



# Koopman-based surrogate models for multi-objective optimization of agent-based systems

Jan-Hendrik Niemann<sup>a,b,\*</sup>, Stefan Klus<sup>c</sup>, Nataša Djurdjevic Conrad<sup>a</sup>, Christof Schütte<sup>a,b</sup>

<sup>a</sup> Department of Mathematics and Computer Science, Freie Universität Berlin, Germany

<sup>b</sup> Zuse Institute Berlin, Germany

<sup>c</sup> School of Mathematical & Computer Sciences, Heriot-Watt University, UK

## ARTICLE INFO

### Keywords:

Multi-objective optimization  
Agent-based models  
Data-driven model reduction  
Koopman operator theory  
Optimal control  
Social dynamics

## ABSTRACT

Agent-based models (ABMs) provide an intuitive and powerful framework for studying social dynamics by modeling the interactions of individuals from the perspective of each individual. In addition to simulating and forecasting the dynamics of ABMs, the demand to solve optimization problems to support, for example, decision-making processes naturally arises. Most ABMs, however, are non-deterministic, high-dimensional dynamical systems, so objectives defined in terms of their behavior are computationally expensive. In particular, if the number of agents is large, evaluating the objective functions often becomes prohibitively time-consuming. We consider data-driven reduced models based on the Koopman generator to enable the efficient solution of multi-objective optimization problems involving ABMs. In a first step, we show how to obtain data-driven reduced models of non-deterministic dynamical systems (such as ABMs) that depend potentially nonlinearly on control inputs. We then use them in the second step as surrogate models to solve multi-objective optimal control problems. We first illustrate our approach using the example of a voter model, where we compute optimal controls to steer the agents to a predetermined majority, and then using the example of an epidemic ABM, where we compute optimal containment strategies in a prototypical situation. We demonstrate that the surrogate models effectively approximate the Pareto-optimal points of the ABM dynamics by comparing the surrogate-based results with test points, where the objectives are evaluated using the ABM. Our results show that when objectives are defined by the dynamic behavior of ABMs, data-driven surrogate models support or even enable the solution of multi-objective optimization problems.

## 1. Introduction

Modeling social dynamics and studying the resulting collective phenomena is an important research problem, for instance, in the field of epidemic modeling, opinion dynamics, mobility, or innovation spreading in ancient times [1–4]. Agent-based models (ABMs) provide an intuitive yet powerful framework for modeling the interactions of individuals, small groups, or entire populations. The high level of modeling flexibility allows users to gain important insights into the complex dynamic patterns that emerge from the interactions of discrete entities called *agents*. These agents often follow simple rules that describe their behavior, making ABMs accessible to both experts and non-experts without extensive mathematical knowledge. Agent-based modeling is particularly useful in modeling complex systems where interactions and environmental factors play a significant role, where heterogeneity among agents is desired or required, or where agents can learn and adapt to new situations. In addition to simulating and forecasting the dynamics of complex ABMs, optimization problems such

as optimal control need to be solved. One important use case is to support decision-making processes.

Optimal control is the process of determining the best set of actions or inputs to a dynamical system over a given time horizon in order to achieve a desired objective while satisfying given constraints. There exist well-established numerical methods to solve complex, nonlinear optimization problems such as line search, conjugate gradient, or trust region methods [5,6]. The optimization of ABMs, however, leads to significant challenges since basic concepts such as derivatives are not well-defined due to the fact that most ABMs are inherently stochastic and discontinuous. Additionally, most ABMs are high-dimensional systems comprising thousands of agents, which is often not only desirable but also necessary, e.g., to be able to correctly represent even small fractions of the total population. Furthermore, many independent simulations are often required to compute statistical properties such as means and variances. This renders the simulation of ABMs expensive

\* Corresponding author at: Department of Mathematics and Computer Science, Freie Universität Berlin, Germany.  
E-mail address: [niemann@zib.de](mailto:niemann@zib.de) (J.-H. Niemann).

and time-consuming and a more thorough analysis including optimization and also sensitivity analysis or uncertainty quantification nearly impossible.

In this work, we will focus on so-called *multi-objective optimization problems* involving ABMs. These problems arise naturally whenever multiple objectives are to be optimized simultaneously, without any predetermined prioritization. There exist various methods and techniques to solve multi-objective optimization problems, such as scalarization,  $\epsilon$ -constraint methods, evolutionary algorithms, particle swarm, agent-based and multi-agent methods, or set-oriented methods, see [7–14].

In the setting of ABMs, one way to solve multi-objective optimization problems is to use heuristic methods such as evolutionary algorithms. These methods do not rely on derivatives and are simple to use, see, e.g., [15] and references therein, and are commonly used for an automatic or guided calibration of ABMs, for example, to fit the model against real-world financial market data [16] or biological observations [17,18]. Multi-objective calibration can also provide insights into parameters or initial conditions that may not be evident through simulations alone, see [19–24] for recent studies focusing on multi-objective calibration of ABMs using heuristic methods. We refer the reader to [25] for a comprehensive review of calibration techniques. However, the lack of (mathematical) convergence properties and the often costly simulation requirements are drawbacks of heuristic methods for solving multi-objective optimization problems involving ABMs.

Surrogates, either for individual objectives or for the complete model, provide a way to reduce the computational effort when optimizing ABMs. In this sense, a *surrogate* refers to a viable substitute for the original model that has a sufficient level of accuracy and can be evaluated significantly faster, often by several orders of magnitude. Interpolation, regression, or machine learning can be used to quickly obtain a surrogate for expensive objectives [10,26]. However, when model parameters are changed, these surrogates need to be recalculated. Instead of replacing expensive objective functions, the entire model itself can be replaced by a surrogate model. In [27], a surrogate model that approximates the mapping between ABM inputs and the corresponding response in the output is learned using non-parametric machine learning for simulation, calibration, and exploration of the parameter space. Further approaches include, e.g., surrogate models based on difference equations [28–30], partial differential equations [31], or ordinary differential equations (ODEs) [32] to solve multi-objective optimization problems involving ABMs. See also [10] for a recent review on surrogate modeling.

In this work, we use a data-driven method to find suitable surrogate models. Data-driven methods prove to be particularly advantageous for this task as they can extract valuable insights about the behavior of dynamical systems solely from data, without relying on prior knowledge about the system. This makes these methods particularly well-suited to address problems where a closed description of the system is not available. Various methods have been developed in recent years, e.g., for approximating transfer operators associated with the system [33], for detecting metastable and coherent sets [34], or for performing stability analyses [35], as well as for deriving the governing equations of the underlying dynamics [36–38], model reduction, optimization, and control [39–43]. Following [44], we use the infinitesimal generator of the Koopman operator, which is the adjoint of the Perron–Frobenius operator associated with the dynamical system and describes the evolution of observable functions representing any kind of measurement, to obtain a reduced model from ABM simulation data, which is then used as a surrogate model to solve a multi-objective optimization problem. The reduction of the numerical effort is achieved in two ways: First, a data-driven method allows us to obtain a dynamical system that models only the required quantities. In our particular case these are the aggregated dynamics of an ABM, i.e., the collective behavior of the agents. This corresponds to a dimension reduction. Second, the reduced dynamical model, which serves as a surrogate model, is given in terms of differential equations, which can in general be simulated more efficiently than ABMs. The main contributions of this work are:

- We show how *generator extended dynamic mode decomposition* (gEDMD) [38], a data-driven method that approximates the infinitesimal generator of the Koopman operator, can be extended to identify the governing equations of non-deterministic dynamical systems with potentially nonlinear dependence on the control inputs.
- We demonstrate how these reduced models, obtained from ABM simulation data, can be used as surrogate models to enable the efficient solution of multi-objective optimal control problems.
- We consider two different use cases, namely linear and nonlinear dependence on the control inputs for two ABMs, and show for both ABMs that the Pareto sets computed using the data-driven surrogate models indeed approximate the Pareto sets of the ABMs.

In this study, we showcase the potential of data-driven surrogate models to efficiently address the challenges posed by multi-objective optimization problems associated with the dynamical behavior of ABMs. In particular when limit models for large numbers of agents are unknown or non-existent, data-driven surrogate models offer a viable solution for tackling optimization problems that would otherwise be computationally infeasible due to the high computational cost of evaluating the objectives. Our method requires a significant amount of data, but in simulation studies where a surrogate model is needed for optimization or control, data availability is typically not an issue.

The remainder of this paper is organized as follows. In Section 2 we introduce the stochastic Koopman operator and its generator as well as multi-objective optimization problems. We then show in Section 3 how gEDMD can be used to obtain data-driven reduced models with linear and nonlinear dependence on the control inputs. We demonstrate in Section 4 how both approaches can be used to find surrogate models to solve multi-objective optimal control problems for high-dimensional ABMs. We consider the voter model and an ABM modeling the transmission dynamics of SARS-CoV-2. Open questions and future work will be discussed in Section 5.

## 2. Preliminaries

We will first introduce the required mathematical tools and provide a brief overview of the stochastic Koopman operator and its generator as well as multi-objective optimization problems.

### 2.1. The Koopman operator and its generator

Let  $\mathbb{X} \subset \mathbb{R}^d$  be the state space. We consider stochastic differential equations (SDEs) with (time-varying) control input  $u: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{d_u}$  of the form

$$dX_t = b(X_t, u) dt + \sigma(X_t, u) dW_t, \quad (1)$$

where  $\{X_t\}_{t \geq 0} \in \mathbb{X}$  is a time-homogeneous stochastic process and  $b: \mathbb{R}^d \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^d$  denotes the drift term,  $\sigma: \mathbb{R}^d \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d \times s}$  the diffusion term, and  $W_t$  an  $s$ -dimensional Wiener process. Let  $\Phi^t$  denote the associated flow map and  $f \in L^\infty(\mathbb{X})$  a real-valued observable of the system representing any kind of measurement. Assuming that  $u$  is constant, the semigroup  $\{\mathcal{K}_u^t\}_{t \geq 0}$  of Koopman operators  $\mathcal{K}_u^t: L^\infty(\mathbb{X}) \rightarrow L^\infty(\mathbb{X})$  is defined by

$$(\mathcal{K}_u^t f)(x) = \mathbb{E}[f(\Phi^t(x, u)) \mid X_t = x],$$

i.e., the conditional expectation of  $f(\Phi^t(x, u))$  given  $X_t = x$ . The Koopman operator is an infinite-dimensional, linear and non-expansive operator forming a contraction semigroup. Provided that  $f$  is twice continuously differentiable, it can be shown using Itô's lemma that the infinitesimal generator  $\mathcal{L}_u$  of the Koopman operator can be characterized by

$$(\mathcal{L}_u f)(x) = \sum_{i=1}^d b_i(x, u) \frac{\partial f}{\partial x_i} + \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d a_{ij}(x, u) \frac{\partial^2 f}{\partial x_i \partial x_j},$$

where  $a = \sigma \sigma^\top$ . We refer to it as the *Koopman generator*, see [45] for details. The function  $v(t, x) := (\mathcal{K}_u^t f)(x)$  satisfies  $\frac{\partial v}{\partial t} = \mathcal{L}_u v$ , which is a second-order partial differential equation commonly known as the *Kolmogorov backward equation*, see [46]. For deterministic dynamical systems, i.e.,  $\sigma \equiv 0$ , we obtain a first-order partial differential equation, the Liouville equation. See [38] for details.

**Remark 2.1.** We denote the Koopman operator and its generator for an uncontrolled system of the form (1), i.e.,  $u \equiv 0$ , by  $\mathcal{K}_0^t$  and  $\mathcal{L}_0$ , respectively.

