

Description of modelling framework per April 29, 2020

Table of Contents

<i>Introduction</i>	2
<i>Layer 1: Stochastic IBM network with SEIR epidemiologic transmission dynamics</i>	5
Structure of the municipality network IBM	5
SEIR-type epidemiologic dynamics.....	7
Implementation of intervention measures.....	9
Base-line testing of IBM with empirical data	10
<i>Layer 2: Nonlinear model predictive control (MPC)</i>	12
Nonlinear MPC applied to control of the IBM.....	13
<i>Layer 3: Region- and nation-level complex network model</i>	18
<i>Table 1. Parameters used</i>	19
<i>References</i>	20

Introduction

After the COVID-19 outbreak hit East Asia, Europe and the US with full force, the need for advanced and more accurate methods for prediction and control of pandemic outbreaks become urgent. The gold-standard in current epidemiological modeling is to use an individual-based model (IBM), where individual agents are assigned a behavior pattern which is meant to simulate the society, while infection between the individuals is based on probabilities based on observations.

One of the first publications concerning the modeling of the COVID-19 epidemic came from Imperial College (Ferguson et al. 2020). This report describes in detail how the IBM from the team at Imperial College predicts the effects of the COVID-19 outbreak in the UK and USA. It also discusses different strategies, such as social distancing, which could be implemented to control and mitigate the outbreak. Moreover, the report shows a way to control the outbreak through utilizing a suppression strategy with an on/off approach to soften the blow on national economy while still avoiding that the healthcare system does not become overloaded with patients. Later, several other studies have been released using the same or similar modeling frameworks and approach to interventions. Using an on/off strategy with repeated actions and random switching times is a **suboptimal** choice from a feedback control point of view since **optimality** is not addressed. Another concern with this *ad hoc* approach is that, simple control strategies can generate closed-loop stability problems and they can also be the root to exponential epidemic growth. In our view, stability and optimality must be explicitly addressed when designing control strategies for that inform policy choices. Consequently, our approach is to optimize the

- Choice and combination of policies, i.e. which different groups should be subjected to social distancing measures; children, students, workers, elderly etc.
- The timing of when to start each policy, and for how long they should stay in effect.

The main mathematical framework to achieve this goal is optimal control theory for dynamic constrained systems, such as model predictive control (MPC) (Borrelli et al. 2017). An MPC controller, when properly designed, can anticipate future epidemic spread and, from this predicted future, compute the optimal control actions. For an epidemic model, the control actions will be the optimal combination of policies and interventions, as well as their optimal starting time and duration. The accuracy of the optimal policy structure is further improved by measuring the number of infectious people through extensive testing. This input of experimental evidence allows the MPC controller to recompute the predicted policies each time new information is available.

When selecting and designing our computational modeling framework, we have explicitly taken into consideration the following aspects: (A) It should be possible to directly test a wide variety of interventions, such as partial workplace and school closings, social distancing, and testing with subsequent isolation of infected individuals. (B) Each infected person should have a disease-state development where the stages have realistic time delays. If a lag in response time is not included, we anticipate that all estimated response times to interventions will be markedly wrong. This is especially evident for the correct estimation of R , and thus, the capacity to determine how quickly R responds to interventions. The mentioned challenges are inherent to most metapopulation-based models, whereas an IBM approach is not hampered with such issues.

Guided by the above reasoning, we have developed a complex system modelling framework to describe the spreading pattern of COVID-19 in mainland Norway. The model design and build features are chosen with the intent that it can be used to carefully assess a multitude of relevant intervention strategies and predict optimal scenarios for disease mitigation and control. Briefly, the NTNU network-control model is based on three interacting layers of description:

1. At the basal layer, we generate a separate IBM using complex network theory for each Norwegian municipality. The features of the model's layers and each individual all adhere to high-resolution demographic data specific for the respective municipality. The disease-state of each individual follows an epidemiological SEIR-type model.
2. The intermediate level consists of a nonlinear model predictive controller (MPC) for state feedback. The NMPC is designed using an adaptive low-fidelity SEIR model to approximate the governing dynamics of the IBM. The model-based controller allows for determination of optimal intervention strategies for each municipality, in addition to allowing for uniform control of a sub-set of municipalities forming a region or even uniform control of all municipalities at a national level.
3. At the highest level, municipalities are connected by the movement of people due to regular activities, such as commuting, and intermittent activities.

