Mathematical Biosciences and Engineering

https://www.aimspress.com/journal/mbe

MBE, 23(7): 2178–2207.
DOI: 10.3934/mbe.2026079
Received: 10 March 2026
Revised: 30 June 2026
Accepted: 02 July 2026
Published: 16 July 2026

Research article

Using active subspaces to explore discrepancies between global and local parameter sensitivities in a Lotka-Volterra system

Huiyan Zou and Allison L. Lewis*

Department of Mathematical Sciences, Lafayette College, 730 High Street, Easton, PA 18042, USA

* **Correspondence:** Email: lewisall@lafayette.edu; Tel: +16103305276.

Abstract: Global sensitivity metrics are essential tools for assessing parameter importance in complex models, particularly when precise information about parameter values is unavailable. In many cases, such metrics are used to provide parameter rankings that allow for necessary dimension reduction in moderate-to-high dimensional systems. However, globally derived sensitivity results may obscure localized variability in parameter sensitivities, resulting in misleading conclusions about parameter importance and ensuing consequences for subsequent tasks such as model calibration and surrogate model construction. In this study, we illustrated how discrepancies between globally and locally based sensitivity information may arise for an emerging sensitivity metric based on active subspace methodology, as well as for other commonly used sensitivity techniques. In response, we outlined a framework that exploits the active subspace to evaluate the stability of parameter sensitivities over the admissible parameter space. This analysis allows one to determine the subregions of the parameter space in which a globally derived sensitivity metric may be considered representative of local behavior. The proposed framework was illustrated on a collection of simple examples for ease of visualization, as well as a variation of the Lotka-Volterra model used to study interactions between competing tumor cell lines, for which we demonstrated how these issues may be exacerbated in higher dimensions. Our findings suggest that globally derived sensitivity information should be treated with caution, and that incorporating analysis on local subregions may improve robustness and accuracy in downstream modeling tasks.

Keywords: active subspaces; global sensitivity analysis; dimension reduction; model calibration; surrogate model construction

1. Introduction

The purpose of sensitivity analysis is to quantify how variations in model inputs affect model outputs. Such analysis is widely used to guide parameter selection, allow for model simplification, or ensure

model identifiability—tasks that are essential within the field of mathematical biology, since available data is often sparse—only observed for some subset of model state variables, and/or subject to large amounts of measurement noise. Global sensitivity analysis measures variations in model output due to parameter perturbations across the entire specified parameter space. Common methods for performing global analysis include screening methods such as Morris elementary effects, variance-decomposition methods such as Sobol' indices, fourier amplitude sensitivity test (FAST), or extended FAST (eFAST), or correlation and regression methods such as partial correlation coefficients, among others [1–5]. In contrast, local sensitivity analysis utilizes linearization at a set of nominal parameter values to measure model behavior more precisely at a particular point of interest.

When prior information about parameter values is limited, researchers often rely on global sensitivity analysis. However, conclusions derived via such a global approach may not accurately reflect model behavior in all regions of the parameter space, especially in nonlinear and/or high-dimensional models [6]. The inherent averaging of sensitivity information from a random sampling of points across the space may overinflate parameter importance due to a small region of out-sized influence, or prioritize parameters that are somewhat sensitive everywhere at the cost of parameters that are generally inert but critically sensitive in certain local regions. For example, there exist safety-critical models such as those used in nuclear engineering for which knowledge of local sensitivities is essential; if such a system is highly sensitive in a specific region, relying solely on global results that may obscure these region-specific sensitivities could yield critical risks [2]. Global methods are also highly dependent upon the imposed parameter distributions and ranges, and incapable of differentiating between regions of the space that might produce implausible model outputs due to parameter-driven bifurcations. Averaging of sensitivity results from regions with drastically different model behavior can result in discrepancies with regard to parameter importance [7–9]. As such, an overreliance on global sensitivity measures is likely to lead to poor performance in downstream model analysis tasks.

Previous works have sought to understand the relationship between localized sensitivity variations and global sensitivity measures. For instance, plots illustrating contributions to the sample mean (CSM) and contributions to the sample variance (CSV) help practitioners to understand which local regions most substantially contribute to mean values of the model output [10], and quantify the potential reduction in model output variance that can be achieved via restrictions of the input ranges to certain localized regions [11]. Wei et al. [12] employed regional sensitivity analysis (RSA) and parametric sensitivity analysis (PSA) techniques to study how global sensitivity measures such as Sobol' indices depend on alterations to the input ranges and/or distributions. Likewise, Pannier and Graf [13] developed a grid-based strategy for deriving *sectional* global sensitivity measures (SGSMs), to allow for an enhanced understanding of the regional characterization of the functional relationship between model inputs and outputs. In each of these investigations, the overarching goal is to provide a more detailed description of the input-output relationship than what can be obtained from a single global sensitivity measure, as is often reported. Moreover, the issues addressed by such analyses are particularly relevant in biological modeling, where input ranges and distributions tend to be imposed based on values that are known to be biologically reasonable, but without consideration of the impact of those choices on the model output variability. In such scenarios, the development of diagnostic tools that can alert a researcher to regions in the parameter space for which global sensitivity measures may not be locally representative is vital.

Traditional global sensitivity methods provide a ranking of parameters in accordance with their overall influence on the model output. However, these methods do not provide a systematic way to evaluate

how well globally derived rankings represent parameter sensitivity in more localized regions. In this investigation, we study how sensitivity information derived from emerging active subspace methodology can provide all of the same information as other commonly used metrics, while also addressing this limitation. The active subspace—computed via an eigendecomposition of a matrix containing gradient vectors sampled at points from around the parameter space—identifies influential directions in the input space, which are defined to be linear combinations of parameters that contribute most to variation in the model output [2, 14]. By prioritizing linear combinations of inputs that need not align with the standard coordinate space, the active subspace addresses not just basic input prioritization but is also ideally suited for examining how parameters may interact to produce large changes in model response. Connections between active subspace methodology and traditional sensitivity methods (specifically the use of the sensitivity Fisher information matrix, or sFIM) have been previously studied [6]. Here, we propose the use of regional subspace distances between locally and globally defined active subspaces to allow the user to quantify the degree to which parameter sensitivity varies across the admissible space. By identifying regions in which global results are not representative of local behavior, researchers are given a roadmap as to the conditions under which global metrics may be considered consistent with localized behavior, and can identify areas of the parameter space that are worthy of further focus.

Specifically, this investigation builds upon previous works addressing the issue of localized sensitivity variability in three ways. First, we propose the systematic use of regional subspace distances—computed by measuring the discrepancy between locally and globally derived active subspaces—as a diagnostic tool for identifying where global sensitivity measures may fail to adequately capture local sensitivity behavior. Second, our use of a grid-based stability map allows for the visualization of regions of the input space where global sensitivity information is inconsistent with its local counterpart, flagging areas of the input space that warrant deeper investigation. Finally, we provide an in-depth examination of the consequences of relying on global sensitivity measures that are not locally representative, with particular attention to how an incomplete picture of the full input-output representation can affect downstream analysis tasks such as model calibration or surrogate model construction.

In Section 2, we introduce the active subspace, define the corresponding metric used to rank parameters in terms of importance, and introduce our strategy to assess the stability of sensitivity results across the parameter space. Section 3 begins with a detailed analysis of several simple, two-dimensional models for visualization purposes. We then provide an in-depth investigation of a variation of the Lotka-Volterra model used to study interactions between competing cell lines in a tumor spheroid. This six-dimensional model—widely used in both mathematical biology and mathematical ecology—serves as an illustration of how local sensitivity variability may exacerbate the disagreement between local and global analyses in moderate-to-high dimensional examples. Using the Lotka-Volterra model, we conclude by demonstrating the potential downstream impacts of an overreliance on global information when performing subsequent tasks such as model calibration and surrogate model construction. Finally, Section 4 concludes with a summary of the findings and suggestions for future work.

2. Methodology

2.1. Constructing the active subspace

Active subspace methodology enables dimension reduction in complex models by identifying linear combinations of input parameters that most significantly influence the model output, in addition to

orthogonal directions that can be safely neglected in ensuing analysis with minimal loss in accuracy. The following provides a concise summary of active subspace construction; for further details, see [2, 14].

Let $q = f(\mathbf{x})$ represent the mapping from the parameter space to a scalar-valued model quantity of interest (QoI), such that $f : \mathbb{R}^m \to \mathbb{R}$. The direction of greatest change in the QoI with respect to the input parameters is captured by the function's gradient, $\nabla_{\mathbf{x}} f(\mathbf{x})$. We introduce the matrix $\mathbf{C}$,

$$\mathbf{C} = \int \nabla_{\mathbf{x}} f(\mathbf{x}) \nabla_{\mathbf{x}} f(\mathbf{x})^T \rho(\mathbf{x})\, d\mathbf{x}, \tag{2.1}$$

which—defined as the average of the outer product of the function's gradient with itself—quantifies how parameters interact with one another to drive changes in the QoI over the parameter density $\rho(\mathbf{x})$. Due to its symmetry, $\mathbf{C}$ admits an eigendecomposition, $\mathbf{C} = \mathbf{W}\Lambda\mathbf{W}^T$, where $\Lambda = \text{diag}(\lambda_1, \ldots, \lambda_m)$ contains the eigenvalues of $\mathbf{C}$ with $\lambda_1 \geq \ldots \geq \lambda_m \geq 0$, and the matrix $\mathbf{W}_{m\times m}$ is an orthonormal matrix containing the normalized eigenvectors of $\mathbf{C}$, $\mathbf{W}_{m\times m} = [\mathbf{w}_1, \ldots, \mathbf{w}_m]$. We note that the matrix $\mathbf{C}$ is closely related to the derivative-based global sensitivity measures (DGSMs) introduced in [15] and the improved Morris importance measure proposed in [16], both of which also employ gradient information averaged over the admissible parameter space. Sobol' and Kucherenko further established a connection between DGSMs and Sobol' indices, bounding the total Sobol' index above by a function of the DGSM [15]. Thus, the use of averaged gradient information is already standard in global sensitivity analysis; what distinguishes the active subspace method is that it deploys the information garnered by the matrix $\mathbf{C}$ to identify a low-dimensional subspace of influential directions, rather than attributing sensitivity strictly to the original input parameters.

The eigendecomposition of $\mathbf{C}$ provides a natural way to define the active subspace. In particular, we separate the eigenvectors of $\mathbf{C}$ into two distinct sets according to a natural split in the magnitude of the eigenvalues, should one occur:

$$\Lambda = \begin{bmatrix} \Lambda_1 & \mathbf{0} \\ \mathbf{0} & \Lambda_2 \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_1 & \mathbf{W}_2 \end{bmatrix},$$

where $\Lambda_1 = \text{diag}(\lambda_1, \ldots, \lambda_n)$ with $n < m$ contains the n eigenvalues that are largest in magnitude, with $\mathbf{W}_1$ containing the corresponding first n eigenvectors. Similarly, $\Lambda_2 = \text{diag}(\lambda_{n+1}, \ldots, \lambda_m)$ contains the smaller eigenvalues corresponding to the eigenvectors in $\mathbf{W}_2 = [\mathbf{w}_{n+1}, \ldots, \mathbf{w}_m]$. The active subspace is then defined as the space spanned by $\mathbf{W}_1$, which captures the most influential directions in the input space with respect to the quantity of interest. By projecting $\mathbf{x}$ into the subspaces defined by the first n eigenvectors (the active subspace) and the remaining $m - n$ eigenvectors (the inactive subspace), respectively, we can define the active variables to be $\mathbf{y} = \mathbf{W}_1^T\mathbf{x}$ and the inactive variables to be $\mathbf{z} = \mathbf{W}_2^T\mathbf{x}$. Thus, any given $\mathbf{x}$ in the parameter space can be expressed in terms of $\mathbf{y}$ and $\mathbf{z}$ as

$$\mathbf{x} = \mathbf{W}\mathbf{W}^T\mathbf{x} = \mathbf{W}_1\mathbf{W}_1^T\mathbf{x} + \mathbf{W}_2\mathbf{W}_2^T\mathbf{x} = \mathbf{W}_1\mathbf{y} + \mathbf{W}_2\mathbf{z}.$$

This decomposition allows us to analyze $f(\mathbf{x})$ in a reduced space by considering its behavior with respect to the components $\mathbf{y}$ and $\mathbf{z}$. Since $\mathbf{W}_2$ contains directions that contribute very little information about the QoI, the function $f(\mathbf{x})$ can be approximated by a function on the active variables only,

$$f(\mathbf{x}) \approx g(\mathbf{y}),$$

which effectively reduces the dimensionality of our model from m inputs to n, where $n < m$.

In moderate-to-high-dimensional models, directly computing the matrix $\mathbf{C}$ is impractical due to the high-dimensional integration required. In practice, $\mathbf{C}$ is approximated using a Monte Carlo algorithm, as outlined in [14] and restated in Algorithm 1.

Algorithm 1 Random sampling to estimate the active subspace [14]

1. Draw M samples $\{\mathbf{x}_j\}$ independently according to the density function $\rho(\mathbf{x})$.
2. For each $\mathbf{x}_j$, compute $\nabla_{\mathbf{x}} f_j = \nabla_{\mathbf{x}} f(\mathbf{x}_j)$.
3. Approximate
$$\mathbf{C} \approx \hat{\mathbf{C}} = \frac{1}{M} \sum_{j=1}^{M} (\nabla_{\mathbf{x}} f_j)(\nabla_{\mathbf{x}} f_j)^T.$$
4. Compute the eigendecomposition $\hat{\mathbf{C}} = \hat{\mathbf{W}} \hat{\Lambda} \hat{\mathbf{W}}^T$.

