFK	0.003	0.0258	62	160	162	United States
BOG	0.0029	0.0248	40	74	74	Colombia
PVG	0.0029	0.0221	56	153	152	China

Additional node statistics and centralities are shown in the other columns, including in-degree (**deg**(+)) and out-degree (**deg**(−))

We continue our exploratory analysis by studying the percolation properties of the network [18]. We percolate the network by sequentially removing the nodes (one at a time) and their connections, and observing how the connectivity of the graph changes. We remove the nodes following both a random order and following a decreasing order of the out-degree and other centrality measures. The results are shown in Fig. 3. The network is highly resilient to random attacks, since the removal of almost all nodes is required in order to disrupt network connectivity. However, we note that the network is moderately more vulnerable to targeted (degree-based) attacks, yet this would still require more than half of all airports to be removed for the single giant component to disappear. Similarly, the network is moderately vulnerable to targeted attacks according to other ranking factors (Page-Rank, betweenness, coreness). We note that these procedures are also averaged over many trials to account for the removal of vertices of equal rankings in different orders, however, we find these results very quite similar to percolation by degree.

In conclusion, the Open Flights network summary statistics show that airports which are large regional destinations, or hubs for connecting flights, tend to have high importance to network connectivity. Furthermore, it is observed that some nodes in the network are extremely well connected, both at regional and global level, with significant geographical community structure. The network is also highly resilient to random or deterministic attacks.
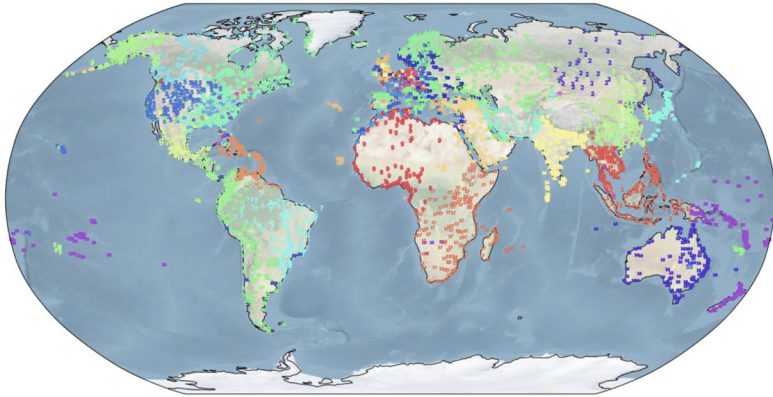
(a) Map of the SBM results, with colours indicating the cluster memberships.
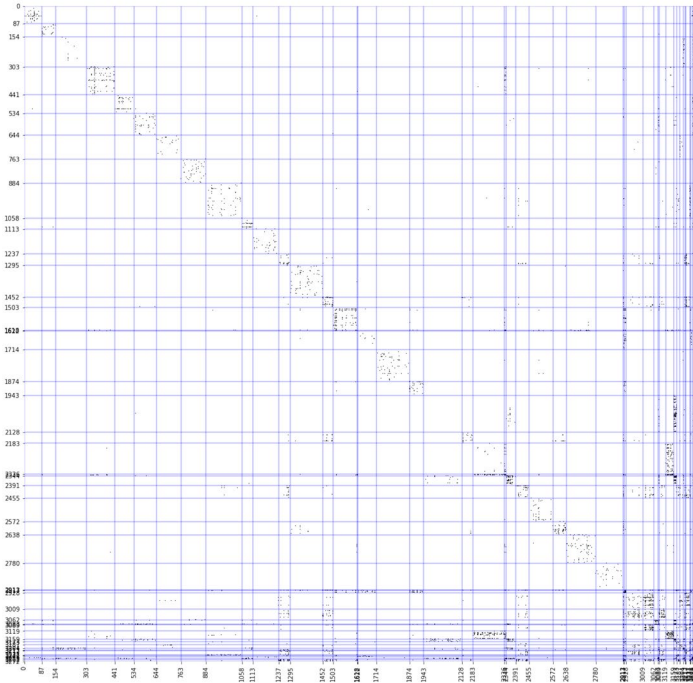
(b) Block matrix visualisation of the adjacency matrix, where the black dots indicate routes between airports and the blue boxes indicate the clusters.

Fig. 2 Fitted SBM results

In the context of epidemics, these initial findings provide a very solid evidence that the flight connections can sadly be a very efficient vehicle to facilitate the spread of diseases, and, more importantly that substantial network "damage" (e.g. airport closures) is required to ensure that an epidemic does not spread to a world-wide scale. This evidence motivates our work, in that we aim at finding optimal
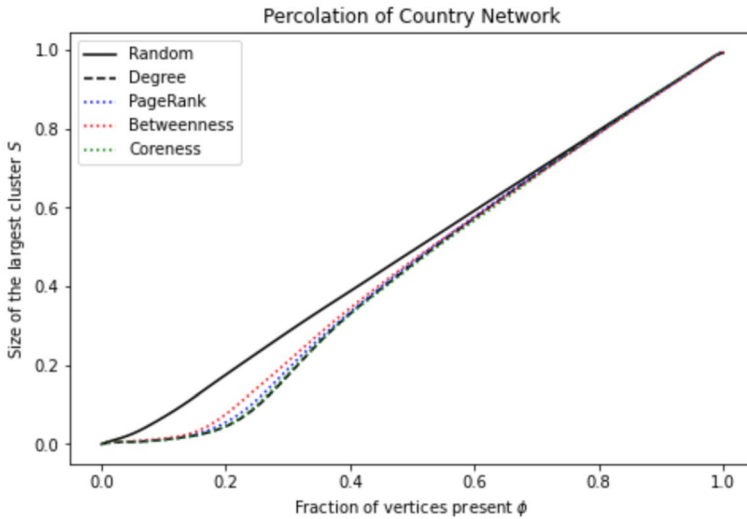
Fig. 3 Percolation on the Open Flights dataset via a variety of ranking criteria (results are averaged over 100 independent simulations)

mitigation strategies that can reduce the pace of the epidemic to a much smaller scale, without causing excessive disruption to economies and to this particular infrastructure network.

3 Model Specification

3.1 Theoretical Underpinnings

Before we develop the main model of this paper, we must first introduce two existing models which can be found in the domains of applied mathematics and epidemiology [18]. First, we specify the graph diffusion model, which describes the flow of a fluid across a network, driven by pressure differences between adjacent nodes. This can be expressed as a vector of differential equations denoting changes of fluid volumes at each node and time step:

$$\frac{d\boldsymbol{\psi}}{dt} = c(\mathbf{A} - \mathbf{D})\boldsymbol{\psi}.$$

We use the notation $\boldsymbol{\psi}$ to represent the vector of fluid volumes at every node, $\boldsymbol{A}$ to denote the adjacency matrix of the network, $\boldsymbol{D}$ to denote a diagonal matrix of containing the degrees of every node, and $\boldsymbol{c}$ is called the diffusion constant. A full derivation of this can be found in [18].

Additionally, we introduce the SEIRS compartmental epidemiology model. Each letter of the model name denotes a compartment of the system (Susceptible ($\boldsymbol{S}$), Exposed ($\boldsymbol{E}$), Infectious ($\boldsymbol{I}$) and Recovered ($\boldsymbol{R}$)), in which some number of individuals from the total population ($\boldsymbol{M}$) reside. Figure 4 illustrates the direction of

Fig. 4 SEIRS compartmental flow diagram
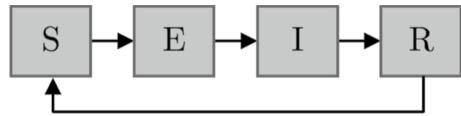
progression from state to state, while (1) indicates the exact rates at which the population in each compartment changes.

$$\frac{dS}{dt} = \delta R - \frac{S\beta I}{M}$$
$$\frac{dE}{dt} = \frac{S\beta I}{M} - \epsilon E$$
$$\frac{dI}{dt} = \epsilon E - \gamma I$$
$$\frac{dR}{dt} = \gamma I - \delta R$$

(1)

The epidemic parameters β, ϵ, γ and δ are non-negative scalars and can be estimated or calibrated to match the characteristics of some observed epidemic. Once the system's initial condition and the epidemic parameters are specified, we can find a numerical approximation to the solution of this system, hence, obtaining the number of individuals in each compartment, at each time $t > 0$.

