

# Sensitivity Analysis for ReaxFF Reparametrization Using the Hilbert–Schmidt Independence Criterion

Michael Freitas Gustavo, Matti Hellström, and Toon Verstraelen\*



Cite This: <https://doi.org/10.1021/acs.jctc.2c01320>



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

**ABSTRACT:** We apply a global sensitivity method, the Hilbert–Schmidt independence criterion (HSIC), to the reparametrization of a Zn/S/H ReaxFF force field to identify the most appropriate parameters for reparametrization. Parameter selection remains a challenge in this context, as high-dimensional optimizations are prone to overfitting and take a long time but selecting too few parameters leads to poor-quality force fields. We show that the HSIC correctly and quickly identifies the most sensitive parameters and that optimizations done using a small number of sensitive parameters outperform those done using a higher-dimensional reasonable-user parameter selection. Optimizations using only sensitive parameters (1) converge faster, (2) have loss values comparable to those found with the naive selection, (3) have similar accuracy in validation tests, and (4) do not suffer from problems of overfitting. We demonstrate that an HSIC global sensitivity is a cheap optimization preprocessing step that has both qualitative and quantitative benefits which can substantially simplify and speed up ReaxFF reparametrizations.



## 1. INTRODUCTION

**1.1. ReaxFF Reparametrization.** The ReaxFF (reactive force field)<sup>1,2</sup> is a potential energy surface (PES) for modeling reactive chemistry at spatial and temporal scales typically unreachable with other more accurate but expensive methods. ReaxFF's speed comes at the cost of replacing accurate formalism with empirical equations that can typically contain hundreds of fitted parameters. Many of these parameters have no easy physical interpretability, and only expert users may have a notion of appropriate values.

This makes the task of fixing the parameters quite difficult. The procedure typically used to do this is the minimization of some cost function which measures the deviations between predictions made by ReaxFF and some training set of values the user would like to replicate.<sup>3–12</sup> However, the question of which parameters should be optimized has always been a difficult one. Optimizing all potentially relevant parameters simultaneously is prone to producing overfitted results, and very high dimensional optimizations are costly and insufficiently exploratory.<sup>4</sup> A key conclusion from our previous work<sup>4,13</sup> is that the careful conditioning of the error function (in terms of parameter selection and training set items) is crucial to obtaining reasonable results.

Typically, researchers apply rudimentary sensitivity analyses and expert knowledge to guide parameter selection. However, these processes are usually inaccurate or laborious or do not sample the space sufficiently. In this work, we aim to address this problem by introducing a systematic global sensitivity analysis

technique to guide parameter selection in a more robust manner. This method has also been coded into our ParAMS package,<sup>14</sup> making it a seamless part of any reparametrization workflow.

**1.2. Global Sensitivity Methods.** Uncertainty analysis concerns itself with the propagation of uncertainty from a model's inputs to its outputs, i.e., estimating an outputs' distribution given all possible changes in the inputs.<sup>15</sup> Sensitivity analysis attributes the uncertainty of the outputs to particular inputs. This is broadly done by measuring how much the unconditioned output distribution,  $\mathbb{P}(Y)$ , differs from the output distribution given certain inputs,  $\mathbb{P}(Y|X_i)$ .<sup>16,17</sup> The type of statistical operator used to make this comparison divides global sensitivity methods into two types; variance-based and distribution-based.

Variance-based methods assume that the variances of these distributions are sufficient to describe them.<sup>16,18</sup> This approach is well-established and the first we tried. One of the most popular variance-based techniques is known as the Sobol method. It is based on the assumption that the total variance on the outcomes

Received: December 31, 2022

of a model  $Y$  can be linearly apportioned to every input and combination of inputs:<sup>19</sup>

$$\sum_i S_i + \sum_i \sum_{j>i} S_{ij} + \sum_i \sum_{j>i} \sum_{k>j} S_{ijk} + \dots + S_{123\dots d} = 1 \quad (1)$$

