



# A Higher-Order Singular Value Decomposition Tensor Emulator for Spatiotemporal Simulators

Giri GOPALAN and Christopher K. WIKLE

We introduce methodology to construct an emulator for environmental and ecological spatiotemporal processes that uses the higher-order singular value decomposition (HOSVD) as an extension of singular value decomposition (SVD) approaches to emulation. Some important advantages of the method are that it allows for the use of a combination of supervised learning methods (e.g., random forests and Gaussian process regression) and also allows for the prediction of process values at spatial locations and time points that were not used in the training sample. The method is demonstrated with two applications: The first is a periodic solution to a shallow ice approximation partial differential equation from glaciology, and second is an agent-based model of collective animal movement. In both cases, we demonstrate the value of combining different machine learning models for accurate emulation. In addition, in the agent-based model case we demonstrate the ability of the tensor emulator to successfully capture individual behavior in space and time. We demonstrate via a real data example the ability to perform Bayesian inference in order to learn parameters governing collective animal behavior.

**Key Words:** Agent-based model; Bayesian; Collective movement; Machine learning; Surrogates; Tensor.

## 1. INTRODUCTION

Multivariate spatiotemporal processes are ubiquitous in agricultural, biological, and environmental science. Modeling such processes is complicated by complex dependencies across processes, time, and space, as well as uncertainty in observations, process specification, and parameters that are contained in the observation and process models. Since the mid-1990s, Bayesian hierarchical models have been the primary tool to consider such processes in statistics (1996). Although there are challenges to each component of such a model, for

---

G. Gopalan (✉) Statistics Department, California Polytechnic State University, San Luis Obispo, CA 93407, USA (E-mail: [gopalan@calpoly.edu](mailto:gopalan@calpoly.edu)).

C. K. Wikle Department of Statistics, University of Missouri, Columbia, MO 65211, USA (E-mail: [wiklec@missouri.edu](mailto:wikle@missouri.edu)).

© 2021 International Biometric Society  
*Journal of Agricultural, Biological, and Environmental Statistics*  
<https://doi.org/10.1007/s13253-021-00459-x>

multivariate spatiotemporal processes, arguably the greatest challenge is the specification of a model for the process that is rich enough, yet parsimonious enough, to capture realistic behavior (e.g., nonlinear interactions, non-stationarity, non-separability, collectiveness, etc.). One approach to this problem has been to specify mechanistically motivated dynamical parameterizations within the hierarchical statistical framework (i.e., “physical-statistical models,” [Berliner 2003](#); [Kuhnert 2014](#)) such that they include the potential for known mechanistic behavior (e.g., diffusion, advection, density-dependent growth, etc.; see [Wikle and Hooten 2010](#)). In situations where one has a greater knowledge of the fundamental process dynamics, it is often preferable to embed a “black box” dynamical model (simulator) within a formal statistical framework, such as is done in the computer model calibration literature ([Kennedy and O’Hagan 2001](#)). Examples of such black box simulators include partial differential equation (PDE) models ([Gopalan et al. 2019](#)) or agent-based models (ABMs) that have multiple individuals interacting across time and space (e.g., [Fadikar et al. 2018](#)).

The intersection of these approaches for the statistical modeling of complex dynamical processes occurs when the black-box models are too expensive to implement within statistical inference procedures (typically Bayesian), and one must build an emulator (i.e., surrogate statistical model) to represent the input–output relationships in the black-box model. Essentially, the emulator is a function that mimics the output of the simulator but is computationally less expensive to evaluate. Emulators are most often specified by Gaussian processes (GPs), but nonlinear surrogate models can also be used (see the overview in [Gramacy 2020](#)).

Although surrogate modeling in statistics is a mature and well-studied topic, there are still challenges with respect to developing efficient surrogates for black-box models with high-dimensional multivariate spatiotemporal output. Perhaps the most common approach to dealing with spatiotemporal model output in the computer model calibration literature has been to project the model output on a reduced rank set of basis functions as in principal components analysis (PCA or singular value decomposition, SVD) ([Higdon et al. 2008](#)). In this case, one typically models the lower-dimensional PC coefficients in terms of GPs. However, different basis functions can be used ([Salter et al. 2019](#)) and one can also model the relationship between the inputs and projection coefficients (right singular vectors) via nonlinear functions in the mean. This is important because often the relationship between the inputs and the right singular vectors is not smooth.

It is much less common to build surrogate models for multivariate or multi-state spatiotemporal data. [Leeds et al. \(2014\)](#) considered a surrogate model for a multivariate spatiotemporal process consisting of phytoplankton, sea surface temperature, and sea level pressure. A more recent exception is [Pratola and Chkrebti \(2018\)](#), who made use of the so-called higher- order singular value decomposition (HOSVD) tensor analog of SVD ([De Lathauwer et al. 2000](#)) to build a GP surrogate model for multi-state spatiotemporal data. However, it is important to realize that input–output structure can be quite different whether one considers space, time, multiple-processes, or parameters. As mentioned above, it is quite possible that such relationships are not smooth and that in some cases it might be better to model the relationships in the mean (first-order) and in other cases in the covariance (second-order), as via a GP. Thus, we extend the first-order SVD modeling approach to the HOSVD tensor framework. In essence, the proposed tensor-based method stores a

tensor of simulator runs and performs a HOSVD; then, machine learning models are trained to learn the analogs of right singular vectors. Importantly, we are free to select different sub-emulators for space, time, processes, and parameters in this framework.

Our goal is to demonstrate this HOSVD approach for generating emulators for complex input–output structures. Such an emulator can then be used either as a pure multivariate spatiotemporal prediction model that uses only simulator runs for training, for model calibration given field data, or to perform inference on model parameters, depending on the goals of the researcher. Thus, the contributions of this methodology are:

1. It provides for a combination of different supervised learning approaches (e.g., random forests (RFs) and GP regression) that can be used to emulate the spatial, temporal, and parameter components of the model output as obtained through the HOSVD.
2. Parameters can be batched together in the tensor-decomposition, which is useful if parameters are correlated *a priori*.
3. Surrogate model predictions can be made at locations different from the set of spatial and temporal locations used in the training runs (in contrast to many SVD approaches for emulating high-dimensional output).
4. Multiple individuals in an ABM with spatiotemporal locations can be emulated efficiently in space and time.
5. Computational savings are afforded by using a reduced rank tensor factorization through the truncated HOSVD (De Lathauwer et al. 2000); a low-rank tensor factorization is usually sufficient to capture much of the variation in simulator output, just as a few principal components are usually sufficient to capture most of the variation in lower-dimensional simulator output.

The remainder of the paper is organized as follows. In Sect. 2, we briefly review aspects of SVD emulation, and tensors. In Sect. 3, we introduce a statistical model in terms of the HOSVD and describe the method for emulation based on this model. In Sect. 4, we provide empirical results of the application of the tensor-based method in two scenarios: (i) a periodic solution to the shallow ice approximation from glaciology (Bueler et al. 2005), which is used to demonstrate the advantage of predicting at locations not considered in the training, and (ii) a simulation of an ABM for collective animal movement to illustrate that the emulator can successfully capture the locations of multiple individuals moving collectively in space and time. These examples demonstrate the tensor method’s ability to apply to predictions off of a grid, flexibility in combining different unsupervised learning methods, and use in Bayesian inference. We present a brief discussion and conclusion in Sect. 5.

## 2. SVD EMULATION AND TENSOR BACKGROUND

Before introducing the statistical model and associated methodology of the HOSVD tensor-based emulator approach that is the subject of the paper, we provide a brief introduction to emulators, a review of the approach from Higdon et al. (2008) and Hooten et al. (2011), and an overview of tensors, as our approach builds on these methods.

## 2.1. EMULATORS/SURROGATES

The use of emulators, or equivalently surrogate models, has appeared widely in the uncertainty quantification and computer experiments literature, as is detailed in [Gramacy \(2020\)](#). Their principal use is to approximate computationally intensive functions at a reduced computational cost. Very broadly, they can be categorized as methods that use GPs, as in [Gu et al. \(2018b\)](#) for instance, and polynomial basis (polynomial chaos) expansions, as in [Sargsyan \(2016\)](#). In addition, others have considered surrogate models that are based on mechanistic models (e.g., [Leeds et al. 2014](#)) and neural network models (e.g., [Tripathy and Bilonis 2018](#)). This is an active area of research. For example, [Kasim et al. \(2020\)](#) recently introduced a method that uses a deep neural network that learns its own architecture for the purposes of emulation. Below, we discuss in more detail emulators that are based on modeling the right singular vectors in an SVD and come back to tensor-based emulators after a brief introduction to tensor decomposition.