The approach of combining an IBM with optimal nonlinear control in studying epidemics allows the fine-tuning of intervention strategies which is impossible with the widely used application of *ad hoc* measures. Our total model contains approximately 5 million people and simulates spreading dynamics through social interactions in layers such as households, schools and workplaces.

Layer 1: Stochastic IBM network with SEIR epidemiologic transmission dynamics

Structure of the municipality network IBM

We generate a high-fidelity IBM for a single municipality by creating a set of households corresponding to the population N_m of that municipality. Each household consists of one or several nodes, to each of which we assign a list of attributes:

1. Age
2. Domicile
3. Layer memberships / group membership within layers (see Fig. 1)
4. Disease state (see Fig. 2), and date of last change in disease state
5. Disease test state

The model has 9 layers. Layer a)-h) each consist of many groups, and each individual is only member of one of these groups. Further, an individual can only be present in a single of the layers b)-h). The layers are: a) Household, b) Day-care, c) Primary school, d) Secondary school, e) High school, f) Workplace, g) Nursing home, h) Hospital, and i) Generic contact network. A group is designed as a k-clique, i.e. where all members of a group are in contact. The assignments that form the IBM contact network are constrained such that known high-resolution demographic data for each municipality is matched:

- 1) The household layer consists of separate households with a size and age distributions that follow demographic data for that municipality.
- 2) The number of schools, type, and their student populations are based on demographic data.
- 3) The number of day care facilities is based on demographic data.
- 4) The number of nursing homes and their population sizes are based on demographic data.

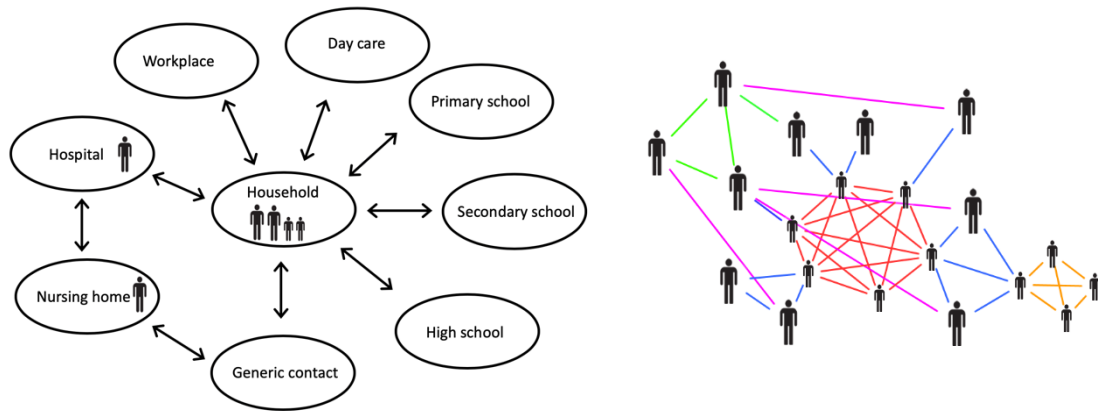


Fig. 1 Structure of network IBM. (Left) Possible domicile indicated by person figure. Named circles show available layers that a person can be member of. (Right) Illustrative example of resulting connection network between individuals caused by shared group membership in different layers: household (blue), primary school (red), day care (orange), work (green) and generic (pink).

- 5) If a household has multiple children in e.g. day-care age, these children are assigned the same day-care. Similar for primary and secondary schools. For high-school age, the assignment to school unit is random.
- 6) In the work layer, the number of companies and their sizes are based on demographic data. Note that this layer is intended to represent spread between co-workers. For professions with large amounts of exposure to the general public, contacts with customers are represented through the generic contact network.
- 7) The generic contact layer is designed to capture heterogeneity in a person's daily contact patterns. We include two different modes: (A) For young (< 20) and elderly (80+) we assign number of contacts following a normal distribution, whereas the remaining age groups are assigned contacts following a combination of a normal and a power law distribution. The latter generates larger contact heterogeneity.
- 8) Using SSB statistics, adults with different work and household municipalities are assigned work locations in the correct municipality.
- 9) When a person is committed to a hospital, they are removed from their domicile (household or nursing home).

