

# Discovering an Active Subspace in a Single-Diode Solar Cell Model

Paul G. Constantine<sup>1\*</sup>, Brian Zaharatos<sup>1</sup> and Mark Campanelli<sup>2</sup>

<sup>1</sup>*Department of Applied Mathematics and Statistics, Colorado School of Mines, 1500 Illinois Street, Golden, CO 80401, USA*

<sup>2</sup>*National Renewable Energy Laboratory, Golden, CO, USA*

Received 30 June 2014; revised 6 March 2015; accepted 1 June 2015

DOI:10.1002/sam.11281

Published online 24 July 2015 in Wiley Online Library (wileyonlinelibrary.com).

**Abstract:** Predictions from science and engineering models depend on the values of the model's input parameters. As the number of parameters increases, algorithmic parameter studies such as optimization and uncertainty quantification require many more model evaluations. One way to combat this curse of dimensionality is to seek an alternative parameterization with fewer variables that produces comparable predictions. The *active subspace* is a low-dimensional linear subspace defined by important directions in the model's input space; input perturbations along these directions change the model's prediction more, on average, than perturbations orthogonal to the important directions. We describe a method for checking if a model admits an exploitable active subspace and apply this method to a single-diode solar cell model with five input parameters. We find that the maximum power of the solar cell has a dominant one-dimensional active subspace, which enables us to perform thorough parameter studies in one dimension instead of five. © 2015 Wiley Periodicals, Inc. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 8: 264–273, 2015

**Keywords:** single-diode solar cell model; active subspaces; dimension reduction; parameterized simulations

## 1. INTRODUCTION

Science and engineering simulations often contain several input parameters—e.g. physical constants, boundary conditions, or geometry descriptions. When presented with such parameterized simulations, the scientist naturally wonders how the simulated predictions depend on the input parameters. Which parameters, when perturbed, create the largest change in predictions? How precisely must the parameters be specified to ensure accurate predictions? And what is the effect of imprecisely prescribed input parameters on the predictions? If the goal is to maximize or minimize a predicted quantity, which combinations of parameters correspond to larger or smaller values of the prediction? In a similar vein, one may ask which parameter values correspond to predictions outside a safe region of operation—or which parameter values yield predictions that are consistent with a set of observations.

Many scientists rely on intuition about the physical system to answer these questions. But intuition becomes less

trustworthy as the simulations become more complex, e.g. when they include several interacting physical components. Algorithms for optimization, uncertainty quantification, and model calibration become more attractive as model complexity increases. If an algorithm can easily interface with the simulation—e.g. by automatically evaluating predictions given the values for the inputs—then applying the algorithm to the simulation becomes relatively easy.

The number of times the algorithm needs to evaluate a prediction increases—sometimes extremely rapidly—as the number of inputs increases. The situation is worse if each evaluation requires significant computational resources. For example, finding the global minimum of a complicated prediction depending on 100 inputs is not tractable if available resources permit only ten model evaluations. In practice, the scientist may choose only the most important parameters to vary so that the study fits within the computational budget. Alternatively, one may seek a low-dimensional description of the prediction as a function of the input parameters. If such a description is sufficiently accurate, then studies such as optimization and calibration can work in the space of fewer variables—potentially

\* Correspondence to: Paul G. Constantine (pconstan@mines.edu)

allowing the desired parameter studies within the given budget.

The *active subspace* is a low-dimensional linear subspace defined by important directions in the model's input space; input perturbations along these directions change the model's prediction more, on average, than perturbations orthogonal to the important directions. Not all models have an active subspace. Some models' predictions respond significantly to input perturbations along all directions. However, if a model does admit an active subspace, then one can exploit it to perform parameter studies in the coordinates of the subspace—i.e. the *active variables*—which are linear combinations of normalized versions of the model's input parameters. Therefore, it can be very advantageous to discover that a model admits an active subspace.

Active subspaces have been studied in a variety of contexts under different names. Cook's excellent text *Regression Graphics* [1] reviews and develops statistical methods for dimension reduction in the context of regression surfaces, and it contains references to the major works in the statistics literature. What we call the *active subspace* is a type of *dimension-reduction subspace* in Cook's parlance [1, Chapter 6]—though we are working with noiseless computer simulations in contrast to general regression surfaces. Russi's 2010 Ph.D. dissertation uses the phrase *active subspace* in a way comparable to our use [2]. He exploits the active subspace to construct quadratic surrogate models for uncertainty quantification in chemical kinetics models. Our prior work develops a theoretical framework for active subspaces including applications to kriging response surfaces [3]. We have applied these methods to several models in aerospace engineering [4–6]. Abdel-Khalik has applied similar methods in nuclear engineering applications [7].

In this paper, we describe how to test a model for an active subspace and apply this test to the maximum power from a single-diode model of a photovoltaic solar cell with five input parameters. In Section 2, we generically describe the active subspace and how to search for it. We then describe the single-diode model, its input parameters, and its predicted *performance parameters* (i.e., model outputs) in Section 3. We apply the tests for the active subspace to the model's maximum power in Section 4 and show that a dominant one-dimensional active subspace is present in the five-dimensional space of input parameters. We conclude in Sections 5 and 6 with a discussion of how one may exploit the low-dimension of the active subspace to further study the parameter dependence in the single-diode model's maximum power.

## 2. ACTIVE SUBSPACES

We consider a generic multivariate function  $f = f(\mathbf{x})$ , where  $\mathbf{x}$  represents the inputs of the model and  $f$  represents

a specific scalar performance parameter that the model predicts. Let  $\mathcal{X} = [-1, 1]^m$  be the domain with  $\mathbf{x} \in \mathcal{X}$  and let  $\rho : \mathcal{X} \rightarrow \mathbb{R}_+$  be a bounded and continuous weight function on  $\mathcal{X}$ ; we assume  $\rho$  is normalized to integrate to 1. Assume  $f$  is differentiable and absolutely continuous and denote the gradient  $\nabla_{\mathbf{x}} f(\mathbf{x}) = [\partial f / \partial x_1, \dots, \partial f / \partial x_m]^T$  oriented as a column vector.