3.2 Model Definition

In order to model the transmission of disease through the international flights network, we opt to use the SEIRS model combined with a graph diffusion model as described in the previous section. We will refer to the airports' adjacency matrix as A and denote airport nodes as v_j, for $j = 1, \ldots, N$. The total number of nodes in the network is N and the associated population at each node is M_j. Let us define the local epidemic state vector as $\theta_j(t) = (S_j \; E_j \; I_j \; R_j)^\top$,, which represent the compartments of the SEIRS model for airport population at any given time. A condition of the SEIRS model constrains the total population of all compartments to equal the total population, at each time t:

$$S_j(t) + E_j(t) + I_j(t) + R_j(t) = M_j(t),$$

We assume that the local population is fully mixed (i.e. everyone has equal chance of being infected), as this is a standard assumption of compartmental epidemic models; however, we assume this only to be true at the individual airport level and not for the entire global system. Additionally, we introduce α_j as the proportion of the population which can travel, and c as the probability that an individual departs from an airport on any given day. The proportions allow us to define an additional variable $\psi_{j(t)} = \alpha_j \theta_j(t)$ which corresponds to the mobile epidemic state.

We can now define, at a high level, our simulation procedure, which we use to generate the SEIRS data conditionally on a specific set of epidemic parameters:

1. At each time step, we must first initiate community spread of the disease through each airport's local population. This is represented by the compartments of its SEIRS model denoted by the local epidemic state vector θ_j. Using a fourth-order Runge-Kutta approximation technique, we update each compartment's value in accordance with the dynamics of the SEIRS model. The new state of the compartments after community spread is denoted θ_j^* as opposed to the previous state θ_j, for all js.

2. Now take this new local epidemic state θ^* and split it into two types of population members: a base population denoted θ_B, who are permanent residents to the local area (i.e. the **locals**); and a second group θ_T, identifying the transient population who are temporary **visitors**, who have arrived at the airport on business or holidays and will return home after a short period. This differentiation is necessary to ensure that the local populations remain stable over time (we assume no permanent migration in our model).

3. Next, we compute the proportions of θ_B and θ_T who can fly (in each compartment), these are $\psi_+ = \alpha_+ \theta_B$ and $\psi_- = \alpha_- \theta_T$, describing the number of outbound passengers and returning travellers, respectively.

4. Use a diffusion model to compute the changes of θ_B and θ_T at each airport. The diffusion model is tuned to take into account several factors including the differences between outbound and returning passengers at every connected airport, the relative importance of the airports in the network, and the airport's number of connections. These guarantee that the diffusion model can affect the SEIRS statuses without changing the distribution of the population across locations.

5. Recombine the updated values of θ_B and θ_T into the aggregate populations θ and repeat the whole procedure for the number of iterations ("days") required.

For completeness, we include both a diagrammatic form of the algorithm (Fig. 5) as well as the full algorithm (Table 2) which reflects the high-level overview above. A full derivation is provided in Appendix A.1.

3.3 Modelling assumptions

In this section, we provide further information on the parameters of our model, on their interpretation, and on how we have used different data sources to identify a range of plausible values for these parameters. We specifically focus on the dataset that have been used and where they come from, and how we have transformed these datasets to obtain the key elements for our study.

Our framework requires information on airports, routes, demographics, wealth, which is difficult to obtain and use. A summary of the model's parameters is provided in Table 3, for convenience.

The international flights network introduces and analysed in Sect. 2, is the foundation for the diffusion part of our model. From this network of route connections, we aim at constructing a migration matrix connecting more than **3,000** locations. One fundamental transformation that we apply is the following: let P be the vector of Page-Rank values obtained from earlier analysis and A be the adjacency matrix of
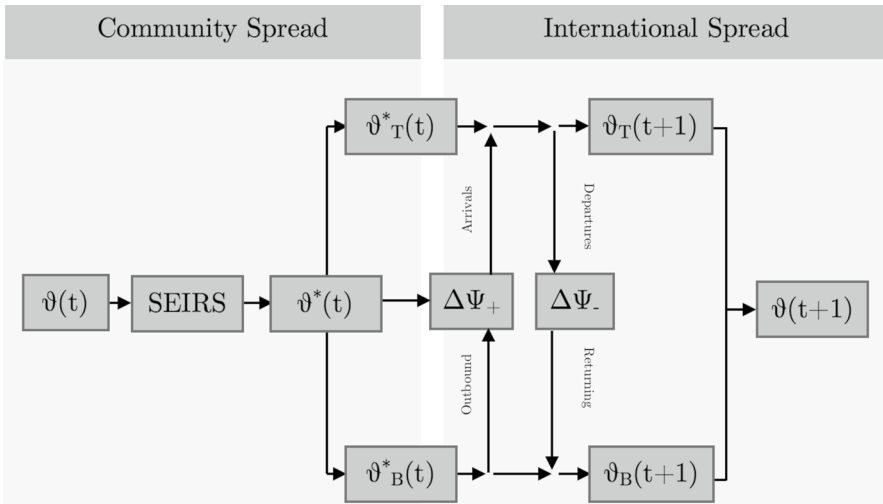
Fig. 5 Diagram of the epidemic diffusion model steps

Table 2 Simulation pseudocode, where with Σ we indicate a matrix representation of the SEIRS process

Algorithm pseudocode for T days

for $t = 1, \dots, T$:

$\theta^*(t-1) = \theta(t-1) + \Sigma\, \theta(t-1)$	Community spread
$\theta^*(t-1) = \theta_B^*(t-1) + \theta_T^*(t-1)$	Preparation for diffusion
$\psi_+^*(t-1) = \theta_B^*(t-1)\, \alpha_+$	
$\psi_-^*(t-1) = \theta_T^*(t-1)\, \alpha_-$	
$\Delta\psi_+(t) = -c_+ \psi_+^*(t-1)\left(I - C^{\mathsf{T}}\right)$	International spread
$\Delta\psi_-(t) = -c_- \psi_-^*(t-1)\left(I - C^{\mathsf{T}}\right)$	
$\theta_B(t) = \theta_B^*(t-1) + Min\left(\Delta\psi_+(t), 0\right) + Max\left(\Delta\psi_-(t), 0\right)$	
$\theta_T(t) = \theta_T^*(t-1) + Max\left(\Delta\psi_+(t), 0\right) + Min\left(\Delta\psi_-(t), 0\right)$	

flight connections. Then, we define and work with the matrix C, which is the relative centrality matrix with elements defined by $C_{ij} = \frac{A_{ij}P_j}{\sum_k A_{ik}P_k}$. The interpretation of this new matrix is that the positions of its non-null entries are identical to those of A, except now the edges have been assigned weights based on the relative importance of adjacent nodes. This transformation ensures that the flows of individuals between airports are scaled so that they reflect higher traffic to major airports and less traffic to smaller airports, thus, preventing the system from diverging towards unrealistic configurations as the simulation progresses.

With regard to the distribution of the population into a metapopulation structure, we must define the number of individuals that have access to and are served by, each of the airports. This corresponds to estimating the initial state of the system, θ, in that we are estimating the counts M for each airport location. For this task, we

Table 3 Model parameters, with [*] indicating the necessary conditions for the escalation of the epidemic

	Parameter	Description	Dimension	Range
Data based	$\theta(\mathbf{0})$	Initial Populations (S, E, I, R) for each airport	$\mathbb{R}^{4 \times N}$	(N/A) Gridded Population of the World (GPW), SEDAC
	C	Relative Centrality	$\mathbb{R}^{N \times N}$	(N/A) Open Flights Data
	α_+	% Base Population who can fly	$\mathbb{R}^{N \times N (Diag.)}$	(N/A) World Bank Data
Model assumption	α_-	% Transient Population who can fly	$\mathbb{R}^{N \times N (Diag.)}$	$\alpha_- = I$, assuming no permanent migration
	c_-	% of people who return at the end of the week	$\mathbb{R}$	$c_- = 1$, assuming no permanent migration
Heuristic based	c_+	% of people who depart on any given day	$\mathbb{R}$	$c_+ \in (0, 1)$
	β	Rate of infection / contact (Daily)	$\mathbb{R}$	$\beta \in (0, \infty); \beta > \epsilon > \gamma$ [*]
	ϵ	Rate of transfer from exposed to infectious state (daily)	$\mathbb{R}$	$\epsilon \in (0, \infty)$
	γ	Rate of transfer from infectious to recovered state (daily)	$\mathbb{R}$	$\gamma \in (0, \infty)$
	δ	Rate of loss of immunity (daily)	$\mathbb{R}$	$\delta \in (0, \infty)$

use the gridded population of the world dataset [20], and we assign a value for total population at each airport location. In order to perform this assignment, there is an immediate problem because many airports are often in close proximity of each other (for example, some cities are served by multiple airports). We assume that the maximal distance that anyone will travel to reach an airport is **240**km (**60 km/h * 4 hours = 240**km). Using this assumption, for each cell of the grid of the population dataset, we search all the airports that are within this radius. Then, we assign a population contribution to each, proportionally to their Page-Rank values. This metric is chosen because it provides a reasonable indication of which airport the travellers will tend to use. Note that this approach will automatically exclude population grid cells which are not within **240**km of any airport. These populations are excluded from the simulation as they are unlikely to be flying regularly, yet we note that this corresponds to less than **3%** of the total population.