For models in which the gradient of the QoI is not available analytically, finite differences may be used to estimate $\nabla_{\mathbf{x}} f(\mathbf{x})$. For a derivation of upper bounds on approximation errors due to Monte Carlo sampling and the use of finite differences, see [14].

Traditionally, the choice of dimension n for the active subspace has been based on locating a natural gap in the eigenvalue spectrum, such that the largest gap in the eigenvalues occurs between λ_n and λ_{n+1}. Such a criterion can be shown to lead to the most accurate estimation of the active subspace when using Algorithm 1 to estimate the eigendecomposition of $\mathbf{C}$ [14]. However, other methods for identifying the dimension of the active subspace have been explored, including the use of response surfaces, energy thresholds such as those used in principal component analysis (PCA), or error-based criteria to identify possible reductions [17–19]. The advantages and disadvantages of each method are discussed in [19], where the authors examine the tradeoff between ensuring accuracy in the approximated QoI and balancing the need for computational efficiency.

2.2. Sensitivity analysis using the active subspace

Traditional global sensitivity analysis methods—including screening methods such as Morris elementary effects [3] and variance-based decomposition methods such as Sobol' indices [4]—use sensitivity information from across the parameter space to rank the parameters in order of their influence to the QoI. Parameters that are found to be uninfluential are often fixed at their nominal values, such that the reduction of the problem to be studied is aligned with the original coordinate space. We refer to this as *subset*-based reduction. This is in opposition to *subspace*-based reduction, in which the discarded directions may be linear combinations of the original model parameters.

In general, the use of a subset-based metric for performing model reduction is beneficial because it allows for easier interpretability of the results; a subset of parameters is removed from consideration, and the remaining parameters maintain their physical meaning. However, subspace-based reduction—while not as easily interpretable due to its removal of linear combinations of parameters—is often able to retain more of the total information in the model when using the same dimension as the subset-based

analysis, due to its optimization of the information content through the rotation of the parameter space.

Construction of the active subspace allows for both subset- and subspace-based dimension reduction. If subset-based reduction is of interest, parameter rankings can be obtained using *activity scores*, derived from the eigendecomposition of the matrix $\mathbf{C}$. The activity score for parameter i, denoted as α_i, quantifies the overall contribution of the parameter to the QoI [20] and is defined as

$$\alpha_i = \sum_{j=1}^{m} \lambda_j \mathbf{w}_{i,j}^2, \quad \text{for} \quad i = 1, ..., m, \tag{2.2}$$

where $\mathbf{w}_{i,j}$ represents the (i, j)th element of the matrix $\mathbf{W}$. This sensitivity metric scales each squared eigenvector component by its corresponding eigenvalue. As such, larger α_i values indicate that the parameter in question is contributing to directions of greater importance, thus exerting greater influence on the model quantity of interest.

Throughout this investigation—in addition to reporting the activity scores derived from the active subspace—we will also include parameter rankings based on Morris elementary effects (mean effect μ_i and absolute mean effect μ_i^*, along with standard deviation σ_i) [3], and Sobol' indices (first-order indices S_i and total-effect indices S_{T_i}) [4], in order to compare the performance of our active subspace methodology to more traditional sensitivity methods. For each metric, rankings will be derived on both global and local scales. When deviations between global and local rankings occur, subset-based reduction based on global metrics should be used with caution, as it indicates that we may be removing parameters that appear uninfluential globally but are highly sensitive in certain local regions.

Subspace-based dimension reduction is also possible using information from the active subspace, a significant benefit to using this method over others such as Morris screening or Sobol' indices. Specifically, rather than discarding non-influential parameters as determined by the activity scores, one may choose to discard non-influential *directions* in the input space, defined by the inactive variables $\mathbf{z}$. Such a procedure is often used for applications such as surrogate model construction, in which a lower-dimensional model approximation may be constructed on a subspace of the original parameter space in order to increase computational efficiency for future model evaluations. Surrogate model construction utilizing subspace-based reduction will be illustrated in Section 3.3.2.

2.3. Analyzing stability of sensitivity information

Metrics such as activity scores, Morris elementary effects, and Sobol' indices are designed to summarize information about parameter influence at a global scale, but may do so at the expense of obfuscating local sensitivity variations. This issue is particularly prominent in models that contain nonlinear parameter interactions. In particular, the matrix $\mathbf{C}$, which defines the active subspace, can be reframed as the integral of the sensitivity Fisher information matrix (sFIM) over the parameter density. When parameters interact in a nonlinear manner to produce changes in the model output, such averaging will produce different eigenvectors than would be obtained using a local sFIM, and thus lead to a different approximation of the model output [6]. Our goal is to develop a tool to quantify the degree to which this issue exists in any given model and highlight the consequences of relying solely on globally derived information when such nonlinearities are present. Here, we present a metric to determine when global and local parameter sensitivities are out of alignment, based on information provided by the active subspace. This framework can help to identify regions of the parameter space in which global

information can be considered locally representative or point to subregions in which further analysis should be conducted. Though we use the phrase "local analysis" throughout the remainder of this investigation, it should be noted that this refers to the computation of global sensitivity metrics on a localized region of the global admissible parameter space, in contrast to the traditional definition of local sensitivity analysis. As this distinction is critical for what follows, we emphasize it once again.

Remark. For the remainder of this study, any reference to "local analysis" will refer to computation of global sensitivity metrics (e.g., activity scores, Morris elementary effects, or Sobol' indices) on a subregion of the admissible parameter space, as opposed to the traditional definition of local sensitivity analysis, in which derivative-based sensitivity information is measured at a single set of nominal values.

Our subspace-based criterion assesses the stability of the active subspace by measuring its consistency across different regions of the parameter space. We consider the matrices $\mathbf{C}$ and $\widetilde{\mathbf{C}}$ defined as in Eq (2.1), where $\mathbf{C}$ is computed using gradient samples from across the global admissible parameter space, while $\widetilde{\mathbf{C}}$ utilizes only gradient samples from within a particular localized region of interest that is a subregion of the global parameter space. Let $\mathbf{W} = \begin{bmatrix}\mathbf{W}_1 & \mathbf{W}_2\end{bmatrix}$ and $\widetilde{\mathbf{W}} = \begin{bmatrix}\widetilde{\mathbf{W}}_1 & \widetilde{\mathbf{W}}_2\end{bmatrix}$ represent two matrices containing the eigenvectors of the globally computed and locally computed matrices $\mathbf{C}$ and $\widetilde{\mathbf{C}}$, respectively. The distance between the two active subspaces defined by the ranges of $\mathbf{W}_1$ and $\widetilde{\mathbf{W}}_1$ can be calculated as

$$\text{dist}(\text{range}(\mathbf{W}_1), \text{range}(\widetilde{\mathbf{W}}_1)) = \left\| \mathbf{W}_1\mathbf{W}_1^T - \widetilde{\mathbf{W}}_1\widetilde{\mathbf{W}}_1^T \right\|_2, \tag{2.3}$$

which quantifies the distance between the projection matrices of the two subspaces using the spectral norm [14, 21]. As detailed in [21], this distance is equivalent to the sine of the maximum canonical angle between the two subspaces. This physical interpretation provides natural bounds of 0 and 1 on our subspace distance metric.

The column dimension n of the matrices $\mathbf{W}_1$ and $\widetilde{\mathbf{W}_1}$ is chosen based on the eigendecomposition of the locally derived matrix $\widetilde{\mathbf{C}}$ using a threshold analogous to a PCA truncation criterion. Specifically, the dimension n of the active subspaces is chosen such that $\geq 95\%$ of the spectral energy is captured in the first n eigenvalues of $\widetilde{\mathbf{C}}$, which ensures that distances are measured across the n directions determined to be locally significant. Thus, the metric of Eq (2.3) is a measure of whether the global sensitivity metric is able to adequately capture information that is important in the local region.

By analyzing the relationship between subspace distance and parameter values, one can determine subregions of the space in which locally and globally computed sensitivity information is in close alignment, or identify certain parameter ranges for which the globally computed active subspace prioritizes different combinations of parameters than would be favored locally. In Section 3, we will illustrate how such regions can be visually identified for both two-dimensional and higher-dimensional models.

3. Results

In what follows, we illustrate the proposed methodology on several two-dimensional examples—chosen in order to facilitate straightforward visualization of how local and global sensitivities may not align—and on the six-dimensional Lotka-Volterra model, where we demonstrate how such issues may be further intensified in higher dimensions. All of the analysis is performed using MATLAB software.

3.1. Two-dimensional models

To illustrate the variation in sensitivity that can occur as one moves around the parameter space, we begin with three simple, two-dimensional examples. Using $\mathbf{x} \in [0, 1] \times [0, 1]$, we define our three models as follows:

$$f_1(\mathbf{x}) = e^{0.7x_1+0.3x_2}, \quad f_2(\mathbf{x}) = x_1 e^{0.7x_1+0.3x_2}, \quad \text{and} \quad f_3(\mathbf{x}) = e^{0.7x_1^2x_2+0.3x_2}.$$

All three of these examples permit a one-dimensional active subspace—as determined by the negligibility of the second eigenvalue when compared to the first—which is defined by the span of the gradient vector. The variability in gradient vector direction for these examples is visualized in Figure 1. Function f_1 exhibits no variation at all over the parameter space; the gradient vector is a scalar multiple of $[0.7,\ 0.3]$ everywhere in the space. For such an example, an active subspace computed globally will be identical to an active subspace computed on any sub-portion of the parameter space, and a ranking of the parameters via activity scores will always prioritize x_1 over x_2. With function f_2, there is a slight amount of variability in the direction of the gradient vector as one moves around the space, though x_1 is clearly always more influential than x_2. The function f_3 is an example of a function for which the direction of the gradient vector is highly dependent on the location in the parameter space where it is computed.

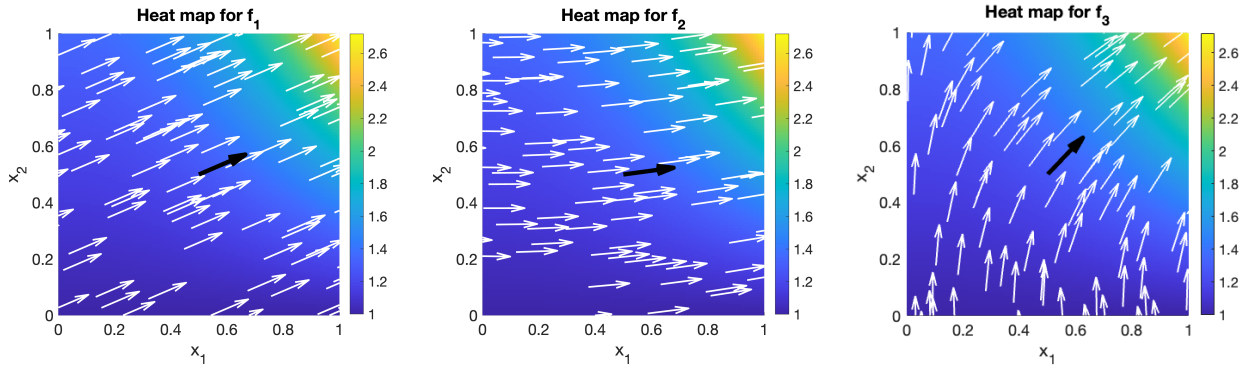

Figure 1. Visualization of locally computed gradient vectors (white) versus a gradient vector computed from global sampling (black).

We want to understand how this variation in gradient vector direction affects the subsequent sensitivity analysis. For each function, we divide the domain into local grids of size 0.01×0.01, and compute the active subspace within each grid square using Algorithm 1 with $M = 10$ samples randomly drawn from a uniform density function using a Latin hypercube design. The locally computed active subspaces are then used to determine activity score rankings for each grid square using Eq (2.2). Finally, we use 100,000 samples from a uniform Latin hypercube design across the full parameter space—corresponding to 10 samples per local region—to compute a global active subspace and global activity score ranking.

As expected, locally computed activity scores for functions f_1 and f_2 indicate that x_1 is more influential than x_2 in every local region, matching their global rankings. For function f_3, the global normalized activity scores of $[0.9412, 1]$ indicate a slight prioritization of x_2; however, in 2187 of the 10,000 local regions—namely, those concentrated in the upper-right corner—x_1 is prioritized over x_2 instead.

Figure 2 illustrates the distances between locally and globally computed active subspaces, computed using Eq (2.3). Again, heat maps for f_1 and f_2 show little or no variation across the space, indicating that the global active subspace is relatively well-aligned with local information everywhere. However, the heat map for f_3 indicates that there are regions in the parameter space where the global active subspace is a very poor representation of local activity, most notably when x_1 and/or x_2 is small: regions where the influence of the parameter x_2 strongly outweighs x_1. Given these large disparities in active subspace alignment, together with the differing activity score rankings from two regions of the space, this constitutes our first example in which a globally computed sensitivity metric may give misleading information about function behavior in a local region.