## 2.2. SVD EMULATORS

The proposed tensor-based emulator is an extension of SVD-based emulators. Here we describe two approaches to SVD emulation—a GP (second-order) approach in [Higdon et al. \(2008\)](#) and a random forest (first-order) approach in [Hooten et al. \(2011\)](#). [Higdon et al. \(2008\)](#) construct Bayesian hierarchical models involving observational data in the manner of [Kennedy and O’Hagan \(2001\)](#) in order to calibrate physically important parameters and to make predictions with uncertainties. Because the physics-based computer simulators they employ are computationally expensive to evaluate, they use a basis of principal components of simulator runs (derived with the SVD) and use GP prior distributions on the weights associated with the linear combination of basis vectors derived from the SVD. [Hooten et al. \(2011\)](#) utilize a similar SVD approach for constructing an emulator, except they relax the constraint that the weights in the linear combination have GP prior distributions, and consider instead other functional forms in the mean response, such as random forests. An extension of this approach to the tensor case is one of the features of our methodology.

First, assume that computer simulator runs have been placed in a matrix,  $\mathbf{C}$ , where each column of  $\mathbf{C}$  corresponds to a distinct computer simulator run. If there are  $N$  computer simulator runs for distinct parameter values  $\boldsymbol{\theta}$ , with simulator output of dimension  $M$ , then the matrix  $\mathbf{C}$  is of dimension  $M$  by  $N$ . The first-order emulator relies on using the singular value decomposition,  $\mathbf{C} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ , where  $\mathbf{U} \in \mathbb{R}^{M \times M}$  and  $\mathbf{V} \in \mathbb{R}^{N \times N}$  are orthonormal and  $\mathbf{D} \in \mathbb{R}^{M \times N}$  is diagonal. In particular, a statistical model for the  $M$ -dimensional computer simulator output,  $\mathbf{c}$ , given a parameter,  $\boldsymbol{\theta}$ , is:

$$\mathbf{c} = \mathbf{U} \mathbf{D} \mathbf{v}(\boldsymbol{\theta}) + \boldsymbol{\epsilon}$$

for a residual term  $\boldsymbol{\epsilon}$  that has mean 0 and possibly non-diagonal covariance matrix  $\boldsymbol{\Sigma}$ . A multivariate normal random variable for  $\boldsymbol{\epsilon}$  is suggested, though is not necessary.

The central notion encapsulated by the above model is that a vector of computer simulator output can be approximately expressed as a linear combination of the columns of  $\mathbf{U} \mathbf{D}$ , and the function  $\mathbf{v}(\boldsymbol{\theta})$  specifies the  $N$  coefficients of the linear combination for a particular value

of  $\theta$ ; as such, the function  $v(\theta)$  has  $N$ -dimensional vector output. The approach of [Hooten et al. \(2011\)](#) is to model  $v(\theta)$  as  $g(\theta, \beta)$ , where  $g$  is a nonlinear regression function that involves  $\beta$  as tuning parameters and takes as input the parameter of inferential interest,  $\theta$ . Each component of the vector  $g(\theta, \beta)$  uses a nonlinear regression model, as for instance with a random forest: The training data come from the  $N$  input values of  $\theta$  for the computer simulator runs, and the outputs are the rows of  $V^T$ . In contrast, [Higdon et al. \(2008\)](#) use a prior distribution with 0 mean GPs in order to learn the coefficients for the singular vectors along with other model parameters.

Computational cost reduction stems from using less than  $N$  columns of  $UD$ , in the way that only a few principal components are used in a principal components analysis. The first few columns of  $UD$  are usually sufficient to capture most of the variation of the simulator output (in previous examples,  $r = 3$  was sufficient to capture more than 99% of the variation). This leads to a reduced computational cost for a few reasons: If the first  $r$  columns ( $< N$ ) of  $UD$  are used, then  $v(\theta)$  is  $r$ -dimensional, so only  $r$  machine learning models are needed for  $v(\theta)$ , reducing the time needed to evaluate  $v(\theta)$  and compute  $UDv(\theta)$ .

### 2.3. TENSOR BACKGROUND

A real-valued tensor is a multidimensional array of real numbers, just as a matrix is a real-valued two-dimensional array of real numbers. In this subsection, our intention is to provide a brief review of aspects of tensors that will enable the reader to understand the tensor emulator that we utilize in Sect. 3. A comprehensive reference for tensors is [Kolda and Bader \(2009\)](#).

We focus on the HOSVD since it is arguably the most natural extension of the SVD. Assume that  $X$  is a  $K$ -tensor with dimensions  $n_1, n_2, \dots, n_K$ . The higher order singular value decomposition of  $X$ , or HOSVD( $X$ ), decomposes  $X$  as  $Z \times U_1 \times U_2 \times \dots \times U_K$ . The properties of the decomposition are:

- $Z$ , referred to as the *core tensor*, is a completely orthogonal tensor whose dimensionality is the same as the dimensionality of  $X$ . Being completely orthogonal means that any two sub-tensors along the same coordinate have a 0 inner product; the inner product between real-valued tensors multiplies coordinates of the same type and then sums the resultant products.
- $U_1, \dots, U_K$  are matrices with dimensions  $n_1$  by  $n_1$ ,  $n_2$  by  $n_2$ ,  $\dots$ , and  $n_K$  by  $n_K$ , respectively. Additionally,  $U_1, \dots, U_K$  are orthogonal, analogous to the usual SVD.

In order to interpret the tensor decomposition discussed previously, one must understand tensor-matrix multiplication, which is described in terms of the familiar matrix-vector multiplication operation. If  $Z$  can be thought of as  $n_2 \times n_3 \times \dots \times n_K$  length  $n_1$  column vectors; then  $Z \times U_1$  is an  $n_1$  by  $n_2 \dots$  by  $n_K$  tensor where each column vector of length  $n_1$  in  $Z$ , denoted as  $z \in \mathbb{R}^{n_1 \times 1}$ , is replaced with  $U_1 z$ . Successive multiplication with  $U_2, \dots, U_K$  proceeds in the same manner.

The specific use of tensors for surrogate modeling appears in the statistics literature in [Pratola and Chkrebtii \(2018\)](#). Building off of this work, the method presented in Sect. 3 allows for spatiotemporal predictions that are off of the “grid” of observation locations and

also allows for the flexibility to use a combination of different supervised learning methods, in addition to GPs. Additionally, a regularized regression approach for a low-rank tensor approximation appears in [Chevreuil et al. \(2015\)](#) and a low-rank polynomial tensor basis expansion appears in [Konakli and Sudret \(2016\)](#); in contrast, the methodology presented here uses the truncated HOSVD ([De Lathauwer et al. 2000](#)) with a combination of existing supervised learning methodology (i.e., GP regression and RFs).

### 3. HOSVD TENSOR EMULATION

In this section, we develop a statistical emulator by using the HOSVD. An important advantage of this approach is being able to emulate each component of a multidimensional spatiotemporal input–output space as well as model parameters by including separate surrogate models in each dimension. Additionally, an important aspect is the selection of different machine-learning models for the different tensor components, which is explained in Sect. 3.3.

#### 3.1. MODEL FORMULATION

Without loss of generality, assume we seek to emulate a nonlinear function  $f(x, y, t, \theta_1, \dots, \theta_p)$ , where  $(x, y, t)$  are spatiotemporal coordinates and  $(\theta_1, \dots, \theta_p)$  are parameters, compactly referred to as  $\theta$ . The parameters can be either vector or scalar quantities. Assume one has evaluated  $f$  for all combinations of  $M$   $x$ -coordinates,  $N$   $y$ -coordinates,  $T$   $t$ -coordinates,  $P_1$  values of  $\theta_1$ ,  $P_2$  values of  $\theta_2$ ,  $\dots$ ,  $P_p$  values of  $\theta_p$ . (So  $M \times N \times \dots \times P_p$  evaluations of the function  $f$ .) These evaluations (e.g., computer simulator runs) can be stored in an  $M$  by  $N$  by  $T$  by  $P_1 \dots$  by  $P_p$ -dimensional tensor,  $X$ , where the  $i, j, k, l_1, \dots, l_p$  cell of the tensor stores the function evaluation under the  $i$ th  $x$  coordinate,  $j$ th  $y$  coordinate,  $k$ th  $t$  coordinate,  $l_1$ th parameter value of  $\theta_1, \dots$ , to  $l_p$ th parameter value of  $\theta_p$ .