Consider the following matrix  $\mathbf{C}$  defined as

$$\mathbf{C} = \int_{\mathcal{X}} (\nabla_{\mathbf{x}} f)(\nabla_{\mathbf{x}} f)^T \rho \, d\mathbf{x}. \quad (1)$$

In the context of dimension reduction for regression functions, Samarov calls this matrix an *average derivative functional* [8]. We are not studying regression functions, per se. Instead,  $f$  represents a noiseless, parameterized computer simulation. The matrix  $\mathbf{C}$  is symmetric and positively semi-definite, so it admits a real eigenvalue decomposition

$$\mathbf{C} = \mathbf{W} \Lambda \mathbf{W}^T, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_m), \\ \lambda_1 \geq \dots \geq \lambda_m \geq 0. \quad (2)$$

We can partition the eigenvectors,

$$\mathbf{W} = [\mathbf{W}_1 \quad \mathbf{W}_2], \quad \Lambda = \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix}, \quad (3)$$

where  $\mathbf{W}_1$  contains the first  $n$  eigenvectors and  $\Lambda_1$  contains the  $n$  largest eigenvalues. We use the two sets of eigenvectors to create new sets of variables  $\mathbf{y} = \mathbf{W}_1^T \mathbf{x}$  and  $\mathbf{z} = \mathbf{W}_2^T \mathbf{x}$ . We call the subspace defined by  $\mathbf{W}_1$  the *active subspace*, and we call the variables  $\mathbf{y}$  the *active variables*; the term *active subspace methods* for this type of analysis was first used in Russi's 2010 PhD dissertation [2]. The following two lemmas justify these labels.

LEMMA 1: The mean-squared directional derivative of  $f$  with respect to the eigenvector  $\mathbf{w}_i$  is equal to the corresponding eigenvalue,

$$\int_{\mathcal{X}} ((\nabla_{\mathbf{x}} f)^T \mathbf{w}_i)^2 \rho \, d\mathbf{x} = \lambda_i. \quad (4)$$

LEMMA 2: The mean-squared gradients of  $f$  with respect to the coordinates  $\mathbf{y}$  and  $\mathbf{z}$  satisfy

$$\int_{\mathcal{X}} (\nabla_{\mathbf{y}} f)^T (\nabla_{\mathbf{y}} f) \rho \, d\mathbf{x} = \lambda_1 + \dots + \lambda_n, \quad (5)$$

$$\int_{\mathcal{X}} (\nabla_{\mathbf{z}} f)^T (\nabla_{\mathbf{z}} f) \rho \, d\mathbf{x} = \lambda_{n+1} + \dots + \lambda_m.$$

The proofs of these lemmas can be found in our prior work [3]. In words, they mean that  $f$  changes more, on average, when its inputs are perturbed along the directions  $\mathbf{W}_1$  than along the directions  $\mathbf{W}_2$ ; the eigenvalues quantify precisely how much more. If an eigenvalue is exactly zero, then  $f(\mathbf{x})$  is constant along the direction defined by the corresponding eigenvector over all of  $\mathcal{X}$ .

To gain some intuition, consider the extreme case where all eigenvalues are precisely zero except  $\lambda_1$ . Then,  $f(\mathbf{x}) = g(\mathbf{w}_1^T \mathbf{x})$ , where  $\mathbf{w}_1$  is the first column of  $\mathbf{W}$ , and  $g$  is a function of one variable. In many applications, the smaller eigenvalues are not precisely zero, but they may be much (e.g., orders of magnitude) smaller so that  $f(\mathbf{x})$  may be reasonably approximated by a function of  $n < m$  linear combinations of  $\mathbf{x}$ .

If a given model  $f$  admits such structure, then certain operations—e.g., response surface modeling or optimization—become much less expensive. In particular, these operations can be performed in the  $n$ -dimensional space of the active variables  $\mathbf{y}$  instead of the full  $m$ -dimensional space. It is therefore extremely valuable to determine if  $f$  admits an active subspace. This analysis assumes that the weight function  $\rho$  on the domain  $\mathcal{X}$  is given. The computed quantities (such as  $\mathbf{W}$  and  $\Lambda$ ) will change if a different  $\rho$  is given.

### 2.1. Identifying an active subspace

To identify the active subspace, we must approximate the matrix  $\mathbf{C}$  from (1). Since  $\mathbf{C}$  is the mean of the outer product of the gradient with itself, we can approximate it with simple Monte Carlo. One could achieve a more accurate approximation of  $\mathbf{C}$  and its eigendecomposition with an integration rule that is more accurate than simple Monte Carlo. However, if  $m$  is greater than two or three, then tensor product constructions of accurate univariate numerical integration rules (e.g., Gaussian quadrature) require too many evaluations of the gradient to be practical—especially, if the gradient is expensive to compute. The simple Monte Carlo proceeds as follows. Draw  $\mathbf{x}_i$  independently according to  $\rho$  with  $i = 1, \dots, M$ . Then,

$$\mathbf{C} \approx \hat{\mathbf{C}} = \frac{1}{M} \sum_{i=1}^M \nabla_{\mathbf{x}} f(\mathbf{x}_i) \nabla_{\mathbf{x}} f(\mathbf{x}_i)^T = \hat{\mathbf{W}} \hat{\Lambda} \hat{\mathbf{W}}^T. \quad (6)$$

We use the eigenvalues of the approximation  $\hat{\mathbf{C}}$  as evidence of an active subspace. In particular, a large gap in the eigenvalues indicates a separation between the corresponding active and inactive subspaces defined by  $\hat{\mathbf{W}}_1$  and  $\hat{\mathbf{W}}_2$ , respectively. The qualification *large* depends on the application. For example, a good low-dimensional approximation of  $f$  over its entire domain may need a much larger

separation than a good approximation of the bounds or the average of  $f$ .