10) When an infected person manifests a symptom, is confirmed COVID-19 positive, or an asymptomatic person is confirmed COVID-19 positive, the model assumes they will self-quarantine from activities in all layers except their domicile.

All individuals participate in the generic contact network, which is designed as a random time-dependent scale-free network to capture heterogeneity in contact patterns. A new instance of this network (per municipality) is generated every day. Except for the generic contact network, each grouping in each layer is represented as a k -clique, where every individual is connected to all others in the group, thus representing a well-mixed group. Fig. 1 shows a schematic of the IBM layers inside which there are smaller groups (left panel) and a simplified example of a possible resulting contact network (right panel).

The probability of infecting depends on which layer the contact is located. Within a layer, all members of a single group have a constant probability of infecting other members of that group. The specific values used for each layer are detailed in Table 1.

SEIR-type epidemiologic dynamics

Each individual in the IBM model is either healthy, in various states of infection, or recovered from the disease: Susceptible (S); Exposed (E); Infected, asymptomatic (I_a); Infected pre-symptomatic (I_p); Infected, symptomatic (I_s); Hospitalized (H); Intensive care (ICU); Recovered (R); or Dead (D).

The different states and their possible transitions are captured by the state-transition schematic of Fig. 2. The direct transition from I_s to D captures the disease trajectory for some people in nursing homes. The transition from S to E is governed by disease transmission probabilities in the IBM contact network. After an individual is infected from a neighbor in the contact network, the SEIR-type dynamics of that individual's disease state (onward from E) are governed by stochastic

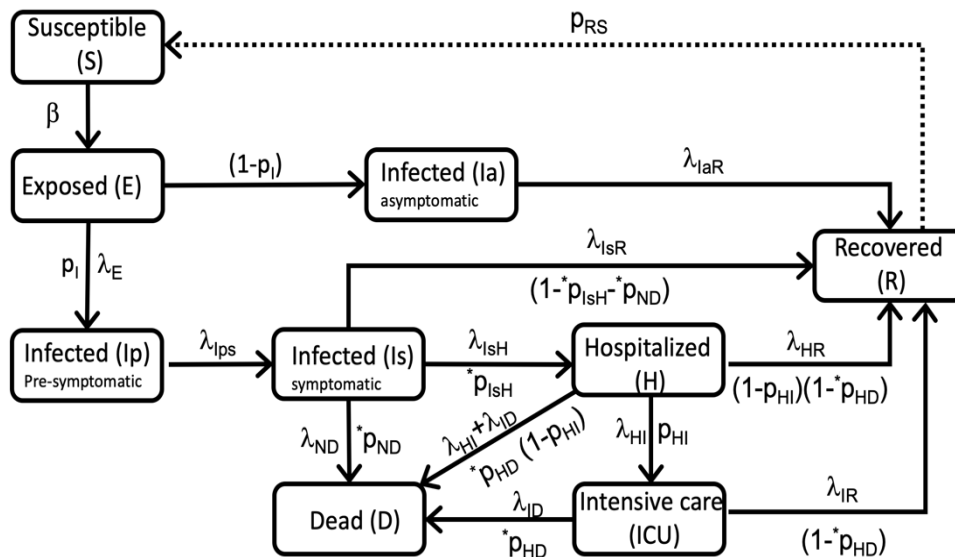


Fig. 2 Block diagram showing schematic of the SEIR-type epidemiological state model for each individual in the IBM. Solid lines show transitions currently active in the model. The dotted line indicates current uncertainty about long-term COVID-19 immunity. Starred parameters indicate age dependent values.

processes with appropriate waiting times according to empirical data for COVID-19. We continually update our parameter estimations given the daily update in disease state numbers for Norway by region.

At each time point, an individual will store four points of information about their progression through the SEIR states: current state, date of last change of state, next state, and date of next change of state. Each day, the model updates the state as needed. Upon entering a new state, the model selects the following state according to the probabilities and waiting times specified in Fig. 2 and Table 1. Once the next state is determined, the duration for which the individual will remain in the recently entered state is generated according to that state's specific λ , and the date of next change is set accordingly.

Note that our SEIR-model uses age-stratified transition rates, where the age of an individual decides which of the age groups that individual is part of. Table 1 shows the rates used for the different transitions in Fig. 2.