$S_i$  is the Sobol index for the first-order effect of parameter  $i$ , calculated as<sup>19</sup>

$$S_i = \frac{\mathbb{V}(\mathbb{E}(Y|X_i))}{\mathbb{V}(Y)} \quad (2)$$

where  $\mathbb{E}$  and  $\mathbb{V}$  indicate the expected value and variance, respectively. Higher-order effects are calculated similarly:

$$S_{ij} = \frac{\mathbb{V}(\mathbb{E}(Y|X_i, X_j))}{\mathbb{V}(Y)} \quad (3)$$

A special case is known as the total effect, which is calculated as

$$S_{Ti} = 1 - \frac{\mathbb{V}(\mathbb{E}(Y|X_{-i}))}{\mathbb{V}(Y)} \quad (4)$$

where  $X_{-i}$  denotes the vector of all parameters except  $X_i$ . A full decomposition of the variance is prohibitive for all but the smallest number of dimensions since a total of  $2^d - 1$  indices need to be calculated. It has, however, been shown that the first effect and total effect are sufficient for identifying the most sensitive parameters.<sup>19</sup>

The main disadvantage of this method is its cost in terms of the number of samples needed. An efficient method to calculate the first-order and total effect Sobol indices requires  $\eta(d + 2)$  function evaluations.<sup>19</sup> There is no prescribed way to select  $\eta$ , but it is typically on the order of several thousand.<sup>20,21</sup> Since the cost of the calculation of the indices themselves is trivial in comparison to the cost of obtaining the samples, the recommended procedure in literature is to iteratively sample and calculate until the indices converge. In this work, we attempt to introduce a preprocessing step to a ReaxFF parametrization, and thus, the cost of this should be significantly less than the optimization; literature sources suggest that this will not be the case.<sup>20,21</sup>

The intractable expense of the Sobol method is well-known, and much effort has gone into decreasing its cost.<sup>19</sup> One popular method which is significantly cheaper than the Sobol one is the elementary effects method, also known as Morris screening.<sup>22</sup> Morris screening uses a careful sampling strategy to calculate an “estimated effect” which has been shown to be a good proxy for the total effect Sobol index.<sup>19,22</sup> This method is a popular screening method for dimensionality reduction in literature, particularly for expensive models. In theory, Morris screening should be a suitable method for our purposes; however, in our early investigations, this proved to not be the case.

Both Sobol indices and Morris screening impose rules on how samples may be generated. In the case of Sobol indices, certain parameter values must be fixed while allowing others to vary. In the case of Morris screening, samples are gathered in trajectories where individual parameters are changed one at a time. In our particular context, these rules make obtaining an adequate number of valid samples very frustrating and time-consuming. This is because the ReaxFF loss function surface is full of undefined points corresponding to parameter values which result in crashed calculations and nonphysical results. If such points are encountered when generating a Morris trajectory, for

example, then the entire trajectory must be abandoned because these methods do not allow the entry of nonfinite or undefined values. It would be simpler to find a method which allowed these values or did not have sampling rules so that such instances could simply be discarded.

Aside from the real, practical issues related to their use, the assumptions made by variance-based techniques are inappropriate in the context of identifying parameter sensitivities in the ReaxFF error function. This is because it can produce extremely large values and minima which are often located in narrow valleys and are unlikely to be captured during a sampling procedure.<sup>4</sup> Cheaper approximation methods are typically insufficiently space-filling or rely on strong assumptions about the structure of the inputs or outputs.<sup>15,23</sup> The main criticism of variance-based techniques, however, is the fact that they limit the amount of information extracted from distributions to a single metric.

Distribution-based sensitivity approaches represent a newer family of methods. Unlike variance-based methods, these techniques do not reduce distributions to only their variance and can, in principle, capture arbitrary dependencies by considering the entirety of the distributions. Different statistics are used to measure the differences between distributions,<sup>16</sup> and various methods of this type have been developed.<sup>18,23,24</sup> We are interested in the Hilbert–Schmidt independence criterion (HSIC)<sup>25</sup> for its simplicity, speed, and accuracy.

The HSIC was first introduced in Gretton et al.<sup>25</sup> Through the use of kernels, it can identify arbitrary nonlinear dependencies between the inputs and outputs.<sup>16</sup> It is simpler and faster to converge than many other distribution-based techniques and does not require explicit regularization.<sup>25</sup> It has since been used in feature selection,<sup>23,24,28–29</sup> event detection,<sup>30,31</sup> machine learning,<sup>32,33</sup> and dimensionality reduction in the context of global optimization,<sup>17</sup> among other applications. This final use-case is of particular interest in our context and, other than Spagnol et al.,<sup>17</sup> we have found no other work that directly applied the HSIC as a preoptimization step.