Other eigenvalue-based dimension-reduction schemes (e.g., principal component analysis) use heuristics based on the magnitude of the eigenvalues, such as choosing the dimension of the subspace such that the ratio  $\sum_{i=1}^n \lambda_i / \sum_{i=1}^m \lambda_i$  is larger than some threshold. For active subspaces, such heuristics are not well-justified. We are not interested in capturing the gradient’s variance; we want to identify directions along which the function changes the most. In our recent work, we show that the distance between the  $n$ -dimensional subspace defined by  $\mathbf{W}_1$  and its approximation defined by  $\hat{\mathbf{W}}_1$  is inversely proportional to  $\lambda_{n+1} - \lambda_n$  [9]. Therefore, the quality of the finite-sample subspace approximation depends more on the gap between the eigenvalues than their magnitude.

### 2.2. Bootstrap to estimate variability

If a gap is present in the approximate eigenvalues  $\hat{\Lambda}$ , one may naturally ask if a comparable gap is present in the true eigenvalues  $\Lambda$ . To address this question, we use nonparametric bootstrap. The bootstrap is most appropriate when the simulation is expensive and the number  $M$  of gradient evaluations is constrained by a computational budget. For  $j = 1, \dots, M'$ , let  $\boldsymbol{\pi}_j = [\pi_1^j, \dots, \pi_M^j]$  be an  $M$ -vector of integers drawn uniformly at random between 1 and  $M$ . The  $j$ th bootstrap replicate of  $\hat{\mathbf{C}}$  is computed as

$$\hat{\mathbf{C}}^j = \frac{1}{M} \sum_{i=1}^M \nabla_{\mathbf{x}} f(\mathbf{x}_{\pi_i^j}) \nabla_{\mathbf{x}} f(\mathbf{x}_{\pi_i^j})^T = \hat{\mathbf{W}}^j \hat{\Lambda}^j (\hat{\mathbf{W}}^j)^T. \quad (7)$$

The collection of eigenvalues  $\hat{\Lambda}^j$  yields a bootstrap distribution for the eigenvalues, which can be used to estimate bootstrap intervals. These intervals are not true confidence intervals, since the estimates are biased. However, this bias decreases as  $M$  increases; see [10, Section 7.2] and [11, Chapter 3] for related discussions of the bias in bootstrap estimates for principal components. Nevertheless, we use the 99% bootstrap intervals as evidence of gaps in the estimated eigenvalues.

We can also use the bootstrap replicates to estimate the error in the estimated subspace. Partition the eigenvectors from the bootstrap replicate as in (3). Define

$$e_j = \|\hat{\mathbf{W}}_1 \hat{\mathbf{W}}_1^T - (\hat{\mathbf{W}}_1^j)(\hat{\mathbf{W}}_1^j)^T\|_2 = \|\hat{\mathbf{W}}_1^T \hat{\mathbf{W}}_2^j\|, \quad (8)$$

which is the distance between the subspace defined by  $\hat{\mathbf{W}}_1$  and bootstrap replicate  $\hat{\mathbf{W}}_1^j$  [12]. The mean and 99% bootstrap intervals of the set  $\{e_j\}$  quantify the variability in the estimated active subspace.

Our current research efforts are focused on rigorously justified, computable metrics for determining the relationship between the approximate eigenpairs  $\hat{\Lambda}$ ,  $\hat{W}$  and the true eigenpairs  $\Lambda$ ,  $W$ . Our recent paper applies non-asymptotic random matrix theory to characterize the error in these approximations [9].

### 2.3. Visualization with the active subspace

Scatter plots are a common way to visualize data sets in search of a trend. Unfortunately, visualization tools can only display scalar responses  $f$  as a function of at most two variables. When the response depends on more than two variables, one can plot responses versus each variable or each pair of variables.

Using the vectors defining the active subspace, we can create scatter plots based on the active variables, which often show discernible trends that can be exploited in further studies (e.g., building response surfaces or optimization). This idea is described in detail in the texts on *regression graphics* [1], where such plots are called *sufficient summary plots*. For example, suppose we have noticed a large gap between the first and second eigenvalue. We can confirm the presence of the active subspace by first sampling  $\mathbf{x}_i$  from  $\rho$  as in (6), computing  $f_i = f(\mathbf{x}_i)$  (often computed along with the gradients in (6)), and plotting the pairs  $(\mathbf{w}_1^T \mathbf{x}_i, f_i)$ —where  $\mathbf{w}_1$  is the first column of  $\hat{W}$ . If a tight, univariate trend is visibly present, then this verifies the active subspace. Figure 4 in Section 4 shows these scatter plots for the single-diode model.

These sufficient summary plots can also be used to examine the variability in the vectors defining the active subspace. In particular, we can use the bootstrap replicates  $\hat{W}^j$  from (7). The first eigenvector from each can be used to create pairs for the scatter plot. Plotting all pairs together creates clusters in the scatter plot that can indicate the spread around a trend. We create these visualizations for the predicted performance parameter from the single-diode model in Section 4.

### 2.4. Gradient approximation

Often the gradient  $\nabla_{\mathbf{x}} f$  is not available or is too complicated to compute. To approximate partial derivatives, one builds a model of  $f(\mathbf{x})$  that is easily differentiable, e.g., a polynomial model. A first-order finite difference approximation of the partial derivative with respect to  $x_i$  at  $\mathbf{x}$  computes the slope of the plane that interpolates  $f$  at  $\mathbf{x}$  and  $\mathbf{x} + \varepsilon \mathbf{e}_i$ , where  $\mathbf{e}_i$  is a vector of zeros with a one in the  $i$ th component. When the variables are properly normalized, a reasonable choice for  $\varepsilon$  to approximate the derivative and avoid round-off issues is the square root of machine precision [13, Section 8.1].

If the simulation output is well behaved as a function of the variables  $\mathbf{x}$ , then finite difference approximations of gradients can be used as a substitute for the true gradient. Each finite difference approximation requires  $m + 1$  evaluations of  $f$ , so the number of evaluations to compute  $\hat{C}$  is  $M(m + 1)$ . Therefore, finite differences are most appropriate when the following are satisfied:

- the simulation output behaves well as a function of the parameters, i.e., there is no noise due to limited iterations of a nonlinear solver or large changes of the function on scales smaller than the finite difference step size,
- the simulation completes in a short enough time on available computing resources to permit  $M(m + 1)$  evaluations.