### 2.1.1. Generator extended dynamic mode decomposition

Due to the infinite-dimensional nature of the Koopman operator, it is common practice to consider projections onto finite-dimensional subspaces. We briefly introduce *generator extended dynamic mode decomposition* (gEDMD), a method which approximates the Koopman generator. For details, see [38]. We will assume for now that  $u \equiv 0$  and omit any subscripts to simplify the notation. Given  $m$  measurements of the system's state  $\{x_i\}_{i=1}^m$ , its drift  $\{b(x_i)\}_{i=1}^m$ , and diffusion  $\{\sigma(x_i)\}_{i=1}^m$ , as well as a set of basis functions  $\{\psi_k\}_{k=1}^\ell$ , which can be written in vector form as  $\psi(x) = [\psi_1(x), \dots, \psi_\ell(x)]^\top$ , we define

$$d\psi_k(x) = (\mathcal{L}\psi_k)(x) = \sum_{i=1}^d b_i(x) \frac{\partial \psi_k}{\partial x_i}(x) + \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d a_{ij}(x) \frac{\partial^2 \psi_k}{\partial x_i \partial x_j}(x)$$

for  $k = 1, \dots, \ell$ . We then construct the matrices  $\Psi_X, d\Psi_X \in \mathbb{R}^{\ell \times m}$  for all measurements and basis functions, i.e.,

$$\Psi_X = \begin{bmatrix} \psi_1(x_1) & \dots & \psi_1(x_m) \\ \vdots & \ddots & \vdots \\ \psi_\ell(x_1) & \dots & \psi_\ell(x_m) \end{bmatrix} \quad \text{and} \quad d\Psi_X = \begin{bmatrix} d\psi_1(x_1) & \dots & d\psi_1(x_m) \\ \vdots & \ddots & \vdots \\ d\psi_\ell(x_1) & \dots & d\psi_\ell(x_m) \end{bmatrix}, \quad (2)$$

and assume that  $d\Psi_X = M\Psi_X$ . In general, this problem cannot be solved exactly such that we solve it in the least-square sense by minimizing  $\|d\Psi_X - M\Psi_X\|_F$ , where  $\|\cdot\|_F$  denotes the Frobenius norm. The solution is given by  $M = d\Psi_X \Psi_X^+$ , where  $\Psi_X^+$  denotes the Moore–Penrose pseudoinverse of  $\Psi_X$ . The matrix  $L = M^\top$  is an empirical estimate of the matrix representation of the infinitesimal generator  $\mathcal{L}$ . The convergence of gEDMD in the infinite data limit to a Galerkin approximation of the generator, i.e., a projection onto the space spanned by the basis functions, was shown in [38].

### 2.1.2. System identification

Assuming that  $\mathbb{X}$  is bounded so that the full-state observable  $g(x) = x$  is contained component-wise in  $L^\infty(\mathbb{X})$ , we can use the observable to reconstruct the governing equations of the underlying dynamical system. We make the assumption that  $g(x) = x$  can be represented by the basis functions  $\psi$ , which can easily be accomplished by adding the observables  $\{x_i\}_{i=1}^d$  to the set of basis functions. The system can directly be represented in terms of the basis functions,

$$(\mathcal{L}g)(x) = b(x) \approx (LB)^\top \psi(x),$$

where we define  $B \in \mathbb{R}^{\ell \times d}$  such that  $g(x) = B^\top \psi(x)$ . For deterministic dynamical systems, this is equivalent to SINDy [47]. For non-deterministic systems and for  $\psi_k(x) = x_i x_j$ , we identify the diffusion term using

$$a_{ij}(x) \approx (\mathcal{L}\psi_k)(x) - b_i(x)x_j - b_j(x)x_i,$$

where we assume that  $b_i$  and  $b_j$  as well as  $b_i(x)x_j$  and  $b_j(x)x_i$  are contained in the space spanned by the basis functions. If it is necessary to obtain the drift term  $\sigma$ , it can be obtained by computing a Cholesky decomposition of  $a$ . For details, see [38].

## 2.2. Multi-objective optimization

*Multi-objective optimization* concerns the simultaneous optimization of  $k$  – often contradictory – objective functions  $f_1, \dots, f_k : \mathbb{R}^n \rightarrow \mathbb{R}$ . For example, one might want to minimize the spread of an infectious disease with curfews while at the same time also trying to minimize the socioeconomic cost of this intervention. A decision maker must find a compromise, i.e., an agreement reached through mutual concessions that often alter or combine the original goals. An optimal compromise cannot be further improved without worsening at least one of the other goals. This is called *Pareto-optimal*. In contrast to single-objective optimization, where conflicting objectives are prioritized a priori, e.g., by a weighted sum of the objectives, multi-objective optimization prioritizes a posteriori. This means that if we know Pareto-optimal points, we can select a solution according to the preference of each objective.

We consider multi-objective optimization problems of the form

$$\begin{aligned} \min_{y \in \mathbb{R}^n} \quad & F(y) = [f_1(y), \dots, f_k(y)]^\top, \\ \text{s.t.} \quad & g_i(y) \leq 0, \quad i = 1, \dots, q, \\ & h_j(y) = 0, \quad j = 1, \dots, p, \end{aligned} \quad (3)$$

where  $F : \mathbb{R}^n \rightarrow \mathbb{R}^k$ ,  $g : \mathbb{R}^n \rightarrow \mathbb{R}^q$  and  $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$  for  $g := [g_1, \dots, g_q]^\top$  and  $h := [h_1, \dots, h_p]^\top$ . The space of all feasible decisions is called the *feasible decision space*  $\mathcal{R}$  and is given by the constraints in (3), i.e.,

$$\mathcal{R} = \{y \in \mathbb{R}^n \mid g(y) \leq 0, h(y) = 0\},$$

where  $\leq$  is defined component-wise. A decision  $y^* \in \mathcal{R}$  is called (*globally*) *Pareto-optimal* or (*global*) *Pareto point* if there is no  $y \in \mathcal{R}$  that dominates  $y^*$ , i.e.,  $F(y) \neq F(y^*)$  and  $F(y) \leq F(y^*)$  hold component-wise. A decision  $y^*$  is called *locally Pareto-optimal* or *local Pareto point* if there exists  $\epsilon > 0$  such that  $y^*$  is Pareto-optimal in a neighborhood  $y \in U_\epsilon(y^*) \subset \mathcal{R}$ . The set of all Pareto-optimal points is called *Pareto set* and its image under the map of the objective functions *Pareto front* [8,14].

**Example 2.2.** Consider the multi-objective optimization problem (3) from [13] with objective functions  $f_i : \mathbb{R} \rightarrow \mathbb{R}$  given by

$$\begin{aligned} f_1(y) &= (y - 1.5)^2, \\ f_2(y) &= y^4 - 4y^3 + 4y^2 \end{aligned}$$

and feasible decision space  $\mathcal{R} = [-0.5, 2.5]$ . Fig. 1(a) shows both objective functions. The set of locally Pareto-optimal points is given by  $[0, 1] \cup [1.5, 2]$ . Fig. 1(b) shows not only the Pareto front as dotted and dashed line segments, but also that only the interval  $[1.5, 2]$  is globally Pareto-optimal, as every point in  $[1.5, 2]$  clearly dominates all points in the other interval.  $\triangle$

We will use a set-oriented method known as *sampling algorithm* [13], which computes a box-covering of the Pareto set. The sampling algorithm is an iterative two-step process, which in a first step subdivides a collection of boxes with respect to one coordinate and then discards every box that does not pass a set-wise non-dominance test, i.e., boxes that only contain dominated points. Fig. 2 shows how the boxes and the Pareto front covering evolve using the optimization problem introduced in Example 2.2. A complete description of the sampling algorithm can be found in Appendix A. For further details, see also [8,13,14,48].

Each box can be efficiently represented by a center and a radius, so that all collections can be stored in a binary tree whose memory consumption grows linearly with  $n$ , see [49] for the MATLAB toolbox GAIO, which contains binary tree data structures and the algorithms for set-oriented calculations.

## 3. Koopman-based surrogate models with control

For the following discussion, we consider non-deterministic dynamical systems with control of the form (1). The aim of this section is to learn the Koopman generator associated with such systems. We consider two cases: (i) we assume that the control  $u$  acts linearly on (1), and (ii) we will consider the more general case, i.e.,  $u$  acts nonlinearly on (1).

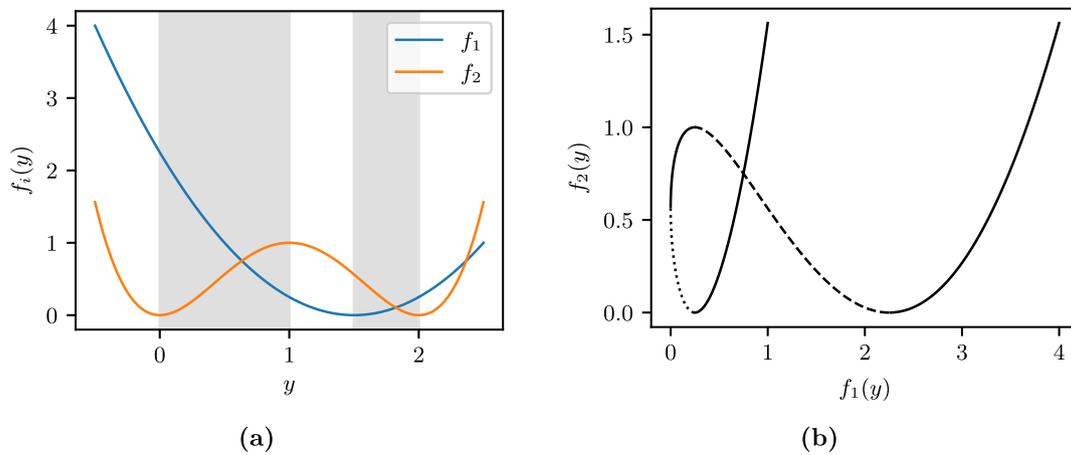


Fig. 1. (a) Objective functions  $f_1$  and  $f_2$  on the interval  $R = [-0.5, 2.5]$  with Pareto set (gray shaded) and (b) Pareto front (dotted and dashed). The dotted line segment shows that only points in  $[1.5, 2]$  are globally Pareto-optimal since they clearly dominate all points in  $[0, 1]$ .