Finally, we estimate the percentage of each population who can fly α_+. It is obvious that this percentage will vary across countries, depending on a number of factors, primarily wealth. To find a reasonable value for this model parameter, we acquire passenger estimates by country as supplied by the World Bank and divide these through by the total airport populations for the given country. We then assume the proportion by airport is the same as at country level[1].

One further fundamental modification that we make regarding α_+, is that the infectious individuals in the *I* compartment are not allowed to travel. This assumption is motivated by the fact that individuals suffering from the disease, especially if symptomatic, may be incapacitated to fly, or would likely be identified as infectious by mandatory testing procedures, hence their contributions to the epidemic would be under some reasonable control. Still, this assumption does not stop the disease from spreading between locations, since the exposed individuals in each of the *E* compartments would still be carrying the disease.

The SEIRS model, as well as the graph diffusion model, relies on several strong assumptions that inevitably impact the results. For clarity, we list and summarise below these assumptions, to highlight the specific features that our approach will exhibit and that should be kept in mind.

1. Fully mixed local populations: within any given node, every member of the population has equal chance of contact, and, thus, equal chance of passing on the disease.
2. Maximal travel distance: we assume that the maximum distance someone will travel is **240**km to get to an airport, and thus, anyone who is based outside of all airport radiuses is assumed to be isolated and excluded from the model.
3. Air transit only: we assume that the only way for the disease to spread between nodes is via air routes and that spread via other means are negligible.

---

[1] Sometimes this proportion is greater than **1** particularly for major hubs, where annual traffic through the airport may exceed the size of the local population assigned. In these cases, we replace our estimate by the average global proportion, clearly this is imperfect but is likely as accurate as we will be able to obtain given the fact that a more correct information is not readily available.

4. Fully mixed wealth: the proportion of population which may fly is the same for all airports in the same country.
5. Heavy symptoms: the infectious individuals do not travel.
6. **No permanent immigration:** outbound and inbound air traffic volume are balanced for each location.
7. Universal rates: we assume that the parameters of the SEIRS model are universal and do not vary between countries.

While it is obvious that each of these assumptions can have a non-negligible effect, we do argue that our simulations can show a remarkable resemblance with a realistic epidemic, such as the recent COVID-19 one. However, as we will show in the next sections, we consider a variety of different epidemic framework to ensure that our results can be generalised and to provide a tool that can be used in an arbitrary epidemic setting.

4 Simulations

4.1 Visualisation of Unmitigated Spread

Now that we have outlined the theory and processes to develop our model, we proceed to simulate and visualise an unmitigated epidemic (i.e., no measures to decrease the infection rate). One crucial decision that we need to make regards the calibration of the epidemic parameters. The previous study by [10] used a SEIRS model and estimated $\epsilon = 0.14$, $\gamma = 0.048$ and $\beta = 0.4$, with $\delta = 1/730$ set conservatively as there is uncertainty as to how long recovered individuals will remain immune to COVID-19. Since we aim at recreating a realistic epidemic resembling the recent COVID-19 spread, we consider parameter ranges around the above values and run a sensitivity analysis to explore the various different outcomes.

The SEIRS model parameters that we consider are as follows:

· $\epsilon = 1/7$: individuals remain in E for an average of **7** days;
· $\delta = 1/730$: recovered individuals remain immune for an average of two years;
· $\gamma \in \{1/5, 1/25, 1/50\}$: three possible scenarios where individuals remain infectious for an average of **5**, **25** or **50** days on average;
· $\beta(t) = 3e^{-\eta t/365}$, for $t = 0, 1, 2, \dots$: each infectious individual has an average of **3** contacts; however, this number is scaled down with an exponential decay as time progresses. This is a reflection of the fact that the response to an epidemic will play a crucial role in slowing down or stopping the disease, and that such response will improve over time, due to a variety of factors including social distancing, testing, better understanding of the disease and treatments.
· $\eta \in \{0, 2.5, 5, 7.5, 10\}$: the exponential decay can have different slopes, the value **0** corresponds to no slope and, thus, a constant β, whereas **10** gives a dramatic decrease, and thus, it corresponds to a very effective response.
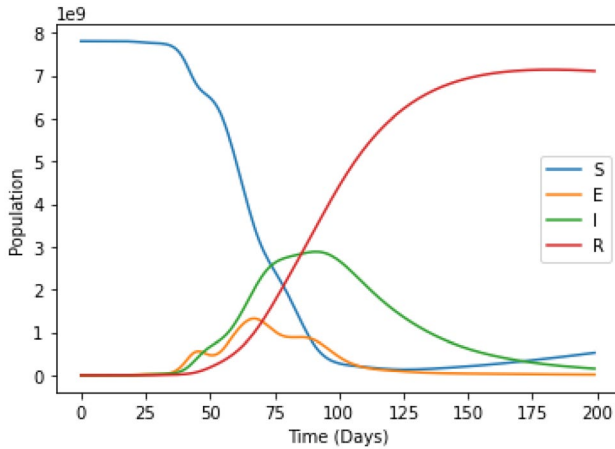
Fig. 6 Aggregated SEIRS compartments' sizes, for the unmitigated scenario under the default configuration

As a **default configuration**, we set $\eta = 5$, $\epsilon = 1/7$, $\gamma = 1/25$ and $\delta = \frac{1}{730}$.

In terms of initial conditions, we setup the model such that the first cases occur in the Wuhan Airport metapopulation, for similarity with the COVID-19 outbreak [1]. In the initial state of the system, all of the metapopulations are in the susceptible compartment, with the exception of Wuhan Airport where we have **0.001**% of the local population in the infectious compartment.

As an initial example, we run the model with the default configuration and show a qualitative analysis of the results. This simulation corresponds to an unmitigated scenario, because all the airports remain fully operational throughout the simulation, and the only response to the epidemic is the one implied by the parameter $\eta = 5$, i.e. a fairly slow exponential decay of the average number of contacts.

The simulation's results, shown in Fig. 6, are clearly tragic, in that the epidemic easily escalates into a pandemic and affects almost all individuals. While unrealistic, this worst case scenario example illustrates how airport closure interventions can potentially play a crucial role in slowing the epidemic, but it also allows us to visualise the spatio-temporal spread of the disease. Figure 7 is an illustrative sketch of the spread through the network which is derived from our simulations. There is evidence of a gradual dispersion of the virus across the world starting in China, moving onward into other areas including South East Asia, Japan, Russia, India, South Africa and Middle East. Some of the last places to be infected are the Americas, Nordic Countries, Alaska and Turkey.

4.2 Sensitivity Analysis of Key Parameters

In this section, in order to obtain a better understanding of the impact of these SEIRS parameters on our results, we conduct a sensitivity analysis. As per the epidemic parameters considered in Sect. 4.1, here we run the simulations for the
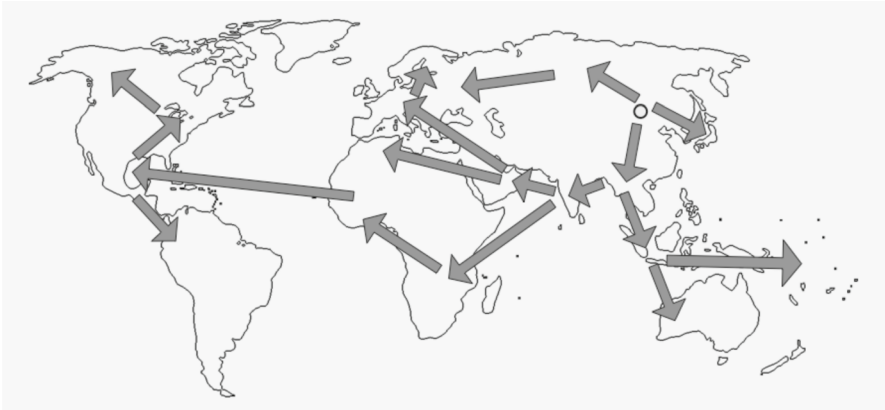
Fig. 7 Illustrative sketch of the epidemic based on the simulations of Fig. 6

epidemic parameters varying in their respective ranges, for the unmitigated scenario. This means that the outcome of the epidemic is solely determined by the epidemic parameters, because there are no airport closures. When extracting the results, we focus on a few key airport locations, and we also consider a benchmark scenario where there is no network structure (i.e. all population is in the same location).