Implementation of intervention measures

A high-fidelity IBM allows for the implementation of detailed interventions, where contact patterns of a selection, or all, of the individuals may be modified. It is thus straightforward to conduct computational experiments where some or all schools are closed or partially closed, various levels of social distancing is included, or testing and vaccination is administered. The most basic intervention strategy consists in locking down various IBM-layers. How this is implemented in detail depends on the layer:

1. For day care facilities the layer is disabled in its entirety. Similarly for secondary schools and high schools.
2. For primary schools, a shutdown is implemented separately for grades 1.-4. and grades 5.-7., allowing for the opening of grades 1.-4. Independent of the other grades.
3. In the work layer, a shutdown is not complete – instead, a fraction of cliques intended to represent workplaces where work-from-home is feasible, are disabled.
4. In the random contact layer, lockdown is implemented by setting an upper bound on the number of possible daily contacts for a fraction of the nodes.
5. The household and nursing home layers are never disabled.

General social distancing and hygiene measures are represented by a change in the base transmission rates of each layer.

Targeted quarantine of individuals

The quarantining of an individual is represented by disabling workplace, school, and random layer spread for asymptomatic and pre-symptomatic individuals put into quarantine, similarly to behavior patterns for individuals who self-quarantine when symptoms manifest. Such quarantines may be imposed on individuals who test positive, as well as other individuals at high risk of contagion from someone who has recently tested positive (such as those in the same household, workplace, or school).

COVID-19 testing of (sub-) population

The testing itself is represented by drawing individuals from the symptom-free population (susceptible, latent, asymptomatic, pre-symptomatic, or recovered), and returning a positive test if they are asymptomatic or pre-symptomatic. Pooled tests are represented similarly, but with multiple people (forming a pool) simultaneously tested and all tests marked as positive if any of the people in the test pool are in the asymptomatic or pre-symptomatic states. Our IBM implementation allows for random selection of individuals for testing, or testing based on a pre-conceived schema.

Base-line testing of IBM with empirical data

The parameters in the stochastic IBM are treated as fixed, and they are either set from the literature or estimated directly from publicly available empirical data, as indicated in Table 1. We determined model parameters by fitting the predicted hospitalization rate of our Oslo model to hospitalizations for Oslo in the period from March 1st to April 20th. We assume a sudden transition between two regimes, with infections following original “unrestricted” probabilities until March 13th. After this, all schools and day care facilities close, as do 50% of workplaces. Random contacts are reduced by 75%. Infection probabilities in nursing homes and within households remain the same in both regimes. Model simulations start in March and continue until June 20th, 2020.

Since the lockdown in Norway occurred in the early stages of the epidemic, the fitting of the IBM results to data depends very much on the number of symptomatic infected on the day of lock down (March 13, 2020). Thus, for a given parameter set and municipality, we have used a large number of simulations to determine the optimal number of infected (N_{lock}) on this day for the model predictions to fit to time series data of number of hospitalized persons. Simulations are subsequently conducted by initiation the network with a small number of infections and running the code until N_{lock} is reached. The date is set to be March 13 at this point. Formal model fitting,

sensitivity analysis and meta modelling are work in progress. We anticipate that several of the IBM parameters can be estimated from other data sources, e.g. from aggregated data from the infection tracking application *Smittestopp* and the large-scale systematic testing regime for COVID-19.

Layer 2: Nonlinear model predictive control (MPC)

Control theory, a subfield of *cybernetics*, deals with the control of dynamical systems in machines, biological processes and engineering. A feedback control system has an output that is controlled to a desired state or trajectory by using measurements or observations as feedback signals. The feedback signals, if properly designed, will stabilize dynamic systems that are poorly stable or unstable, as long as they are controllable.

MPC is an advanced feedback control method that is used to control a dynamic system while satisfying a set of constraints on the inputs and the states (Johansen 2011, Borrelli et al. 2017). The objective is to satisfy all constraints using a control action algorithm in an optimal manner and to ensure that the closed-loop system is stable. MPC relies on a dynamical model of the system, e.g. an SEIR model, which is used for dynamic optimization by minimizing a user-specified cost function. This corresponds to Layer 2 in our model framework. The parameters of the SEIR model are continuously adjusted (adaptive control) by fitting observed time series data to the model using system identification methods or other methods for data analysis.