Both of these conditions are satisfied for the performance parameter of interest from the single-diode model as a function of its input parameters. Therefore, we use finite differences in place of the true gradient to search for an active subspace.

If these conditions are not satisfied, then one must pursue other options for estimating the gradients. One idea is to use local linear models within a budget-constrained set of runs. In short, assume one has  $B$  pairs  $(\mathbf{x}_j, f(\mathbf{x}_j))$ . Given a point  $\mathbf{x}$  in the parameter space, select  $M$  points in  $\{\mathbf{x}_j\}$  near  $\mathbf{x}$ , fit a linear model to the selected pairs, and estimate the gradient as the gradient of the linear model.  $M$  must be large enough and the points must be chosen to enable fitting the linear model. We have used this approach in practice, but we are still analyzing its properties. The advantage of this approach over finite differences is that it is robust to noise in  $\{f_j\}$  when the number of points used to fit the linear model is greater than the number of variables ( $M > m + 1$ ). However, if the points  $\{\mathbf{x}_j\}$  are not sufficiently dense in the parameter space, then the gradient estimates can be very poor, which lead to poor estimates of the eigenvalues and eigenvectors. We continue to analyze this approach.

## 3. A SINGLE-DIODE SOLAR CELL MODEL

In this section we introduce the single-diode solar cell model, the input parameters, and the predicted quantity of interest. Over the past decade, a growing demand for clean energy has caused rapid growth in the photovoltaic (PV) industry. The overall health of the industry depends on proper characterization of the risks associated with a PV system. One can characterize such risks using a mathematical model to describe the performance of the PV device, such as a single-diode lumped-parameter

equivalent-circuit model defined in Section 3.1. For many series-wired PV devices, the single-diode model accurately describes the current–voltage ( $I$ – $V$ ) characteristics of the device at given irradiance and temperature conditions. The single-diode model contains several input parameters that must be estimated—e.g., using methods described in [14] or [15]—to properly pose the model. With the inputs fixed, practitioners can estimate the *key performance parameters*, such as the maximum power output  $P_{\max}$  and energy conversion efficiency  $\eta$ .

Estimating the model’s input parameters presents challenges. For example, several single-diode model parameter estimation methods have been proposed in the PV literature, and it is not always clear if any particular estimation yields a unique result [16]. Further, since each key performance parameter is a function of the model parameters, errors in estimating the latter will cause issues in estimating the former.

Given these issues, it is desirable to know how the variability or uncertainty in each of the input parameters affects variability in the outputs. In particular, it would be useful to know whether the model admits an active subspace. If an active subspace is present, then estimation of the performance parameters is greatly simplified. For example, rather than conducting current–voltage measurements in a way that yields accurate estimates of all model parameters, practitioners could instead focus on measurements that accurately estimate parameters in the active subspace.

### 3.1. Calculating the key performance parameters

The single-diode model describes a relationship between the current ( $I$ ) and voltage ( $V$ ) in the single diode. There are multiple variants of the single-diode model, each of which has accompanying auxiliary equations that describe how components depend on irradiance and temperature. In this work, we use the auxiliary equations derived in [17; see [18] for an example of others. These relationships are defined in the following equations.

$$I = I_L - I_S \left( \exp \left( \frac{V + I R_S}{N_S n V_{th}} \right) - 1 \right) - \frac{V + I R_S}{R_P}. \quad (9)$$

$N_S$  is the number of cells connected in series, which we set to 1. The thermal voltage  $V_{th}$  is a given, fixed constant for fixed temperature  $T = 25^\circ\text{C}$ . An auxiliary equation defines the photocurrent  $I_L$  as

$$I_L = I_{SC} + I_S \left( \exp \left( \frac{I_{SC} R_S}{N_S n V_{th}} \right) - 1 \right) + \frac{I_{SC} R_S}{R_P}. \quad (10)$$

The remaining terms are input parameters to the single-diode model; their names, units, and ranges are given in

**Table 1.** The ranges of the input parameters for the single-diode model in (9) and (10). These ranges correspond to typical  $2\text{cm}^2$  crystalline silicon PV cells.

Parameter	Name	Lower bound	Upper bound	Units
$I_{SC}$	Short-circuit current	0.05989	0.23958	amps
$I_S$	Diode reverse saturation current	$2.2e-11$	$2.2e-7$	amps
$n$	Ideality factor	1	2	unit-less
$R_S$	Series resistance	0.16625	0.66500	ohms
$R_P$	Parallel (shunt) resistance	93.75	375.00	ohms

Table 1. The ranges correspond to typical  $2\text{cm}^2$  crystalline silicon PV cells.

Given a device and  $I$ – $V$  data, one can estimate the precise values for the input parameters ( $I_{SC}$ ,  $I_S$ ,  $n$ ,  $R_S$ ,  $R_P$ ); see [14,15,17]. With these values fixed, one can compute the maximum power of the device as

$$P_{\max} = \max_{I,V} I V, \quad (11)$$

where current  $I$  and voltage  $V$  are constrained by (9) and (10). Changing the inputs ( $I_{SC}$ ,  $I_S$ ,  $n$ ,  $R_S$ ,  $R_P$ ) changes the nonlinear constraints in the optimization (11), thus changing  $P_{\max}$ . Therefore, we can write

$$P_{\max} = P_{\max}(I_{SC}, I_S, n, R_S, R_P). \quad (12)$$

We use MATLAB to solve the optimization (11) subject to the constraints (9) and (10). Figure 1 shows ten  $I$ – $V$  curves with the single-diode model parameters chosen uniformly at random from the ranges in Table 1. The black dots show the values of  $I_{\max}$  and  $V_{\max}$  that produce  $P_{\max}$ . A wrapper function computes  $P_{\max}$  as a function of the input parameters. Requests for the MATLAB code may be sent to the third author Campanelli.

## 4. AN ACTIVE SUBSPACE IN THE SINGLE-DIODE MODEL

We apply the techniques from Section 2 to search for an active subspace of the  $P_{\max}$  performance parameter as a function of the parameters defined in Table 1. A different performance parameter (e.g., efficiency) would require an independent and identical analysis. The relationship between the respective active subspaces depends on the relationship between the performance parameters and their inputs. In some cases, the active subspaces might be similar; in other cases, they may differ substantially.