In this work, we apply an HSIC analysis to a ReaxFF reparametrization problem. Based on the sensitivity results, we reduce the original high-dimensional optimization to a lower-dimensional one. We then demonstrate that the optimization in this reduced space is able to produce ReaxFF force fields which are competitive with those produced by the higher-dimensional optimization but converge more quickly and reduce the risk of overfitting.

We focus on ReaxFF because parametrizations of this force field occur regularly in literature.<sup>10,34–37</sup> However, the techniques we introduce here could conceivably be applied to any high-dimensional parametrization problem of the type that occurs frequently in many fields<sup>38–46</sup> because they all share the same underlying sloppy mathematical behavior.<sup>47</sup>

In the next section we will introduce ReaxFF, the HSIC metric, the kernels used, and the reparametrization loss function. In Section 4 we introduce a test case on which we apply the sensitivity technique, the results of which are described in Section 5, and future research paths are discussed in Section 6. Section 3 details our implementation of the calculation, and conclusions are drawn in Section 7.

## 2. MATHEMATICAL DERIVATIONS

**2.1. ReaxFF Energy Potential.** The elucidation of the full ReaxFF energy potential is beyond the scope of this work, but we

provide a brief introduction here for context. Interested readers can consult Chenoweth et al.<sup>5</sup> and Senftle et al.<sup>2</sup> for more details.

The ReaxFF energy potential is a summation of different energy contributions:

$$E_{\text{sys}} = E_{\text{bond}} + E_{\text{lp}} + E_{\text{over}} + E_{\text{under}} + E_{\text{val}} + E_{\text{pen}} + E_{\text{coa}} \\ + E_{\text{C}_2} + E_{\text{triple}} + E_{\text{tors}} + E_{\text{conj}} + E_{\text{H-bond}} + E_{\text{vdW}} \\ + E_{\text{charges}} + \dots \quad (5)$$

The main contribution is the bonding energy,  $E_{\text{bond}}$ , which calculates an energy based on bond orders, which are, in turn, calculated from atomic positions (the most fundamental inputs to any PES). The other contributions can be thought of as corrections to the bonding energy. Some corrections are general, for example  $E_{\text{over}}$  and  $E_{\text{under}}$ , which penalize atoms which are over- or undercoordinated by their bonding. Some corrections are extremely specific, for example  $E_{\text{C}_2}$ , which corrects the energy for the  $\text{C}_2$  molecule. There is also an energy term for fluctuating charges ( $E_{\text{charges}}$ ). ReaxFF typically uses either the EEM<sup>48</sup> or ACKS2<sup>49</sup> charge models. In this work we use EEM charges.

Each of the energy contributions is typically a complicated function of bond orders and empirical parameters. For example:<sup>5</sup>

$$E_{\text{bond}} = -D_e^\sigma \cdot \text{BO}_{ij}^\sigma \exp[p_{\text{be}1} (1 - (\text{BO}_{ij}^\sigma)^{p_{\text{be}2}})] - D_e^\pi \cdot \text{BO}_{ij}^\pi \\ - D_e^{\pi\pi} \cdot \text{BO}_{ij}^{\pi\pi} \quad (6)$$

where  $D_e^\sigma$ ,  $D_e^\pi$ ,  $D_e^{\pi\pi}$ ,  $p_{\text{be}1}$ , and  $p_{\text{be}2}$  are parameters and  $\text{BO}_{ij}$  are bond orders.

The complexity and number of contributions quickly leads the ReaxFF PES to have a very large number of parameters, particularly because most parameters are a function of one or more atom types. For example, there is a  $D_e^\sigma$  parameter for every combination of atom types being modeled.

In an effort to make the parameters easier to understand, the original authors organized them into the following groups or blocks: general, atoms, bonds, off-diagonal, angles, torsions, and hydrogen bonds. These groups have no bearing on the calculation of the PES but serve as a helpful organizational tool and appear in the formatting rules for the inputs of most ReaxFF implementations. We will return to these groups when we introduce a new Zn/S/H force field in Section 4.1.

We use the term “force field” in this work to refer to a full set of ReaxFF parameters and their values. Force fields are functions of the atom types they include and are sometimes only appropriate for specific conditions. For example, ReaxFF force fields containing hydrogen and oxygen are generally classed as appropriate for aqueous or combustion chemistry.