In Table 4, we report the max and the argmax of the I compartment, for each of the locations, and for the various combinations of η and γ. The max of the I compartment is a crucial measure of the impact of an epidemic, since it is directly related to the stress caused to the local health institutions. As regards the argmax of the I compartment, this can be used as a measure of the speed of the disease: if the disease progresses slowly, the epidemic curve is flat and it is easier to keep the situation under control until better solutions or treatments are found.

We can see in the results in Table 4 that, as η increases, the max of I tends to decrease and the argmax of I tends to increase. This means that actioning an effective response makes the epidemic curve flatter, in that its peak is reduced and delayed.

The different values of γ correspond instead to a different duration of the disease in individuals. In most cases, we see that a smaller value of γ, which signals a longer duration of the disease, causes higher peaks of infections indicating a more difficult situation. The relation of γ with the timing of the peak seems less clear, possibly because if the peak decreases then it might also be reached sooner, but this would not indicate a more difficult situation, in general.

The table also shows that the two parameters jointly affect the results: the impact of η on the results is stronger when γ is large, and, vice versa, η has less of an effect when γ is small. If the duration of the disease is short (large γ), then we see a relatively low number of cases, which (for most locations) drops down to **0.0** when β decays at a fast rate. The situation is different when the duration of the disease is long (small γ): in this case, we see more infections overall, and the number of infections does not decrease for faster decays of β. In fact, for the major airports that we consider, we observe a counter intuitive scenario where the peak of infections

Table 4 Sensitivity analysis. The tables on the left column show the time until the peak of *I* is attained, at that particular location. The tables on the right column show the max value of *I* (in millions) at that particular location. The simulations were run for **200** days. Local populations (in millions) served by these airports are **8.34** for London Heathrow Airport, **3.97** for Dubai International Airport, **28.39** for Hong Kong International, **8.08** for JFK International

London Heathrow Airport:

Days until peak of infections

γ	η				
	0	2.5	5	7.5	10
0.2	57	65	78	105	105
0.04	60	62	70	80	101
0.02	65	65	75	85	100

London Heathrow Airport:

Peak of infections (millions)

γ	η				
	0	2.5	5	7.5	10
0.2	3.1	3.0	2.6	1.2	0.0
0.04	6.8	6.9	7.1	7.2	6.3
0.02	8.4	8.6	8.9	9.1	9.1

Dubai International Airport:

Days until peak of infections

γ	η				
	0	2.5	5	7.5	10
0.2	57	65	79	106	105
0.04	60	65	70	81	105
0.02	65	70	80	85	105

Dubai International Airport:

Peak of infections (millions)

γ	η				
	0	2.5	5	7.5	10
0.2	3.9	4.0	3.7	1.8	0.0
0.04	8.7	9.1	9.7	10.2	9.6
0.02	10.9	11.6	12.3	13.3	14.0

Hong Kong International:

Days until peak of infections

γ	η				
	0	2.5	5	7.5	10
0.2	46	51	58	70	87
0.04	47	50	54	60	68
0.02	51	55	57	62	70

Hong Kong International:

Peak of infections (millions)

γ	η				
	0	2.5	5	7.5	10
0.2	9.4	9.0	8.2	6.3	2.1
0.04	20.4	20.5	20.6	20.6	19.9
0.02	24.8	25.0	25.2	25.4	25.5

JFK International:

Days until peak of infections

γ	η				
	0	2.5	5	7.5	10
0.2	59	67	81	109	105
0.04	60	63	70	82	105
0.02	61	66	75	85	105

JFK International:

Peak of infections (millions)

γ	η				
	0	2.5	5	7.5	10
0.2	2.8	2.7	2.2	0.8	0.0
0.04	6.1	6.2	6.2	6.1	5.0
0.02	7.4	7.6	7.7	7.8	7.4

Benchmark (no network):

Days until peak of infections

γ	η				
	0	2.5	5	7.5	10
0.2	30	32	34	37	41
0.04	33	34	35	37	39

Benchmark (no network):

Peak of infections (millions)

γ	η				
	0	2.5	5	7.5	10
0.2	0.7	0.6	0.6	0.6	0.5
0.04	1.4	1.4	1.4	1.4	1.4

Table 4 (continued)

Benchmark (no network):					Benchmark (no network):				
Days until peak of infections					Peak of infections (millions)				
η					η				
0.02	36	37	38	40	41	0.02	1.7	1.7	1.7	1.7	1.7

increases when β decays at a faster rate. This is motivated by travelling: the airports that we consider are hubs in the network, meaning that they have a substantial and constant in-flow of exposed individuals, which then switch to infected while staying in the major city. Clearly this situation does not occur in the no-travelling benchmark framework; however, even in this case, we can notice that the peak of infection is essentially unaffected by the decay rate of β.

5 Mitigation Strategies

In this section, we examine potential mitigation strategies in the context of epidemics on international flight networks. We introduce several mitigation strategies and test them on our SEIRS model. The strategies that we consider are the following:

· Nth day rule: close-all airports after a prefixed number of days.
· Threshold rule: close an airport if the number of infected reaches a prefixed threshold.
· Limited Nth day rule: sort the airports based on some nodal attribute, and then close a number of them after a prefixed number of days.
· Optimisation approach: use a genetic algorithm to determine which airport is best to close after a prefixed number of days.

In order to evaluate these strategies, we must first select some performance metrics. We decide to select metrics which are easy to interpret by policy makers and by the general public, while also being useful in the context of managing hospital intensive care unit capacity and overall impact of the epidemic. Specifically, we will measure the peak number of infections and the total cases (measured as the peak of recovered individuals), as these are transparent and can easily be measured from our simulations. We will report these values in relation to a baseline, which is provided by the unmitigated scenario of Sect. 4.1, whereby no airports are closed.

In all simulations, we use the default configuration of the SEIRS parameters, but for some of the models, we also vary η to provide more complete results. We argue that the results can be generalised and that the same qualitative results are obtained with alternative realistic configurations of the SEIRS parameters.
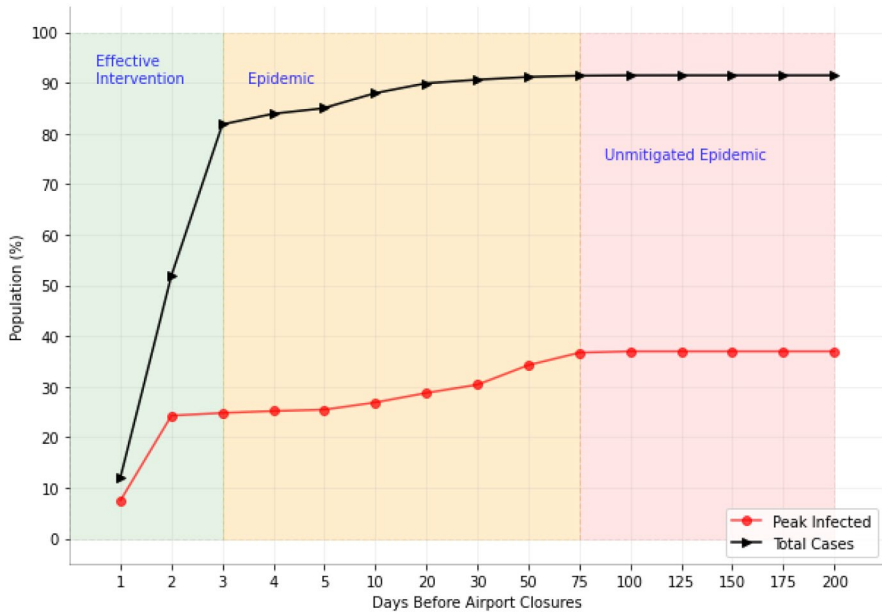
Fig. 8 Impact of worldwide permanent airport closures from nth day since first infection

5.1 Nth Day Rule

The first mitigation strategy that we consider is defined by the permanent closure of air routes from the nth day after the initial outbreak. Our simulations in Fig. 8 demonstrate that closing routes early reduce the peak number of infections but does not significantly reduce the total number of cases, unless all airports are closed by the end of day **2**. While a substantial mitigation can be achieved through airport closures, the time to action these changes is extremely limited, highlighting the importance of timely interventions. In Table 5, we further study some of the aspects of the mitigation strategy from a global perspective, in conjunction with different decay values of β. We note that closing airports immediately after day **1** would have a significant impact, reducing the total number of cases by almost **80%** for the default configuration. This demonstrates the high level of connectivity within the network, with cases of infections proliferating across every continent within just the first **3** days of the outbreak.