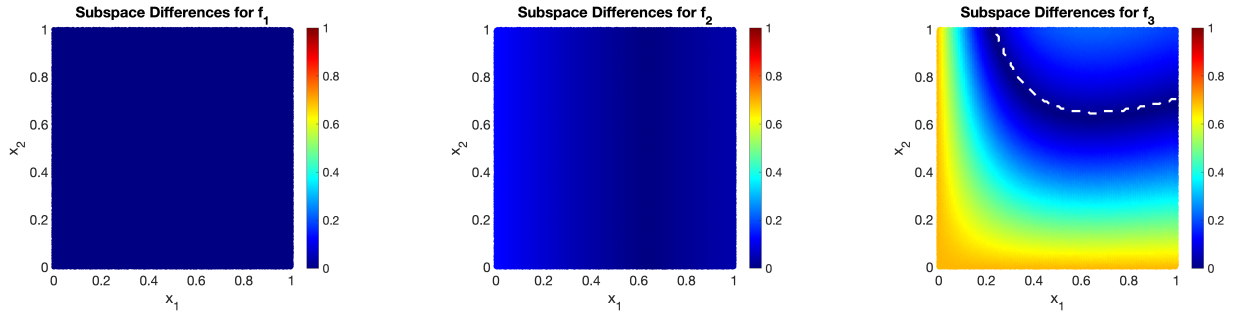

(a) **(b)** **(c)**

Figure 2. Subspace distance visualization. In Figure 2(c), the dashed curve indicates the separation of regions where x_1 is ranked first (above) and where x_2 is prioritized (below), according to locally computed activity scores.

3.2. Lotka-Volterra model

To demonstrate how the disparity in globally versus locally derived sensitivity information may be exacerbated in models that are of higher dimension and/or are highly nonlinear, we turn our attention to the six-dimensional Lotka-Volterra model.

3.2.1. Model description

The Lotka-Volterra model is a system of differential equations originally developed to describe predator–prey interactions in ecological systems [22, 23]. It has also been widely adapted across disciplines, including physics [24–26] and chemistry [27, 28], to describe dynamics in systems such as lasers, reaction kinetics, and phase transitions. The general form of the model captures nonlinear interactions between two or more species or components, making it a flexible structure for exploring the effects of parameter variation in dynamical systems. In this study, we apply our analysis to a variant of the Lotka-Volterra model describing the growth of a heterogeneous tumor spheroid comprised of two cancer cell types, Type S and Type R [29]:

$$\frac{dS}{dt} = r_S S\left(1 - \frac{S}{K_S} - \frac{\gamma_R R}{K_S}\right), \tag{3.1}$$

$$\frac{dR}{dt} = r_R R\left(1 - \frac{R}{K_R} - \frac{\gamma_S S}{K_R}\right), \tag{3.2}$$

where each population is considered to undergo logistic growth in the absence of the other cell line, with growth rates of r_S and r_R and carrying capacities of K_S and K_R, respectively. The signs and magnitudes of the parameters γ_S and γ_R determine interaction strength between the two populations, and $S(t)$ and $R(t)$ represent the volume (mm^3) of Type S and Type R cancer cells at time t. For this investigation, we initialize the system using a total tumor volume of 0.02 mm^3 with a 9:1 initial ratio of Type S to Type R cells. Additionally, we set parameter ranges of [0, 1] for all parameters $[r_S, r_R, K_S, K_R, \gamma_S, \gamma_R]$, which restricts us to positive growth scenarios with competitive interactions only. Figure 3 showcases some of the variability in behavioral patterns that can occur for this system, using different sets of parameter values drawn from within these ranges. We define our primary scalar-valued QoI to be the integral of the total tumor volume over an eight-week (56-day) growth observance period,

$$QoI_{pri} = \int_0^{56} (S(t) + R(t))\ dt, \tag{3.3}$$

which we approximate computationally using the trapezoidal rule with bin sizes of one day. This QoI can be considered a representation of the total tumor burden for a patient. It is worth noting that sensitivity measures are highly dependent upon the chosen quantity of interest. While Eq (3.3) will serve as our primary quantity of interest for this investigation, we also present results for a secondary QoI in Section 3.2.5 to illustrate that the issue of variability in localized sensitivity information is not unique to this choice.

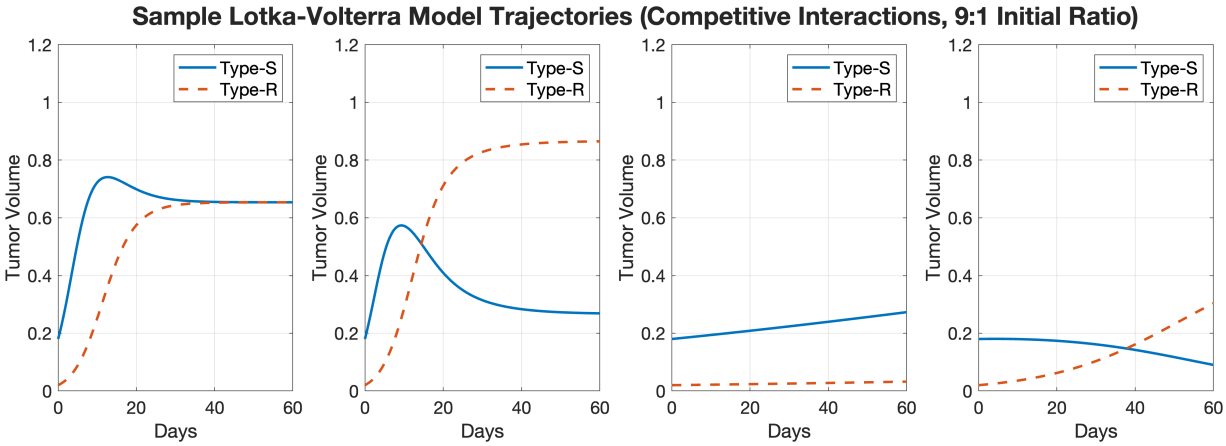

Figure 3. Competitive Lotka–Volterra model trajectories with a 9:1 initial Type S:Type R ratio. Trajectories for Type S (blue, solid) and Type R (orange, dashed) cell volumes illustrate how different choices of growth rates, carrying capacities, and interaction strengths can impact the behavior of the system. From left to right, the parameter values are: $[r_S, r_R, K_S, K_R, \gamma_S, \gamma_R]$ = [0.36, 0.36, 0.85, 0.85, 0.30, 0.30]; [0.35, 0.50, 0.80, 0.95, 0.20, 0.45]; [0.01, 0.01, 0.80, 0.80, 0.60, 0.75]; [0.05, 0.08, 0.20, 0.50, 0.60, 0.75].

3.2.2. Sensitivity analysis

For the localized sensitivity analysis, we divide each parameter's range into eight segments of length 0.125, resulting in a total of $8^6 = 262{,}144$ local regions for consideration. The number of segments is chosen to showcase regional sensitivity variations while still maintaining a feasible computational

budget. Within each region, we generate $M = 25$ sample points from a uniform Latin hypercube design to compute the local active subspace, a number chosen to satisfy the bound derived in [14], as discussed further in Section 3.2.3. Global analysis is conducted using 25×8^6 sample points from a Latin hypercube design, corresponding to 25 sample points per local region. Gradient vectors required for active subspace computations are computed using forward finite differences with step size $h = 10^{-5}$. For comparison with other traditional sensitivity methods, we also conduct Morris screening with six levels and $r = 25$ trajectories per local grid, and estimate total Sobol' indices using 25 samples per local grid. Both the Morris elementary effects and Sobol' indices are computed using the algorithms presented in [2].

Beginning with subset-based analysis, we first compute the globally determined sensitivity metrics for each parameter in the model, presented in Table 1. Using the normalized activity scores derived from the active subspace analysis, rankings indicate that the Type S growth rate, r_S, and carrying capacity, K_S, have the most significant influence on the QoI (with these two outranking their r_R and K_R counterparts due to the initial cell ratio of 9:1), while the interaction parameters γ_S and γ_R contribute the least. The overall global ranking is given by $[r_S,\ K_S,\ K_R,\ r_R,\ \gamma_S,\ \gamma_R]$. In contrast, both Morris screening and Sobol' indices prioritize the two carrying capacities, K_S and K_R, when analyzed across the full parameter space, with global ranking $[K_S,\ K_R,\ r_R,\ r_S,\ \gamma_S,\ \gamma_R]$.

Table 1. Global sensitivity metrics for all methods.

		r_S	r_R	K_S	K_R	γ_S	γ_R
Activity Scores	$\alpha_i/\alpha_{\max}$	1.0000	0.6715	0.9364	0.6942	0.0746	0.0567
Morris Screening	μ	1.0877	1.1604	4.4625	3.9410	-0.7857	-0.6296
	μ^*	1.1050	1.1727	4.5086	3.9669	0.8030	0.6438
	σ	0.9532	0.8820	3.9881	3.7662	2.6908	0.8389
Sobol' Indices	S_i	0.0837	0.0922	0.3707	0.2439	0.0134	0.0089
	S_{Ti}	0.1543	0.1776	0.4747	0.3482	0.0351	0.0274

To understand whether global results can be accurately applied to analyze model behavior on a more localized scale, we conduct our analyses individually on the set of 8^6 locally defined subregions and look for patterns in which parameters are prioritized. For the active subspace analysis, we observe 304 unique parameter rankings (out of 720 possible rank orders) among the 8^6 regions, revealing that sensitivity information varies drastically around the space. The Morris screening and Sobol' indices methods reveal the observance of 336 and 720 unique rankings, respectively, illustrating that this issue is not unique to the activity score metric.

The top ten most frequently observed local rankings for each of the activity score, Morris screening, and Sobol' indices methods are tallied in Table 2. Notably, the global activity score ranking of $[r_S,\ K_S,\ K_R,\ r_R,\ \gamma_S,\ \gamma_R]$ obtained from the active subspace method is only the 53rd most frequently observed local ranking, occurring in only 1525 out of the total 262,144 regions. The two most globally favored parameters, r_S and K_S, are ranked as the top-two parameters in only three of the ten most commonly observed local rankings. While K_S is regularly observed among the three most influential parameters in the most frequently observed rankings, r_S is actually ranked as the *least* influential parameter in five of the top ten observed local rankings, casting doubt upon the reliability of the global active subspace findings when applied to local analyses. Similarly, the global ranking of

$[K_S,\ K_R,\ r_R,\ r_S,\ \gamma_S,\ \gamma_R]$ obtained via the Morris screening and Sobol' indices methods appears as the 35th and 51st most frequent ranking in those two methods, respectively.

Table 2. Ten most frequently observed local rankings for the Lotka-Volterra model, using each of the three sensitivity methods.

Activity Scores - Top 10 Rankings					
Order	Ranking	Frequency	Order	Ranking	Frequency
1	$K_R, r_R, K_S, \gamma_S, \gamma_R, r_S$	12,311	6	$K_R, K_S, \gamma_R, r_R, \gamma_S, r_S$	5676
2	$K_S, r_S, K_R, \gamma_S, \gamma_R, r_R$	10,381	7	$K_S, K_R, \gamma_S, r_S, \gamma_R, r_R$	5019
3	$K_S, r_S, K_R, \gamma_S, r_R, \gamma_R$	9336	8	$K_S, K_R, \gamma_R, \gamma_S, r_R, r_S$	4778
4	$K_S, r_S, K_R, \gamma_R, \gamma_S, r_R$	8971	9	$K_R, r_R, K_S, \gamma_S, r_S, \gamma_R$	4626
5	$K_R, K_S, r_R, \gamma_R, \gamma_S, r_S$	5921	10	$r_R, K_R, K_S, \gamma_S, \gamma_R, r_S$	4577
Morris Screening - Top 10 Rankings					
Order	Ranking	Frequency	Order	Ranking	Frequency
1	$K_S, K_R, r_R, \gamma_S, \gamma_R, r_S$	9006	6	$K_R, K_S, \gamma_R, \gamma_S, r_R, r_S$	5448
2	$K_S, r_S, K_R, \gamma_S, r_R, \gamma_R$	5995	7	$K_R, r_R, K_S, \gamma_S, \gamma_R, r_S,$	5398
3	$K_S, K_R, r_R, \gamma_S, r_S, \gamma_R$	5920	8	$K_S, r_S, K_R, \gamma_S, \gamma_R, r_R$	5250
4	$K_R, K_S, r_S, \gamma_R, \gamma_S, r_R$	5710	9	$K_S, K_R, \gamma_S, \gamma_R, r_S, r_R$	5009
5	$K_R, K_S, r_S, \gamma_S, \gamma_R, r_R$	5573	10	$K_S, r_R, K_R, \gamma_S, \gamma_R, r_S$	4983
Sobol' Indices - Top 10 Rankings					
Order	Ranking	Frequency	Order	Ranking	Frequency
1	$K_S, K_R, \gamma_S, \gamma_R, r_R, r_S$	2013	6	$K_S, r_S, K_R, \gamma_R, r_R, \gamma_S$	1662
2	$K_S, r_S, r_R, \gamma_R, \gamma_S, K_R$	1888	7	$K_S, K_R, \gamma_S, r_S, r_R, \gamma_R$	1660
3	$K_S, r_S, \gamma_R, r_R, \gamma_S, K_R$	1885	8	$K_S, r_S, r_R, \gamma_S, \gamma_R, K_R$	1615
4	$K_S, K_R, \gamma_S, r_R, \gamma_R, r_S$	1737	9	$K_R, K_S, \gamma_S, \gamma_R, r_S, r_R$	1611
5	$K_S, K_R, \gamma_R, \gamma_S, r_R, r_S$	1706	10	$K_S, r_S, \gamma_S, r_R, \gamma_R, K_R$	1608

Figure 4 displays the percentage of cases in which each parameter appears in the top 1, 2, or 3 parameter ranking spots among all of the localized regions, for each of the activity score, Morris screening, and Sobol' indices methods. Despite the differing global rankings, all methods are in agreement about which parameters are most often flagged as influential among the locally defined regions. For all three methods, one of the two carrying capacities is most often chosen as the most sensitive parameter. The growth rates are chosen as most influential far less frequently but often show up in the top 2 or 3 of the ranking list. Finally, the interaction parameters are rarely (if ever) chosen as the most influential, but do occasionally make appearances in the top 2 or 3 rankings—especially in the case of Sobol' indices, where parameter ranking distributions are more evenly dispersed.