The MATLAB scripts for generating the figures in this section can be found at <https://bitbucket.org/paulcon/>

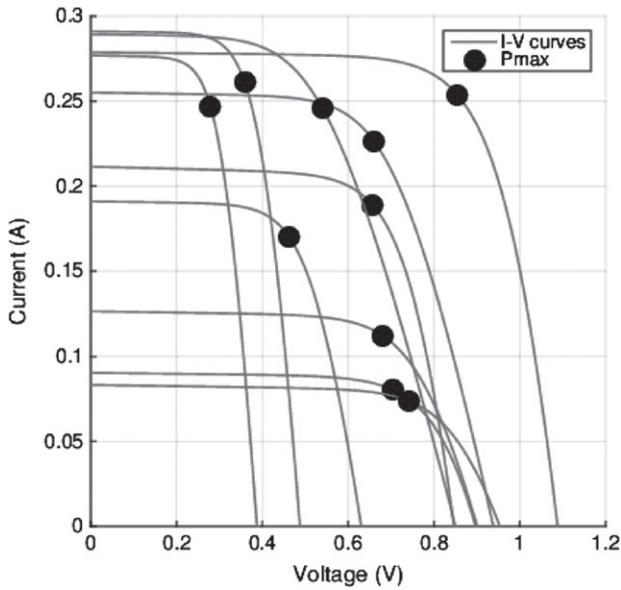


Fig. 1 Ten current–voltage curves with parameters chosen uniformly at random from the ranges in Table 1. The black circles plot the location of  $P_{\max} = I_{\max} V_{\max}$ .

active-subspaces-in-a-single-diode-model. To generate the figures and results, we evaluated the model  $(m + 1)M = 6000$  times, which took approximately 3 minutes on a MacBook Air. A more expensive model would permit fewer evaluations, and we would choose  $M$  differently—or perhaps choose a different method.

#### 4.1. Normalizing the input parameters

We first normalize the domain of  $P_{\max}$  to the hypercube; denote the normalized variables by  $\mathbf{x} = [x_1, x_2, x_3, x_4, x_5]^T$ . The variable  $I_S$  varies over several orders of magnitude, and preliminary tests show that  $P_{\max}$  changes rapidly near smaller values of  $I_S$ . To address this, we work with  $\log(I_S)$ . The range of  $\log(I_S)$  is bounded below by  $-24.54$  and above by  $-15.32$ .

Let  $\mathcal{L}(\cdot)$  and  $\mathcal{U}(\cdot)$  return the upper and lower bounds, respectively, of the argument. We define the normalized input parameters as

$$\begin{aligned} x_1 &= 2 \left( \frac{I_{SC} - \mathcal{L}(I_{SC})}{\mathcal{U}(I_{SC}) - \mathcal{L}(I_{SC})} \right) - 1, \\ x_2 &= 2 \left( \frac{\log(I_S) - \mathcal{L}(\log(I_S))}{\mathcal{U}(\log(I_S)) - \mathcal{L}(\log(I_S))} \right) - 1, \\ x_3 &= 2 \left( \frac{n - \mathcal{L}(n)}{\mathcal{U}(n) - \mathcal{L}(n)} \right) - 1, \\ x_4 &= 2 \left( \frac{R_S - \mathcal{L}(R_S)}{\mathcal{U}(R_S) - \mathcal{L}(R_S)} \right) - 1, \\ x_5 &= 2 \left( \frac{R_P - \mathcal{L}(R_P)}{\mathcal{U}(R_P) - \mathcal{L}(R_P)} \right) - 1. \end{aligned} \quad (13)$$

Then,  $\mathbf{x} \in [-1, 1]^5$  and has no units. This map is invertible. In other words, given a value  $\mathbf{x} \in [-1, 1]^5$ , one can shift and scale the components of  $\mathbf{x}$ —and transform the second component with the exponential—to produce a valid input for the single-diode model. We take the weight function  $\rho$  in (1) to be a constant  $2^{-5}$  inside  $[-1, 1]^5$  and zero elsewhere. One can interpret this weight function as a uniform probability density on the space of input parameters.

Proper scaling is important. The results of the active subspaces analysis are not scale invariant, so the scientist must carefully choose input parameter bounds that are appropriate for the application. Changing the bounds can change the results—sometimes dramatically. For example, the quantity of interest may be irregular and badly behaved over a wide range of inputs. But reducing the range of interest might focus on a small region in the parameter space where  $f$  is smooth—even nearly linear. A linear function has a one-dimensional active subspace. Thus, loosely speaking, if the range of parameters is sufficiently small, then the active subspace will be nearly one-dimensional for continuous and differentiable quantities of interest.

#### 4.2. Estimating the eigenvalues and eigenvectors

We choose  $M = 1000$  points  $\mathbf{x}_i$  uniformly at random from  $[-1, 1]^5$ , and for each  $\mathbf{x}_i$ , we compute both  $P_{\max}$  and a first-order finite difference approximation of the gradient. We use a finite difference step size of  $10^{-6}$  in the normalized domain, which was small enough to ensure that gradient approximation errors were substantially smaller than the variance due to random sampling.

To study the variability in the computed components of  $\hat{\mathbf{C}}$  from (6), we use the bootstrap described in Section 2.2 with  $M' = 1000$  bootstrap replicates. The 99% bootstrap intervals for the eigenvalue estimates are shown in Figure 2a. The tight bootstrap ranges suggest that there is little eigenvalue variability in the bootstrapped data set. There is a gap of nearly an order of magnitude between both the first and second and the second and third eigenvalues. This suggests a very dominant one-dimensional active subspace. Figure 2b shows the estimated error in the approximated subspace using (8). This error is actually the cosine of the principal angle between the subspaces, so it is bounded above by 1.

Figure 3 displays the components of the first and second eigenvectors of  $\hat{\mathbf{C}}$  computed with  $M = 1000$  samples—along with the  $M' = 1000$  bootstrap replicates. They are normalized so that the first component is positive. The small ranges suggest that the one- and two-dimensional active subspaces are stable within the gradient samples.