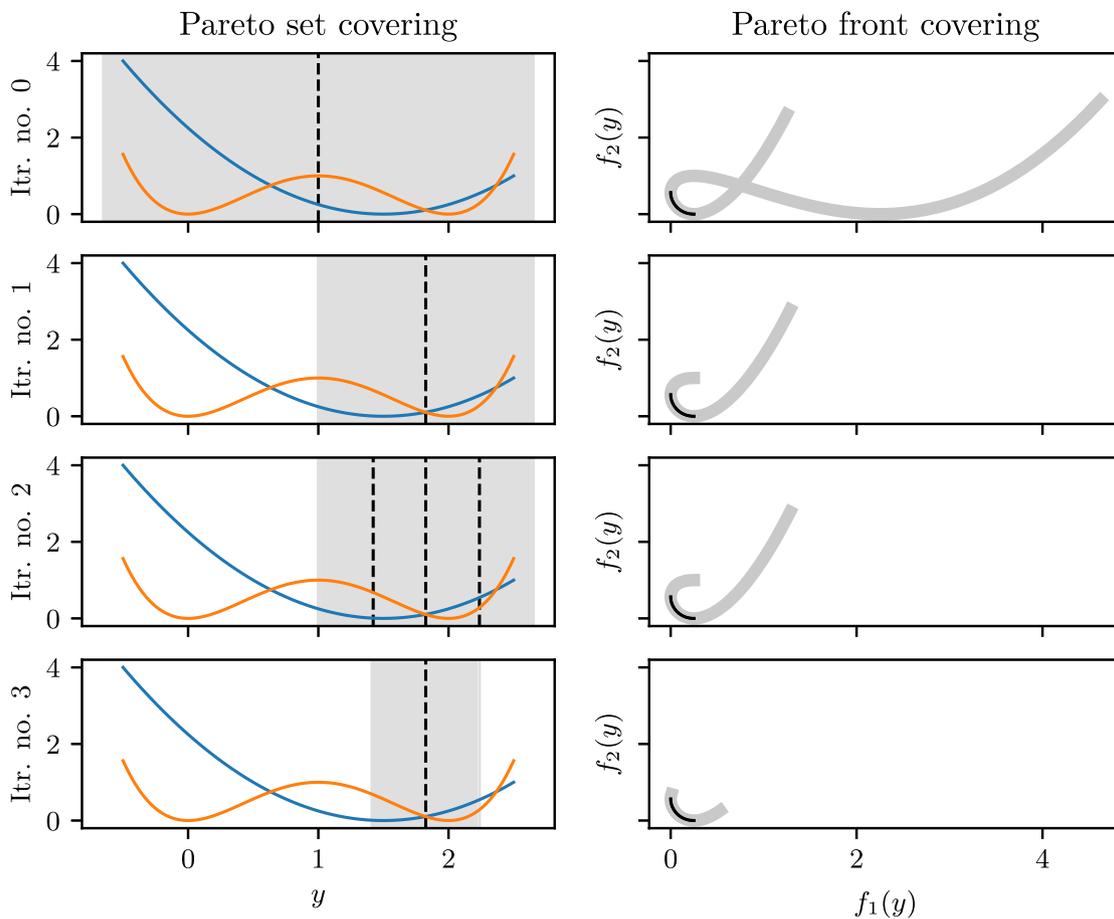


Fig. 2. First four iterations of the sampling algorithm demonstrated using Example 2.2. The dashed lines represent the boundaries of each box. In this example, each box is split in half before a non-dominance test is performed. The plots on the right show the images of the current box collections (gray/shaded) covering the true global Pareto front (black/solid).

### 3.1. Linear case: Generator interpolation

We extend the recent result in [43] for deterministic control-affine dynamical systems to *non-deterministic* control-affine dynamical systems.

**Definition 3.1.** A non-deterministic dynamical system of the form (1) is called *control-affine* if

$$b(X_t, u) = b_0(X_t) + \sum_{i=1}^{d_u} b_i(X_t) u_i, \tag{4}$$

$$a(X_t, u) = a_0(X_t) + \sum_{i=1}^{d_u} a_i(X_t) u_i, \quad (5)$$

where  $u_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$  and  $a = \sigma \sigma^\top$ .

**Theorem 3.2.** *Given a space of twice continuously differentiable functions and controls  $u_1, u_2$ , if the dynamics (1) are control-affine, then the Koopman generators are control-affine, i.e.,*

$$\mathcal{L}_{\alpha_1 u_1 + \alpha_2 u_2} = \mathcal{L}_0 + \alpha_1 \mathcal{A}_{u_1} + \alpha_2 \mathcal{A}_{u_2},$$

where  $\mathcal{A}_u = \mathcal{L}_u - \mathcal{L}_0$  and  $\alpha_1, \alpha_2 \in \mathbb{R}$ .

The proof of Theorem 3.2 is analogous to the proof for deterministic control-affine systems, which can be found in, e.g., [43]. For the sake of completeness it is included in Appendix B.

**Example 3.3.** Consider the nonlinear control-affine system

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} (\gamma + g(u))x_1 \\ \delta(x_2 - x_1^2) \end{bmatrix}, \quad (6)$$

where in this example  $u \equiv \text{const}$  and  $g$  is a function. Defining  $f_1 := x_1$ ,  $f_2 := x_2$ , and  $f_3 := x_1^2$ , the system has a finite-dimensional, linear representation

$$\frac{d}{dt} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} = \underbrace{\begin{bmatrix} \gamma + g(u) & 0 & 0 \\ 0 & \delta & -\delta \\ 0 & 0 & 2(\gamma + g(u)) \end{bmatrix}}_{=L_u} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}, \quad (7)$$

where  $L_u$  is a finite-dimensional representation of the generator  $\mathcal{L}_u$ . Following Theorem 3.2, if  $g$  is linear,  $L_u$  can be split into

$$L_u = \underbrace{\begin{bmatrix} \gamma & 0 & 0 \\ 0 & \delta & -\delta \\ 0 & 0 & 2\gamma \end{bmatrix}}_{=L_0} + \underbrace{\begin{bmatrix} g(u) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 2g(u) \end{bmatrix}}_{=A_u},$$

where  $A_u$  is a finite-dimensional representation of  $\mathcal{A}_u$  and linear with respect to  $u$ . Thus,  $L_u$  is control-affine.  $\triangle$

### 3.2. Nonlinear case: State augmentation

To motivate what follows, take again a look at Example 3.3 and assume that  $g$  is nonlinear. The system (6) can still be represented as in (7), however, nonlinearity of  $g$  prohibits to use Theorem 3.2 as the operator  $A_u$  is not linear with respect to  $u$ . A workaround is to augment the system so that the control  $u$  is represented as an additional state, i.e., the augmented system is given by

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ u \end{bmatrix} = \begin{bmatrix} (\gamma + g(u))x_1 \\ \delta(x_2 - x_1^2) \\ 0 \end{bmatrix}.$$

A closed representation of the augmented system as in (7) for the non-augmented system (6), however, is not possible any longer. We will now show how gEDMD can be extended to control inputs. The idea of augmenting the state space has also been used in other data-driven methods such as DMDc [50] or EDMDC [39], both of which provide approximations of the Koopman operator using linear and nonlinear basis functions, respectively, or SINDYc [51], which can be used to find the governing equations of a dynamical system with control.

For a general system (1), we assume that we have a set of  $m$  measurements of the augmented system state  $\{[x_t, u_t]\}_{t=1}^m$  as well as  $\{\dot{u}_t\}_{t=1}^m$ , which is trivial as we have full access to  $u$ , i.e., we can choose its values and thus know its derivative. Let  $[x_t, u_t] = [x_{t1}, \dots, x_{td}, u_{t1}, \dots, u_{td_u}] \in \mathbb{R}^{d+d_u}$ . To simplify the notation, let  $\bar{x}_t := [x_t, u_t]$ . Additionally, assume that we have a set of the augmented drift

$\{[b(\bar{x}_t), \dot{u}_t]\}_{t=1}^m$  as well as the augmented diffusion  $\{\sigma(\bar{x}_t, u_t)\}_{t=1}^m$ , which is defined as

$$\sigma(\bar{x}_t, u_t) := \begin{bmatrix} \sigma_{12}(u_t) & \sigma_{22}(u_t) \\ \sigma_{21}(u_t) & \sigma_{22}(u_t) \end{bmatrix} \in \mathbb{R}^{(d+d_u) \times (d+d_u)},$$

where  $\sigma_{12}(u_t) \in \mathbb{R}^{d \times d_u}$ ,  $\sigma_{21}(u_t) \in \mathbb{R}^{d_u \times d}$  and  $\sigma_{22}(u_t) \in \mathbb{R}^{d_u \times d_u}$  denote the diffusion terms of the control  $u$  (typically zero). We assume that either both drift  $b(\bar{x}_t)$  and diffusion  $\sigma(\bar{x}_t)$  are known or that they can be estimated pointwise. Let  $\{\psi_k\}_{k=1}^{\ell}$  denote the set of basis functions and define

$$d\psi_k(\bar{x}) := (\mathcal{L}\psi_k)(\bar{x}) = \sum_{i=1}^{d+d_u} b_i(\bar{x}) \frac{\partial \psi_k}{\partial \bar{x}_i}(\bar{x}) + \frac{1}{2} \sum_{i=1}^{d+d_u} \sum_{j=1}^{d+d_u} a_{ij}(\bar{x}) \frac{\partial^2 \psi_k}{\partial \bar{x}_i \partial \bar{x}_j}(\bar{x}),$$

where  $a = \sigma(\bar{x}_t, u_t) \sigma(\bar{x}_t, u_t)^\top$ . The partial derivatives of the basis functions can be precomputed analytically. We can then compute the gEDMD matrices  $\Psi_X$  and  $d\Psi_X$  in (2) for  $\bar{x}_1, \dots, \bar{x}_m$  and solve the minimization problem  $\min \|d\Psi_X - M\Psi_X\|_F$  to obtain a finite-dimensional empirical estimate of the generator  $L = M^\top$ .

Note that we make three important restrictions on the feasible controls: (i) the controls are given by families of solutions of ODEs, (ii) the controls act on long time scales only (e.g., bang-bang controls are not allowed), and (iii) the controls are representable in terms of the basis functions. However, these constraints do not limit our approach for the following reasons: (i) in most scenarios, controls are typically given in such a form, (ii) rapidly changing controls are unrealistic in the context of ABMs, and (iii) insufficient and inappropriately chosen basis functions lead to inaccurate results.

**Example 3.4.** To illustrate how the state augmentation works, consider a stochastic SIR model to simulate an infectious disease and its containment with non-pharmaceutical interventions. For simplicity, we consider a population of size  $N$  and assume that an infection occurs at rate  $\beta(u)$ , which depends on the disease itself and the containment strategy  $u(t)$ . Once infected, individuals recover at rate  $\gamma$  and cannot be re-infected. Defining the system state  $X_t := [S_t, I_t]^\top$ , where  $S_t$  and  $I_t$  denote the numbers of susceptible and infected individuals, respectively, the stochastic SIR model is given by

$$\begin{aligned} dS_t &= -\beta(u) \frac{S_t I_t}{N} dt - \frac{1}{N} \sqrt{\beta(u) S_t I_t} dW_1(t), \\ dI_t &= \left[ \beta(u) \frac{S_t I_t}{N} - \gamma I_t \right] dt + \frac{1}{N} \sqrt{\beta(u) S_t I_t} dW_2(t) - \sqrt{\gamma N^{-1} I_t} dW_3(t). \end{aligned}$$

The number of recovered individuals is given by  $R_t = N - S_t - I_t$  since the size of the population is assumed to be constant.

It has been known – not only since the COVID-19 pandemic – that non-pharmaceutical interventions such as social distancing, wearing face masks, or lockdowns can slow or prevent the spread of infectious diseases. Such interventions can be expected to have a nonlinear effect on the dynamics because the probability of infection changes nonlinearly with the viral load in the air to which a person is exposed [52]. For this toy model we assume that  $u$  acts quadratically via  $\beta(u) = \tilde{\beta}(1 - u(t))^2$  for time  $t \geq 0$ . The system is not control-affine. Thus, Theorem 3.2 cannot be applied. State augmentation, on the other hand, can be applied. We assume that the control  $u$  is a continuously relaxed intervention of the form  $u(t) = A Q / (1 + e^{Bt})$ . We set  $\tilde{\beta} = 0.5$ ,  $\gamma = 0.05$ ,  $A = 0.5$ ,  $Q = 1000$ , and  $B = 0.1$ . Fig. 3(a) shows the trajectories for susceptible and infected individuals as a fraction of the population without (solid) and with control  $u$  (dashed). The control  $u(t)$  is shown in Fig. 3(b).