Our subset-based analysis reveals that many different rankings may appear in localized regions that differ from the global results and that this issue persists across multiple established sensitivity methods. To determine in which local regions the global results can be considered representative of local behavior, we conduct a subspace-based analysis to analyze the stability of the active subspace as we move around the admissible parameter space. The subspace distance is evaluated for each local region by analyzing the deviation in the local active subspace from the global subspace using the dimension of the local active subspace, calculated as described in Section 2.3. To understand how these distances change as we

move around the space, we report the distributions of subspace distances for each parameter-subregion combination in Figure 5. Analyzing these distributions allows us to identify trends in subspace distance as a function of a particular parameter, while all other parameters are allowed to vary. For example, subspace distances tend to decrease on average as K_S increases in value; that is, global measures are more likely to properly represent local behavior in regions where K_S is large. Similarly, subspace distances appear to increase on average as both γ_S and γ_R increase, such that global measures are more indicative of local behavior for smaller values of the interaction parameters. The non-monotonic nature of the distribution and means progression for parameters r_S, r_R, and K_R suggests that there are additional interactions at play that need to be understood.

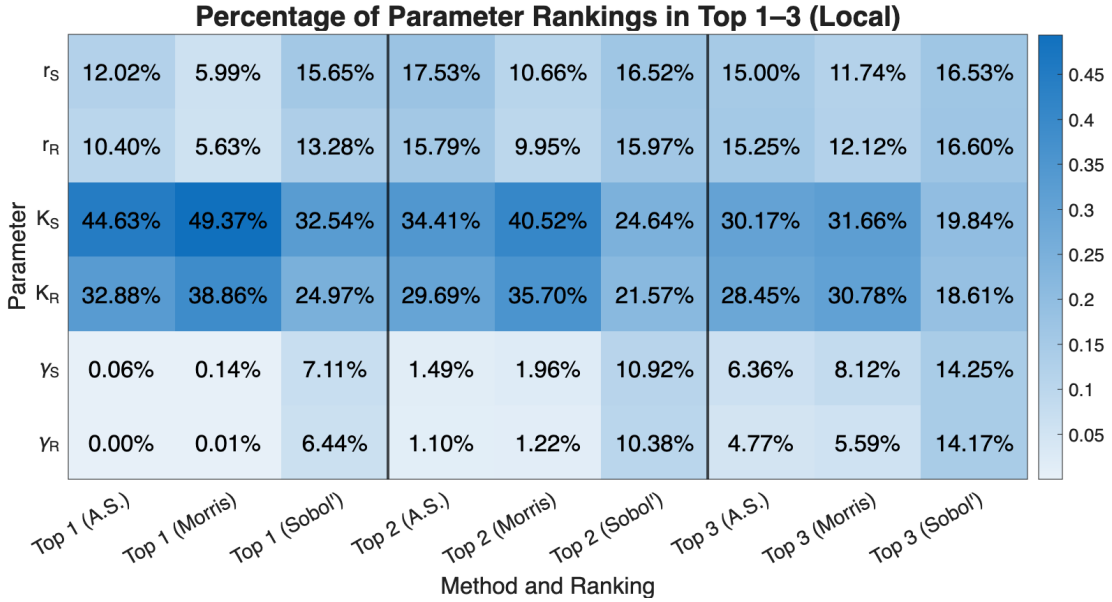

Figure 4. Percentage of local regions in which each parameter is ranked among the top 1, 2, or 3 of the ranking list.

To better understand how parameters may interact to drive changes in the quantity of interest across local regions, we include interaction heatmaps in Figure 6, which quantify the average subspace distance across all parameter combinations contained within a particular pairwise-defined region. These heatmaps enable us to identify pairs of parameter values for which subspace distances are particularly problematic. For instance, large values of either γ_S or γ_R paired with small values of K_R are likely to lead to misaligned subspaces, such that global measures are not indicative of local behavior.

To gain additional insight into what magnitudes of subspace distances are likely to cause contradictory sensitivity analysis results on global versus local scales, we examine the level of agreement between parameter rankings obtained at both scales as a function of the distance between their subspaces. Let ϕ_k represent the set of k most influential parameters as defined by the global analysis. For a given local region, let ξ_k represent the k most influential parameters as defined by the local analysis on that region. We define the g-score to be the cardinality of their intersection divided by the number of factors considered:

$$g = \frac{\text{card}(\phi_k \cap \xi_k)}{k}. \tag{3.4}$$

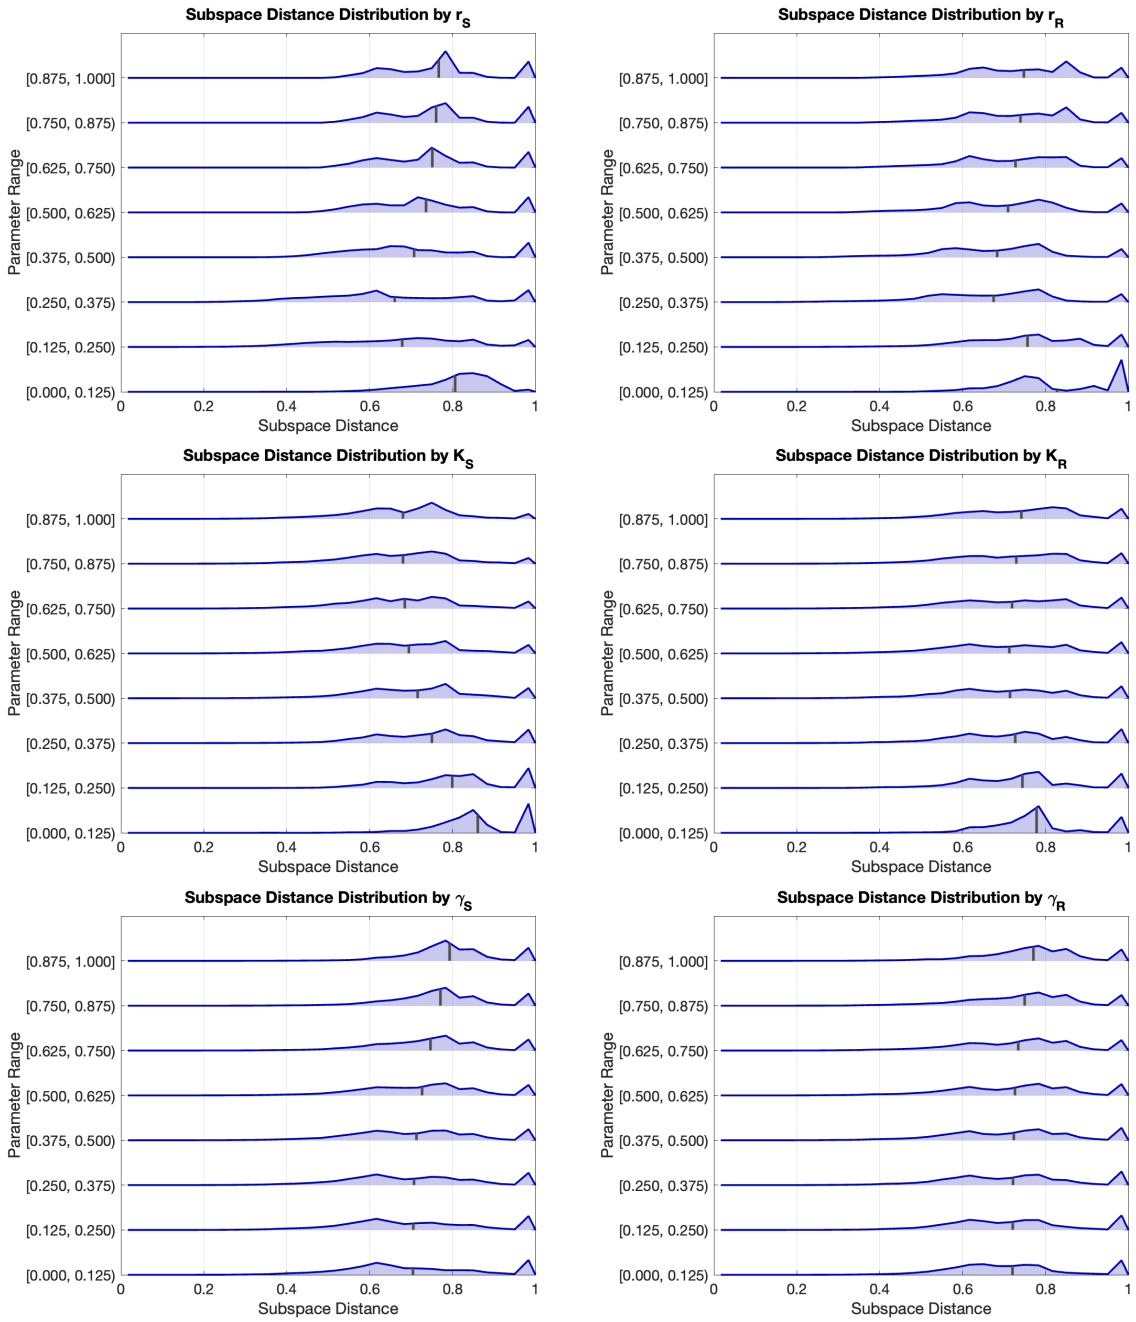

Figure 5. Distributions of local–global subspace distances across eight parameter bins for each model parameter. Each ridge represents a bin of width 0.125, and gray vertical bars denote bin-wise mean subspace distances.

The metric of Eq (3.4)—used in [31] to compare rank-based screening methods to Sobol' indices across an array of test functions—is employed here to capture the level of agreement between the local and global methods with regard to the top k parameters. Though such a metric treats all parameter substitutions equally, such that exchanging the top-two most influential parameters is penalized in the

same manner as exchanging the most and least influential parameters, the g-score provides a quick assessment of the hazards of using misleading sensitivity rankings in a parameter subset selection context where only the top k parameters are retained, a common usage of sensitivity analyses.

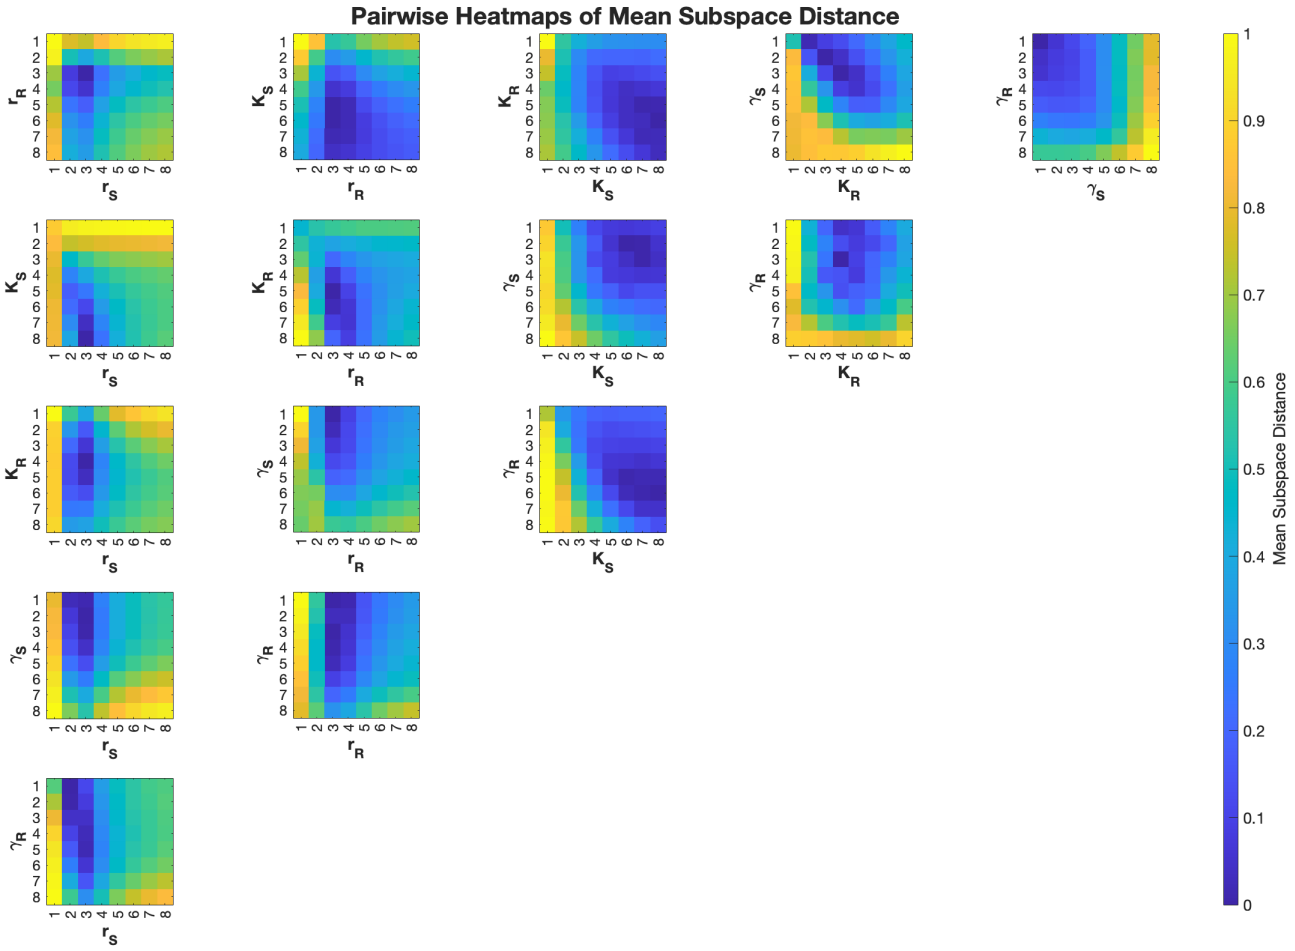

Figure 6. Pairwise heatmaps of the mean subspace distance between local and global active subspaces for the Lotka–Volterra model. Each heatmap corresponds to a pair of parameters, with colors indicating the average subspace distance over the associated two-dimensional parameter grid. Axis labels 1–8 denote the eight parameter bins of width 0.125, corresponding to the intervals [0,0.125), [0.125,0.25), ..., [0.875,1]. Larger values indicate greater deviation between local and global active subspaces.