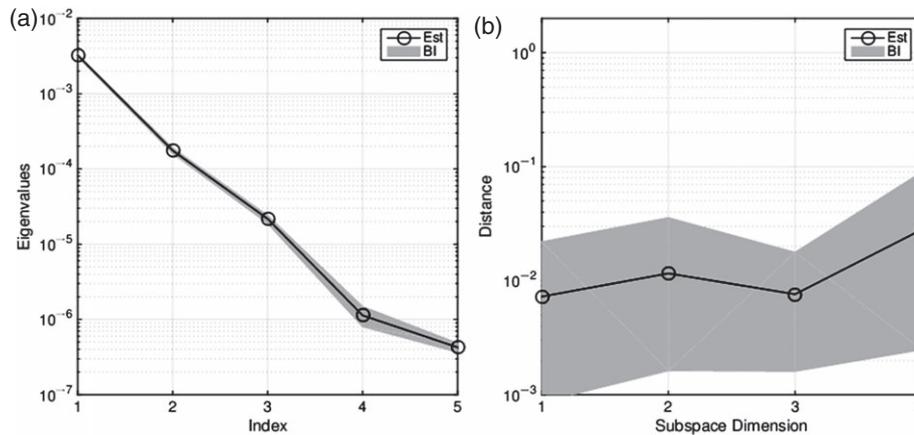


Fig. 2 Figure (a) shows the eigenvalues of the matrix  $\hat{C}$  from (6) with  $M = 1000$  samples and their bootstrap ranges with  $M' = 1000$  bootstrap replicates. Figure shows the mean and ranges of the bootstrap estimates of the subspace error (see (8)) with  $M' = 1000$  bootstrap replicates.

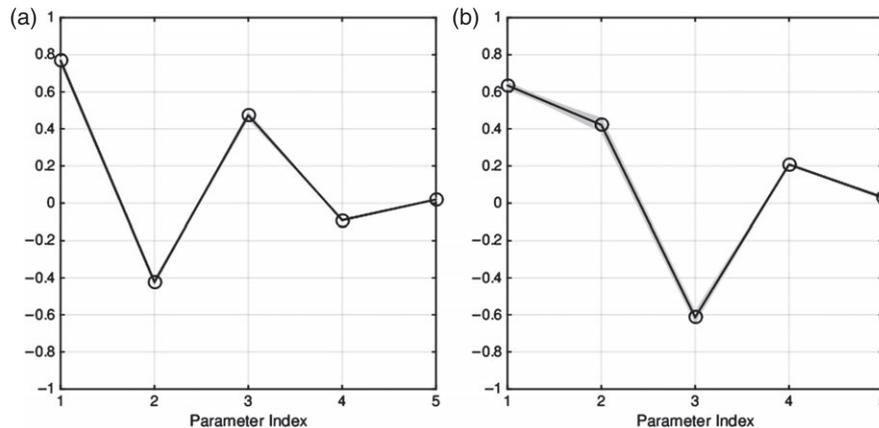


Fig. 3 These figures show the components of the first (a) and second (b) eigenvectors of  $\hat{C}$  with  $M = 1000$  gradient samples. The gray-shaded region superimposes the  $M' = 1000$  bootstrap replicates of the eigenvector components; the tight ranges suggest stability in the gradient samples. These values can be used as measures of sensitivity of  $P_{\max}$  with respect to the input variables in the model, whose indices are given along the horizontal axis of the figures. See Table 1 for a description of the variables.

### 4.3. Identifying the active subspace

The order-of-magnitude gap between the first and second eigenvalues suggests a dominant one-dimensional active subspace. Recall that gaps between eigenvalues are more important than the magnitudes of the eigenvalues. Figure 4a plots 100 values of  $P_{\max}$  against a linear combination of the corresponding normalized inputs; the weights of the linear combination are the components of the first eigenvector  $\hat{w}_1$  from (6). Figure 4(a) is a one-dimensional sufficient summary plot, as described in [1]. The plot shows a potentially exploitable low-dimensional relationship between the input parameters and the quantity of interest,  $P_{\max}$ .

Figure 4(b) shows the two-dimensional sufficient summary plot. The grayscale corresponds to the value of  $P_{\max}$ . The black dot in the center of each cluster has a horizontal

and vertical location determined by linear combinations of  $P_{\max}$ 's corresponding input parameters; the weights of the linear combinations are the eigenvectors  $\hat{w}_1$  and  $\hat{w}_2$ , respectively. Each cluster corresponds to the value of  $P_{\max}$  with horizontal and vertical positions determined by 20 bootstrap replicates of  $\hat{w}_1$  and  $\hat{w}_2$ . The relatively small spread in the clusters indicates that the estimated subspace is stable.

In Figure 4(b),  $P_{\max}$  increases primarily as a function of the first active variable. This provides more evidence of the dominance of the one-dimensional active subspace.

### 4.4. Eigenvectors for sensitivity analysis

The eigenvector components connect the active subspace to the (normalized) variables in the model; Figure 3

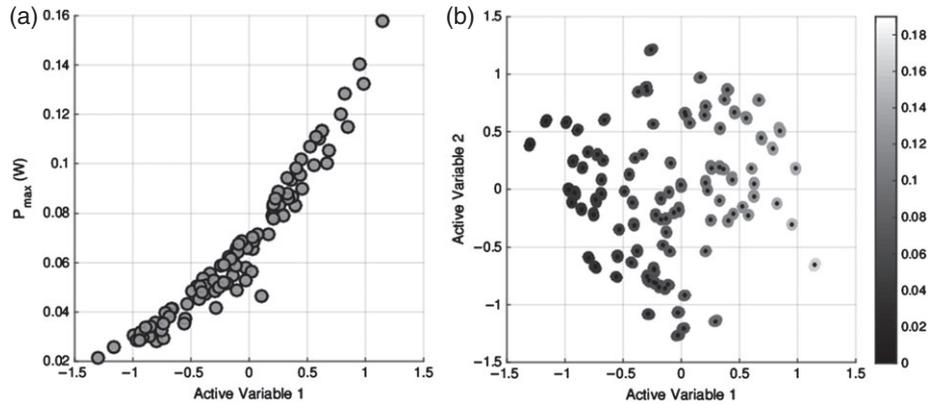


Fig. 4 One-dimensional (a) and two-dimensional (b) sufficient summary plots of 100 realizations of  $P_{\max}$  using  $\hat{\mathbf{w}}_1$  (a) and both  $\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2$  (b). The clusters of points on the right correspond to 20 bootstrap replicates of  $\hat{\mathbf{w}}_1$  and  $\hat{\mathbf{w}}_2$ . The strong trend in 4a is verified by looking at the variation along *Active Variable 1* (the horizontal axis) in 4b, and the combination of these provides evidence of a one-dimensional active subspace.