One fundamental aspect of the epidemic, which is especially apparent in the unmitigated scenario of Table 4, is that the actual spread of the disease only becomes noticeable after several weeks, when the number of infections starts peaking in some of the locations. This reflects well how the recent COVID-19 was able to spread across the world unnoticed for months before it was identified and treated as a pandemic.

Our interpretation of these results is that the graph diffusion model spreads tiny portions of the epidemic in each of the seed's neighbours, thus, seeding also these

Table 5 Nth day rule. Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for various decay parameters (on the rows) and for permanent airport closures happening on different days since the start of the epidemic (over the columns)

Nth Day Rule : peak infections						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Day					
	1	2	3	4	5	200
	η					
0	84.1	38.3	19.1	17.5	16.5	48.3
5	79.8	34.4	33.0	32.0	31.2	37.0
10	60.2	58.8	57.9	57.2	56.7	15.1
Nth Day Rule : total cases						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Day					
	1	2	3	4	5	200
	η					
0	86.9	43.3	4.8	0.9	0.4	94.5
5	86.8	43.2	10.5	8.2	7.0	91.5
10	77.1	76.3	75.1	73.7	72.1	51.0

new locations. These "seeds" could be interpreted as local patient zeros. Then, if the SEIRS parameters guarantee a local escalation of the epidemic, we inevitably observe an increasing number of cases even a long time after the airport closures. Once the airports are closed, the escalation of local epidemics becomes independent of one another and solely dependent on the local epidemic parameters, i.e. on the decay η. This is in fact a very positive message, since in these cases, the local escalations can be mitigated or even reversed if the number of infectious individuals is small enough, and if β decays rapidly. We can see this clearly in 5, where faster decay rates lead to much greater mitigation, with smaller numbers of total infections and recoveries.

We also notice that the impact of closures is more relevant when there is no decay, whereas the effect is less noticeable for faster decays. This suggests that local interventions aimed at reducing the parameter β can play a fundamental role and potentially replace airport closures.

5.2 Threshold Rule

A further modification to the previous results involves dynamically closing airports whenever the total number of cases exceeds a given percentage of the local population. This differs from the previous method in which we implemented blanket global closures. The results of this new experiment, which are shown in Table 6, indicate

Table 6 Threshold rule

Threshold Rule : peak infections						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Thresh					
	1e-7	**1e-5**	**1e-3**	**1e-2**	**1e-1**	**no threshold**
	η					
0	2.2	1.5	1.0	0.6	0.3	48.3
5	15.5	10.3	5.5	3.6	1.9	37.0
10	56.8	52.2	32.8	19.8	10.1	15.1
Threshold Rule : total cases						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Thresh					
	1e-7	**1e-5**	**1e-3**	**1e-2**	**1e-1**	**no threshold**
	η					
0	0.0	0.0	0.0	0.0	0.0	94.5
5	0.9	0.5	0.2	0.1	0.0	91.5
10	65.8	48.4	26.8	16.3	8.5	51.0

Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for various decay parameters (on the rows) and for airport closures happening at different threshold levels (over the columns)

that this "wait and see" strategy is totally ineffective under our modelling framework. Even considering a highly unrealistic version of this strategy in where we suppose it is possible to detect cases up to a fineness of **1** in every **10** million people, it is already too late to close airports, whereby mitigation is mostly irrelevant unless the decay is really fast.

5.3 Limited Nth Day Rule

In the previous results, we have shown that it is far more effective to close airports preemptively than it is to wait on some threshold level on infections to be attained within the local population. However, one could argue that it is impractical to close-all airports globally, both from an economic and political point of view. This motivates the two following strategies that we propose, where we proceed to examine what performance we can achieve by only closing a subset of key airports, where we can carefully choose the particular subset using different strategies. Here, we rank the airports using several metrics: population, PageRank and betweenness; then, we propose to close a number of airports following the order given by one of these ranking. The resulting strategy is similar to the nth day rule; however, in this case, we selectively close fewer airports based on their metric score. The main argument and motivation behind this limited nth day rule is that closing all airports will inevitably carry a substantial damage, so

Table 7 Limited nth day rule by local population

Nth Day Rule by population : peak infections						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	30.1	9.4	8.4	7.8	7.5	37.0
5%	64.0	21.1	20.1	19.4	19.0	37.0
10%	79.7	29.6	28.2	27.2	26.5	37.0
25%	79.8	33.7	32.3	31.3	30.5	37.0
100%	79.8	34.4	33.0	32.0	31.2	37.0
Nth Day Rule by population : total cases						% affected population
% reduction wrt to worst scenario						in the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	24.2	9.5	2.9	2.1	1.7	91.5
5%	51.7	21.1	6.1	4.6	3.8	91.5
10%	72.8	28.1	8.2	6.3	5.4	91.5
25%	86.8	40.5	9.9	7.7	6.6	91.5
100%	86.8	43.2	10.5	8.2	7.0	91.5

we are interested in checking whether good results can also be achieved by closing fewer airports.

Crucially, we observe in Tables 7, 8 and 9 that even closing just the top **1%** of airports, we still obtain significantly greater reductions than the threshold rule. This further emphasises the point that early intervention is far more important in the network than attempting to detect when certain infection thresholds have been breached. We also note that the fall-off in the strategy's performance is quite nonlinear both across the rows or columns.

With regard to comparisons between the different metrics, it appears that ranking the airports by local population is most effective for first day closures; however, the two other centrality-based procedures tend to perform better when airports are closed after day **1**.

5.4 Optimisation Approach

In Sect. 4, we explored several mitigation strategies, ultimately finding greatest flexibility and mitigation in the limited nth day rule. This strongly suggests that we could find a very effective airport closure strategy that will achieve excellent mitigation results, without totally disrupting the network connectivity. In this section, we employ a genetic algorithm (GA) to search for the optimal combination

Table 8 Limited nth day rule by Page-Rank centrality. Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for various levels of airport closures happening on different days since the start of the epidemic (over the columns)

Nth Day Rule by Page − Rank : **peak infections** % Reduction wrt to worst scenario						% affected population in the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	19.3	11.0	10.0	9.4	8.9	37.0
5%	74.0	31.7	29.8	28.6	27.8	37.0
10%	79.5	34.1	32.5	31.4	30.5	37.0
25%	79.8	34.3	32.9	31.9	31.1	37.0
100%	79.8	34.4	33.0	32.0	31.2	37.0
Nth Day Rule by Page − Rank : **total cases** % Reduction wrt to worst scenario						% affected population in the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	5.1	1.5	0.8	0.7	0.6	91.5
5%	52.2	13.2	4.0	3.0	2.5	91.5
10%	81.5	31.4	8.3	6.2	5.1	91.5
25%	86.7	42.1	9.9	7.5	6.3	91.5
100%	86.8	43.2	10.5	8.2	7.0	91.5

Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for various levels of airport closures happening on different days since the start of the epidemic (over the columns)

of airport closures that maximises the utility of the nth Day Rule. GAs, in their simplest form, operate on binary strings called chromosomes which are an encoding of the parameters of interest, commonly referred to as genes. Any particular instance of a chromosome has a genotype which refers to a specific string of bits each with values **1** or **0** representing a particular gene's allele. Once the problem is formulated within this framework, the GA procedure is as follows:

1. Define a fitness function $F(X)$ which evaluates the optimality of a given genotype.
2. Initialise a population of chromosomes with randomly assigned genotypes.
3. Evaluate the fitness of all members of the population. Individuals will then be selected for breeding at a frequency proportional to their fitness (this emulates the "survival of the fittest" mechanism).
4. In a process known as *crossover*, pairs of *selected* genotypes are split uniformly at some locus along the chromosome and then *recombined*, to form new chromosomes.

Table 9 Limited nth day rule by betweenness centrality

Nth Day Rule by betweenness : peak infections % Reduction wrt to worst scenario						% affected population in the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	17.8	9.4	8.7	8.2	7.8	37.0
5%	70.9	36.0	33.8	32.3	31.2	37.0
10%	74.6	36.2	34.3	33.0	32.0	37.0
25%	79.7	34.9	33.3	32.2	31.4	37.0
100%	79.8	34.4	33.0	32.0	31.2	37.0
Nth Day Rule by betweenness : total cases % Reduction wrt to worst scenario						% Affected population In the worst scenario
	Day					
	1	2	3	4	5	200
Closed airports						
1%	5.6	1.5	0.7	0.6	0.5	91.5
5%	55.9	19.2	6.5	4.9	4.0	91.5
10%	78.2	29.3	8.7	6.7	5.6	91.5
25%	84.8	39.5	10.3	8.1	6.9	91.5
100%	86.8	43.2	10.5	8.2	7.0	91.5

Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for various levels of airport closures happening on different days since the start of the epidemic (over the columns)

5. *Mutations* are then applied at random to alleles of the recombined chromosomes (simply by bit flipping) in order to prevent irrevocable loss of any characteristic.
6. The process (3–5) is then repeated so that the fittest members of the population are selected.