To compute a representation of the Koopman generator, we augment the system state by the control  $u$  and generate  $m = 1000$  uniformly distributed training points  $\{[x_t, u_t]\}$ , where  $x_t := [S_t/N, I_t/N]^\top$ . We use exact values for the augmented drift  $\{[b(x_t, u_t), \dot{u}_t]\}_{t=1}^m$  as well as the augmented diffusion  $\{\sigma(x_t, u_t)\}_{t=1}^m$ . For the basis functions we choose

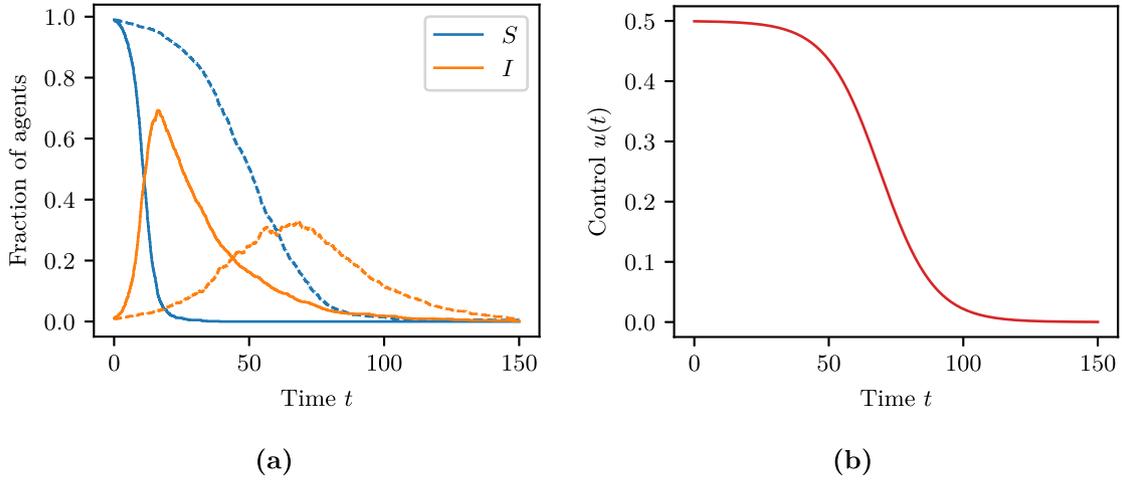


Fig. 3. (a) Trajectories of susceptible and infected agents as a fraction of the population for the SIR model with (dashed) and without (solid) control  $u$  and (b) control  $u(t)$  applied to SIR model.

monomials up to degree 5. The system is correctly identified with drift and diffusion terms given by

$$b(x, u) = \begin{bmatrix} -\tilde{\beta} x_1 x_2 (1-u)^2 \\ \tilde{\beta} x_1 x_2 (1-u)^2 - \gamma x_2 \\ B u + \frac{B}{A} u^2 \end{bmatrix}$$

and

$$a(x, u) = \frac{1}{N} \begin{bmatrix} \tilde{\beta} x_1 x_2 (1-u)^2 & -\tilde{\beta} x_1 x_2 (1-u)^2 & 0 \\ -\tilde{\beta} x_1 x_2 (1-u)^2 & \tilde{\beta} x_1 x_2 (1-u)^2 + \gamma x_2 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Since in general we are only interested in identifying the dynamics of  $X_t = [S_t, I_t]^T$  and not the dynamics of the control  $u$  (which is actually known), we could relax the third restriction concerning the representation in terms of the basis function and neglect the columns in  $L$  corresponding to  $u$ . For further details on system identification and model reduction using gEDMD, see [38].  $\Delta$

#### 4. A multi-objective optimization approach for ABMs

The aim in this section is to find the Pareto set of  $n$  controls  $u_i$  that optimize  $k$  objective functions  $f_i : \mathcal{R} \rightarrow \mathbb{R}$  in (3) depending on the state of an ABM. We assume that each function  $f_i$  can be written as

$$f_i(X, u) = \mathbb{E} \left[ \int_{t_0}^{t_1} r_i(X(t; u), u(t)) dt + s_i(X(t_1, (u(t_1)))) \right], \quad t_0 < t_1,$$

where  $r_i$  and  $s_i$  are the running and terminal cost functions depending on state  $X(t; u)$  of an ABM and  $u$  is the control. An optimal control  $u^*$  is given by a solution of the multi-objective optimization problem

$$\min_{u \in U} F(X, u),$$

where  $F(X, u) = [f_1(X, u), \dots, f_k(X, u)]^T$  and  $U = \{u : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{d_u} \mid u(\cdot) \text{ measurable}\}$  the space of admissible controls. Three major issues occur when optimizing ABMs:

- (i) **Computational complexity:** Most ABMs involve thousands of agents and simulate complex interactions between agents, often at high time resolution, making the calculation of single trajectories computationally expensive.
- (ii) **Stochasticity:** Most ABMs are inherently stochastic so that many independent simulations are required to compute statistical properties such as  $\mathbb{E}[\cdot]$ .
- (iii) **Discontinuity:** Most ABMs are inherently discontinuous since agents make discrete decisions, and thus single trajectories are not differentiable with respect to  $u$ .

By replacing the ABM with a surrogate model that is less expensive to evaluate and directly approximates expectations, we address the problems of computational complexity and stochasticity. It has recently been shown in [53] that under certain conditions it is possible to obtain reduced models, given by differential equations, from ABM data that accurately approximate the aggregated dynamics of ABMs, i.e., the collective behavior of the agents. That is, instead of dealing with the ABM state space whose size grows combinatorially as  $d^N$ , where  $N$  denotes the number of agents and  $d$  the number of states each agent can have, the space of the surrogate model is at most  $d$ -dimensional (with  $d \ll N$ ). Together with the results presented in Section 3, we now construct computationally cheap surrogate models that preserve features relevant for optimizing ABMs.

The sampling algorithm introduced in Section 2.2 avoids the problem of discontinuity and the resulting nonexistence of derivatives since it is derivative-free, i.e., it relies only on function evaluations. We will demonstrate in Sections 4.1 and 4.2 how data-driven reduced models of ABMs can be effectively used as surrogate models in multi-objective optimization problems.

**Remark 4.1.** For the sake of illustration, we will only consider constant controls. Nevertheless, the approach can also be applied to time-varying controls such as different levels of curfews that change, e.g., every week. Additionally, having a reduced dynamical model based on the Koopman formalism opens up the possibility to use existing methods for controlling linear problems.

We will now introduce two ABMs which will be used as examples to demonstrate the different approaches described in Section 3.

*The voter model.* The voter model was first introduced in [54] and has since not only been used in the social context but also in many other contexts, such as chemical systems or ant colonies [55–57]. In its basic definition,  $N$  identical agents interact with each other at any time in a given network and change their opinions based on (stochastic) transition rules. In this work, we consider the model with two opinions and two transition rules. Given two agents with opinions  $S_i \neq S_j$ , the first transition rule is given by

$$R_{ij} : S_i + S_j \xrightarrow{\gamma_{ij}} 2S_j,$$

which is a second-order transition, meaning that an agent with opinion  $S_i$  adopts another agent's opinion  $S_j$  at rate  $\gamma_{ij}$ . The second transition rule is a first-order transition of the form

$$R'_{ij} : S_i \xrightarrow{\gamma'_{ij}} S_j,$$

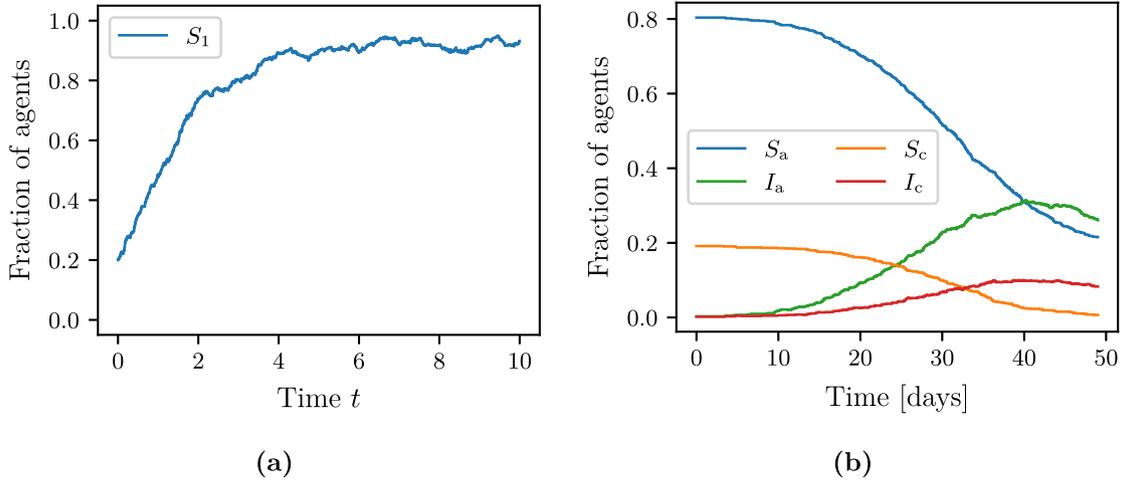


Fig. 4. Aggregated trajectories of (a) the voter model and (b) the GERDA model for the parameters given in Table 2.

where an agent changes its opinion independently of all other agents at rate  $\gamma'_{ij}$ . Gillespie's stochastic simulation algorithm [58], which constructs exact realizations in continuous time, can be used to simulate the ABM. We choose  $N = 500$  agents and set  $\gamma_{12} = 1$ ,  $\gamma_{21} = 2$ , and  $\gamma'_{12} = \gamma'_{21} = 0.1$ . A single aggregated simulation is shown in Fig. 4(a). That is, we visualize the number of agents with opinion  $S_1$  at time  $t$ . The trajectory for  $S_2$  is omitted due to conservation of  $N$ , i.e.,  $S_2 = N - S_1$ . For further details, especially with respect to convergence to ordinary or stochastic differential equations, we refer the reader to [59].

*The Georeferenced Demographic Agent-based model.* The second example is the individualized **GEo**Referenced **Demographic Agent-based model** (GERDA) for the transmission of SARS-CoV-2 and the disease dynamics of COVID-19 [1]. Using detailed location data, a demographically matched population, and realistic, hourly schedules for each agent, this model simulates contacts between people in the given locations and the resulting infection events. Data from multiple towns, e.g., Tepoztlán (Mexico), Zikhron Ya'akov (Israel), or Gangelt (Germany) have been used to calibrate the GERDA model and analyze the respective infection dynamics. In this paper, the calibrated data of the German municipality of Gangelt is used. For more information, see the authors' original publication. The model is implemented in Python and is available at [https://ford.biologie.hu-berlin.de/jwodke/corona\\_model](https://ford.biologie.hu-berlin.de/jwodke/corona_model).

To simplify the disease dynamics, which originally distinguishes between agents that are susceptible, diagnosed, hospitalized, in intensive care, and recovered or deceased, we summarize and consider only the classical three compartments  $S$  (susceptible),  $I$  (infected), and  $R$  (recovered) for the aggregated dynamics. Since in the subsequent analysis we consider homeschooling and working from home as controls, and since a separation by age is also justified by different infection rates (see, e.g., [60]), we group the agents by children (age 0 to 18 years) and adults (19+) and use the subscripts  $c$  and  $a$ , respectively. Fig. 4(b) shows an aggregated trajectory of a single simulation of GERDA in terms of numbers of susceptible, infected, and recovered adults and children for the parameters given in Table 2. The trajectories for  $R_a$  and  $R_c$  are omitted due to the conservation of the number of adults  $N_a$  and children  $N_c$ .