To understand how this agreement changes as a function of subspace distance, we sort the subspace distances into intervals of width 0.2, and average the g-scores for all regions that fall into each bin. The results are presented in Figure 7, with error bars provided by bootstrapped confidence intervals with $N = 1000$ samples taken from among the regions in each bin. Note that errors are largest for the 0–0.2 distance case, since there are fewer regions in this category to be averaged; the bin counts are given by 63, 4362, 44,939, 133,127, and 79,653 regions, respectively. As expected, agreement between the two parameter rankings is highest when the subspace distance between them is close to zero, and generally declines as the two subspaces move out of alignment. Moreover, while the two methods frequently agree on the *least* influential parameters (as evidenced by high average g-scores for

$k = 3, 4, 5$, even for larger subspace distances), the two are very likely to disagree on the *most* influential parameter, as illustrated by the $k = 1$ scenario. For instance, when selecting the single most influential parameter, the average g-score is only 54% even for subspace distances in the 0–0.2 range, and declines sharply for larger subspace distances. When selecting the top-two most influential parameters, there is moderately strong agreement (an average of $\approx$ 76%) for the smallest observed subspace distances between 0–0.2, but that agreement score declines to an average of 36% for subspace distances between 0.8–1.0. Overall, the results of Figure 7 lend credence to the importance of considering local sensitivity variability when conducting parameter subset selection. Moreover, producing an analysis akin to that in Figure 7 provides a researcher with a way to quantify what subspace distances are acceptable for their specific model and purpose.

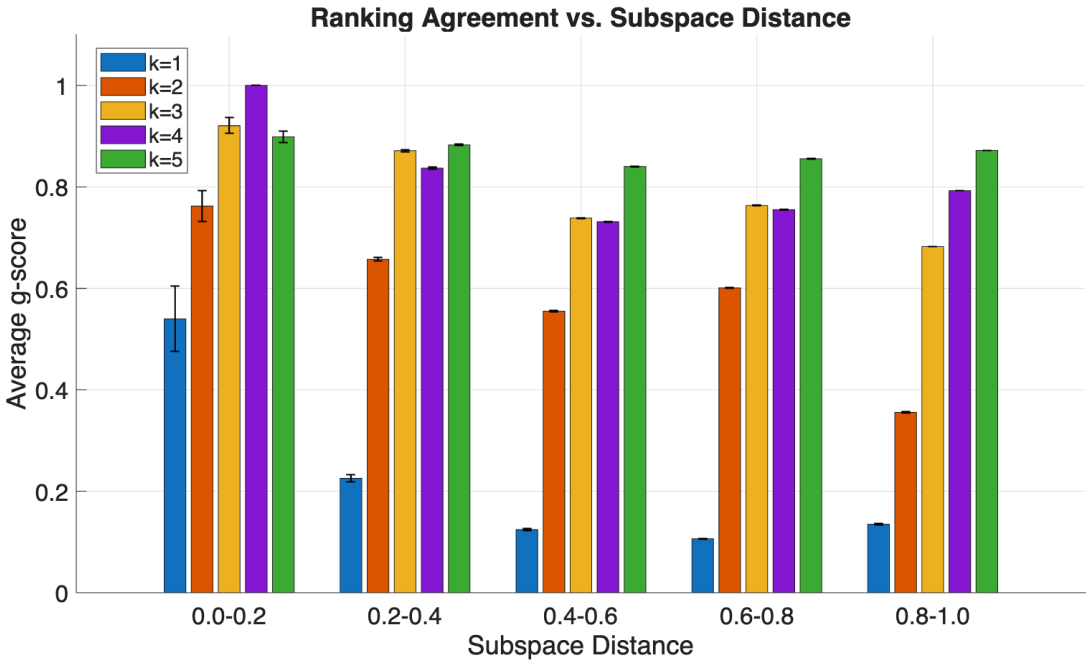

Figure 7. Average parameter ranking agreement between local and global analyses as a function of distance between locally and globally derived active subspaces. Error bars represent bootstrapped errors with $N = 1000$ samples.

3.2.3. Statistical reliability

The number of Monte Carlo samples used for the active subspace calculations throughout this investigation ($M = 25$ for each local region) is motivated by the bound derived in [14]. Specifically, Constantine recommends the use of $M = \alpha k \log(m)$ samples to ensure accurate estimation of the first k eigenvalues of the matrix $\mathbf{C}$, where m represents the number of input parameters and $\alpha \in [2, 10]$ is an oversampling factor. For our model, even the conservative choice of $k = m$ yields a suggested range of M between 11 and 54 across the range of α, placing our choice of $M = 25$ comfortably within this range.

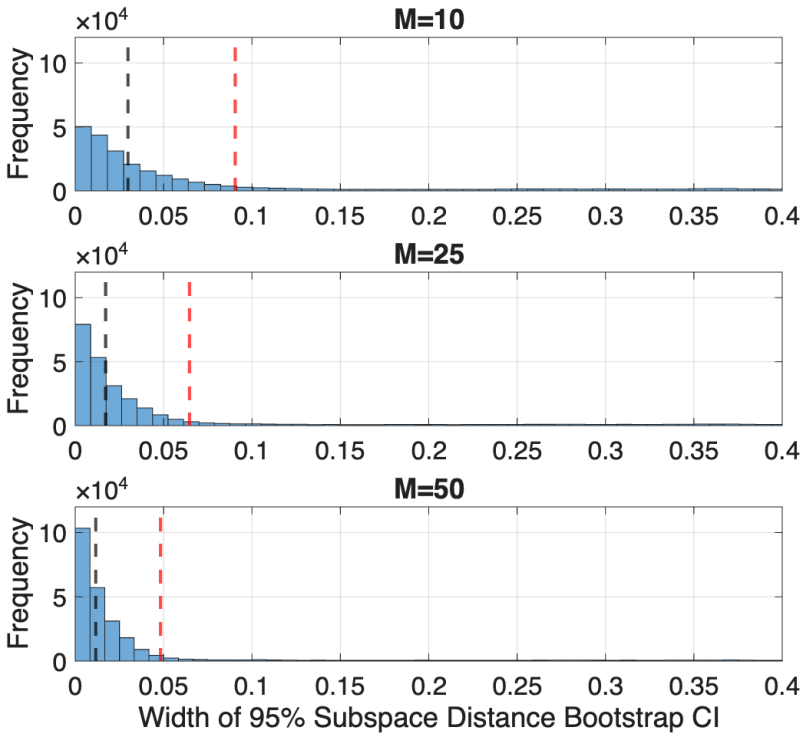

Figure 8. Subspace distance bootstrap confidence interval widths for $M = 10$, $M = 25$, and $M = 50$ samples. Mean CI widths are shown in red and median CI widths in black.

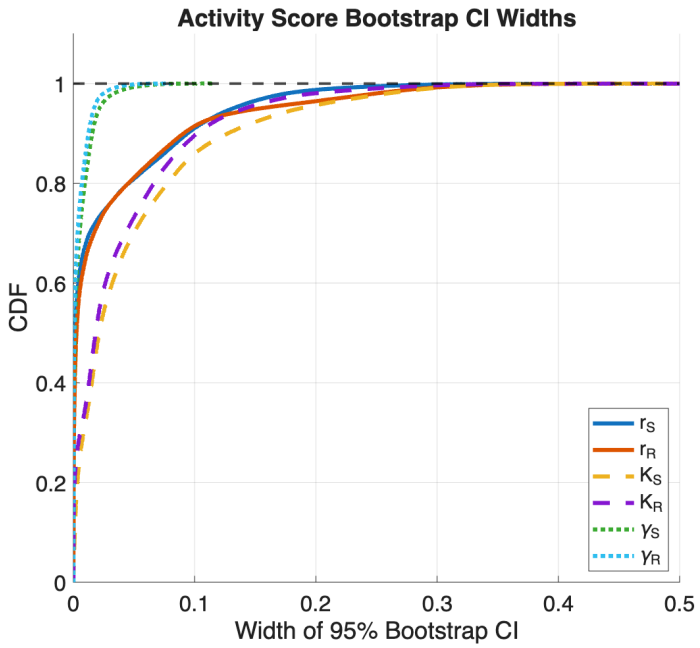

Figure 9. CDF plots of activity score bootstrap confidence interval widths for $M = 25$ samples.

To further verify that $M = 25$ is sufficient, we examine bootstrapped confidence intervals on the subspace distance estimates for the set $M \in \{10, 25, 50\}$ in Figure 8. If the eigendecomposition were poorly estimated for small M, we would expect the confidence intervals to narrow substantially as M increases. Instead, the reduction in average interval width as we move from $M = 10$ to $M = 25$ is moderate, and from $M = 25$ to $M = 50$ is only minimal, suggesting that there is little to be gained by increasing the sample size beyond $M = 25$. Moreover, the number of unique local rankings observed decreases only slightly, from 306 unique rankings with $M = 10$ to 304 rankings with $M = 25$ to 301 rankings using $M = 50$; this provides further evidence that the existence of so many unique rankings is in fact due to differences in local sensitivity and not just due to noise in Monte Carlo sampling.

Finally, Figure 9 shows the cumulative distribution functions of bootstrapped confidence interval widths—one confidence interval per local region—for each of the six activity scores across all 262,144 regions, using $M = 25$ Monte Carlo samples. Parameters with low overall influence, such as γ_S and γ_R, exhibit narrow confidence intervals throughout. Even the more influential parameters, such as K_S and K_R, show distribution functions that rise steeply toward 1, indicating that the confidence interval widths are small across the vast majority of regions. Together, these diagnostics support the adequacy of $M = 25$ samples for the analysis contained in this investigation.

3.2.4. Bifurcation analysis

One major drawback of global sensitivity analysis is the dependence of the obtained measures upon the imposed parameter ranges and distributions. Because global sensitivity measures cannot identify regions of the input space that produce implausible model outputs or differentiate between regions for which model outputs vary drastically in behavior due to parameter-driven bifurcations [7–9], one might ask whether the substantial subspace distances observed previously reflect the comparison of local sensitivities to a global measure averaged across multiple different bifurcation regions. In what follows, we illustrate that while restricting comparison to a global region consisting solely of parameter values leading to the same long-term outcome as the local region in question may reduce the subspace distance slightly, the general issue persists.

The Lotka-Volterra system as defined in Eqs (3.1) and (3.2) has four equilibria: the mutual extinction equilibrium, two competitive exclusion equilibria, and one coexistence equilibrium. Using the Jacobian matrix,

$$J(S,R) = \begin{bmatrix} r_S - \frac{2r_S}{K_s}S - \frac{r_S\gamma_R}{K_S}R & -\frac{r_S\gamma_R}{K_S}S \\ -\frac{r_R\gamma_S}{K_R}R & r_R - \frac{2r_R}{K_R}R - \frac{r_R\gamma_S}{K_R}S \end{bmatrix},$$

the stability of each equilibrium can be determined for a given parameter set $[r_S, r_R, K_S, K_R, \gamma_S, \gamma_R]$. In particular, the extinction equilibrium $(0, 0)$ is always unstable, while exactly one of the three remaining equilibria is asymptotically stable for any given sample. The breakdown of samples corresponding to each possible long-term outcome is provided in Table 3.

By grouping parameter samples according to the long-term behavior of their corresponding model trajectories, we recompute subspace distances comparing local sensitivity information to global sensitivity measures derived solely from samples with matching outcomes. This allows us to assess both whether sensitivity variability differs across outcome types, and whether controlling for bifurcation region meaningfully reduces the misalignment between local and global sensitivity information.

Table 3. Equilibria of the Lotka-Volterra system (3.1) and (3.2), and the number of parameter samples for which that equilibrium is asymptotically stable.