The main advantage of MPC when applied to epidemic models is that, the method allows the current timeslot to be optimized while taking future timeslots into account. This is achieved by optimizing a finite time-horizon, but only implementing the current timeslot and then optimizing future timeslots again, repeatedly. Also, MPC has the ability to predict the future output of the system and take optimal control actions accordingly. As an analogy for the MPC, we can look to board games. Just as with MPC, board games (for instance chess) often have the players plan out a series of moves, their strategy, that to them seems optimal based on their understanding of the game and predictions of the opponent's possible moves. This series of planned moves are your actions in the *prediction horizon*. However, as you cannot account for all possible outcomes, you might need to rethink your strategy. For every move your opponent takes, you will have to adapt to their moves, especially if a move was not incorporated

in your current strategy. It is the same with MPC, which only takes one action before it generates a new action strategy based upon new information gained. This behavior is effectively a feedback control, where future timeslots are used to predict the system behavior or an opponent in chess. A major benefit of this approach is that the closed-loop system dynamics will be more robust to model inaccuracies and prediction errors as compared to an open-loop strategy without feedback (Nowzari et al. 2016).

Nonlinear MPC applied to control of the IBM

In order to apply MPC to control the epidemic outbreak in an optimal manner (Bussell *et al.* 2019), a high-fidelity individual-based model is used to generate observations, as shown in Fig. 3 (see description of Layer 1 for details). The model is adaptive in the sense that observations are used to adjust the model parameters by data-analysis methods. The next step is to fit a low-fidelity model of SEIR-type to the IBM. This is the Layer 2 modeling approach. Note that the low-fidelity model is only used to compute the MPC feedback control law using the governing dynamics of the system. Hence, model inaccuracies are acceptable up to a certain level. Moreover, it is well known that feedback controllers are robust to model and parameter uncertainties, which gives us confidence in using a low-fidelity model to implement the MPC controller.

The low-fidelity SEIR model is also used to predict the future states of the system (see Fig. 3), and each time new observations are available, the predicted horizon and control actions are updated accordingly using the latest information. This process is illustrated in Fig. 4. The computed control signal is an optimal decision strategy in the form of a series of policies, which is fed back to the high-

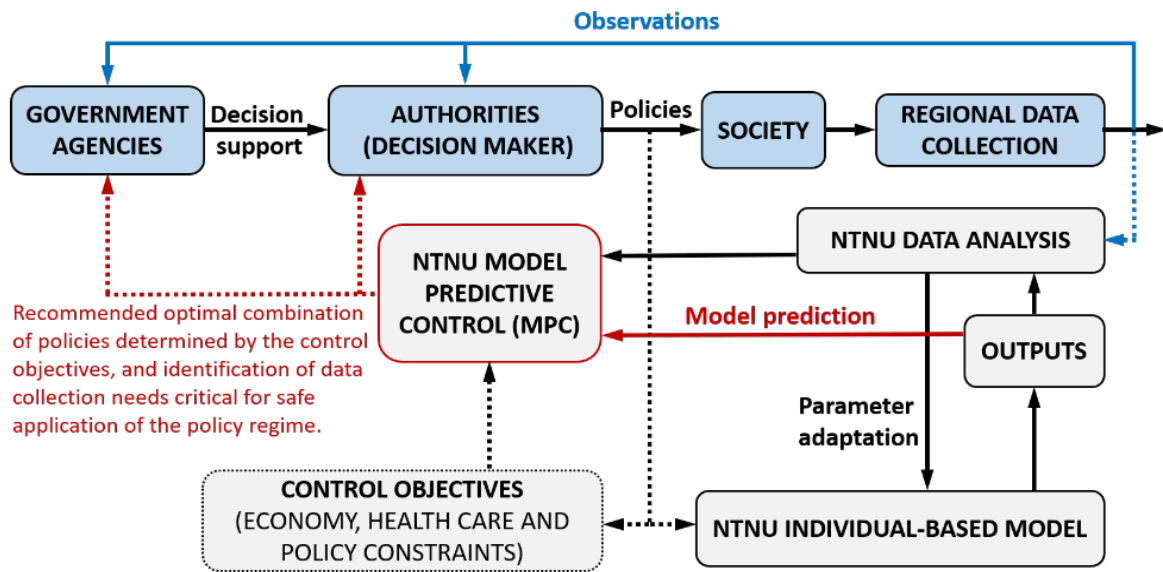


Fig. 3 Block diagram showing the MPC control and control objectives used to specify the optimization problem. The data analysis block is used to fit the model in the MPC controller to the “real world” by parameter adaptation.