It is important to note that there are multiple variations possible for the initial function evaluations of  $f$ . For example, the  $x$ -coordinates,  $y$ -coordinates, and  $t$ -coordinates can be considered jointly (e.g., as  $(x, y)$  tuples or  $(x, y, t)$  tuples). Furthermore, an additional spatial dimension can be considered (i.e.,  $(x, y, z)$ ) if desired. Additionally, scalar parameters  $\theta_i$  and  $\theta_j$  can be considered jointly as  $(\theta_i, \theta_j)$  tuples. Such an approach is illustrated in the first example of Sect. 4, where  $(x, y)$  tuples and amplitude and period parameters are sampled jointly. The initial values chosen to populate  $X$  may vary from application to application; for instance, a predefined lattice of  $(x, y)$  coordinates may be used in the context of a numerical PDE solver. In the absence of prior distributions from which to sample parameters, we suggest using a latin hypercube design or some hybrid variant due to its prevalence in the computer experiments literature ([Gramacy 2020](#)).

Another important variation is a scenario where one is emulating a spatiotemporal ABM, where, without loss of generality, the function  $f$  has both  $x$  positions ( $f_x$ ) and  $y$  positions ( $f_y$ ). In that scenario,  $f_x$  and  $f_y$  take the form  $f_x(i, t, \theta_1, \dots, \theta_p)$  and  $f_y(i, t, \theta_1, \dots, \theta_p)$ , where  $i$  indexes over the agents,  $t$  stands for time, and  $\theta_1, \dots, \theta_p$  represent the parameters. Tensor emulators are then constructed separately for *both*  $f_x$  (the  $x$  values of the trajectories)

and  $f_y$  (the  $y$  values of the trajectories). An example of this variation is illustrated in the second application. The tensor-based methodology proceeds in the same essential manner as follows.

Let  $\mathbf{Z} \times \mathbf{U}_1 \times \mathbf{U}_2 \times \cdots \times \mathbf{U}_{p+3}$  be the HOSVD of  $\mathbf{X}$ . The statistical model for an evaluation of  $f$  at  $(x^*, y^*, t^*, \theta^*_1, \dots, \theta^*_p)$  is:

$$f(x^*, y^*, t^*, \theta^*_1, \dots, \theta^*_p) = \mathbf{Z} \times \mathbf{u}_1(x^*) \times \mathbf{u}_2(y^*) \times \mathbf{u}_3(t^*) \times \cdots \times \mathbf{u}_{p+3}(\theta^*_p) + \epsilon,$$

for a residual term  $\epsilon$  where  $E[\epsilon] = 0$  and  $\mathbf{u}_1, \dots, \mathbf{u}_{p+3}$  are nonlinear, vector-valued functions. That is, the functions  $\mathbf{u}_1$  through  $\mathbf{u}_{p+3}$  behave like the function  $\mathbf{v}$  from the SVD approach described in Sect. 2.2. In particular,  $\mathbf{u}_1(\cdot) : \mathbb{R} \rightarrow \mathbb{R}^{1 \times M}$ ,  $\mathbf{u}_2(\cdot) : \mathbb{R} \rightarrow \mathbb{R}^{1 \times N}$ , and so on. The most basic model assumes that the error term  $\epsilon$  is independent and identically Gaussian, though that is not a necessary requirement. For instance, [Gopalan et al. \(2019\)](#) consider a multivariate random walk with spatiotemporal correlation.

### 3.2. MODEL IMPLEMENTATION

Our purpose in this section is to describe in detail how to construct an emulator (i.e., how to obtain an estimator  $\hat{f}$ ) using the HOSVD components of the tensor simulator runs.

### 3.3. EMULATOR CONSTRUCTION

Assume that  $\mathbf{X}$  is a tensor of function evaluations as in the previous section, with HOSVD that is given by  $\mathbf{Z} \times \mathbf{U}_1 \times \mathbf{U}_2 \times \mathbf{U}_3 \times \cdots \times \mathbf{U}_{p+3}$  for core tensor  $\mathbf{Z}$  and matrices  $\mathbf{U}_1$  through  $\mathbf{U}_{p+3}$ . The methodology presented is to develop estimators for functions  $\mathbf{u}_1, \dots, \mathbf{u}_{p+3}$ ; these are denoted as  $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_{p+3}$ . Then, the estimated value for  $f(x^*, y^*, t^*, \theta^*_1, \dots, \theta^*_p)$  is

$$\hat{f}(x^*, y^*, t^*, \theta^*_1, \dots, \theta^*_p) = \mathbf{Z} \times \hat{\mathbf{u}}_1(x^*) \times \hat{\mathbf{u}}_2(y^*) \times \hat{\mathbf{u}}_3(t^*) \times \cdots \times \hat{\mathbf{u}}_{p+3}(\theta^*_p).$$

First, consider the function  $\mathbf{u}_1(x)$ , which when expanded as a vector is  $(u_{11}(x), \dots, u_{1M}(x))$ , an element of  $\mathbb{R}^{1 \times M}$ . Specifically,  $\hat{u}_{11}(x)$  is obtained by using a standard supervised learning approach, such as GP regression or a random forest. The training data for learning  $u_{11}$  are the  $M$  values of  $x$  (i.e., the features) and the first column of  $\mathbf{U}_1$  (i.e., the responses). The training data for learning  $u_{12}$  are the  $M$  values of  $x$  and the second column of  $\mathbf{U}_1$  and so on, up to  $u_{1M}$ , which uses the  $M$ th column of  $\mathbf{U}_1$ . The same estimation procedure is applied for  $\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \dots$ , and  $\hat{\mathbf{u}}_{p+3}$ .

An important aspect of the above procedure is the selection of the machine learning models for learning the  $\mathbf{u}$  functions. While in practice this must be achieved with experimentation in a real-data scenario (as in the application examples that follow), one reason why this is potentially advantageous is because not all components of the  $\mathbf{u}$  functions are going to be very smooth (i.e., continuously differentiable), so considering alternate machine learning models besides a GP regression is sensible. (An example is presented in the supplementary figures.)



The package *rTensor* (Li et al. 2018) provides an option for a reduced rank tensor factorization through the truncated HOSVD (De Lathauwer et al. 2000). A reduced rank factorization does not completely recover the input tensor,  $X$ , but it approximates  $X$  as  $Z \times V_1 \times V_2 \times V_3 \times \cdots \times V_{p+3}$  where the number of columns (i.e., *ranks*) of  $V_1, V_2, \dots, V_{p+3}$  are no more than  $M, N, T, \dots, P_p$ , respectively. The number of rows, however, remain to be  $M, N, T, \dots, P_p$  of  $V_1, V_2, V_3, \dots, V_{p+3}$ , respectively. The core tensor  $Z$  has dimensions equal to the column numbers of  $V_1, V_2, \dots, V_{p+3}$ , respectively. Using the low-rank approximation is computationally advantageous for the same reason that one uses only the top few of the columns of  $UD$  for emulation in the first-order SVD case; the number of machine learning models needed is reduced based on the rank, which in turn speeds up the emulator evaluation and tensor-matrix multiplication.

### 3.4. SELECTION OF RANKS

The starting point of the procedure outlined above is to take the HOSVD decomposition of the tensor of simulator runs ( $X$ ), for which it may be computationally beneficial to use a low-rank representation via the truncated HOSVD. Here we provide guidance for selecting the ranks for the truncated HOSVD in a principled manner, if the user chooses to use a low-rank approximation. The method iteratively performs rank selection via SVD scree plots for each mode of the tensor after tensor unfolding. Specifically, the suggested method is as follows:

1. For each mode of the tensor  $X$ , unfold the tensor. For example, if the tensor has dimensions  $n_1, n_2, n_3$ , and  $n_4$ , unfolding the tensor by the first mode yields a matrix with  $n_1$  rows and  $n_2 n_3 n_4$  columns, where row  $i$  groups all of the tensor elements where the first mode is  $i$  (of  $n_1$ ). The tensor unfolding operation, along with a computational implementation, is described in Li et al. (2018).
2. For each mode of  $X$ , take the SVD of the unfolded and centered tensor and use a standard procedure to determine the total number of singular vectors to keep. (As is usual, centering refers to ensuring column sample means are 0.) Specifically, we use a scree plot of singular values along with a comparison to singular values of a randomized matrix (each column shuffled) to determine a singular value threshold for calculating the total number of singular vectors to keep, as suggested in Chapter 14 of Hastie et al. (2009) for PCA.
3. For each mode, set the rank in the truncated HOSVD to the number of retained singular vectors from the previous step, except adding 1 to account for the mean that is subtracted when centering the data.