#### 4.1. Multi-objective optimization of the voter model

In order to win over as many voters as possible to opinion  $S_2$ , we want to find combinations of push and pull arguments that simultaneously move voters away from opinion  $S_1$  and towards opinion  $S_2$ . We represent this behavior in terms of the control  $u = [u_{\text{push}}, u_{\text{pull}}]^T \in \mathcal{R}$ , where  $u_{\text{push}}$  acts on the transition rate  $\gamma_{12}$  from  $S_1$  to  $S_2$  and  $u_{\text{pull}}$  acts analogously on  $\gamma_{21}$ . We assume that  $u$  acts linearly on the dynamics

given by the ABM, i.e.,  $\gamma_{12}(u) = \bar{\gamma}_{12} + u_{\text{push}}$  and  $\gamma_{21}(u) = \bar{\gamma}_{21} + u_{\text{pull}}$ , and set up a multi-objective optimization problem of the form (3) with feasible decision space  $\mathcal{R} := [-1, 5] \times [-2, 5]$  and objective function  $F : \mathcal{R} \rightarrow \mathbb{R}^2$  with

$$f_1(u) = \mathbb{E}[X_1(t; u)/N], \quad (8)$$

$$f_2(u) = u_{\text{push}}^2 + u_{\text{pull}}^2, \quad (9)$$

where  $X_1(t; u)$  denotes the number of agents believing in opinion  $S_1$  at time  $t$  depending on control  $u \in \mathcal{R}$ . The second objective is chosen such that any effort to influence the agent's opinion is costly. We evaluate the objective (8) at time  $t = 10$ . Figs. 5(a) and (b) show the objective functions (8) and (9), respectively.

Even though this ABM is computationally inexpensive in comparison to other ABMs, it is still inefficient to optimize it directly. Therefore, we first construct a surrogate model using gEDMD, and then solve the multi-objective optimization problem. As the control  $u$  acts linearly on the dynamics, Theorem 3.2 guarantees that we can obtain a surrogate model for varying control  $u$  using interpolated Koopman generators. In fact, for sufficiently large  $N$  and due to conservation of  $N$ , the drift and diffusion terms identified by gEDMD correspond to the well-known SDE approximation by Kurtz [59] (also known as the chemical Langevin equation [61]) satisfying

$$\begin{aligned} dC_t = & \left[ ((\gamma_{21} + u_{\text{pull}}) - (\gamma_{12} + u_{\text{push}})) C_t (1 - C_t) - \gamma'_{12} C_t + \gamma'_{21} (1 - C_t) \right] dt \\ & + \frac{1}{\sqrt{N}} \left[ -\sqrt{(\gamma_{12} + u_{\text{push}}) C_t (1 - C_t)} dW_1(t) \right. \\ & \left. + \sqrt{(\gamma_{21} + u_{\text{pull}}) C_t (1 - C_t)} dW_2(t) \right] \\ & + \frac{1}{\sqrt{N}} \left[ -\sqrt{\gamma'_{12} C_t} dW_3(t) + \sqrt{\gamma'_{21} (1 - C_t)} dW_4(t) \right], \end{aligned}$$

where  $C_0 = \lim_{N \rightarrow \infty} X_1(0; u)/N$ . Drift and diffusion terms can be written as (4) and (5). Thus, Theorem 3.2 holds and we learn four generator approximations to cover the feasible decision space  $\mathcal{R}$ . We use the four controls marking the vertices of  $\mathcal{R}$  to construct the surrogate model. The matrix representation of the Koopman generator approximation for the surrogate is then given by  $L_u = \sum_{i=1}^4 \alpha_i L(u_i)$ , where  $\sum_{i=1}^4 \alpha_i = 1$  with  $\alpha_i \in [0, 1]$  and  $L(u_1) = L([0, 0])$ ,  $L(u_2) = L([1, 0])$ ,  $L(u_3) = L([0, 1])$  and  $L(u_4) = L([1, 1])$  denote the matrix representations of the Koopman generator approximation for the dynamics with control  $u_i$ . To train the reduced models at the vertices, i.e., learn the matrix representations  $L(u_i)$ , we sample the ABM at 100 uniformly distributed points, each having 100 Monte Carlo simulations to calculate pointwise drift and diffusion estimates using the Kramers–Moyal formulae. For details on Kramers–Moyal expansions, see [62].

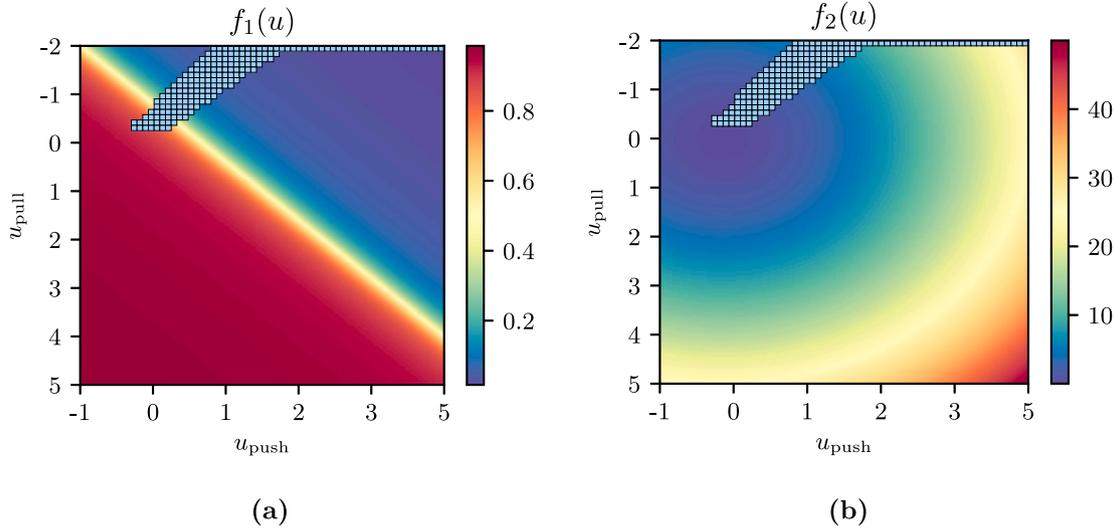


Fig. 5. Pareto set covering (light blue boxes) after 12 iterations plotted on (a) objective function (8) and (b) objective function (9) at time  $t = 10$  depending on the control  $[u_{\text{push}}, u_{\text{pull}}]^T \in \mathcal{R}$ . The yellow line in (a) marks the area of interest where the majority switches from  $S_1$  to  $S_2$ .

Fig. 5 shows the Pareto set covering obtained after 12 iterations for the feasible decision space  $\mathcal{R}$  together with the objectives (8) and (9). In order to achieve an absolute majority for opinion  $S_2$ , we refine the decision space to  $\mathcal{R}^* = [0.25, 0.75] \times [-0.75, -0.25]$ . Fig. 6 shows in (a) the computed covering of the Pareto set for  $\mathcal{R}^*$  and in (b) the image of  $\mathcal{R}^*$  under objective function  $F$  as a light blue area as well as the approximated Pareto front (red/solid), which is obtained by mapping the center of each box covering the Pareto set in (a). To verify that the surrogate model approximates the dynamics of the ABM sufficiently well and consequently is also capable of approximating the Pareto set, we randomly choose some test points (blue/dots) for which we evaluate the objective (8) from 100 Monte Carlo simulations using the ABM. Fig. 6(c) shows close-ups of these test points. We observe that the test points that are covered by the boxes in Fig. 6(a) (and are therefore not visible) are non-dominated points. These points are mapped to the Pareto front computed via the surrogate model. The visible (blue marked) test points (in general, points that are not covered by boxes in Fig. 6(a)) are on the right side of the Pareto front and thus dominated. Thus, the box covering computed via the data-driven surrogate model actually approximates the Pareto-optimal set of the ABM as it captures the dynamical behavior of the ABM under control  $u$  sufficiently well. The error bars indicate a confidence level of 99.9%.

#### 4.2. Multi-objective optimization of the GERDA model

The design of optimal containment strategies for a complex model like GERDA has recently been discussed in [63]. Although the considered objective functions are only prototypical, they show that the design of optimal containment strategies using ABMs crucially depends on the level of detail of the model. However, even though the main objective in [63] consists of three sub-objectives (i.e., health, society, and economy), multi-objective optimization is not considered. We now show how a multi-objective optimal containment strategy can be found using a suitable reduced model of GERDA.

Following [63], we consider *constant* controls  $u = [u_s, u_w]^T \in \mathcal{R}$  and set the feasible decision space for our multi-objective optimization problem (3) to  $\mathcal{R} = [0, 1] \times [0, 0.8]$  and, for better comparability, choose the same objectives and parameters as in [63], i.e.,  $F : \mathcal{R} \rightarrow \mathbb{R}^2$  with