For a more detailed explanation of the entire process we refer to [16]. We opt to use a genetic algorithm for this problem since the solutions of our combinatorial problem can be easily represented as binary strings. By contrast, problems of this nature are ill suited to classical optimisation methods such as Gradient Descent, as it is not possible to analytically compute gradients and our search space is too large for approximation methods. Additionally, a fitness function can be easily defined from the metrics that we have previously described. In the following paragraph, we provide more details regarding the setup of the GA into our framework.

GA details. We define the key information required to formulate our problem within the GA framework as follows:

– Chromosome: instead of optimising *which airports to close on each day*, we choose to optimise *which countries should close their airports on each day*.

This reduces the GA search space by **15** times which is highly desirable for convergence and global optimum. This re-framing is also reasonable because it reflects the likely course of action that we would observe in epidemics, where decisions are often taken at country level.

We define the chromosome as a **195**-bit string where **1** in the i-th entry indicates that the i-th country's airports are closed, and vice versa **0**s indicate closed airports.

– Fitness function: in designing the fitness function, we seek to balance the economic impacts of airport closures along with the reduction in peak infections and peak recoveries, relative to an unmitigated epidemic. Since we have no prior biases between the following **3** criteria:

$$T = \% \text{ in reduction of peak infections}$$
$$P = \% \text{ in reduction of peak recoveries}$$
$$A = \% \text{ airports remaining open}$$

we aggregate them into the objective functions as follows:

$$F(X) = T(X) * P(X) * \sin\left(\frac{\pi}{2} * A(X)\right)$$

where X represents a simulated SEIRS environment that is obtained for a given chromosome.

In order to evaluate the fitness function, we must compute the values of T and P (computing A is trivial). This clearly involves inputting the parameters encoded in the chromosomes genotype into our simulation developed in previous sections. We perform this by first using a lookup table to convert between the **195**-bit country closures string specified in the genotype to a **3,425** bit string airport closures vector required by our model. Next we set to zero the rows and columns of the closed airports within the adjacency matrix at the appropriate time steps (to disconnect an airport from the network). Finally, we proceed to run the algorithm for **200** days, which is sufficient to capture the first wave of the epidemic under all conditions (Table 11).

5.5 Optimisation Results and Comparisons

The results of the algorithm are shown in Table 10 in the usual format for consistency. Similarly to the other methods, we see a substantial difference between a day **1** closure and a day **5** closure, both with respect to the peak of infections and peak of recoveries. The mitigation effect, measured as a percentage reduction of the number of cases with respect to the unmitigated scenario, is approximately at the same level as the one obtained with the limited nth day rule. Differently from the limited nth day rule, the GA approach does not consider airport closures but country closures, which is clearly more realistic but also less efficient. However, we note that the objective function can be calculated in all of the solutions, and it is maximised by the GA solution, for each day.

Table 10 Genetic algorithm strategy

GA Rule :						% Affected population
% Reduction wrt to worst scenario						In the worst scenario
	Day					
	1	**2**	**3**	**4**	**5**	**200**
Measure						
Max(I)	79.8	36.9	29.0	29.8	31.1	37.0
Max(R)	86.8	34.1	8.9	7.2	5.6	91.5

Mitigation level of the disease measured as a percentage reduction with respect to the baseline which has no airport closures (reported on the last column). The results are reported for permanent airport closures happening on different days since the start of the epidemic (over the columns). Different airports are closed on different days

Table 11 Genetic algorithm strategy

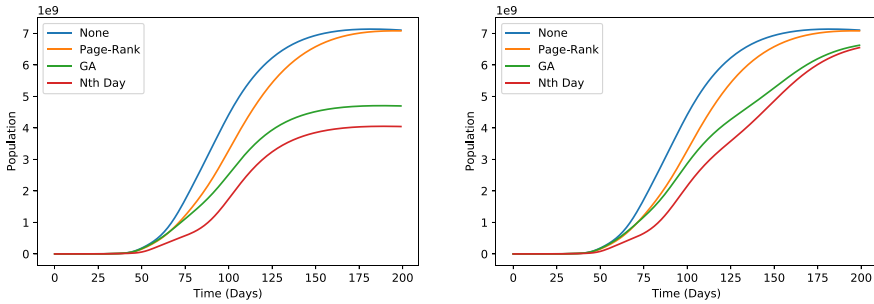| GA strategy	
Day	Number of closed countries
1	83
2	113
3	88
4	109
5	103

Number of countries that close they airports on each day, according to the GA's solutions. The total number of countries is **220**.
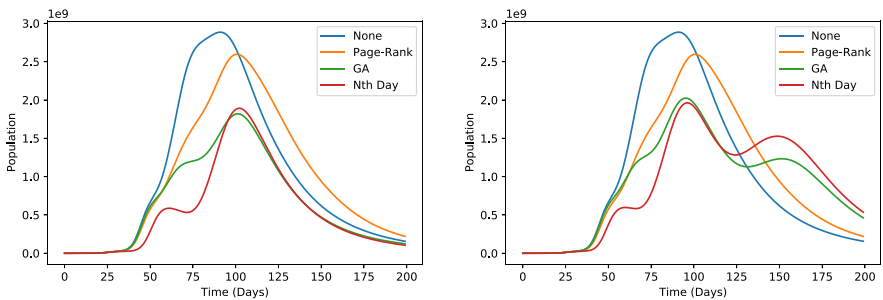
(a) I compartment for day **2** closures. (b) I compartment for day **4** closures.

Fig. 9 A comparison of GA strategy versus naive strategies and total inaction

We offer a more detailed overview of the comparisons between the different mitigation strategies in Figs. 9 and 10. In these figures, we represent the evolution of the two main compartments of interest, I and R. We notice that the closure strategies

(a) **R** *compartment for day **2** closures.　　(b) **R** *compartment for day **4** closures.*

Fig. 10 GA results. A comparison of SEIRS compartments for the GA strategy versus the other strategies considered. The limited nth day rule uses Page-Rank centrality and closes **1**% of the airports

affect these compartments in a rather non-trivial way, in that, at the global level, they can delay the epidemic and fragment it into various waves (i.e. peaks of infections).

In Fig. 9, the line labelled as `None` refers to the unmitigated scenario: we see that in this case, the epidemic peaks very quickly and that the *I* compartment reaches the highest value across all methods. This is the worst scenario that we observe, when compared to the other strategies.

For day **2** closures, we see that the GA strategy strictly dominates the centrality-based strategy because the infectious cases are fewer at all times. The nth day rule, which refers to the closure of all airports from the nth day, naturally gives the best results because this also corresponds to the strictest of measures.

The peak of infections is in fact higher than that of the GA strategy: this is perfectly reasonable because the different strategies not only affect the peaks but also the distribution of infections over time. For day **4** closures, the results are not as clear; however, we argue that also in this case, the GA algorithm attains better results than the Page-Rank strategy, getting close to the nth day rule but without as many closures.

Figure 10 highlights in a more clear way the same message, providing a neat ranking of the strategies. For both day **2** and day **4** closures, we can see that the GA achieves excellent results without having to resort to full closures; on the other hand, the centrality-based strategy provides little improvement over the unmitigated scenario.