To check the quality of the low rank tensor decomposition that is found with the previous steps, we suggest calculating the proportion of variance explained by the low rank decomposition with the familiar formula:

$$1 - RSS/TSS,$$



where RSS stands for residual sum of squares and TSS stands for total sum of squares. Here, RSS is the square of the Frobenius norm of the residual tensor, and TSS is the squared Frobenius norm of the training tensor minus the grand mean. That is – let  $\mathbf{X}$  be the training tensor as usual, and let  $\mathbf{R}$  be the low rank approximation via the truncated HOSVD. Finally, let  $\mathbf{X}_c$  be the tensor formed by subtracting every element of  $\mathbf{X}$  by  $\bar{X}$ , a scalar that is the mean of all elements of  $\mathbf{X}$  (i.e., the grand mean). Then

$$\begin{aligned} RSS &= \|\mathbf{X} - \mathbf{R}\|_F^2, \\ TSS &= \|\mathbf{X}_c\|_F^2, \end{aligned}$$

where the subscript  $F$  denotes the Frobenius norm. The proportion of variance explained by the low-rank approximation is more interpretable than the Frobenius norm of the residual.

## 4. EXAMPLES

We consider two examples to demonstrate some of the advantages of the HOSVD tensor emulator. Specifically, in Sect. 4.1 we consider an emulator for a periodic exact solution to the shallow ice approximation partial differential equation (SIA PDE) from [Bueler et al. \(2005\)](#), and in Sect. 4.2 we consider an emulator for an ABM simulation of collective animal movement. Both examples show the advantages of considering different surrogates for spatial components and parameters. In addition, the SIA PDE example demonstrates how our emulator approach can generate spatial predictions at spatial locations and time points that were not observed, and the collective movement ABM example demonstrates that our approach can reasonably emulate the locations in space and time of multiple individuals. Both experiments make use of the *rTensor* ([Li et al. 2018](#)) R package to find the HOSVD and to perform tensor matrix multiplication.

### 4.1. OSCILLATING GLACIER EXAMPLE

The SIA PDE is a commonly used mathematical model that describes the time evolution of glacier thickness in glaciers that are shallow; e.g., see [Aðalgeirsdóttir \(2003\)](#), [Flowers et al. \(2005\)](#), [Jarosch et al. \(2013\)](#), [Guan et al. \(2016\)](#), [Werder et al. \(2020\)](#). The main physical principle utilized in this model is mass conservation, and the shallowness approximation allows one to ignore stress terms in formulating the PDEs.

Here we use a HOSVD tensor-based emulator for a periodic solution to the SIA PDE described in [Bueler et al. \(2005\)](#). The periodic solution is a dome that oscillates according to a particular period and amplitude. This periodic function is used in [Gopalan et al. \(2019\)](#) and is repeated here for completeness:

$$\begin{aligned} H(r, t) &= H_s(r) + P(r, t), \\ P(r, t) &= C_p \sin(2\pi t/T_p) \cos^2 \left[ \frac{\pi(r - 0.6L)}{.6L} \right]; \text{ if } 0.3L < r < .9L, \\ P(r, t) &= 0; \text{ if } 0 \leq r \leq 0.3L \text{ or if } r \geq 0.9L. \end{aligned}$$

In this formulation,  $H$  stands for glacier thickness and is a function of distance from the origin,  $r$ , and time,  $t$ ; in particular, this is a radially symmetric solution. The term  $H_s$  (Eq. 21 of [Bueler et al. \(2005\)](#)) corresponds to a static profile that is added to a time-varying periodic function,  $P$ , that has an amplitude  $C_p$  and period  $T_p$ . Here,  $L$  refers to the length of the glacier profile. Overall, the solution looks like a dome with a periodic oscillation in thickness within the domain  $0.3L < r < .9L$ . The parameters that we focus on for this example are  $C_p$ , the amplitude of the periodic perturbation, and  $T_p$ , the period of the perturbation. For more details of this model and the periodic solution, see [Bueler et al. \(2005\)](#).

#### 4.1.1. Glacier Example Experimental Design

To demonstrate the HOSVD tensor emulator for the glacier dynamics problem, we consider the following experiment. For  $s = 10, 20, 30$ :

1.  $s^2$  2-tuples of  $(x, y)$  are drawn using latin hypercube sampling on the range  $-5 \times 10^5 m$  to  $5 \times 10^5 m$  for the  $x$  and  $y$  location values.
2.  $s$  values of  $t$  are sampled uniformly from 0 to 10000 years.
3.  $s^2$  2-tuples of (period, amplitude) are drawn using latin hypercube sampling on the range (1000 to 5000) years for period and (100 to 400) meters for amplitude.
4. For each combination  $(x, y)$ , (period, amplitude), and  $t$ , the glacier thickness function that is to be emulated is evaluated and the result is stored in a tensor,  $\mathbf{X}$ . The dimensions of the tensor are  $s^2$  by  $s$  by  $s^2$ , so there are  $s^5$  total elements in the tensor.
5. The methodology from Sect. 3 is applied to the tensor  $\mathbf{X}$ . The combinations of machine learning models tried are all random forests for the spatial, temporal, and parameter components, all GPs and RFs for the parameter component, but GPs for the spatial and temporal components. We implemented these machine learning models with the *kernlab* R package ([Karatzoglou et al. 2004](#)) and the *randomForest* R package ([Liaw and Wiener 2002](#)).
6. The predictive accuracy of the emulator is assessed by calculating the mean absolute relative error on a test set. Absolute relative error is  $|(H_{true} - H_{pred})/H_{true}|$ . The test set is randomly selected with the same design steps as above, and a test set of size 100 is used; because the test  $(x, y)$  and  $t$  values are not necessarily the same as the training  $(x, y)$  and  $t$  values, this example illustrates the capability of the tensor-based method to emulate locations and times that were not in the training sample.

Note that since  $(x, y)$  and (period, amplitude) are considered jointly, we end up with only three  $\mathbf{U}$  matrices,  $\mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3$ , where  $\mathbf{U}_1$  corresponds to  $(x, y)$ ,  $\mathbf{U}_2$  corresponds to  $t$ , and  $\mathbf{U}_3$  corresponds to (period, amplitude). Additionally, we used the truncated HOSVD from *rTensor* and the rank selection procedure detailed in Section 3.4. The supplementary figures show the scree plots used for determining ranks, and the proportion of variance explained by the low rank approximation is .99.

#### 4.1.2. Glacier Example Results

The results of the previously described experiment are illustrated in Fig. 1. The  $x$ -axis of this plot represents a comparison of the mean absolute relative error over all sizes, which indicates that the absolute relative error decreases as the size of the training tensor increases. In addition, this plot shows that the HOSVD approach that considers mixed machine learning models yields the smallest error. Specifically, the emulator that included a GP surrogate for  $(x, y)$ , GP surrogate for  $t$ , and random forest surrogate for (period, amplitude) performed better than pure random forest or GP approaches in terms of error. Additionally, we repeated the experimental design ten times in order to assess the variability in the predictions with respect to the initial input  $(x, y)$ ,  $t$ , and (period, amplitude) values. Boxplots comparing the absolute relative errors are shown in Fig. 2, again showing that for this example, the mixture of machine learning surrogates tends to perform best. In addition, a comparison of the actual and emulated values of glacial thickness using the mixture of machine learning models tensor emulator is shown in Fig. 3. This shows that there is generally a good agreement between the actual thickness and predicted thickness across thickness values, with smallest errors towards the higher glacier thickness values; one possible explanation for this behavior is that the magnitude of periodic movement for this particular example is largest within the interior of the glacier, at lower thickness values.

Finally, for an assessment of how the HOSVD tensor-based emulator compares to simply using a standard off-the-shelf supervised learning approach, we compared the absolute relative errors of a GP regression (*kernlab* R package (Karatzoglou et al. 2004)) and a




Figure 1. Comparison of tensor size and mean absolute relative error, for  $s = 10, 20$ , and  $30$  (see Section 4.1.1)—generally, the error decreases with a larger tensor and also is smallest with a combination of machine learning models.




Figure 2. Comparison of mean absolute relative error for a pure random forest, pure GP, and a GP for  $(x, y, t)$  and random forest for (period, amplitude). The mixture of machine learning models generally performs the best.




Figure 3. Comparison of emulated and actual glacial thickness values, using the tensor-based emulator approach with a mixture of machine learning models.

random forest (*randomForest* R package (Liaw and Wiener 2002)). In each of these cases, we did not consider a tensor decomposition. Specifically,  $10^4$  5-tuples of  $(x, y, t, \text{period}, \text{amplitude})$  values were sampled with a latin hypercube design to create the training dataset. In all cases, the errors were larger than using the tensor decomposition; specifically, a random forest produced a mean of 0.0488 absolute relative error (SD: 0.0015), and a GP produced a mean of 0.0387 absolute relative error (SD: 0.000604). Additionally, training data of at least  $2 \times 10^4$  training examples resulted in a memory error, suggesting that the tensor emulator method is a more memory efficient procedure.