Equilibrium	(S^*, R^*)	# Samples	% Samples
Mutual extinction	$(0, 0)$	0	0
Competitive exclusion	$(0, K_R)$	1,638,454	25%
Competitive exclusion	$(K_S, 0)$	1,638,696	25%
Coexistence	$\left(\frac{K_S - \gamma_R K_R}{1-\gamma_S \gamma_R}, K_R - \gamma_S \frac{K_S - \gamma_R K_R}{1-\gamma_S \gamma_R}\right)$	3,276,450	50%

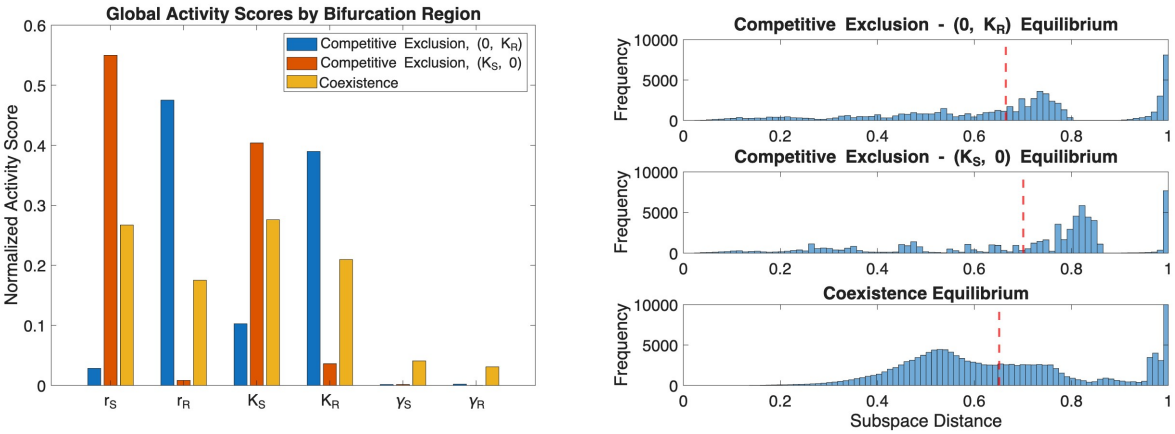

Figure 10. Left: Global activity scores for each of the three bifurcation regions. Right: Comparison of subspace distances when controlling for bifurcation region among regions with different long-term outcomes. Mean subspace distances are shown in red.

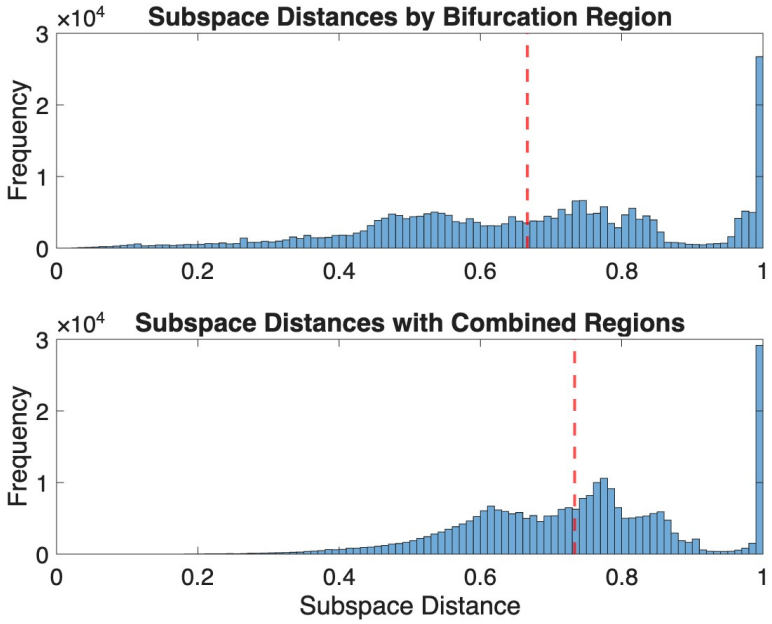

Figure 11. Comparison between subspace distances when distances are computed according to bifurcation region (top) versus a single global average (bottom). Mean subspace distance values are shown in red.

Figure 10 displays the globally computed activity scores and subspace distances computed according to bifurcation region, broken down by the long-term outcome. The global activity scores reveal stark differences in sensitivity rankings across bifurcation regions. For the regions characterized by the two competitive exclusion equilibria, the total tumor burden QoI is driven almost entirely by the growth rate and carrying capacity of the dominating cell type, while the region characterized by an asymptotically stable coexistence equilibrium yields more balanced scores with meaningful contributions from parameters of both cell types. However, even when restricting comparison to samples that share the same asymptotic behavior, subspace distances remain substantial across all three outcome types, with mean distances of 0.6650, 0.7010, and 0.6513, respectively.

Collectively, Figure 11 compares subspace distances obtained when each local region is matched to its corresponding bifurcation-specific global active subspace versus the single global active subspace computed across the full admissible parameter space. Controlling for bifurcation region shifts the overall mean subspace distance modestly downward (from 0.7334 to 0.6665), but the misalignment between local and global sensitivity information remains substantial.

3.2.5. Secondary QoI

Because sensitivity measures are derived using gradients of the quantity of interest computed with respect to each of the model inputs, the sensitivity landscape is highly dependent upon the choice of QoI. Here, we present results for a second quantity of interest in order to illustrate that the issue of sensitivity information misalignment persists even when the QoI is altered.

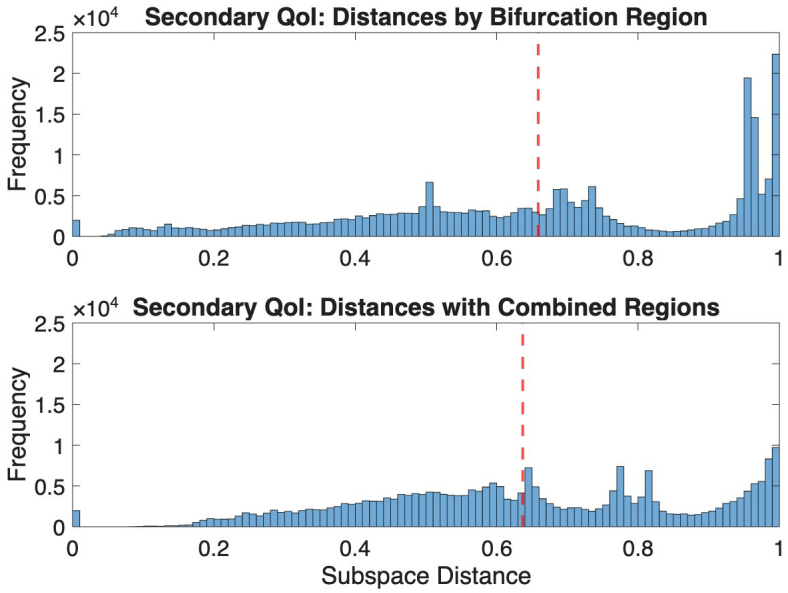

Figure 12. Using the final ratio of Type S cell volume to total tumor volume as the QoI: Comparison between subspace distances when distances are computed according to bifurcation region (top) versus a single global average (bottom). Mean subspace distance values are shown in red.

Our secondary quantity of interest,

$$QoI_{sec} = \frac{S(56)}{S(56) + R(56)}, \tag{3.5}$$

measures the final ratio of Type S cell volume to the total tumor volume after eight weeks. This ratio is clinically meaningful in the context of the Lotka-Volterra model, which has been used to analyze interactions between radiosensitive (Type S) and radioresistant (Type R) cell types [29, 30]. In particular, knowing the cellular composition of the tumor at a given point in the treatment period allows clinicians to reassess the predicted efficacy of radiation therapy and—if resistance has become prevalent—pursue alternative treatments.

Figure 12 shows the collective subspace distances for this quantity of interest, both when broken down by bifurcation region (top) and when compared to a single global average (bottom). Overall, subspace distances remain substantial, with means of 0.6588 and 0.6369, respectively. Of note, controlling for bifurcation region actually slightly increases the average subspace distance for this QoI, illustrating that this misalignment issue is due to factors other than just parameter-driven bifurcations.

3.3. Implications for downstream tasks

In general, the discrepancy between the global and local metrics reveals an important pitfall of relying solely on global sensitivity rankings for model reduction. For instance, a global analysis of this model using the active subspace method on the specified parameter space suggests that attention should be focused primarily on the growth rate and carrying capacity of Type S cells in subsequent tasks, while local analyses reveal that in many subregions, the Type S growth rate may actually be one of the least influential parameters. Sections 3.3.1 and 3.3.2 will illustrate how the use of global sensitivity measures in the presence of such discrepancies may lead to issues with downstream tasks that rely upon model dimension reduction: specifically, model calibration and surrogate model construction. These two procedures illustrate complementary uses of global sensitivity analysis: identifying a *subset*-based reduction of the input space consisting of influential parameters to be estimated (model calibration) and identifying a *subspace*-based reduction of the input space consisting of influential combinations of parameters onto which the model response may be projected for dimension reduction (surrogate model construction).

3.3.1. Model calibration illustration

For models in which the full parameter set is not uniquely identifiable given the available data, it is common practice to conduct a sensitivity analysis and fix any uninfluential parameters at their nominal values prior to fitting the model to data. Here, we illustrate how using global information to select which parameters are fixed may lead to a less-effective model calibration result. This is an example of a *subset*-based reduction task, for which we are applying the activity score metric derived from the active subspace methodology.

For each of the 8^6 local regions defined previously for the Lotka-Volterra model, we generate synthetic data by computing the primary QoI of Eq (3.3), q, for a parameter set drawn at random from that region and adding up to 10% noise sampled uniformly,

$$q_{data} = (1 + \varepsilon)q, \quad \varepsilon \sim \mathcal{U}(-0.1, 0.1).$$

We then fit the model to our synthetic data—using the `lsqnonlin` function in MATLAB—in two ways: once estimating only the top k parameters identified by the global activity scores and once estimating only the top k parameters from the corresponding local activity score analysis, with the remaining $m - k$ parameters fixed at the center of their specified ranges. The performance of the resulting model fits (assessed using $(q_{data} - q_{model})^2$ as the optimized metric) is compared. This procedure is outlined in Figure 13.

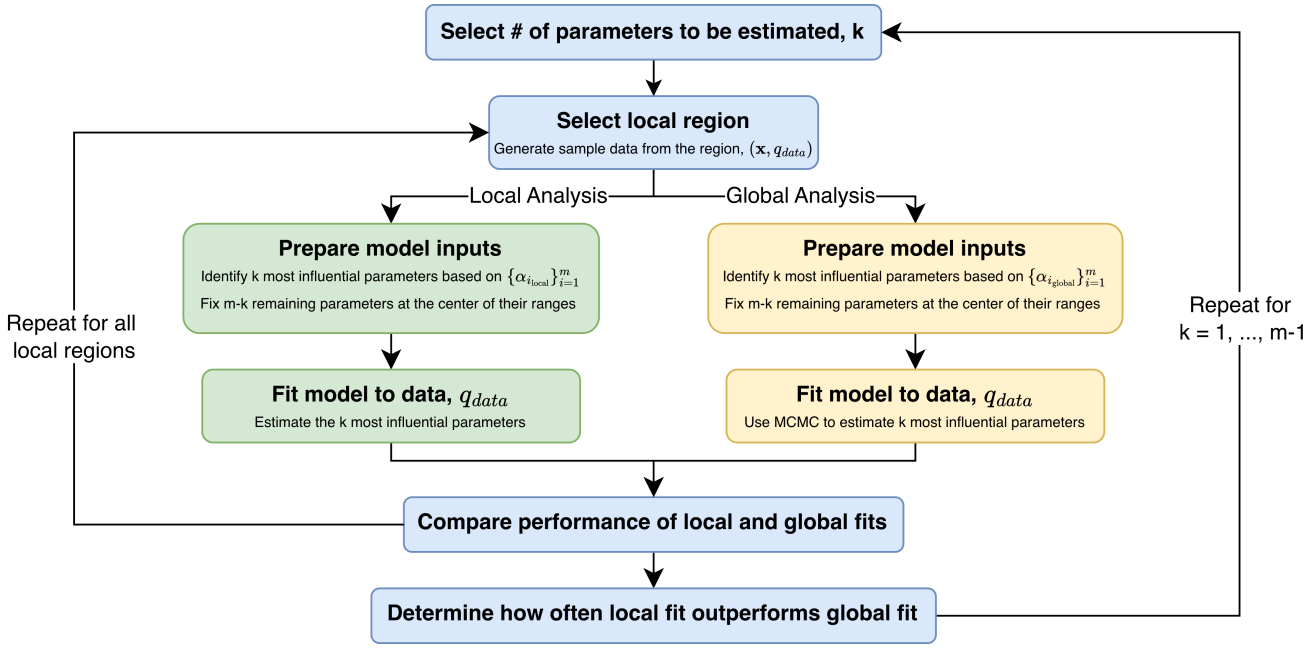

Figure 13. Model calibration experiment procedure.

Figures 14 and 15 show the results of our model calibration experiment. In Figure 14, we record the percentage of scenarios for which the local calibration results in the better model fit when compared to calibration using the global parameter ranking order. Specifically, calibration using local rankings dominates for $k = 1$ with a 79% win rate; when only one parameter is estimated, it is crucial to focus attention on that which renders the model most flexible in the region from which the data is produced to achieve the best possible fit, which requires the use of a locally sensitive parameter. For $k = 2$ through $k = 5$, the use of locally sensitive parameters still presents a clear advantage, with local fits being superior to global fits between 58% and 64% of the time. Figure 15 illustrates the difference between the global fit errors and local fit errors. From this figure, we observe that in scenarios where the local calibration outperforms the global calibration, it tends to do so by a wider margin than in scenarios for which the global calibration edges out the local approach, particularly for the $k = 1$ and $k = 2$ scenarios. As k increases, the distinction between locally and globally sensitive parameters becomes less important, with model fits performing almost equally well.

This experiment illustrates how the use of global sensitivity metrics may lead researchers to focus their attention on the wrong subset of parameters, reducing the efficacy of the calibration procedure. We suggest that if researchers observe large discrepancies between local and global active subspaces around the parameter space, they conduct further research to narrow the plausible ranges of their parameters prior to performing model calibration.