fidelity model for verification of the strategy. Consequently, the resulting system is a feedback loop where we allow for human-in-the-loop (HITL) actions represented by an on/off approval switch. The green dotted line in Fig. 4 are the future observed number of patients in intensive care. The MPC algorithm does not know the green line in advance so it uses the SEIR model to predict the number of patients in intensive care, illustrated by the red line. Hence, based on the red line the MPC computes the optimal polices for the present and the future times. This process is repeated at each time instance and the result is a nonlinear feedback controller, which take the future states into consideration when computing the optimal policies. The optimal policies are illustrated by the black line in Fig. 4

MPC control objectives and constraints

The control objective decides how much the optimization algorithm is penalizing the states and control inputs. A wide range of realistic epidemic control objectives can be combined, such as limitations on maximum number of hospital beds in simultaneous use, and the level and duration of societal disruptions.

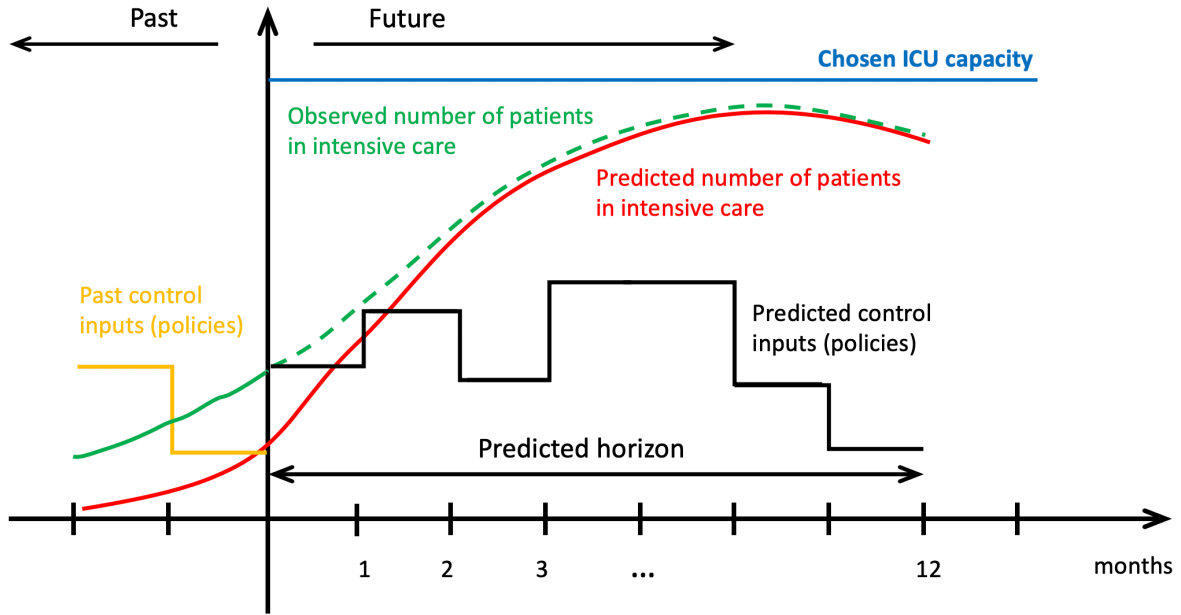


Fig. 4 Example of discrete MPC scheme showing the predicted (red) and measured (green) outputs along with the predicted (black) and past (yellow) control input policies for a 12-month horizon. Our ability to adjust policies according to external constraints, such as ICU capacity is indicated by the blue line.