## 4.2. COLLECTIVE ANIMAL MOVEMENT ABM SIMULATION

This example considers a collective animal movement ABM based on a simplified version of the model in Couzin et al. (2002) in which the animals (agents) can exhibit a variety of realistic behaviors such as grouping and repulsion. We first describe the agent-based model for collective behavior, followed by the associated tensor-based emulator, and conclude with an example of using this emulator to learn collective behavior on a real-world dataset within a Bayesian inferential paradigm.

### 4.2.1. Collective Movement Agent-Based Model Simulator

Assume we have observed location vectors at time  $t = 1, \dots, T$  for the  $i$ th individual ( $i = 1, \dots, N$ ), denoted by  $\mathbf{s}_{i,t} = (x_{i,t}, y_{i,t})'$ . Our basic model decomposes movement into a unit direction vector and scalar speed for each individual at each discrete time. Specifically, we model the evolution of an agent's location by

$$\mathbf{s}_{i,t} = \mathbf{s}_{i,t-1} + \mathbf{d}_{i,t} v_{i,t} + \boldsymbol{\eta}_{i,t}, \quad \boldsymbol{\eta}_{i,t} \sim N(0, \sigma_{\eta}^2 \mathbf{I}_2),$$

where  $\mathbf{d}_{i,t}$  is the unit direction vector, which is a deterministic (black box) function of the most recent past location of all agents and associated parameters (see below), and  $v_{i,t}$  is the speed of the agent. In this formulation there is a stochastic component to the location update represented by the  $\boldsymbol{\eta}_{i,t}$  error term. For the simple example presented here, we further assume that the speed is constant for all agents across time,  $v_{i,t} = v$ , and the primary focus is on the unit direction vector,  $\mathbf{d}_{i,t}$ .

Couzin et al. (2002) note that an important consideration for animal movement is that the animals do not collide (or, rather, that they have some desired personal space). In addition, animals may favor orientation and attraction to other animals within their perception. That is, there is an intermediate zone beyond the no-collision zone in which animals seek to orient with each other (move in the same direction), and beyond this zone they are attracted to other animals. In the simple implementation presented here, we consider a no-collision zone and an orientation zone and assume that individual animals outside of these zones continue moving in the direction they were going at the previous time.

Let  $\delta_{i,j} = \|\mathbf{s}_{j,t-1} - \mathbf{s}_{i,t-1}\|$  be the Euclidean distance between agent  $i$  and agent  $j$  at time  $t - 1$ . Then, we define the deterministic function for the (un-normalized) direction

vector for the  $i$ th agent at time  $t$  by one of the following:

$$\tilde{\mathbf{d}}_{i,t} = \begin{cases} -\sum_{j \in \mathcal{N}_{\alpha,-i}} \frac{s_{j,t-1} - s_{i,t-1}}{\delta_{i,j}}, & \text{if } \exists j \neq i : \delta_{i,j} < \alpha \\ \sum_{j \in \mathcal{N}_{\alpha,\rho_o,i}} \frac{\mathbf{d}_{j,t-1}}{\|\mathbf{d}_{j,t-1}\|} & \alpha \leq \delta_{i,j} \leq \rho_o \\ \mathbf{d}_{i,t-1} & \delta_{i,j} > \rho_o, \end{cases}$$

where  $\alpha$  is the no-collision radius,  $\rho_o$  is the orientation zone radius,  $\mathcal{N}_{\alpha,-i}$  corresponds to all neighbors of  $i$  (not including the  $i$ th individual) within a distance  $\alpha$  of  $s_{i,t-1}$ , and  $\mathcal{N}_{\alpha,\rho_o,i}$  corresponds to the set of neighbors of  $i$  within the distance  $(\alpha, \rho_o]$  of the  $i$ th individual (including the  $i$ th individual). Thus, the first term corresponds to the no-collision term and applies if *any* individual agent is within a distance of  $\alpha$  of the  $i$ th individual. The second term corresponds to orientation with individuals in a zone between  $(\alpha, \rho_o]$  of the  $i$ th individual (assuming the first condition does not hold). If neither of these conditions hold, the  $i$ th individual keeps the same direction as before (the third term). Lastly, we convert this direction vector to a unit vector

$$\mathbf{d}_{i,t} = \frac{\tilde{\mathbf{d}}_{i,t}}{\|\tilde{\mathbf{d}}_{i,t}\|}.$$

In the simulations considered here, we fixed the no-collision radius ( $\alpha$ ) and the location variance ( $\sigma_\eta^2$ ) and focus on the speed ( $v$ ) and radius of orientation ( $\rho_o$ ), which are the primary drivers of collective behavior.

#### 4.2.2. Tensor-Based Emulator

We train a tensor-based emulator in the following manner:

1.  $2^5$  values of  $v$  are sampled uniformly in the range  $[.1,1]$ , and  $2^5$  values of  $\rho_o$  are sampled uniformly in the range  $[5,50]$ . The parameter  $\alpha$  is fixed at  $.5$  and  $\sigma_\eta^2$  is fixed at  $.025$ .
2. For each combination of  $v$  and  $\rho_o$ , and for each time point, the simulator is evaluated for both the  $x$  position and  $y$  position of the 20 animals. The results of the  $x$  positions are stored in a 20 by 101 by 32 by 32 tensor,  $\mathbf{X}$ ; likewise, the results of the  $y$  positions are stored in a 20 by 101 by 32 by 32 tensor,  $\mathbf{Y}$ . The coordinates of these tensors are individual index, time index,  $v$  parameter, and  $\rho_o$  parameter.
3. A tensor emulator is derived using the methodology from Sect. 3. In particular,  $HOSVD(\mathbf{X}) = \mathbf{Z} \times \mathbf{U}_1 \times \mathbf{U}_2 \times \mathbf{U}_3 \times \mathbf{U}_4$ , and the predicted  $x$  positions for the 20 animals over the 101 times points for parameters  $v^*$  and  $\rho_o^*$  is then  $\mathbf{Z} \times \mathbf{U}_1 \times \mathbf{U}_2 \times \hat{\mathbf{u}}_3(v^*) \times \hat{\mathbf{u}}_4(\rho_o^*)$ , where  $\hat{\mathbf{u}}_3(\cdot)$  and  $\hat{\mathbf{u}}_4(\cdot)$  are functions learned with a machine learning model (see Sect. 3.3). As in the previous experiment, we considered random forests and GPs, using the default implementations from [Liaw and Wiener \(2002\)](#) and [Karatzoglou et al. \(2004\)](#), respectively. The  $y$  positions are emulated with the same procedure, starting with the tensor  $\mathbf{Y}$ .

Two test cases are considered to assess the emulator output for the  $x$  and  $y$  trajectories. For Case 1, we let  $v = .5$  and  $\rho_o = 35$ , which forces the animals to move together in a northeast direction given the high degree of collectiveness implied by the larger  $\rho_o$  parameter. For test Case 2, we let  $v = .5$  and  $\rho_o = 5$ , which implies less collectiveness due to the smaller  $\rho_o$  parameter, and two groups of animals move in opposite directions. In our assessment of the emulator approaches, we plot the trajectories over time and compare to the original simulated trajectories. Additionally, we compare the simulated and emulated trajectories with measures of collective animal behavior to gain a more nuanced comparison. Specifically, the two quantitative measures used to compare simulated and emulated animal movements through time are the troop spread and troop elongation metrics (Strandburg-Peshkin et al. 2017). Troop spread measures how spread apart the group is over time, and troop elongation measures how ellipsoidal the group is over time.

Figure 4 shows a comparison between the simulated and emulated trajectories for test Case 1. We can see that the animals move in a collective group and that the emulated trajectories are quite similar to the true trajectories. This is further demonstrated in Fig. 5, which shows a close match for both troop spread and troop elongation between both emulated trajectories and the actual, simulated trajectories. We also compare two emulator approaches with this example. The first is a pure RF for all components, and the second is a multi-model emulator with a RF for the  $y$  positions, but a mix of RF for the  $\rho_0$  component and GP for the  $v$  component of the  $x$  positions. That is, the function  $\hat{u}_3(v)$  is learned with a GP and  $\hat{u}_4(\rho_o)$  is learned with RF regression (for the  $x$  positions); we do not restrict these functions to be learned with the same machine learning approach (i.e., RF). Both the location plots in Fig. 4 and comparison of troop spread and troop elongation metrics in Fig. 5 show that there is generally close agreement between both emulator approaches and the simulated trajectories.

In Case 2, simple visual inspection of both emulator trajectories in Fig. 6 makes it difficult to conclude which emulator (i.e., pure RF or combined RF/GP) emulates the true trajectories closest. However, comparisons of troop spread and troop elongation in Fig. 7 generally show that the combined RF/GP emulator better captures the simulator than the pure RF surrogate emulator, further illustrating the potential utility of a combined RF/GP approach over using a single machine learning model.