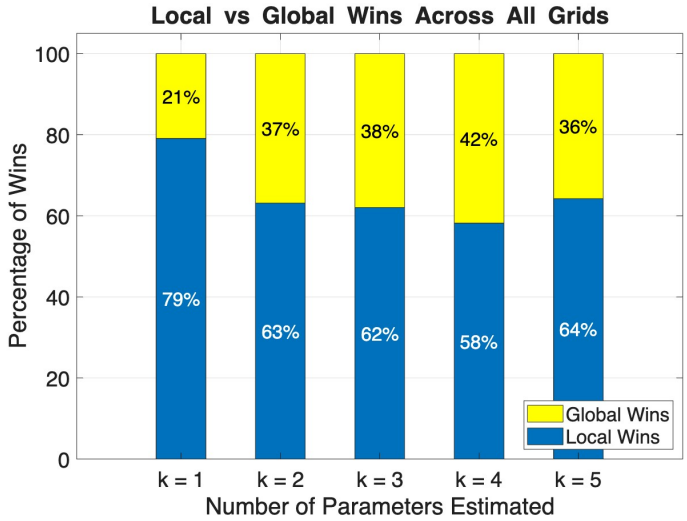

Figure 14. Comparison of local versus global fit wins by dimension.

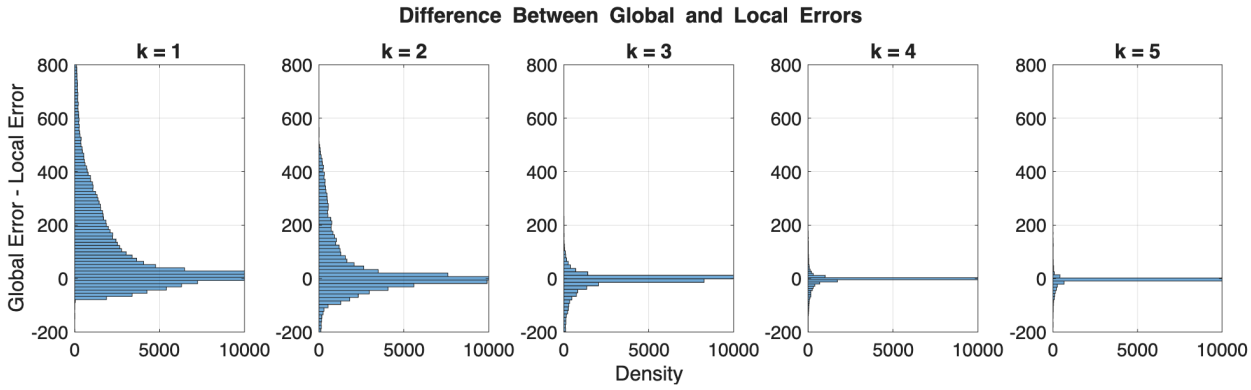

Figure 15. Comparison of model fit errors resulting from global and local approaches by dimension.

3.3.2. Surrogate model construction illustration

Another important task that relies heavily on sensitivity information is the construction of surrogate models to approximate the behavior of complex models, which may be too computationally expensive to evaluate repeatedly. Though the Lotka-Volterra model is not particularly expensive to evaluate, it serves as a useful illustrative example here. This experiment complements the earlier model calibration study by focusing on a *subspace*-based reduction application, in which the directions to be disregarded are linear combinations of the original model parameters as defined by the inactive subspace. In particular, such information cannot be derived from our comparison methods (Morris elementary effects and Sobol' indices), but is readily available using active subspace methodology.

For this investigation, we construct polynomial response surface surrogate models on a subspace of n dimensions, where $n < m$ (with m representing the total number of input parameters). For a given local region, we build two surrogate models to approximate the QoI on that region: one using the globally calculated active subspace and one using the local active subspace. In each case, polynomial response surfaces of orders 1–3 are trained on points generated from the local region in question and projected

into the active subspace. The "best" response surface model is chosen using the Akaike information criterion (AIC) [32], where model performance is determined using a set of testing points generated independently from those used to train the models. Response surface performance is then compared between the two resulting models, with each being evaluated for its ability to reproduce the testing point outputs. The procedure is outlined in Figure 16, and discussed in further detail in [2, 19, 33].

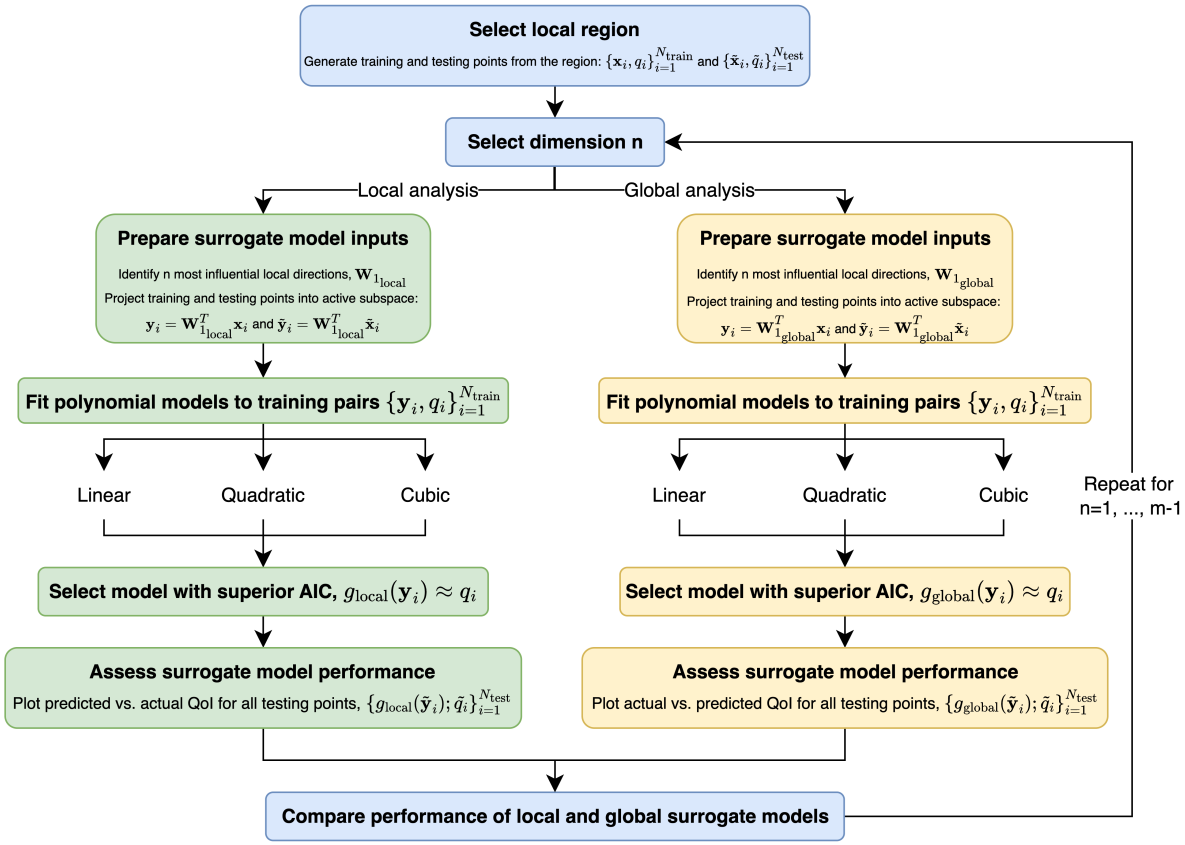

Figure 16. Surrogate modeling experiment procedure.

Specifically, we expect regions with small subspace distances to yield global and local surrogate models of comparable accuracy, whereas regions with large subspace distances should produce global surrogate models that perform markedly worse than their local counterparts. We illustrate this contrast by examining two representative regions: the region represented by

$$\begin{aligned} r_S, \gamma_S &\in [0.125, 0.250], \\ r_R &\in [0.250, 0.375], \\ K_S &\in [0.375, 0.500], \\ K_R, \gamma_R &\in [0.500, 0.625], \end{aligned}$$

which has a small subspace distance of 0.1355, and the topmost region, $[0.875, 1]^6$, with a much more significant subspace distance of 0.9996.

These comparisons are shown in Figures 17 and 18, respectively. On the first region of interest (see Figure 17), we observe that both the locally and globally constructed surrogate models exhibit strong

agreement between actual and predicted values of the quantity of interest across all reduced dimensions, $n = 1, \ldots, 5$. Although the local model achieves marginally tighter agreement, both perform accurately, confirming that for regions with small subspace distances, reliance on global sensitivity measures for dimension reduction does not substantially compromise surrogate model accuracy.

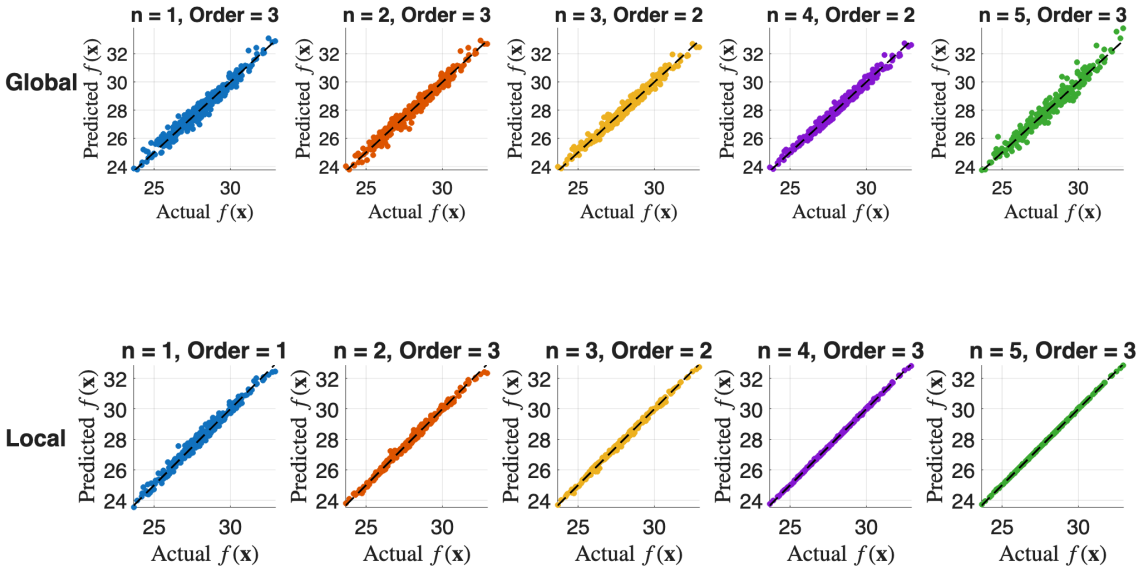

Figure 17. Predictive capability of global (top) versus local (bottom) response surfaces on a region with subspace distance 0.1355.

In Figure 18, which shows the results for the topmost region $[0.875, 1]^6$—a region in which the discrepancy between local and global sensitivity information is more significant with a subspace distance of 0.9996—the benefit of using local sensitivity information is much more apparent. The local model achieves in a single dimension a level of accuracy that the global surrogate requires at least four dimensions to match. Moreover, the local model built on this single dimension (with a quadratic response surface) requires the estimation of only three parameters (parameters c_1–c_3 for the model form $c_1 + c_2 y_1 + c_3 y_1^2$), while the global model built using a quadratic polynomial on $n = 4$ active variables requires the estimation of 15 parameters. The latter clearly no longer constitutes a simplification of the original model as would be desired.

In closing, we confirm that surrogate models built using global information with the goal of approximating local behavior yield less accuracy and tend to require larger dimensions to perform well. While this is generally not surprising—after all, local behavior will always be better approximated using local information—the decrease in global model accuracy (or increase in dimensions required to perform comparably to the local model) is compounded as the alignment in sensitivity information between local and global scales worsens. Thus, for models in which sensitivity varies widely across the parameter space, it is perhaps more beneficial to construct and piece together multiple surrogate models that are tuned to the local information from the regions on which they are constructed, rather than constructing a single global model that will suffer from reduced accuracy throughout the space.

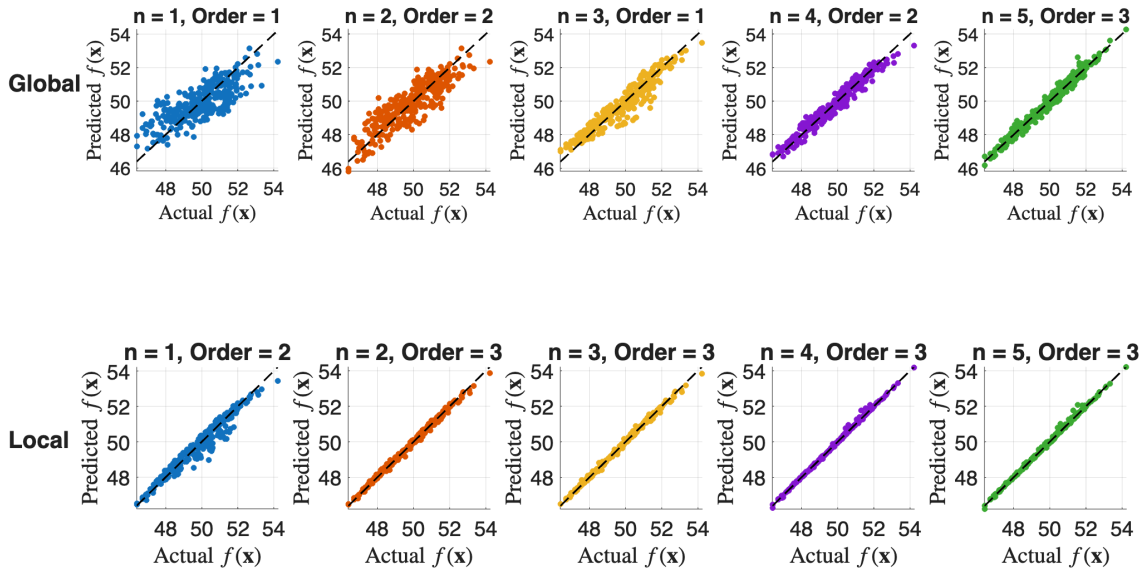

Figure 18. Predictive capability of global (top) versus local (bottom) response surfaces on the region $[0.875, 1]^6$ with subspace distance 0.9996.