Our results suggests that the GA learns to leverage the network structure via closures in such a way as to accelerate the initial infection rate but achieving a lower point of equilibrium. To gain further intuition into the GA behaviour, it is best to look at Fig. 11 which exhibits the evolution of the GA strategy as we alter the day at which closures occur. What we observe is that it is initially optimal to close China and certain other countries such as France which are very well connected to China via air routes. However, as governments delay airport closures, we find that the GA shifts focus away from China and starts to establish "firebreaks" in other regions such as the Middle East, South America, Europe and Afr ica.
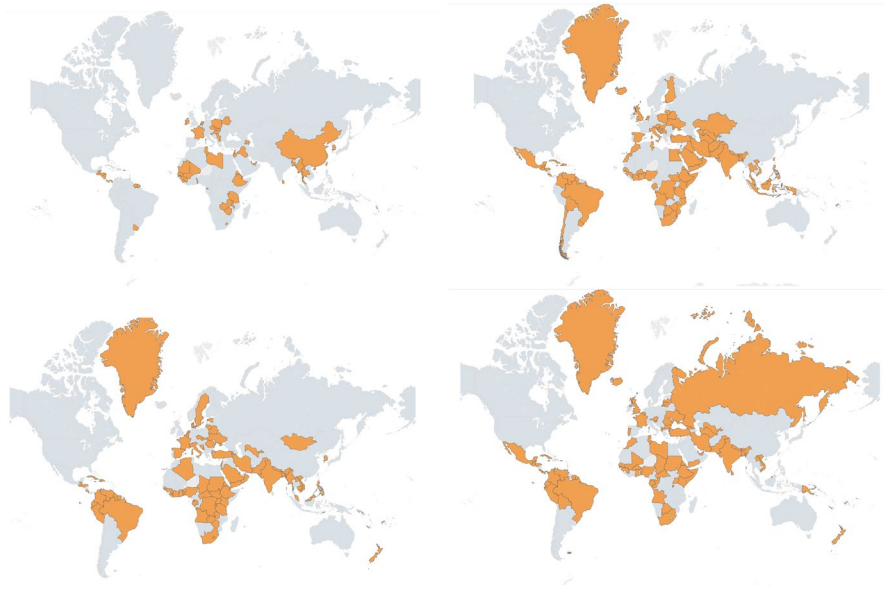
Fig. 11 GA solutions for days **1** to **4**. The countries in orange are the ones that should be closed on day **1** (top left), day **2** (top right), day **3** (bottom left) and day **4** (bottom right)
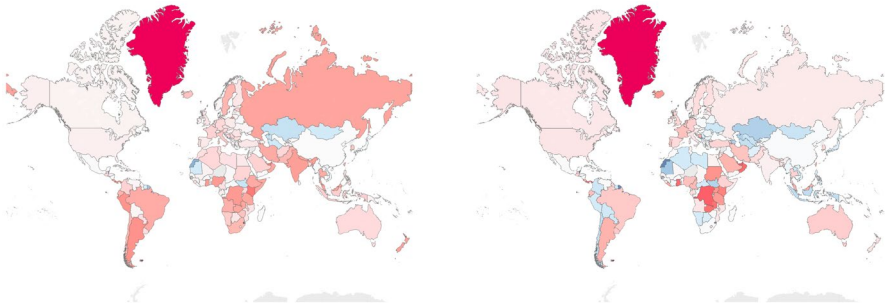
Fig. 12 GA method. Cases reductions achieved by the GA solution for day **3** closures, with respect to the unmitigated scenario. The peak of infections is shown on the left panel, whereas the right panel shows the peak of recoveries. The colours of countries fade from red (GA strategy is favourable) to blue (unmitigated strategy is favourable)

In Figs. 12 and 13, we focus on day **3** closures, and we examine the percentage change in peak infections and peak recoveries under the GA strategy, compared with the unmitigated strategy (i.e. the worst strategy) and nth day rule (i.e. the best strategy), respectively. Despite less than half of the airports being closed under the GA strategy, the vast majority of countries see a reduction in peak infections and peak recoveries relative to the unmitigated case. By contrast, fewer countries see a reduction in peak infections and peak recoveries relative to the nth day rule. We note
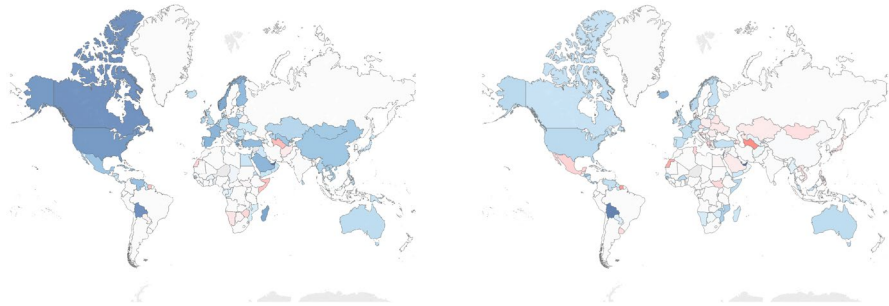
Fig. 13 GA method. Cases reductions achieved by the GA solution for day **3** closures, with respect to the nth day rule. The peak of infections is shown on the left panel, whereas the right panel shows the peak of recoveries. The colours of countries fade from red (GA strategy is favourable) to blue (nth day rule is favourable)

that the countries that are closed in day **3** roughly correspond to the countries that will see the best reductions for the number of cases. In particular, we highlight that some highly connected countries such as China and the United States, do not benefit much from the day **3** GA closure strategy and see little reductions in the number of infections and recoveries (see Fig. 12). However, these countries would benefit more from a close-all strategy, as per the nth day rule (see Fig. 13). Our results provide evidence of the effectiveness of the GA approach over other naive strategies, offering a useful tool can be used to obtain a model-based perspective in this decision problem.

6 Conclusion

In this paper, we have shown that human mobility infrastructure networks are highly complex and highly resilient to node removals. This has important consequences in epidemiology, since the high connectivity facilitates the spread of pandemic diseases to a worldwide scale. These observations motivated us to seek and test nontrivial airport closure strategies that could maintain a good network connectivity while slowing down the spread of a potential epidemic.

The main contribution of our work is that we illustrate a useful methodology which can be used to study epidemics, predict and analyse possible realistic scenarios, and, thus, support decision making with regard to crucial interventions that could help in confining the epidemic. We provide an application of our method to a variety of scenarios which, to some extent, also resemble recent COVID-19 epidemic. Our results do not directly relate to COVID-19, rather we take a more general approach and explore a number of realistic circumstances. Our tool is designed in a flexible way and can accommodate a variety of epidemic settings.

The framework that we propose is based on a metapopulation SEIRS model, where the subpopulations are located in the airports' locations, and they are composed of the individuals living in those nearby areas. Our simulations allowed us to

explore and study a variety of realistic scenarios to understand the dynamics of the spread of the disease. We considered several different airport closure strategies, and we compared them using measures extracted from our simulations.

One main message that arises from our analysis is that, due to travelling, the disease is seeded in many locations worldwide with impressive speed. If conditions are met for isolated local escalations of the epidemic, then most countries will be hit by the seeded disease after a variable delay, regardless of any late interventions on travelling restrictions.

Our findings suggest that the first week of dispersion of the disease through the network is a critical time period for effective intervention, however interventions in the network, such as airport closures, still provide some reductions to peak infections and total cases if implemented after the initial week. Furthermore, we show that policies which reduce community spread can be combined with our proposed airport closure strategies to provide greater benefits than if either policy had been used separately.

Finally, we explored the application of an optimisation approach to identify optimal airport closures within the critical first week of disease spread, in order to reduce the global impact of the epidemic while keeping as many airports open as possible. This optimisation approach, based on a genetic algorithm, improved upon all of the other methods, hence providing a new model-based perspective on the decision-making process that leads to the travel restrictions. One very interesting aspect of our results is that the algorithm leverages the complex structure of the network to place strategic "fire-breaks," which drastically reduce peak infections and total cases.

Appendix

Algorithm Derivation

In this section, we break down the various components of our diffusion and SEIRS models to describe in detail the mechanics of the epidemic simulator. The pseudocode for the full procedure is provided in Table 2.

SEIRS model The algorithm has been vectorised so that θ is a $4 \times N$ matrix containing the states for all airports. The rows of θ are the N dimensional vectors S, E, I and R representing the SEIRS compartments (for each of the airports' locations) with non-negative real numbers.

The fundamental equation for the local SEIRS is given by

$$\frac{dS}{dt} = \delta R - \frac{S\beta I}{M}$$
$$\frac{dE}{dt} = \frac{S\beta I}{M} - \epsilon E$$
$$\frac{dI}{dt} = \epsilon E - \gamma I$$
$$\frac{dR}{dt} = \gamma I - \delta R$$

where β, ϵ, γ and δ are the epidemic parameters. In the pseudocode, we make use of matrix notation and, thus, work with

$$\Sigma = \begin{pmatrix} -\beta & 0 & 0 & \delta \\ \beta & -\epsilon & 0 & 0 \\ 0 & \epsilon & -\gamma & 0 \\ 0 & 0 & \gamma & -\delta \end{pmatrix}.$$

With this notation, one-step forward for the local SEIRS epidemic can be simply obtained with the matrix multiplication $\Sigma\theta$, for a suitable distribution matrix θ.

Diffusion model The mechanism for the diffusion model requires more steps and transformations which we describe in details here below. One main reference for the techniques used here is [18].