Additionally, we use Case 1 and Case 2 to make a comparison to a widely-used GP emulator, `ppgasp` from the `RobustGaSP` R package (Gu et al. 2018a). Specifically, for both synthetic test cases we calculate the mean distances between the emulated trajectories and the simulated trajectories for all of the animals and over all time points, and refer to this as mean error. It appears that the tensor emulator is more accurate in both Case 1 and Case 2, though with most improvement in Case 2. However, `ppgasp` has a much faster emulator evaluation time in both test cases. These results are summarized in Table 1. A faster compute time for emulator evaluation could be achieved with a lower rank HOSVD decomposition, though that would likely come with some accuracy trade-off. Additionally, an important consideration is that output is restricted to be on a temporal grid for all of the animals for the sake of comparison, though in contrast to `ppgasp`, the HOSVD emulator handles the scenario where time is not assumed to be on a discrete grid.



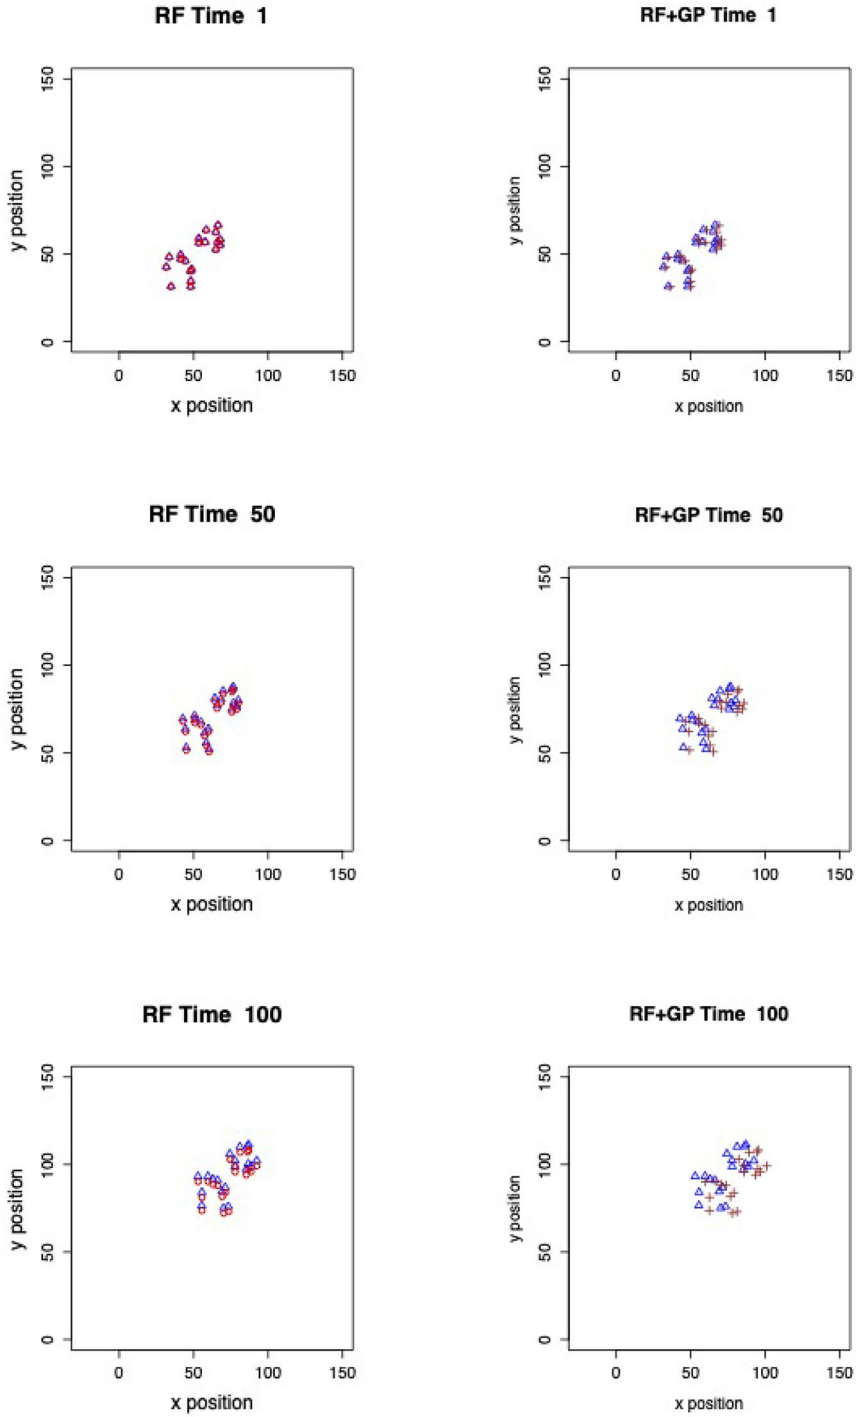


Figure 4. Comparison of simulated (triangles) and emulated trajectories for test Case 1, with the pure RF emulator on the left (circles) and the combined model(RF+GP for  $x$  positions) emulator on the right (crosses). Both emulator approaches appear to match the simulated positions.




Figure 5. Case 1 comparison of troop spread and troop elongation metrics amongst the simulator, pure RF emulator, and multi-model emulator. There is close agreement for both emulators with the simulated trajectories.

#### 4.2.3. Bayesian Learning with the Tensor-Based Emulator

In order to test the emulator on real data in a Bayesian inferential setting, we used the emulator version of our simplified animal movement ABM on guppy (*Poecilia reticulata*) data; specifically, experiment fm7 from [Bode et al. \(2012\)](#). These data come from an experiment using a captive population of guppies. In this experiment, groups of ten guppies of the same sex were filmed from above in a square tank in which one corner contained gravel and shade, which is presumed to be attractive to the guppies because it provides shelter. The guppies were released in the tank in the opposite (lower-right) corner. The data consist of movement trajectories truncated to the time points when all individuals were moving until one guppy reached the shaded target area.

In particular, we assume the following data model for the guppy trajectories:

$$\begin{aligned} O_{itx} &= \hat{f}_x(i, t, v, \rho_0) + \epsilon_{itx}, \quad \epsilon_{itx} \sim iid N(0, \sigma_x^2) \\ O_{ity} &= \hat{f}_y(i, t, v, \rho_0) + \epsilon_{ity}, \quad \epsilon_{ity} \sim iid N(0, \sigma_y^2), \end{aligned}$$

where  $O_{itx}$  and  $O_{ity}$  are observed  $x$  and  $y$  locations for the  $i$ th guppy at time  $t$ , respectively, and  $\hat{f}_x$  and  $\hat{f}_y$  are tensor emulators constructed in the same manner as the previous section (trained with runs from the animal movement simulator described