**Table 2.** The components of the eigenvector  $\hat{\mathbf{w}}_1$ , the Sobol' first-order indices, and the Sobol' total sensitivity indices for the single-diode model parameters. The metrics provide same ranking of importance for the parameters.

Parameter	$\hat{\mathbf{w}}_1$	Sobol' first-order index	Sobol' total sens. index
$I_{SC}$	0.77	0.56	0.61
$I_S$	-0.42	0.17	0.19
$n$	0.47	0.21	0.25
$R_S$	-0.09	0.01	0.01
$R_P$	0.02	0.00	0.00

includes the corresponding indices on the horizontal axis. The magnitudes of the eigenvector components can be used as measures of relative sensitivity for each of the parameters in the model. A large absolute value of the eigenvector component implies that this variable is important in defining the direction along which input perturbations produce the most change, on average, in  $P_{\max}$ . The components corresponding the parallel resistance  $R_P$  and series resistance  $R_S$  are close to zero, which implies that (normalized) changes in  $R_P$  and  $R_S$  do not change  $P_{\max}$  as much as changes in the other parameters.

We compare the components of  $\hat{\mathbf{w}}_1$ , which defines the dominant one-dimensional active subspace, to the Sobol' first-order sensitivity indices and total sensitivity indices. We implemented the method to estimate these indices from Sudret [19] using a tensor product Gaussian quadrature rule with 32768 points in five dimensions. Table 2 displays the eigenvector components and the Sobol' indices.

There are several things to note in Table 2. First, the importance ranking from the Sobol' indices and the ranking from the eigenvector component magnitudes are the same. However, the signs of eigenvector components indicate whether  $P_{\max}$  will increase or decrease, on average,

with changes in the corresponding parameter. (Eigenvectors are unique up to a constant, so these signs should be considered relative to one another.) There is no such interpretation in the Sobol' indices. In other words, the eigenvector components provide more information about the relationship between the corresponding parameters and  $P_{\max}$ .

The numeric values are difficult to compare, since they are normalized differently. The Sobol' indices are divided by the estimated total variance in  $P_{\max}$ , whereas the eigenvector components are normalized to have Euclidean norm equal to 1. In this case, they are the same order of magnitude.

We emphasize the different interpretations of these numbers. The eigenvector  $\hat{\mathbf{w}}_1$  is a single direction in the input space. Input perturbations along this direction change  $P_{\max}$  more, on average, than perturbations orthogonal to this direction. The Sobol' indices indicate the proportion of  $P_{\max}$ 's total variance attributable to the factors of a variance-based decomposition of  $P_{\max}$ . In terms of dimension reduction, one might (i) use the Sobol' indices to conclude that only three of the five parameters were important and (ii) approximate  $P_{\max}$  as a function of those three parameters. In contrast, the active subspace is used to approximate  $P_{\max}$  as with a univariate function of a linear combination of all five parameters. In other words, the variance-based approach reduces the dimension from five to three, and the active subspace approach reduces the dimension from five to one.

There are other challenges facing the active subspace style of dimension that do not affect the variance-based approach. The variance-based approach uses a subset of the model's parameters—as opposed to a subspace. Fixing a subset of parameters at nominal values and allowing the others to vary is relatively straightforward. Exploiting one

important linear combination of the normalized parameters is not as straightforward. In the next section, we discuss some possibilities for exploiting the active subspace.

#### 4.5. Interpreting the sensitivity analyses

We can interpret the consistent importance rankings in the sensitivity metrics in terms of the physics of the single-diode model. First,  $P_{\max} = I_{\max} V_{\max}$ , where  $(I_{\max}, V_{\max})$  is the point on the  $I$ - $V$  curve that maximizes (11).  $I_{\max}$  typically scales linearly with  $I_{SC}$ . The voltage at  $I = 0$ —denoted  $V_{oc}$  for *voltage at open circuit*—typically scales logarithmically with  $I_{SC}$ .  $V_{\max}$  behaves like  $V_{oc}$ . Thus, increases in  $I_{SC}$  affect  $I_{\max}$  linearly and  $V_{\max}$  logarithmically, which changes  $P_{\max}$  as the product of a linear term and a logarithmic term. Since  $I_{SC}$  affects both  $I_{\max}$  and  $V_{\max}$ , it is natural for it to be the driving parameter in  $P_{\max}$ .

Recall that these sensitivity analyses are not scale invariant; the input parameter ranges and the weight function affect the results. The ranges for  $n$  and  $I_S$  are at about their maximum extent for a crystalline silicon PV cell, while the ranges for  $R_S$  and  $R_P$  are somewhat restricted. All of these parameters affect the shape and/or onset of the *knee* of the  $I$ - $V$  curve, which determines the location of  $(V_{\max}, I_{\max})$ , and, thus,  $P_{\max}$ . Here,  $R_S$  and  $R_P$  are relatively unimportant because their ranges are relatively smaller than the ranges for  $n$  and  $I_S$ .

### 5. USING THE ACTIVE VARIABLES

The estimated eigenvalues in Figure 2a and the plots in Figures 4a and 4b provide strong evidence for the presence of a dominant one-dimensional active subspace in the  $P_{\max}$  performance parameter computed from the single-diode model as a function of the model's five input parameters. The natural question is how one can exploit the dimension reduction afforded by the active subspace. There are five types of studies that benefit greatly from fewer input parameters. To keep the scope of this paper limited, we do not address the details of any of these studies and we prefer to reserve them for future exploration.