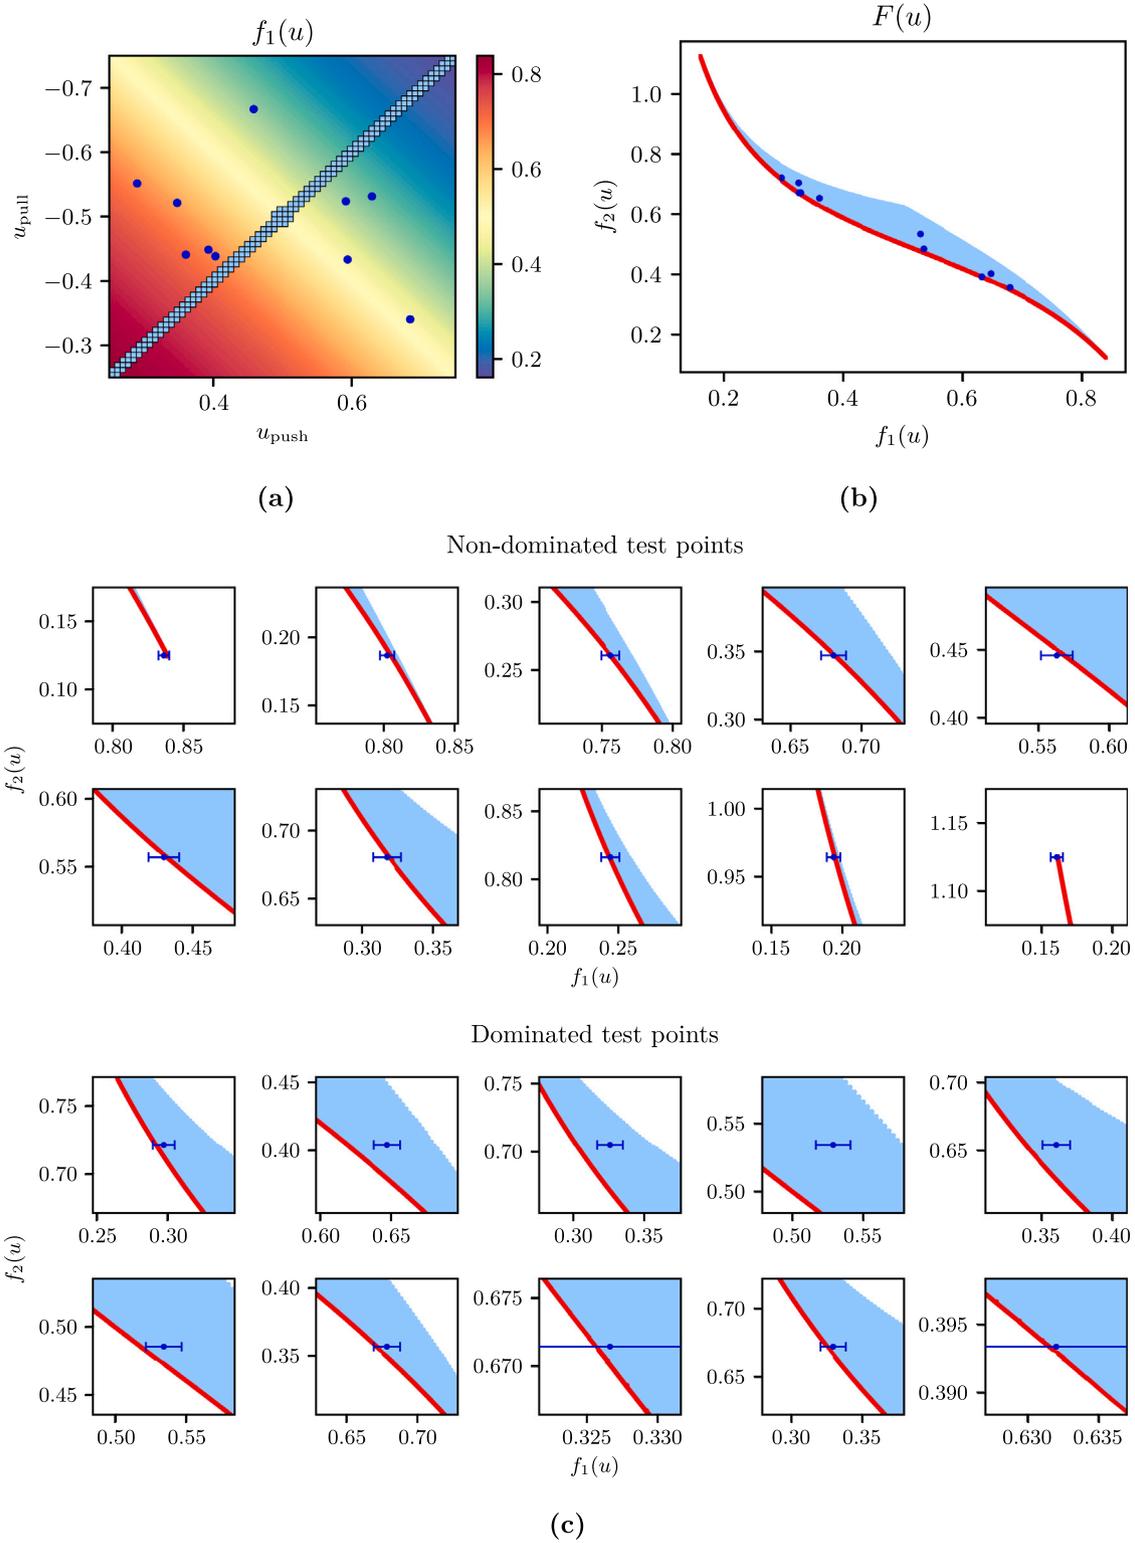
$$f_1(u) = \mathbb{E} \left[ \int_0^T (I(t; u)/N + \exp(10(I(t; u) - I_{\max})/N)) dt \right], \quad (10)$$

$$f_2(u) = \int_0^T [u_s(t)^2 - \log(u_w^{\max} - u_w(t))] dt, \quad (11)$$

where  $f_1$  and  $f_2$  denote the objectives related to health and society/economy, respectively. Further, let  $I(t; u)$  denote the number of infected agents,  $I_{\max}$  and  $u_w^{\max}$  threshold values,  $u_s$  and  $u_w$  the controls representing the fraction of school and work closures (where  $u_s = 0$  is interpreted as no schools being closed and  $u_s = 1$  as all being closed; analogously for  $u_w$ ). Note that we do not consider the objectives for society and economy individually as this would result in a trivial multi-objective optimization problem. In fact, the whole domain  $\mathcal{R}$  would be Pareto-optimal since in this case each objective increases in a different direction.

**Remark 4.2.** Eq. (10) accounts for the direct negative consequences of infections, which is assumed to exhibit an approximate linear relationship with the number of infected individuals, as well as for the social impact, which increases significantly when the capacity of the health care system to treat severely ill individuals is exceeded. This is quantified by a threshold of infected individuals, denoted as  $I_{\max}$ . Eq. (11) represents the direct costs of both controls  $u_s$  and  $u_w$ . The economic costs become  $+\infty$  when the homeworking rate approaches an upper bound  $u_w^{\max} < 1$  in case that too many workplaces are closed or employees are working from home.

Since even for smaller geographic regions such as Gangelte, which has roughly 13000 inhabitants (December 2021), solving the multi-objective optimization problem directly for GERDA is not feasible due to the enormous computational effort and inherent stochasticity of the model. Thus, we construct a surrogate model. Following the arguments in Example 3.4, interpolation between Koopman generators should not be used. In fact, it can be verified easily that gEDMD with control leads to better results (see Table 1). To train the reduced model, which serves as the surrogate model later on, we sample GERDA (for a fixed world) at 49 points and 225 different control inputs along 7-week trajectories for 24 hourly time steps, i.e., one data point every 24 steps is added to the training data set. The space of feasible controls  $\mathcal{R}$  is discretized using an equidistant grid with  $15 \times 15$  grid points. This results in a total of 11025 training data points, each having 1000 Monte Carlo simulations to calculate pointwise drift and diffusion estimates using the Kramers–Moyal formulae. We define the augmented state as  $x := [S_a, S_c, I_a, I_c, u_s, u_w]^T$ , where we distinguish between susceptible and infected adults and children. Note that  $R_a$  and  $R_c$  are omitted due to conservation of  $N_a$  and  $N_c$ . The set of basis functions consists of monomials up to degree 4 following the same arguments as in Example 3.4. Fig. 7 shows the expected aggregated trajectories (dashed) and its standard deviation (shaded) of GERDA as well as the



**Fig. 6.** (a) Pareto set covering (light blue boxes) and test points (blue/dots) for refined feasible decision space  $\mathcal{R}^*$ . (b) Pareto front (red/solid), image of  $\mathcal{R}^*$  (light blue area) and test points under the objective function  $F$ . (c) Close-ups of randomly chosen test points, evaluated via the ABM. While the non-dominated test points are covered by the boxes in (a) and thus are Pareto-optimal points, the dominated test points are not. Error bars indicate a 99.9% confidence level.

solution of the reduced model (solid) in terms of an ODE for different controls. Table 1 compares the true expected trajectories of GERDA with the reduced model in terms of the root mean squared error for the 6-dimensional state  $x$  as well as a further reduced 4-dimensional state  $x := [S, I, u_s, u_w]^T$ . We also include a comparison between the true expected trajectories of GERDA and the generator interpolation

method whose use is invalid for the GERDA model. Even on the training data set, state augmentation yields better results. The 6-dimensional reduced model yields the best approximation. Although all models can be sparsified (e.g., using iterative thresholding [47]), a meaningful relation of the non-zero terms to the classical SIR structure cannot be made. This is, however, also expected as GERDA implicitly contains a

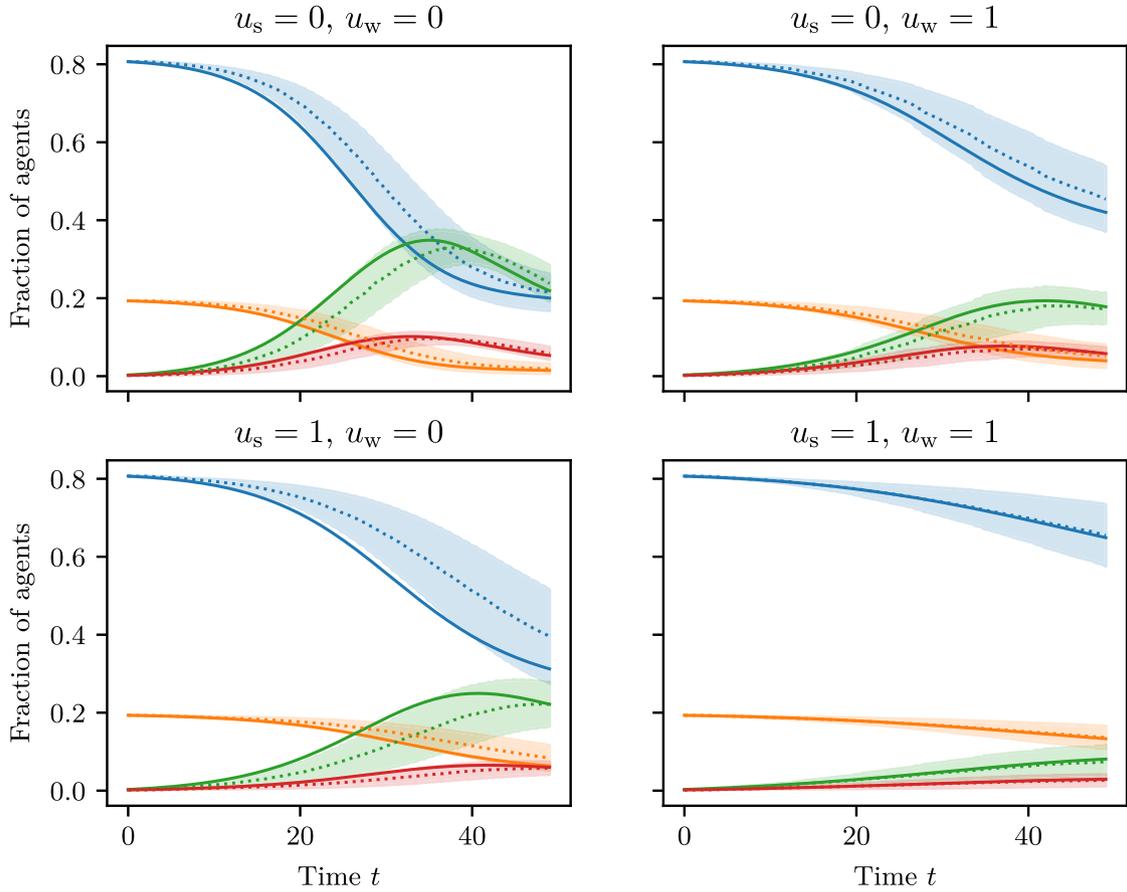


Fig. 7. Mean aggregated trajectories for GERDA (dotted) and the reduced model (solid) of the fraction of susceptible (blue for adults, orange for children) and infected agents (green for adults, red for children) under different controls  $u = [u_s, u_w]^T$ . The shaded areas indicate the standard deviations of the GERDA model. The ABM data is estimated from 1 000 independent simulations using a fixed world and the parameters in Table 2.

Table 1

Root mean square errors of the trajectories for the reduced models obtained by state augmentation and generator interpolation. The ground truth is a Monte Carlo estimate of GERDA with 1 000 simulations for each control  $u$ . The complete parameterization can be found in Table 2.

	Dimension	Control $[u_s, u_w]^T$			
		$[0, 0]^T$	$[0, 1]^T$	$[1, 0]^T$	$[1, 1]^T$
State augmentation	6	0.0325	0.0186	0.0448	0.0023
	4	0.0551	0.0325	0.0617	0.0177
Generator interpolation	6	0.1869	0.1109	0.1172	0.4122
	4	0.0167	0.0347	0.0714	0.0716

network in which agents change locations according to their schedule and interact only with certain agents, i.e., agents present at those locations at the same time. Nevertheless, the reduced model is suitable to solve the multi-objective optimization problem efficiently. Fig. 8 shows the computed Pareto set and front coverings after 14 iterations. In both figures the green cross marks the optimal control  $u^* = [0, 0.63]^T$  and the corresponding objective values computed in [63], which was computed without using a surrogate model. We see that this point is included in the Pareto set covering. This finding further highlights the effectiveness of the data-driven surrogate model in approximating the Pareto-optimal set of the ABM. Moreover, solving the multi-objective optimization problem reveals further optimal control inputs belonging to different combinations of weights in [63].