The mathematical equivalence of this is the cost function

$$J(\mathbf{x}, u) = \int_{t_0}^{t_0+T} (\mathbf{x}(t)^\top \mathbf{Q} \mathbf{x}(t) + r u^2(t)) dt + \mathbf{x}^\top(t_0 + T) \mathbf{Q}_{\text{end}} \mathbf{x}(t_0 + T)$$

where $\mathbf{Q} > 0$, $\mathbf{Q}_{\text{end}} > 0$, and $r > 0$ are the cost function weights, T is the prediction horizon, and

$$\mathbf{x} = [p^S, p^I, p^H, p^C, p^R]^\top$$

denotes the MPC state vector and u is the control input (optimal policy to be computed). The compartments are the number of individuals classified as susceptible (S), infectious (I), hospitalized (H), intensive care (C) and recovered (R). Consequently, the optimal policy can be computed by solving the nonlinear optimization problem

$$u^* = \arg \min J(\mathbf{x}, u)$$

subject to the ODEs

$$\begin{aligned}
\dot{p}^S &= f^S(p^S, p^I, p^H, p^C, p^R) \\
\dot{p}^I &= f^I(p^S, p^I, p^H, p^C, p^R) \\
\dot{p}^H &= f^H(p^S, p^I, p^H, p^C, p^R) \\
\dot{p}^C &= f^C(p^S, p^I, p^H, p^C, p^R) \\
p^R &= 1 - p^S - p^I - p^H - p^C
\end{aligned}$$

and the constraints

$$\begin{aligned}
p^H &\leq K_{\text{cap}}^H \\
p^C &\leq K_{\text{cap}}^C \\
0 &\leq p^i \leq 1, \quad i \in \{S, I, H, C, R\} \\
u_{\min} &\leq u \leq u_{\max}
\end{aligned}$$

with $K_{\text{cap}}^H < 1$ and $K_{\text{cap}}^C < 1$ as constants chosen such that the maximum hospital and intensive care capacities not are exceeded. Since the IBM is stochastic and the nonlinear MPC controller is derived using a deterministic approach, it is necessary to tune the values of K_{cap}^H and K_{cap}^C such that all solutions of the stochastic model satisfy the constraints.

Uncertainty compensation

The SEIR model in the MPC controller is used for dynamic optimization and prediction of future states. A neat feature of feedback control is that it makes the closed-loop system robust for disturbances, parametric uncertainty and structural uncertainty for instance caused by stochastic processes. However, it is important to keep the uncertainty at a minimum by parameter adaptation. The main tool for this is system identification (SI) where data obtained from hospitals and testing of individuals are used to estimate the parameters soon as they are available. The result is an adaptive low-fidelity SEIR model which can be used for optimal feedback control.

In addition to parametric uncertainty it is important to compensate for drift due to the stochastic behavior and model uncertainty. This is done by adding an integral state to the feedback controller,

which minimizes the output error, i.e. the difference of the low- and high-fidelity models (Morari *et al* 2012).

Numerical solvers for MPC

The continuous constrained optimal control problem defined above is turned into a nonlinear program (NLP) with discrete decision variables. This is done using direct multiple shooting and a fourth-order Runge-Kutta method (RK4). The NLP is formulated and solved every timestep using the CasADi library (Andersson et al. 2019) with the Bonmin solver for mixed integer nonlinear programming (MINLP).

Layer 3: Region- and nation-level complex network model

We integrate the municipality models through transport of people: either by regular travel (commuting) or by intermittent travel. We identify commuters as the number of people who have registered work location in a municipality different from their household municipality, according to SSB demographic data. These individuals contribute to the household and generic contact layers in one municipality and in the work layer and generic contact layer of the other municipality.

Intermittent travelers contribute only to the generic contact network in the municipality they are visiting. This structure allows us to predict the effects of export and import of infected individuals across municipalities or larger regions, as well as the regional effects of individuals having been infected abroad. It also allows us to test the effects and optimization of heterogeneous social distancing policy regimes, such that for a given time period different sets of policies apply to different regions depending on infection pressure and other factors influencing the infection dynamics.

Our network IBM allows for integration with actual transit data between different municipalities. At the current time, we do not have access to transit data from major Norwegian mobile phone network providers. Since many travel restrictions have been in place and some are still present, we do not believe this will change the patterns we have uncovered. However, this is an extension we are looking into for the near future. When travel restrictions or recommendations are becoming less constraining, we will be dependent on such transit data to be able to make robust recommendations.