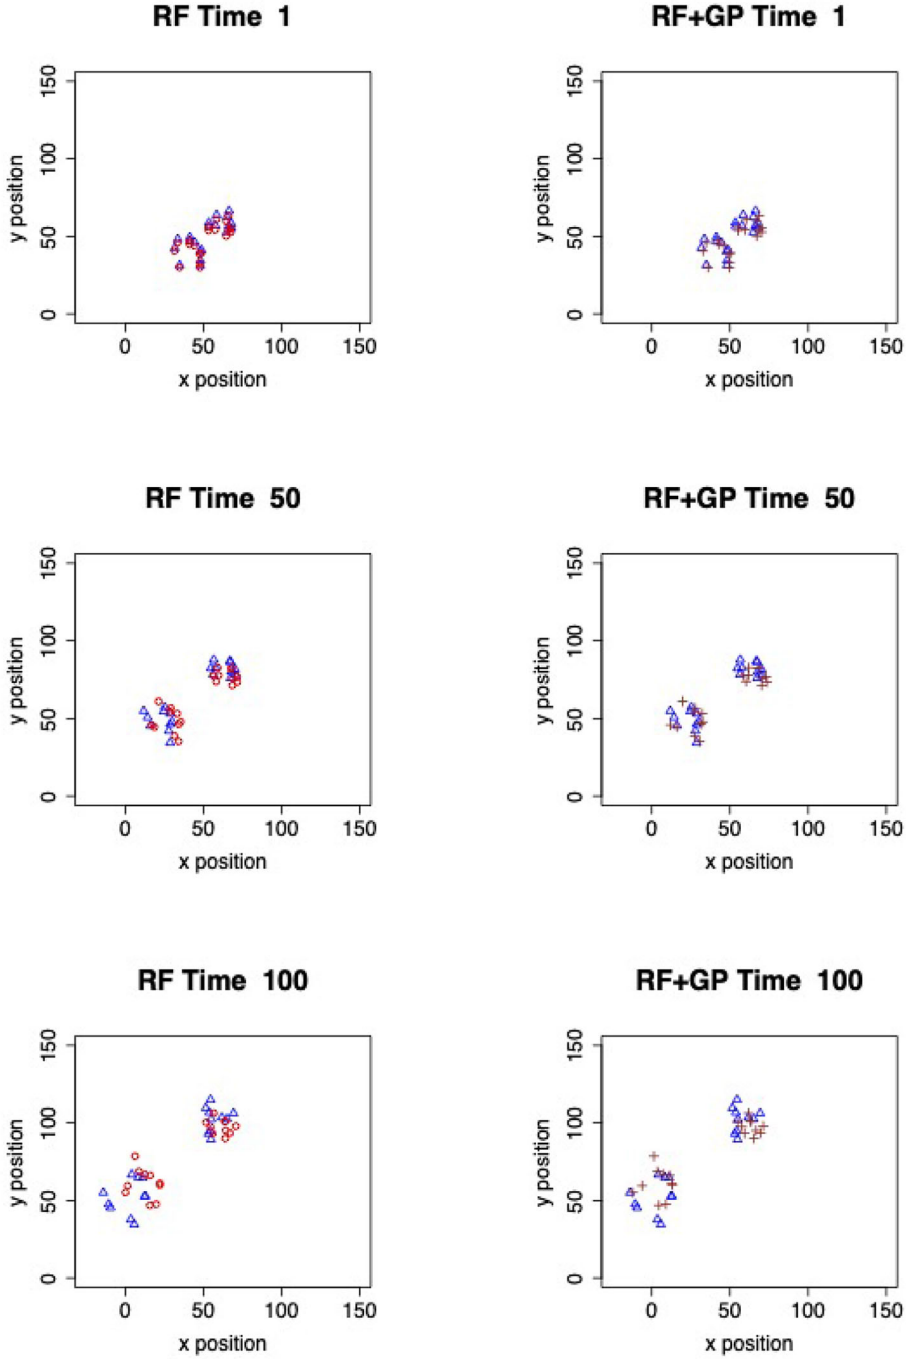


Figure 6. Comparison of simulated (triangles) and emulated trajectories for test Case 2, with the pure RF emulator on the left (circles) and the combined model(RF+GP for  $x$  positions) emulator on the right (crosses); the tensor emulator matches the animal trajectories closely. Visually, it is clear that neither emulator matches the simulator as closely as in Case 1, though it is still a good match in terms of capturing the essential collective behavior.



**Case 2: Troop spread for simulator and emulator****Case 2: Troop elongation for simulator and emulator**

Figure 7. Case 2 comparison of troop spread and troop elongation metrics amongst the simulator, pure RF emulator, and multi-model emulator. The troop spread and troop elongation profiles are best represented by the multi-model emulator.

Table 1. Comparison of run-time and accuracy of the ppgasp and HOSVD emulators on two synthetic ABM test cases. Run-times are listed for evaluation of emulator

Case	Emulator	Mean Error	Run-time
1	HOSVD	1.678	0.245
1	ppgasp	2.762	0.0072
2	HOSVD	5.969	0.2423
2	ppgasp	22.33	0.0077
	Units	m	s

previously). Additionally, we assume the following prior distributions on parameters:  $\rho_o \sim \text{Uniform}(5, 60)$ ,  $v \sim \text{Uniform}(1, 5)$ ,  $\sigma_x^2 \sim \text{InverseGamma}(\text{shape}=3, \text{scale}=4)$ , and  $\sigma_y^2 \sim \text{InverseGamma}(\text{shape}=3, \text{scale}=4)$ .

We implement a Gibbs sampler for sampling from the posterior using inverse-gamma draws (due to conjugacy) for  $\sigma_x^2$  and  $\sigma_y^2$  and grid sampling steps for  $\rho_o$  and  $v$ . Traceplots, given in the supplementary materials, show good mixing, and the resultant posterior dis-




Figure 8. Posterior densities for both  $\rho_o$  and  $v$  are illustrated, which are inferred using guppy data (fm7) from [Bode et al. \(2012\)](#). The posterior for both parameters concentrates in a much more narrow range than the uniform prior, suggesting Bayesian learning.

tributions for  $\rho_o$  and  $v$  are illustrated in Fig. 8. The relatively narrow regions where mass concentrates (in comparison with the uniform priors) suggests that there is definite Bayesian learning of these parameters, though since this is real data, there are no ground-truth parameters to compare to. However, the posterior mean of the  $\rho_o$  orientation distance parameter suggests that there is a substantial amount of collective behavior by these guppies in this experiment. A plot of the simulated trajectory in comparison with the emulated trajectory (using the posterior mean of  $\rho_o$  and  $v$ ) is included in Fig. 9 and indicates overall agreement in the patterns of collective behavior.

## 5. CONCLUSION

The objective of this paper has been to introduce a method to construct an emulator for a complex spatiotemporal function (e.g., computer simulator) using tensor decomposition with potentially different surrogate models for each tensor dimension. In particular, the method is essentially a variant of a SVD-based emulator approach ([Higdon et al. 2008](#)) that uses the HOSVD instead of an SVD. The distinguishing features of the tensor-based method for emulation are:

- The method can predict the function output at spatial locations and time points not considered in the training sample, in contrast to many existing emulator methods.
- A variety of supervised learning approaches (e.g., GP regression and random forests) can be combined for the spatial, temporal, and parameter components. This allows greater flexibility and the ability to combine the strengths (and disadvantages) of different machine learning approaches in constructing an emulator, i.e., a “mix and match” approach.

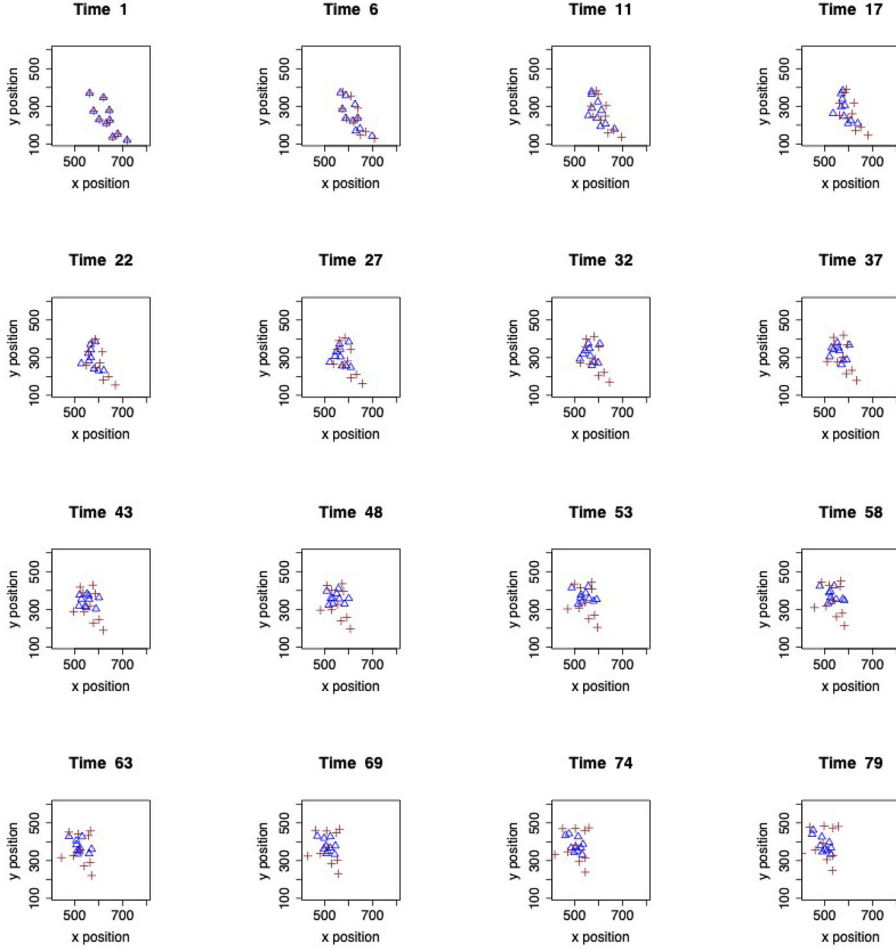


Figure 9. Simulated trajectories for guppy data (fm7 from [Bode et al. \(2012\)](#)) using posterior means of  $\rho_0$  and  $v$  compared to the real data, which indicate a similar pattern in collective behavior. The triangles are the actual trajectories, and the crosses are the simulated trajectories using the posterior means of parameters.

- The method can provide efficient emulation of the spatiotemporal locations of multiple individuals in an ABM.
- By using a reduced rank tensor factorization with the truncated HOSVD ([De Lathauwer et al. 2000](#)), computational savings are achieved.