**Remark 4.3.** Note that our primary goal is not to solve multi-objective optimization problems efficiently, but rather to show that surrogate

Table 2

Parameters used for the simulations of GERDA.

Parameter	
Time $T$ [hours]	1176
Number of agents $N$	1045
Initially infected $I_a$	3
Initially infected $I_c$	2
World	Gangelt (reduced)
General infectivity	0.175
General interaction frequency	1
Health care system's capacity threshold $I_{max}$	0.005 $N$
Threshold economic impact $u_w^{max}$	0.81

models based on the Koopman generator are well-suited for this purpose. Especially in the case of the voter model, it is more efficient to use the well-known ODE or SDE approximations as surrogate models. However, such approximations are not always known or, in the case of GERDA, might not exist. Thus, constructing surrogate models using the Koopman generator offers a viable approach to solving multi-objective optimization problems involving complex ABMs.

## 5. Conclusions

The construction of data-driven surrogate models to speed up or even enable the solution of multi-objective optimization problems, where objectives are defined by the dynamical behavior of ABMs, is of great importance as ABMs become more widely used. Especially in the case when limit models are unknown or non-existent, data-driven reduced models and surrogate models allow to solve optimization problems that would otherwise be computationally infeasible due to

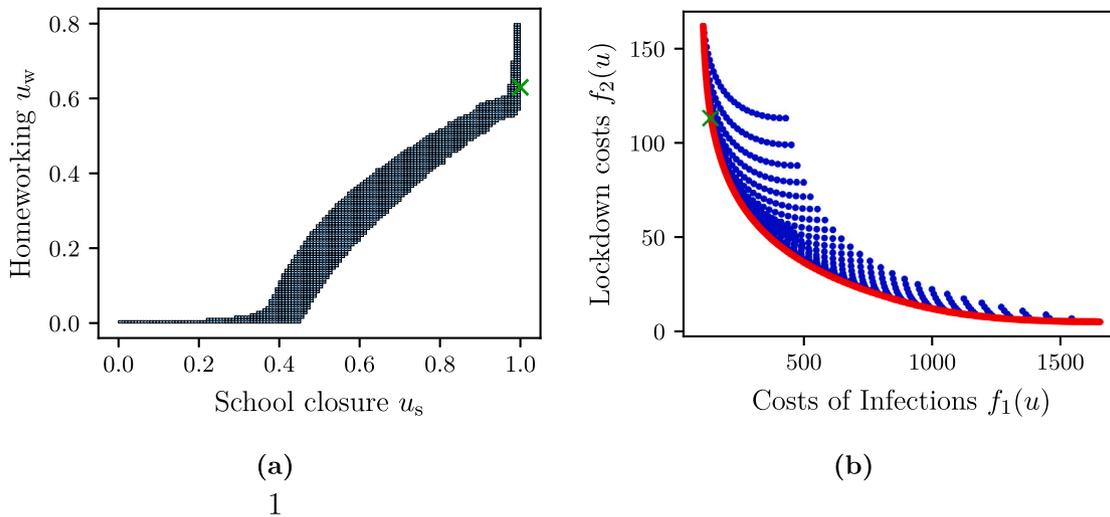


Fig. 8. (a) Pareto set covering for feasible decision space  $\mathcal{R}$  after 14 subdivision steps obtained with the 6-dimensional reduced model for GERDA. (b) Pareto front (red) and dominated points (blue). The green cross indicates the optimal control found in [63].

the very expensive evaluation of the objectives. In this work, we demonstrated how data-driven reduced models based on the Koopman generator enable the efficient solution of multi-objective optimization problems associated with ABMs. For this purpose we constructed surrogate models with varying decision variables, which in this work are controls applied to a non-deterministic dynamical system. We introduced two different methods of constructing these surrogate models. The first one is applicable if the control acts linearly on the system. In this case, linear interpolation between Koopman generators for different control inputs is feasible and can be used to construct surrogate models. This method was illustrated using the voter model, where the goal was to achieve an absolute majority in one opinion with a linear control affecting the second-order transitions. This can be interpreted as repulsion and attraction arguments in an election campaign. The second method overcomes the restriction to linearly acting controls. We showed that a straightforward extension of gEDMD that considers the control as an additional state allows for nonlinear dependence on the controls. The method was demonstrated first using a simple and analytical model. We then used this method to compute the Pareto set of optimal containment strategies for the GERDA model in a prototypical situation and compared it with the solution computed in [63]. We showed that the solution computed for the single-objective optimization problem is contained in our Pareto set.

In this work, we considered multi-objective optimization problems with  $k \ll N$  globally defined objectives, i.e., objectives defined by the collective behavior of the agents, and  $n \ll N$  decision variables. An extension is to consider both objectives defined by the individual and collective behavior of agents, which allows for more general problems without considering aggregated trajectories. An important question then is how to deal with the large computational costs when  $k \approx N$  or  $n \approx N$  or both. Future research will address these questions.

#### CRediT authorship contribution statement

**Jan-Hendrik Niemann:** Conceptualization, Formal analysis, Investigation, Methodology, Project administration, Validation, Visualization, Writing – original draft, Writing – review & editing. **Stefan Klus:** Conceptualization, Investigation, Supervision, Writing – review & editing. **Nataša Djurdjevac Conrad:** Investigation, Supervision, Writing – review & editing. **Christof Schütte:** Conceptualization, Supervision, Writing – review & editing.

#### Declaration of competing interest

There is no conflict of interest between the authors and other persons organizations.

#### Data availability

No data was used for the research described in the article.

#### Acknowledgments

This research has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy MATH+: The Berlin Mathematics Research Center (EXC-2046/1, project ID: 390685689).

#### Appendix A. Sampling algorithm

For  $s = 0$ , let  $\mathcal{B}^0$  denote a collection of finitely many subsets of  $\mathcal{R} = [a_1, b_1] \times \dots \times [a_n, b_n] \subset \mathbb{R}^n$ ,  $a_i \leq y_i \leq b_i$  for  $i = 1, \dots, n$ , such that  $\bigcup_{B \in \mathcal{B}^0} B = \mathcal{R}$ . Then, in each iteration, each box  $B \in \mathcal{B}^s$ ,  $s \geq 0$ , is first subdivided with respect to one coordinate, resulting in a new collection of boxes  $\mathcal{B}^{s+1}$ . In the second step, a set-wise non-dominance test, which can be carried out heuristically by a finite number of test points, is performed on each box  $B \in \mathcal{B}^{s+1}$ , discarding all boxes containing only dominated points. This process is repeated with the new collection of boxes  $\mathcal{B}^{s+1}$  and the next coordinate. For more details see [8,13,14,48]. The sampling algorithm is summarized in Algorithm 1.

#### Appendix B. Proof of Theorem 3.2

Given a space of twice continuously differentiable functions and controls  $u_1, u_2$ , if the dynamics (1) are control-affine, then the Koopman generators are control-affine, i.e.,

$$\mathcal{L}_{\alpha_1 u_1 + \alpha_2 u_2} = \mathcal{L}_0 + \alpha_1 \mathcal{A}_{u_1} + \alpha_2 \mathcal{A}_{u_2},$$

where  $\mathcal{A}_u = \mathcal{L}_u - \mathcal{L}_0$  and  $\alpha_1, \alpha_2 \in \mathbb{R}$ .

**Proof.** Consider a control-affine system of the form (1) and let  $u$  be any linear combination of controls, i.e.,  $u = \sum_{i=1}^{d_u} \alpha_i u_i$  for  $\alpha_i \in \mathbb{R}$ . Given some twice differentiable function  $f$ , the Koopman generator  $\mathcal{L}_u$  depending on  $u$  applied to  $f$  yields

$$\begin{aligned} \mathcal{L}_u f &= b \cdot \nabla_x f + \frac{1}{2} a : \nabla_x^2 f \\ &= b_0 \cdot \nabla_x f + \sum_{i=1}^{d_u} b_i u_i \cdot \nabla_x f + \frac{1}{2} a_0 : \nabla_x^2 f + \frac{1}{2} \sum_{i=1}^{d_u} a_i u_i : \nabla_x^2 f, \end{aligned}$$

**Algorithm 1:** [13,48]

Let  $\mathcal{B}^0$  be a collection of finitely many subsets of  $\mathcal{R}$  such that  $\bigcup_{B \in \mathcal{B}^0} B = \mathcal{R}$ . Then, obtain the new collection  $\mathcal{B}^{s+1}$ ,  $s \geq 0$ , iteratively from  $\mathcal{B}^s$  in two steps:

1. Construct a new collection  $\hat{\mathcal{B}}^{s+1}$  from  $\mathcal{B}^s$  by subdividing each subset  $B \in \mathcal{B}^s$  such that

$$\bigcup_{B \in \hat{\mathcal{B}}^{s+1}} B = \bigcup_{B \in \mathcal{B}^s} B,$$

$$\text{diam}(\hat{\mathcal{B}}^{s+1}) = \theta^{s+1} \text{diam}(\mathcal{B}^s)$$

for  $0 < \theta_{\min} \leq \theta^{s+1} \leq \theta_{\max} < 1$ .

2. Define the new collection  $\mathcal{B}^{s+1}$  by

$$\mathcal{B}^{s+1} := \{B \in \hat{\mathcal{B}}^{s+1} \mid \nexists \hat{B} \in \hat{\mathcal{B}}^{s+1} \text{ such that } \hat{B} \text{ dominates } B\}.$$

where  $\nabla_x^2$  denotes the Hessian. Defining  $\mathcal{A}_u := \mathcal{L}_u - \mathcal{L}_0$ , we obtain

$$\mathcal{A}_u f = \sum_{i=1}^{d_u} u_i \left[ b_i \cdot \nabla_x f + \frac{1}{2} a_i : \nabla_x^2 f \right].$$

The operators  $\mathcal{A}_u$  are linear with respect to the control  $u$ . Moreover, the generators of the Koopman operators are control-affine, that is,

$$\mathcal{L}_u f = \left( \mathcal{L}_0 + \sum_{i=1}^{d_u} \alpha_i \mathcal{A}_{u_i} \right) f. \quad \square$$

**References**

[1] B. Goldenbogen, S.O. Adler, O. Bodeit, J.A.H. Wodke, X. Escalera-Fanjul, A. Korman, M. Krantz, L. Bonn, R. Morán-Torres, J.E.L. Haffner, M. Karnetzki, I. Maintz, L. Mallis, H. Prawitz, P.S. Segelitz, M. Seeger, R. Linding, E. Klipp, Control of COVID-19 outbreaks under stochastic community dynamics, bimodality, or limited vaccination, *Adv. Sci.* 9 (23) (2022) 2200088.

[2] S. Banisch, R. Lima, T. Araújo, Agent based models and opinion dynamics as Markov chains, *Social Networks* 34 (4) (2012) 549–561.

[3] J. Mielke, A. Geiges, Model-Stakeholder Interactions for a Sustainable Mobility Transition, *Global Climate Forum Berlin*, 2018.

[4] N. Djurdjevic Conrad, L. Helfmann, J. Zonker, S. Winkelmann, C. Schütte, Human mobility and innovation spreading in ancient times: A stochastic agent-based simulation approach, *EPJ Data Sci.* 7 (1) (2018) 24.