If consider our network as a dynamic system, we have N nodes, each with an associated amount of fluid $\psi_j \geq 0$. The fluid is transferred from one node to another along the weighted edges of the network. The fluid flows from node i to node j at a rate proportional to the difference in the amount of fluid at each node $c(\psi_i - \psi_j)$, where c is the constant of proportionality or more commonly referred to as the *diffusion constant*. This can be translated into matrix notation with

$$\frac{d\psi_j}{dt} = c \sum_{i=1}^{N} (\psi_i - \psi_j) A_{ij}, \tag{2}$$

where A is a suitable adjacency matrix. Equation 2 can be rewritten as follows:

$$\begin{aligned} \frac{d\psi_j}{dt} &= c \sum_{i=1}^{N} \psi_i A_{ij} - c\psi_j \deg(v_j) \\ &= c(A - D)\psi \\ &= -c(D - A)\psi \end{aligned} \tag{3}$$

where $\deg(v_j)$ is the degree of node j, and D is the $N \times N$ diagonal matrix with degrees as entries.

In our context, the fluid ψ represents the local population available for travel, or, more precisely, a the size of a SEIRS compartment which is available for travel. In fact, we have **4** contemporaneous "fluids" that are moving through the system, corresponding to individuals for each of the compartments travelling

between locations. In order to accommodate this, we define a system through $\psi \in \mathbb{R}^{4 \times N}$ and redefine the equation as follows:

$$\frac{d\psi}{dt} = -c\psi(D - A^\top) \in \mathbb{R}^{4 \times N}. \tag{4}$$

This completes a first step into defining the diffusion equation provided in the algorithm's pseudocode.

Equation 4 simply defines how the local populations that are available for travelling *should* travel, by defining the changes in size for each of the SEIRS compartments, at each of the locations. However, the equation cannot be used as is, since it would create unrealistic travelling patterns. The limitations that we address are the following:

1. The migration process will have a stationary distribution, which likely does not correspond to the initial state of the system. This means that the populations that we observe at each locations will change dramatically over time, which is unrealistic. By contrast, we would like to have constant local populations M_j, which would be in line with our assumption of "no permanent migration".
2. The total out-flow from a node may exceed the node's availability. Although unrealistic, this becomes possible because the continuous diffusion model becomes discretised when implemented. As a consequence, the number of outbound passengers may exceed the number of individuals that can travel from a location, so we need to add a condition to ensure that this does not happen.
3. The number of available travellers should not be the only driving factor behind migration flows. As an example, Ethiopia has a very large population, which corresponds to a large number of potential travellers; however, few routes go through this country.

As regards the first problem, we separate θ into θ_B and θ_T, representing the **locals** and **visitors**, respectively. This allows us to disentangle outbound and return travels, and monitor both. As initial condition, all local populations are in θ_B, i.e. there are no visitors. We proceed with the first step of the diffusion process by applying (4) on the matrix $\psi_+ = \theta_B \alpha$ which gives us the changes of local populations denoted by $\frac{d\psi_+}{dt}$. For those local populations that decrease due to *outbound* individuals, we can write $\theta_B = \theta_B + \min(\frac{d\psi_+}{dt}, 0)$. As regards the positive entries of $\frac{d\psi_+}{dt}$, which correspond to those locations that are supposed to see new individuals *arriving*, we do not add these entries to θ_B but to θ_T. This is in agreement with the fact that these travellers will join the populations of visitors in their new locations, and not the locals. In mathematical terms: $\theta_T = \theta_T + \max(\frac{d\psi_+}{dt}, 0)$.

Analogously, we apply the diffusion process (4) to the matrix $\psi_- = \theta_T \alpha$, which gives us the changes of visitor populations denoted by $\frac{d\psi_-}{dt}$. Then, the negative entries of this term will correspond to the visitors that are *departing*, whereas the positive entries are the visitors that are *returning* at the end of their travels, thus, flowing back into the local populations. The complete updates are then

$$\theta_B = \theta_B + \min\left(\frac{d\psi_+}{dt}, 0\right) + \max\left(\frac{d\psi_-}{dt}, 0\right)$$
$$\theta_T = \theta_T + \max\left(\frac{d\psi_+}{dt}, 0\right) + \min\left(\frac{d\psi_-}{dt}, 0\right).$$

Note that we have carefully chosen the wording *outbound*, *arriving*, *departing*, *returning* to reflect the sequence of steps all travellers pass through in order. This is crucial to ensure logical behaviour of travellers and also to prevent leakage of fluid/people from the network.

In order to address the second issue, we can simply divide ψ_i by the degree of the node which ensures that the outflows computed to every other node will be at most ψ_i. This results in the diffusion equation (4) being changed to

$$\frac{d\psi}{dt} = -c\psi D^{-1}(D - A^\top) \in \mathbb{R}^{4 \times N}. \tag{5}$$

This arbitrary operation simply guarantees a rescaling of each individual value, so that we can ensure that the entries of ψ remain non-negative throughout the procedure.

As regards the third issue, this cannot be addressed with a simple rescaling as we need to ensure that the importance of airport is properly captured by our approach. For this reason, we replace the adjacency matrix A with a weighted matrix C to encourage flows to more central airports. We construct C by first amplifying the entries of A by the centrality of the receiving node, and then we rescale the entries to obtain a row stochastic matrix. The generic element of C is defined as follows:

$$C_{ij} = \frac{A_{ij}P_j}{\sum_{k=1}^{N} A_{ik}P_k},$$

where P_j is the page-rank centrality of node v_j. This makes the migration matrix more realistic since it captures better the role played by each airport, while maintaining exactly the same adjacency structure as informed by A. We then simply make the matrix replacement in (4), noting that, since C is row stochastic, the matrix D becomes the identity matrix and, thus, disappears from the equation. Hence, the diffusion (4) is replaced by

$$\frac{d\psi}{dt} = -c\psi(I - C^\top) \in \mathbb{R}^{4 \times N}$$

which is equal to the formula provided in the pseudocode.

References

1. Shereen MA et al (2020) COVID-19 infection: origin, transmission, and characteristics of human corona-viruses. J Adv Res 24:91
2. Chinazzi M et al (2020) The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. Science 368(6489):395–400
3. OpenFlights: Airport And Airline Data. In: https://openflights.org/data.html. Accessed Jan 2017
4. Wang YJ, Wong GY (1987) Stochastic blockmodels for directed graphs. J Am Stat Assoc 82(397):8–19
5. Nowicki K, Snijders TAB (2001) Estimation and prediction for stochastic blockstructures. J Am Stat Assoc 96(455):1077–1087
6. Daudin J-J, Picard Franck, Robin Stéphane (2008) A mixture model for random graphs. Stat Comput 18(2):173–183
7. Balcan D et al (2010) Modeling the spatial spread of infectious diseases: the GLobal Epidemic and Mobility computational model. J Comput Sci 1(3):132–145
8. Balcan D et al (2009) Multiscale mobility networks and the spatial spreading of infectious diseases. Proc Natl Acad Sci USA 106(51):21484–21489
9. Hethcote HW (2000) The mathematics of infectious diseases. SIAM Rev 42(4):599–653
10. Hou C et al (2020) The effectiveness of quarantine of Wuhan city against the Corona Virus Disease 2019 (COVID-19): a well-mixed SEIR model analysis. J Med Virol 92(7):841–848
11. López L, Rodo X (2020) A modified SEIR model to predict the COVID-19 outbreak in Spain and Italy: simulating control scenarios and multi-scale epidemics. Available at SSRN 3576802
12. Peng L et al (2020) Epidemic analysis of COVID-19 in China by dynamical modeling. In: arXiv preprint arXiv:2002.06563
13. Radulescu A, Cavanagh K (2020) Management strategies in a SEIR model of COVID 19 community spread. In: arXiv preprint arXiv:2003.11150
14. Apolloni A et al (2014) Metapopulation epidemic models with heterogeneous mixing and travel behaviour. Theor Biol Med Model 11(1):3
15. Tizzoni M et al (2014) On the use of human mobility proxies for modeling epidemics. PLoS Comput Biol 10(7):e1003716
16. Spall JC (2005) Introduction to stochastic search and optimization: estimation, simulation, and control, vol 65. Wiley, New York
17. A-Model-Based-Approach-To-Assess-Epidemic-Risk. In: https://github.com/hugo1005/A-Model-Based-Approach-To-Assess-Epidemic-Risk. Accessed Aug 2020
18. Newman Mark (2010). Networks. OUP Oxford, Oxford. https://doi.org/9780191500701
19. Peixoto TP (2014) Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models. Phys Rev E 89(1) . ISSN: 1550-2376. https://doi.org/10.1103/PhysRevE.89.012804
20. Gridded Population of the World (GPW), v4. In: https://sedac.ciesin.columbia.edu/data/set/gpw-v4-admin-unit-center-points-population-estimates-rev11. Accessed Jan 2020