**Visualization.** The active subspace enables one to view the model output's dependence on its inputs with standard computer graphics tools when the active subspace is not more than two-dimensional. Sufficient summary plots, such as those in Figures 4a and 4b, can provide insights to modelers seeking to improve their models.

**Optimization.** Suppose one wanted to maximize  $P_{\max}$  over the input variables. The plot in Figure 4 shows a monotonic trend in the univariate function of the first active variable.

For such functions, maximization is trivial; simply make the first active variable  $\mathbf{w}_1^T \mathbf{x}$  as large as possible subject to the constraints  $\mathbf{x} \in [-1, 1]^m$ . In general, a five-dimensional global optimization where the input/output relationship is not well understood is a difficult problem. For the  $P_{\max}$  output from the single-diode model, the active subspace provides a way to discover the location of the global optimum with ease.

**Response surfaces.** If the single-diode model were expensive to evaluate, then one may wish to construct a response surface that approximates the map from inputs to outputs. A one-dimensional active subspace allows one to build a response surface on only the active variable instead of the five model input parameters. Constructing a response surface in one variable is certainly preferred to constructing one in five variables. Real-time control systems such as maximum power point trackers might benefit from the characterization of a simple, low-dimensional response surface. In addition, the visualization tools give one confidence that the response is sufficiently smooth with respect to the active variable to permit an accurate response surface.

**Averages.** One may wish to compute an average of  $P_{\max}$  over all five input variables. Since the average of  $P_{\max}$  over its inputs is the same as the average of the conditional expectation given the active variable, we need only to approximate a marginal density of the active variable to compute the average. We are working on the specifics of this idea, but the goal would be to transform a five-dimensional integral into a one-dimensional integral requiring many fewer evaluations of  $P_{\max}$ .

**Design.** Suppose a modeler wanted to design a solar cell with  $P_{\max}$  in a specified range and sought an appropriate range of values for the model inputs. A one-dimensional active subspace makes this query much easier—especially if  $P_{\max}$  is monotonic with respect to the active variable as in Figure 4b.

### 6. CONCLUSIONS

We have discussed methods for discovering the directions in a model's input space that change the model's prediction the most, on average. The active subspace is the span of these directions. For global sensitivity analysis, the active subspace offers an alternative to variance decomposition techniques such as Sobol' indices. We have applied this procedure to a single-diode solar cell model with key performance parameter  $P_{\max}$ , and we discovered a dominant one-dimensional active subspace. We offered several possible ways to exploit the knowledge of this low-dimensional parametric dependence to gain greater insight into the model.

## ACKNOWLEDGMENTS

The first author's work is supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC-0011077.

## REFERENCES

- [1] R. D. Cook, *Regression Graphics: Ideas for Studying Regressions through Graphics*, New York, John Wiley & Sons, 1998.
- [2] T. M. Russi, *Uncertainty Quantification with Experimental Data and Complex System Models*. Ph.D. Thesis, University of California, Berkeley, 2010.
- [3] P. G. Constantine, E. Dow, and Q. Wang, Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM J Sci Comput* 36(2014), A1500-A1524.
- [4] P. G. Constantine, Q. Wang, A. Doostan, and G. Iaccarino, *A Surrogate Accelerated Bayesian Inverse Analysis of the HyShot II Flight Data*, American Institute of Aeronautics and Astronautics, Reston, VA, 2011.
- [5] T. W. Lukaczyk, P. G. Constantine, F. Palacios, and J. J. Alonso, *Active Subspaces for Shape Optimization*. American Institute of Aeronautics and Astronautics, Reston, VA, 2014.
- [6] Q. Wang, H. Chen, R. Hu, and P. G. Constantine, *Conditional Sampling and Experiment Design for Quantifying Manufacturing Error of Transonic Airfoil*, American Institute of Aeronautics and Astronautics, Reston, VA, 2011.
- [7] H. S. Abdel-Khalik, Y. Bang, and C. Wang, Overview of hybrid subspace methods for uncertainty quantification, sensitivity analysis. *Ann Nucl Energy* 52 (2013), 28–46.
- [8] A. M. Samarov, Exploring regression structure using nonparametric functional estimation, *J Am Stat Assoc* 88(423) (1993), 836–847.
- [9] P. Constantine and D. Gleich, *Computing Active Subspaces*, arXiv preprint arXiv:1408.0545, 2014.
- [10] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*, Boca Raton, FL, Chapman & Hall/CRC, 1994.
- [11] I. T. Jolliffe, *Principal Component Analysis* (2nd ed.), New York, Springer-Verlag, 2002.
- [12] G. H. Golub and C. F. Van Loan, *Matrix Computations* (3rd ed.), Baltimore, MD, The Johns Hopkins University Press, 1996.
- [13] J. Nocedal and S. J. Wright, *Numerical Optimization* (2nd ed.), New York, Springer, 2006.
- [14] C. Hansen, A. Luketa-Hanlin, and J. Stein, Estimation of parameters for single diode models using measured IV curves. In *39th IEEE Photovoltaics Specialists Conference*. Tampa, FL, 2013.
- [15] G. Farivar and B. Asaei, Photovoltaic module single diode model parameters extraction based on manufacturer data sheet parameters, In *2010 IEEE International Conference on Power Energy (PECon)*, Kuala Lumpur, 2010, 929–934.
- [16] C. Hansen, A. Luketa-Hanlin, and J. Stein, Sensitivity of single diode models for photovoltaic modules to method used for parameter estimation, In *28th European Photovoltaic Solar Energy Conference and Exhibition*, 2013b, 3258–3264.
- [17] B. Zaharatos, M. Campanelli, C. Hansen, K. Emery, and L. Tenorio, Likelihood methods for single-diode model parameter estimation from noisy I-V curve data, In *40th IEEE Photovoltaics Specialists Conference*, Denver, CO, 2014, 2850–2855.
- [18] W. De Soto, S. A. Klein, and W. A. Beckman, Improvement and validation of a model for photovoltaic array performance, *Solar Energy* 80 (2006), 78–88.
- [19] B. Sudret, Global sensitivity analysis using polynomial chaos expansions, *Reliab Eng Syst Safe* 93(7) (2008), 964–979.