[5] J. Nocedal, S. Wright, *Numerical Optimization*, Springer, 2006.

[6] A.R. Conn, N.I.M. Gould, P.L. Toint, *Trust-Region Methods*, SIAM, 2000.

[7] K. Miettinen, *Nonlinear multiobjective optimization*, in: *International Series in Operations Research & Management Science*, Springer, Boston, MA, 1998.

[8] O. Schütze, A. Dell’Aere, M. Dellnitz, On continuation methods for the numerical treatment of multi-objective optimization problems, in: *Dagstuhl Seminar Proceedings*, 2005.

[9] C.A. Coello Coello, G.B. Lamont, D.A. Van Veldhuizen, *Evolutionary algorithms for solving multi-objective problems*, in: *Genetic and Evolutionary Computation Series*, second ed., Springer, Boston, MA, 2007.

[10] S. Peitz, M. Dellnitz, A survey of recent trends in multiobjective optimal control – surrogate models, feedback control and objective reduction, *Math. Comput. Appl.* 23 (2) (2018).

[11] B. Afsar, D. Podkopaev, K. Miettinen, Data-driven interactive multiobjective optimization: Challenges and a generic multi-agent architecture, *Procedia Comput. Sci.* 176 (2020) 281–290, *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 24th International Conference KES2020*.

[12] M.J. Blondin, M. Hale, An algorithm for multi-objective multi-agent optimization, in: *2020 American Control Conference, ACC, IEEE*, 2020.

[13] M. Dellnitz, O. Schütze, T. Hestermeyer, Covering Pareto sets by multilevel subdivision techniques, *J. Optim. Theory Appl.* 124 (1) (2005) 113–136.

[14] O. Schütze, K. Witting, S. Ober-Blöbaum, M. Dellnitz, *Set Oriented Methods for the Numerical Treatment of Multiobjective Optimization Problems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 187–219.

[15] G. An, B.G. Fitzpatrick, S. Christley, P. Federico, A. Kanarek, R. Miller Neilan, M. Oremland, R. Salinas, R. Laubenbacher, S. Lenhart, Optimization and control of agent-based models in biology: A perspective, *Bull. Math. Biol.* 79 (1) (2017) 63–87.

[16] A. Rogers, P. Tessin, Multi-objective calibration for agent-based models, in: *Agent-Based Simulation*, vol. 5, 2004.

[17] B. Calvez, G. Hutzler, Automatic tuning of agent-based models using genetic algorithms, in: *Multi-Agent-Based Simulation VI*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 41–57.

[18] M.N. Read, K. Alden, L.M. Rose, J. Timmis, Automated multi-objective calibration of biological agent-based simulations, *J. R. Soc. Interface* 13 (122) (2016) 20160543.

[19] Z. Liu, D. Rexachs, F. Epelde, E. Luque, A simulation and optimization based method for calibrating agent-based emergency department models under data scarcity, *Comput. Ind. Eng.* 103 (2017) 300–309.

[20] I. Moya, M. Chica, Ó. Cordón, A multicriteria integral framework for agent-based model calibration using evolutionary multiobjective optimization and network-based visualization, *Decis. Support Syst.* 124 (2019) 113111.

[21] J. Párraga-Álava, G.M. Garzón, R.V. Valarezo, Multi-objective genetic algorithms: Are they useful for tuning parameters in agent-based simulation? *Rev. Iber. Sist. Tecnol. Inf.* 19 (2019) 172–184.

[22] R. Korkmaz Tan, Ş. Bora, Adaptive parameter tuning for agent-based modeling and simulation, *Simulation* 95 (9) (2019) 771–796.

[23] M. D’Auria, E.O. Scott, R.S. Lather, J. Hilty, S. Luke, Assisted parameter and behavior calibration in agent-based models with distributed optimization, in: Y. Demazeau, T. Holvoet, J.M. Corchado, S. Costantini (Eds.), *Advances in Practical Applications of Agents, Multi-Agent Systems, and Trustworthiness, The PAAMS Collection*, Springer International Publishing, Cham, 2020, pp. 93–105.

[24] B. Bartkowski, M. Beckmann, M. Drechsler, A. Kaim, V. Liebelt, B. Müller, F. Witing, M. Strauch, Aligning agent-based modeling with multi-objective land-use allocation: Identification of policy gaps and feasible pathways to biophysically optimal landscapes, *Front. Environ. Sci.* 8 (July) (2020) 1–15.

[25] D. Platt, A comparison of economic agent-based model calibration methods, *J. Econom. Dynam. Control* 113 (2020) 103859.

[26] M.B. Berkemeier, S. Peitz, Derivative-free Multiobjective Trust Region descent method using radial basis function surrogate models, *Math. Comput. Appl.* 26 (2) (2021).

[27] F. Lamperti, A. Roventini, A. Sani, Agent-based model calibration using machine learning surrogates, *J. Econom. Dynam. Control* 90 (2018) 366–389.

[28] M. Oremland, R. Laubenbacher, Using difference equations to find optimal tax structures on the SugarScape, *J. Econ. Interact. Coord.* (2014).

[29] M. Oremland, R. Laubenbacher, Optimal harvesting for a predator-prey agent-based model using difference equations, *Bull. Math. Biol.* 77 (3) (2015) 434–459.

[30] S. Koshy-Chenthittayil, P. Mendes, R. Laubenbacher, Optimization of agent-based models through coarse-graining, *Lett. Biomath.* 8 (1) (2021) 167–178.

[31] S. Christley, R. Miller Neilan, M. Oremland, R. Salinas, S. Lenhart, Optimal control of sugarscape agent-based model via a PDE approximation model, *Optim. Control Appl. Methods* 38 (4) (2017) 473–497.

[32] H. Wulkow, T.O.F. Conrad, N. Djurdjevic Conrad, S.A. Müller, K. Nagel, C. Schütte, Prediction of Covid-19 spreading and optimal coordination of countermeasures: From microscopic to macroscopic models to Pareto fronts, *PLoS One* 16 (4) (2021) 1–29.

[33] S. Klus, P. Koltai, C. Schütte, On the numerical approximation of the Perron–Frobenius and koopman operator, *J. Comput. Dyn.* 3 (1) (2016) 51–79.

[34] C. Schütte, S. Klus, C. Hartmann, Overcoming the timescale barrier in molecular dynamics: Transfer operators, variational principles and machine learning, *Acta Numer.* 32 (2023) 517–673.

[35] A. Mauroy, I. Mezić, Global stability analysis using the eigenfunctions of the koopman operator, *IEEE Trans. Automat. Control* 61 (11) (2016) 3356–3369.

[36] A. Mauroy, J. Goncalves, Linear identification of nonlinear systems: A lifting technique based on the koopman operator, in: *2016 IEEE 55th Conference on Decision and Control, CDC*, 2016, pp. 6500–6505.

[37] J.N. Kutz, S.L. Brunton, B.W. Brunton, J.L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, SIAM, Philadelphia, 2016.

[38] S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, C. Schütte, Data-driven approximation of the koopman generator: Model reduction, system identification, and control, *Physica D* 406 (2020) 132416.

- [39] M. Korda, I. Mezić, Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control, *Automatica* 93 (2018) 149–160.
- [40] H. Arbabi, M. Korda, I. Mezić, A data-driven koopman model predictive control framework for nonlinear flows, in: 2018 IEEE Conference on Decision and Control, CDC, 2018, pp. 6409–6414.
- [41] M. Korda, I. Mezić, Optimal construction of koopman eigenfunctions for prediction and control, *IEEE Trans. Automat. Control* 65 (12) (2020) 5114–5129.
- [42] S. Peitz, S. Klus, Koopman operator-based model reduction for switched-system control of PDEs, *Automatica* 106 (2019) 184–191.
- [43] S. Peitz, S.E. Otto, C.W. Rowley, Data-driven model predictive control using interpolated Koopman generators, *SIAM J. Appl. Dyn. Syst.* 19 (3) (2020) 2162–2193.
- [44] J.-H. Niemann, S. Klus, C. Schütte, Data-driven model reduction of agent-based systems using the koopman generator, *PLoS One* 16 (5) (2021) 1–23.
- [45] A. Lasota, M.C. Mackey, Chaos, fractals, and noise: Stochastic aspects of dynamics, second ed., in: *Applied Mathematical Sciences*, vol. 97, Springer, New York, 1994.
- [46] P. Metzner, *Transition Path Theory for Markov Processes: Application to Molecular Dynamics* (Ph.D. thesis), Freie Universität Berlin, 2007.
- [47] S. Brunton, J. Proctor, J. Kutz, Discovering governing equations from data: Sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci.* 113 (2015) 3932–3937.
- [48] S. Peitz, *Exploiting Structure in Multiobjective Optimization and Optimal Control* (Ph.D. thesis), Universität Paderborn, 2017.
- [49] M. Dellnitz, G. Froyland, O. Junge, The algorithms behind GAIO — Set oriented numerical methods for dynamical systems, in: B. Fiedler (Ed.), *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 145–174.
- [50] J.L. Proctor, S.L. Brunton, J.N. Kutz, Dynamic mode decomposition with control, *SIAM J. Appl. Dyn. Syst.* 15 (1) (2016) 142–161.
- [51] S.L. Brunton, J.L. Proctor, J.N. Kutz, Sparse identification of nonlinear dynamics with control (SINDYc), *IFAC-PapersOnLine* 49 (18) (2016) 710–715, 10th IFAC Symposium on Nonlinear Control Systems NOLCOS 2016.
- [52] Y. Cheng, N. Ma, C. Witt, S. Rapp, P.S. Wild, M.O. Andreae, U. Pöschl, H. Su, Face masks effectively limit the probability of SARS-CoV-2 transmission, *Science* 372 (6549) (2021) 1439–1443.
- [53] J.-H. Niemann, *Learning Reduced Models for Large-Scale Agent-Based Systems* (Ph.D. thesis), Freie Universität Berlin, 2022.
- [54] R.A. Holley, T.M. Liggett, Ergodic theorems for weakly interacting infinite systems and the voter model, *Ann. Probab.* 3 (4) (1975) 643–663.
- [55] F. Herreriás-Azcué, T. Galla, Consensus and diversity in multistate noisy voter models, *Phys. Rev. E* 100 (2) (2019).
- [56] T. Biancalani, L. Dyson, A.J. McKane, Noise-induced bistable states and their mean switching time in foraging colonies, *Phys. Rev. Lett.* 112 (3) (2014) 1–5.
- [57] J. Ohkubo, N. Shnerb, D.A. Kessler, Transition phenomena induced by internal noise and quasi-absorbing state, *J. Phys. Soc. Japan* 77 (4) (2008) 044002.
- [58] D.T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, *J. Comput. Phys.* 22 (4) (1976) 403–434.
- [59] T.G. Kurtz, Strong approximation theorems for density dependent Markov chains, *Stochastic Process. Appl.* 6 (3) (1978) 223–240.
- [60] J. She, L. Liu, W. Liu, COVID-19 epidemic: Disease characteristics in children, *J. Med. Virol.* 92 (7) (2020) 747–754.
- [61] D.T. Gillespie, The chemical Langevin equation, *J. Chem. Phys.* 113 (1) (2000) 297–306.
- [62] H. Risken, T. Frank, *The Fokker–Planck Equation: Methods of Solutions and Applications*, Springer Series in Synergetics, Springer-Verlag, Berlin Heidelberg, 1996.
- [63] J.-H. Niemann, S. Uram, S. Wolf, N. Djurdjevic Conrad, M. Weiser, Multilevel optimization for policy design with agent-based epidemic models, 2023, arXiv: 2304.02281.