The method has been demonstrated via application in two scenarios: a periodic solution to a shallow ice approximation partial differential equation from glaciology and a collective animal movement ABM simulator. In both scenarios, combining supervised learning approaches yields an accurate emulator, demonstrating the advantage of model flexibility when different machine learning approaches are combined. It is conceivable the same results will hold for more complex supervised learning methods such as deep learning; in fact, an advantage of the proposed approach is that it can handle the use of a variety of machine learning models, and so is adaptable to the rapid advances being made in the field. While

we conducted some preliminary experiments using the `neuralnetwork` R package with default options, the results were generally poorer than the `randomForest` and `kernlab` defaults for random forests and GP regression; presumably more advanced neural network architectures will show neural networks to be much more competitive. An additional strength of the methodology presented is that it has been implemented with R packages available on the Comprehensive R Archive Network (CRAN), so a tensor-based emulator can be directly included in an application without much difficulty, for instance in a spatiotemporal Bayesian hierarchical model as in [Gopalan et al. \(2019\)](#). This was illustrated in the collective animal movement example with an experiment that documented the movement of guppies in a tank. The Bayesian implementation with the tensor-based emulator suggested that the guppies exhibit a high-degree of collectiveness in this experiment. While we have not covered emulator uncertainty in this treatment, we believe the literature on model discrepancy ([Kennedy and O'Hagan 2001](#); [Brynjarsdóttir and O'Hagan 2014](#); [Gopalan et al. 2019](#)) can be used to construct hierarchical models involving an emulator that accounts for emulator uncertainty and inaccuracy.

## ACKNOWLEDGEMENTS

This work was partially supported by the U.S. National Science Foundation (NSF) through NSF Award DMS-1811745.

*[Received October 2020. Revised June 2021. Accepted June 2021.]*

## REFERENCES

- Aðalgeirsdóttir G (2003) Flow dynamics of Vatnajökull ice cap, Iceland, *Mitteilungen der Versuchsanstalt für Wasserbau, Hydrologie und Glaziologie an der Eidgenössischen Technischen Hochschule Zurich*
- Berliner LM (1996) Hierarchical Bayesian time series models. In: Hanson KM, Silver RN (eds) *Maximum entropy and Bayesian methods*. Springer, Dordrecht
- Berliner LM (2003) Physical-statistical modeling in geophysics. *J Geophys Res Atmos* 108:2003. <https://doi.org/10.1029/2002JD002865>
- Bode NW, Franks DW, Wood AJ, Piercy JJ, Croft DP, Codling EA (2012) Distinguishing social from nonsocial navigation in moving animal groups. *Am Nat* 179:621–632
- Brynjarsdóttir J, O'Hagan A (2014) Learning about physical parameters: the importance of model discrepancy. *Inverse Probl* 30:114007. <http://stacks.iop.org/0266-5611/30/i=11/a=114007>
- Bueler E, Lingle CS, Kallen-Brown JA, Covey DN, Bowman LN (2005) Exact solutions and verification of numerical models for isothermal ice sheets. *J Glaciol* 51:291–306
- Chevreuril M, Lebrun R, Nouy A, Rai P (2015) A least-squares method for sparse low rank approximation of multivariate functions. *SIAM/ASA J Uncertain Quantif* 3:897–921
- Couzin ID, Krause J, James R, Ruxton GD, Franks NR (2002) Collective memory and spatial sorting in animal groups. *J Theor Biol* 218:1–11
- De Lathauwer L, De Moor B, Vandewalle J (2000) A multilinear singular value decomposition. *SIAM J Matrix Anal Appl* 21:1253–1278
- Fadikar A, Higdon D, Chen J, Lewis B, Venkatramanan S, Marathe M (2018) Calibrating a stochastic, agent-based model using quantile-based emulation. *SIAM/ASA J Uncertain Quantif* 6:1685–1706



- Flowers GE, Marshall SJ, Björnsson H, Clarke GK (2005) Sensitivity of Vatnajökull ice cap hydrology and dynamics to climate warming over the next 2 centuries. *J Geophys Res Earth Surface* 110:2005
- Gopalan G, Hrafnkelsson B, Wikle CK, Rue H, Aðalgeirsdóttir G, Jarosch AH, Pálsson F, (2019) A hierarchical spatiotemporal statistical model motivated by glaciology. *J Agric Biol Environ Stat* 24:669–692
- Gramacy RB (2020) Surrogates: Gaussian process modeling, design and optimization for the applied sciences. Chapman Hall/CRC, Boca Raton, Florida. <http://bobby.gramacy.com/surrogates/>
- Gu M, Palomo J, Berger JO (2018a) RobustGaSP: Robust Gaussian stochastic process emulation in R, arXiv preprint [arXiv:1801.01874](https://arxiv.org/abs/1801.01874)
- Gu M, Wang X, Berger JO (2018b) Robust Gaussian stochastic process emulation. *Ann Stat* 46:3038–3066. <https://doi.org/10.1214/17-AOS1648>
- Guan Y, Haran M, Pollard D (2016) Inferring ice thickness from a glacier dynamics model and multiple surface datasets, arXiv e-prints
- Hastie T, Tibshirani R, Friedman J (2009) The elements of statistical learning: data mining, inference, and prediction. Springer, Berlin
- Higdon D, Gattiker J, Williams B, Rightley M (2008) Computer model calibration using high-dimensional output. *J Am Stat Assoc* 103:570–583
- Hooten MB, Leeds WB, Fiechter J, Wikle CK (2011) Assessing first-order emulator inference for physical parameters in nonlinear mechanistic models. *J Agric Biol Environ Stat* 16:475–494. <https://doi.org/10.1007/s13253-011-0073-7>
- Jarosch AH, Schoof CG, Anslow FS (2013) Restoring mass conservation to shallow ice flow models over complex terrain. *Cryosphere* 7:229–240, <https://www.the-cryosphere.net/7/229/2013/>
- Karatzoglou A, Smola A, Hornik K, Zeileis A (2004) kernlab—An S4 Package for Kernel Methods in R. *J Stat Softw* 11:1–20, <http://www.jstatsoft.org/v11/i09/>
- Kasim M, Watson-Parris D, Deaconu L, Oliver S, Hatfield P, Froula D, Gregori G, Jarvis M, Khatiwala S, Korenaga J, et al (2020) Up to two billion times acceleration of scientific simulations with deep neural architecture search, arXiv preprint [arXiv:2001.08055](https://arxiv.org/abs/2001.08055)
- Kennedy MC, O'Hagan A (2001) Bayesian calibration of computer models. *J R Stat Soc Ser B (Stat Methodol)* 63:425–464
- Kolda TG, Bader BW (2009) Tensor decompositions and applications. *SIAM Rev* 51:455–500
- Konakli K, Sudret B (2016) Polynomial meta-models with canonical low-rank approximations: Numerical insights and comparison to sparse polynomial chaos expansions. *J Comput Phys* 321:1144–1169
- Kuhnert PM (2014) Physical-statistical modeling. *Statistics Reference Online, Wiley StatsRef*, pp 1–5
- Leeds WB, Wikle CK, Fiechter J (2014) Emulator-assisted reduced-rank ecological data assimilation for nonlinear multivariate dynamical spatio-temporal processes. *Stat Methodol* 17:126–138
- Li J, Bien J, Wells M (2018) rTensor: An R Package for multidimensional array (tensor) unfolding, multiplication, and decomposition. *J Stat Softw* 87:1–31, <https://www.jstatsoft.org/v087/i10>
- Liaw A, Wiener M (2002) Classification and Regression by random Forest, *R News*, 2, 18–22, <https://CRAN.R-project.org/doc/Rnews/>
- Pratola MT, Chkrebtii O (2018) Bayesian calibration of multistate stochastic simulators. *Stat Sin* 2018:693–719
- Salter JM, Williamson DB, Scinocca J, Kharin V (2019) Uncertainty quantification for computer models with spatial output using calibration-optimal bases. *J Am Stat Assoc* 114:306:1–24. <https://doi.org/10.1080/01621459.2018.1514306>
- Sargsyan K (2016) Surrogate models for uncertainty propagation and sensitivity analysis. Springer, Cham, pp 1–26
- Strandburg-Peshkin A, Farine DR, Crofoot MC, Couzin ID (2017) Habitat and social factors shape individual decisions and emergent group structure during baboon collective movement. *Elife* 6:e19505
- Tripathy RK, Bilionis I (2018) Deep UQ: learning deep neural network surrogate models for high dimensional uncertainty quantification. *J Comput Phys* 375:565–588
- Werder MA, Huss M, Paul F, Dehecq A, Farinotti D (2020) A Bayesian ice thickness estimation model for large-scale applications. *J Glaciol* 66:137–152

Wikle CK, Hooten MB (2010) A general science-based framework for dynamical spatio-temporal models. *Test* 19:417–451

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.