Table 1. Parameters used

Model parameters	Symbol	Value	Function	Source
SEIR-epidemics				
Probability of infection	β	----	Network effect	
Days incubation time	λ_E	1	Fixed	
Days spent pre-symptomatic	λ_{ps}	5	Poisson	FHI model, adjusted for reduced incubation time
Days symptomatic before recovery	λ_{sR}	5	Poisson	FHI model
Days symptomatic before hospitalization	λ_{sH}	6	Poisson	Data from HSØ
Days symptomatic in nursing home before death	λ_{ND}	10	Poisson	Data from HSØ
Days in hospital before recovery (no ICU)	λ_{HR}	8	Poisson	Data from HSØ
Days in hospital before ICU	λ_{HI}	4	Poisson	Data from HSØ
Days in ICU before recovery	λ_{IR}	12	Poisson	Data from HSØ
Days in ICU before death	λ_{ID}	12	Poisson	Data from HSØ
Days asymptomatic before recovery	λ_{aR}	8	Poisson	FHI model adjusted for reduced incubation time
% exposed developing symptoms	P_i	50	Bernoulli	
% symptomatic dying outside of hospital	P_{ND}		Bernoulli	
nursing home residents 70-79 years		26		Adjusted to Norwegian hosp. death rates, Verity et al
nursing home residents 80-89 years		42		Adjusted to Norwegian hosp. death rates, Verity et al
All others		0		
% hospitalized dying	P_{HD}		Bernoulli	Verity et al, Lancet, 2020
0-9 years		1.61 e-3		
10-19 years		6.95 e-3		
20-29 years		3.09 e-2		
30-39 years		8.44 e-2		
40-49 years		0,161		
50-59 years		0,595		
60-69 years		1,93		
70-79 years		4,28		
80+ years		7,8		
% symptomatic being hospitalized	P_{sH}		Bernoulli	Verity et al, Lancet, 2020
0-9 years		0		
10-19 years		0,048		
20-29 years		1,04		
30-39 years		3,43		
40-49 years		4,25		
50-59 years		8,16		
60-69 years		11,8		
70-79 years		16,6		
80+ years		18,4		
% hospitalized needing ICU	P_{HI}	30	Bernoulli	Fitted to FHI ICU numbers
% not developing immunity	P_{RS}	0	Bernoulli	
Individual-based network model				
Infectiousness in Day Care		0,015%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness in Primary School		0,005%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness Secondary School		0,015%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness High School		0,015%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness Household		30%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness Work		0,015%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness Nursing Home		20,0%	Bernoulli	estimated IBM fit to Norwegian clinical data
Infectiousness Generic Contact		0,75 %	Bernoulli	estimated IBM fit to Norwegian clinical data

References

- J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings and M. Diehl. CasADi: A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*. Volume 11, Issue 1, pp. 1-36, Springer, 2019.
- Casadi library for optimal control. URL: <https://web.casadi.org/get/>
- Bonmin (Basic Open-source Nonlinear Mixed INteger programming). Solver for MIO-problems. URL: <https://projects.coin-or.org/Bonmin>
- F. Borrelli, A. Bemporad and M. Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.
- E. H. Bussell, C. E. Dangerfield, C. A. Gilligan and N. J. Cunniffe. Applying optimal control theory to complex epidemiological models to inform real-world disease management. *Philosophical Transactions of the Royal Society B: Biological Sciences*. Volume 374, Issue 1776, July 2019.
- N. M. Ferguson et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. *Imperial College COVID-19 Response Team*, 16 March 2020.
URL: <https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf>
- T. A. Johansen. Introduction to nonlinear model predictive control and moving horizon estimation. In: *Selected Topics on Constrained and Nonlinear Control*. STU/NTNU, ISBN 978-80-968627-4-0, 2011.
- M. Morari and U. Maeder. Nonlinear offset-free model predictive control. *Automatica*, Vol. 48, Issue 9, pp. 2059-2067, September 2012.

- C. Nowzari, V. M. Preciado and G. J. Pappas. Analysis and control of epidemics: a survey of spreading processes on complex networks, *IEEE Control Systems*, Vol. 36, pp. 26-46, February 2016.
- N. J. Watkins, C. Nowzari and G. J. Pappas. Robust economic model predictive control of continuous-time epidemic processes. *IEEE Transactions on Automatic Control*, Vol. 65, Issue 3, pp. 1116-1131, May 2019.