4. Conclusions

Global sensitivity metrics are commonly used in moderate- to high-dimensional models to assess parameter importance. However, these metrics may misrepresent sensitivity when local variability is significant, leading to consequences in subsequent model analysis tasks. In this study, we utilize activity scores derived from the active subspace to demonstrate how parameter rankings may shift throughout the space, and quantify the divergence of local behavior from global results using the distance between locally and globally computed active subspaces. By employing a grid-based approach to studying the stability of the active subspace, we provide a means by which to identify regions of the parameter space for which global sensitivity measures can be considered locally representative.

Employing a variation of the Lotka-Volterra model used to study interactions between competing cell lines in a tumor spheroid, we find that large discrepancies exist between global and local parameter rankings across multiple established sensitivity methodologies, indicating that global summaries cannot adequately capture the full behavior of the system. These inconsistencies have numerous practical implications. A model calibration experiment shows that fitting the model using a subset of influential parameters chosen according to globally derived sensitivity information often results in a poorer model fit, as compared to using local sensitivity information to select which parameters are to be estimated. A second experiment in surrogate model construction reveals that surrogate models that utilize global subspace information generally require higher-dimensional representations while yielding lower accuracy when compared to surrogate models constructed on locally defined active subspaces. Such issues can result in flawed conclusions about parameter importance and may reduce the subsequent reliability of models in sensitive applications such as biological systems or engineering design. Moreover,

the level of severity of such discrepancies even for a simple Lotka-Volterra system—a model used widely across the fields of mathematical biology and ecology, as well as others—raises real concerns about the use of global sensitivity analyses to draw meaningful conclusions about more complex biological models.

Although conducting a thorough local analysis to accompany global measures may be computationally expensive—particularly for models without adjoint capabilities—this investigation illustrates that such analysis is a worthwhile endeavor to avoid misrepresentation of parameter importance in subsequent analysis. The computational expense for our proposed approach—assuming the use of forward finite differences for gradient estimation—scales according to $(m + 1)Mb^m$, where b represents the number of bins employed for each of the m parameters and M is the number of Monte Carlo samples used for each local region. While this may be feasible for moderately sized input spaces such as that of the Lotka-Volterra model employed in this study, it becomes cost-prohibitive as m increases, particularly since M itself grows with m, compounding the curse of dimensionality. Thus, the scalability of such approaches for investigating local sensitivity variability is of concern.

Notably, this scalability concern is not unique to our approach. For instance, Pannier and Graf [13] similarly identified computational inefficiency as a limitation of sectional global sensitivity measures (SGSMs) in high-dimensional settings, and significantly reduced the number of regions required for analysis by adopting a grid-based heuristic that segments the parameter space one input at a time while allowing all others to retain their full ranges. This segmentation could be similarly adopted in our approach if the computational budget is a concern; however, we note that this may come at the cost of obscuring interesting parameter interaction dynamics. An alternative approach could employ adaptive techniques such as those proposed in [19, 34], in which the number of function evaluations required to approximate the active subspace is substantially reduced by taking finite difference steps aligned with the active variables and avoiding model evaluations in directions where the gradient is relatively flat; gradients sampled in this rotated space can then be transformed back into the original parameter space prior to the calculation of the $\mathbf{C}$ matrix. More broadly, a growing body of literature addresses the computational challenge of gradient estimation in high-dimensional settings through the use of surrogate models. For instance, Wycoff et al. [35] demonstrated that performing finite differencing on a trained surrogate model reduces the cost of gradient estimation, and further showed that when the surrogate is a Gaussian process, a closed-form expression for the active subspace exists, rendering finite differencing unnecessary altogether. In Tripathy et al. [36], the authors took a complementary approach, replacing high-fidelity models with surrogates such as polynomial chaos expansions, neural networks (also used as surrogates in [13]), or Gaussian processes, and developed a gradient-free probabilistic formulation of the active subspace that is robust to both observational noise and stochasticity. Together, these works demonstrate that surrogate modeling presents a viable path toward extending the proposed framework to higher-dimensional settings. For practitioners seeking to do so, [37] provides a broad overview of surrogate modeling strategies for global sensitivity analysis, including guidance on selecting an appropriate surrogate for a given problem.

In closing, global sensitivity analysis is often heralded as a means by which to address a multitude of challenges related to model sensitivity and identifiability; however, we caution researchers to investigate the level of variability in local sensitivities for their model before using global sensitivity results to make consequential decisions. Active subspace methodology is well-suited for this purpose, since the active subspace can identify not only which parameters are most important in any given region, but

also regions in which local and global sensitivity patterns are in alignment via subspace-based analysis. Such analysis can then prompt suggestions for refinement of the parameter space to reduce variability prior to completing downstream tasks.

Author contributions

H. Zou: Methodology, software, formal analysis, visualization, writing. A. L. Lewis: Conceptualization, methodology, software, formal analysis, writing, supervision.

Code availability statement

Sample MATLAB code for the two-dimensional and Lotka-Volterra examples can be found at https://github.com/lewisallison10/Global-vs.-Local-Sensitivity-Using-the-Active-Subspace.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The authors declare no conflicts of interest. Allison Lewis is a guest editor for *Mathematical Biosciences and Engineering* and was not involved in the editorial review or the decision to publish this article.

References

1. A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, et al., *Global Sensitivity Analysis: The Primer*, John Wiley & Sons, Chichester, 2008.
2. R. C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, 2nd edition, Society for Industrial and Applied Mathematics, Philadelphia, 2024.
3. M. D. Morris, Factorial sampling plans for preliminary computational experiments, *Technometrics*, **33** (1991), 161–174. https://doi.org/10.1080/00401706.1991.10484804
4. I. M. Sobol', Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates, *Math. Comput. Simul.*, **55** (2001), 271–280. https://doi.org/10.1016/S0378-4754(00)00270-6
5. R. I. Cukier, C. M. Fortuin, K. E. Shuler, A. G. Petschek, J. H. Schaibly, Study of the sensitivity of coupled reaction systems to uncertainties in rate coefficients, *J. Chem. Phys.*, **59** (1973), 3873–3878. https://doi.org/10.1063/1.1680571
6. A. F. Brouwer, M. C. Eisenberg, The underlying connections between identifiability, active subspaces, and parameter space dimension reduction, preprint, arXiv:1802.05641v1.
7. G. Qian, A. Mahdi, Sensitivity analysis methods in the biomedical sciences, *Math. Biosci.*, **323** (2020), 108306. https://doi.org/10.1016/j.mbs.2020.108306

8. G. A. T. Broeke, A. K. Van Voorn, B. W. Kooi, J. Molenaar, Detecting tipping points in ecological models with sensitivity analysis, *Math. Model. Nat. Phenom.*, **11** (2016), 47–72. https://doi.org/10.1051/mmnp/201611405
9. M. Shin, J. H. A. Guillaume, B. F. W. Croke, A. J. Jakeman, Addressing ten questions about conceptual Rainfall-Runoff models with global sensitivity analyses in R, *J. Hydrol.*, **503** (2013), 135–152. https://doi.org/10.1016/j.jhydrol.2013.08.047
10. R. Bolado-Lavin, W. Castaings, S. Tarantola, Contribution to the sample mean plot for graphical and numerical sensitivity analysis, *Reliab. Eng. Syst. Saf.*, **94** (2009), 1041–1049. https://doi.org/10.1016/j.ress.2008.11.012
11. S. Tarantola, V. Kopustinskas, R. Bolado-Lavin, A. Kaliatka, E. Ušpuras, M. Vaišnoras, Sensitivity analysis using contribution to sample variance plot: Application to a water Hammer model, *Reliab. Eng. Syst. Saf.*, **99** (2012), 62–73. https://doi.org/10.1525/bio.2012.62.1.18
12. P. Wei, Z. Lu, J. Song, Regional and parametric sensitivity analysis of Sobol' indices, *Reliab. Eng. Syst. Saf.*, **137** (2015), 87–100. https://doi.org/10.1016/j.ress.2014.12.012
13. S. Pannier, W. Graf, Sectional global sensitivity measures, *Reliab. Eng. Syst. Saf.*, **134** (2015), 110–117. https://doi.org/10.1016/j.ress.2014.09.009
14. P. G. Constantine, *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*, Society for Industrial and Applied Mathematics, Philadelphia, 2015.
15. I. M. Sobol', S. Kucherenko, Derivative based global sensitivity measures and their link with global sensitivity indices, *Math. Comput. Simul.*, **79** (2009), 3009–3017. https://doi.org/10.1016/j.matcom.2009.01.023
16. F. Campolongo, J. Cariboni, A. Saltelli, An effective screening design for sensitivity analysis of large models, *Environ. Modell. Software*, **22** (2007), 1509–1518. https://doi.org/10.1016/j.envsoft.2006.10.004
17. N. Halko, P. G. Martinsson, J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.*, **53** (2011), 217–288. https://doi.org/10.1137/090771806
18. I. T. Jolliffe, *Principal Component Analysis*, 2nd edition, Springer, New York, 2002.
19. K. D. Coleman, A. L. Lewis, R. C. Smith, B. J. Williams, M. Morris, B. Khuwaileh, Gradient-free construction of active subspaces for dimension reduction in complex models with applications to neutronics, *SIAM/ASA J. Uncertainty Quantif.*, **7** (2019), 117–142. https://doi.org/10.1137/16M1075119
20. P. G. Constantine, P. Diaz, Global sensitivity metrics from active subspaces, *Reliab. Eng. Syst. Saf.*, **188** (2018), 243–254. https://doi.org/10.1016/j.ress.2017.01.013
21. G. Stewart, Error and perturbation bounds for subspaces associated with certain eigenvalue problems, *SIAM Rev.*, **15** (1973), 727–764. https://doi.org/10.1137/1015095
22. A. J. Lotka, *Elements of Physical Biology*, Williams & Wilkins, Baltimore, 1925.
23. V. Volterra, Fluctuations in the abundance of a species considered mathematically, *Nature*, **118** (1926), 558–560. https://doi.org/10.1038/118558a0

24. T. Zhang, Z. Cai, Emergent non-Hermitian physics in a generalized Lotka–Volterra model, *Phys. Rev. B*, **108** (2023), 104304.
25. V. V. Usenko, Interpretation and dynamics of the Lotka–Volterra model in the description of a three-level laser, preprint, arXiv:2112.12045.
26. K. W. Fornalski, Application of Lotka–Volterra equations for homeostatic response to an ionizing radiation stressor, *Appl. Sci.*, **13** (2023), 11077. https://doi.org/10.3390/app131911077
27. A. J. Lotka, Contribution to the theory of periodic reactions, *J. Phys. Chem.*, **14** (1910), 271–274. https://doi.org/10.1021/j150111a004
28. A. J. Lotka, Analytical note on certain rhythmic relations in organic systems, *J. Phys. Chem.*, **24** (1920), 239–244. https://doi.org/10.1073/pnas.6.7.410
29. H. Cho, A. L. Lewis, K. M. Storey, H. M. Byrne, Designing experimental conditions to use the Lotka–Volterra model to infer tumor cell line interaction types, *J. Theor. Biol.*, **559** (2023), 111377. https://doi.org/10.1016/j.jtbi.2022.111377
30. M. Paczkowski, W. Kretzschmar, B. Markelc, S. K. Liu, L. A. Kunz-Schughart, A. L. Harris, et al., Reciprocal interactions between tumour cell populations enhance growth and reduce radiation sensitivity in prostate cancer, *Commun. Biol.*, **4** (2021), 6. https://doi.org/10.1038/s42003-020-01529-5
31. F. Campolongo, A. Saltelli, J. Cariboni, From screening to quantitative sensitivity analysis: A unified approach, *Comput. Phys. Commun.*, **182** (2011), 978–988. https://doi.org/10.1016/j.cpc.2010.12.039
32. S. Konishi, G. Kitagawa, *Information Criteria and Statistical Modeling*, Springer, New York, 2008.
33. A. L. Lewis, R. C. Smith, B. J. Williams, Bayesian model calibration on active subspaces, in *2017 American Control Conference (ACC)*, (2017), 3165–3170. https://doi.org/10.23919/ACC.2017.7963434
34. A. L. Lewis, R. C. Smith, B. J. Williams, Gradient-Free active subspace construction using morris screening elementary effects, *Comput. Math. Appl.*, **72** (2016), 1603–1615. https://doi.org/10.1016/j.camwa.2016.07.022
35. N. Wycoff, M. Binois, S. M. Wild, Sequential learning of active subspaces, *J. Comput. Graphical Stat.*, **30** (2021), 1224–1237. https://doi.org/10.1080/10618600.2021.1874962
36. R. Tripathy, I. Bilionis, M. Gonzalez, Gaussian processes with built-in dimensionality reduction: Applications to high-dimensional uncertainty propagation, *J. Comput. Phys.*, **321** (2016), 191–223. https://doi.org/10.1016/j.jcp.2016.05.039
37. K. Cheng, Z. Lu, C. Ling, S. Zhou, Surrogate-assisted global sensitivity analysis: An overview, *Struct. Multidiscip. Optim.*, **61** (2020), 1187–1213. https://doi.org/10.1007/s00158-019-02413-5

AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)