**2.2. Loss Function.** The reparametrization of a force field requires the construction of a training set ( $\mathbf{b}_{\text{ref}} \in \mathbb{R}^m$ ), which is a vector of  $m$  scalar training point values of chemical properties that one would like ReaxFF to replicate. The goal of the reparametrization is the minimization of a loss function which quantifies the overall difference between the ReaxFF model and the training data in a single number. This loss function can take various forms, but the most common is the sum of square errors (SSE). We define the SSE as

$$\text{SSE}(\mathbf{x}) := \sum_{i=1}^m w_i \left( \frac{b_{\text{ref},i} - b(\mathbf{x}, \mathbf{G}_i)}{\sigma_i} \right)^2 \quad (7)$$

Each item in the training set ( $b_{\text{ref},i}$ ) is compared to the result produced by the ReaxFF model ( $b(\mathbf{x}, \mathbf{G}_i)$ ) using a vector of parameters which we allow to be adjusted ( $\mathbf{x}$ ) and a set of system inputs and fixed parameters ( $\mathbf{G}_i$ ). We will call the parameters we have chosen to optimize *active* parameters and denote the set of active parameters as  $\mathcal{D}$ . Items in the training set can be of different types (energies, forces, angles, etc.). To make them comparable and independent of units, each term in the loss function is divided by an appropriate  $\sigma_i$  with the same unit as  $b_{\text{ref},i}$ . The  $\sigma_i$  value is often understood as the acceptable error for that item. Finally, different training values can be weighted differently based on their importance via the weighting vector ( $\mathbf{w}$ ). For example, it is more important to get energy minima correct than it is to replicate energies for very exaggerated geometries; one might weight the former higher than the latter in the loss function. In the case of multidimensional data, like forces, the  $n \times 3$  matrix is simply flattened into the vector of training data. In other words, a single element of  $\mathbf{b}_{\text{ref}}$  would not be “the forces on molecule A” for example, but rather “the  $x$  component of the forces on atom 1 in molecule A”. For more details regarding the implementation itself, readers are directed to the ParAMS documentation.<sup>14</sup>

The vectors  $\sigma$  and  $\mathbf{w}$  perform essentially similar functions since they are both constants which could conceivably be combined. Using one or the other or both is a question of preference for different practitioners. Traditionally, only  $\sigma$  was used, but for some users weights provide a more intuitive value. We have chosen to use a fixed vector of values for  $\sigma$  and adjust the weights to balance the training set.

**2.3. Hilbert–Schmidt Independence Criterion.** This section introduces the HSIC from a practitioner’s perspective. Our presentation is not as general as in other works,<sup>17,25,28</sup> and we only strive to explain the basic idea to readers interested in parameter selection. We will assume that the space of a single parameter is  $\mathcal{X} \subseteq \mathbb{R}$  and the space of loss values is  $\mathcal{Y} \subseteq \mathbb{R}$ . Typically, parameters are bounded and loss values are positive, but these are not strict requirements for HSIC.

Let the Hilbert space of  $\mathbb{R}$ -valued functions on  $\mathcal{X}$  be  $\mathcal{H}_x: \mathcal{X} \rightarrow \mathbb{R}$ , for which an inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}_x}$  is defined. This is called a *reproducing kernel* Hilbert space (RKHS) when a kernel  $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  exists such that  $\forall x \in \mathcal{X}$  and  $\forall f \in \mathcal{H}_x$ :

$$k(x, \cdot) \in \mathcal{H}_x \\ \langle f, k(x, \cdot) \rangle_{\mathcal{H}_x} = f(x)$$

The second condition is known as the *reproducing* property: a function evaluation  $f(x)$  can be reproduced by taking an inner product of  $f$  with a partially evaluated kernel. We similarly define the RKHS  $\mathcal{H}_y$  with kernel  $l$  for for the loss-value space  $\mathcal{Y}$ .

The kernel  $k$  can be defined in terms of a feature map  $\phi: \mathcal{X} \rightarrow \mathcal{F}_x$ , where  $\mathcal{F}_x$  is a new Hilbert space, often called the feature space of  $x$ . This means that each value  $x$  is uniquely represented by a function  $\phi(x) \in \mathcal{F}_x$ . Any positive-definite and symmetric kernel  $k$  can always be described as an inner product in feature space:

$$k(x, x') := \langle \phi(x), \phi(x') \rangle_{\mathcal{F}_x} \quad (8)$$

In practice, direct use of  $\phi$  is avoided, and the kernel  $k$  is used instead to keep calculations manageable. Kernels are generally cheap to evaluate, even when the feature map would be

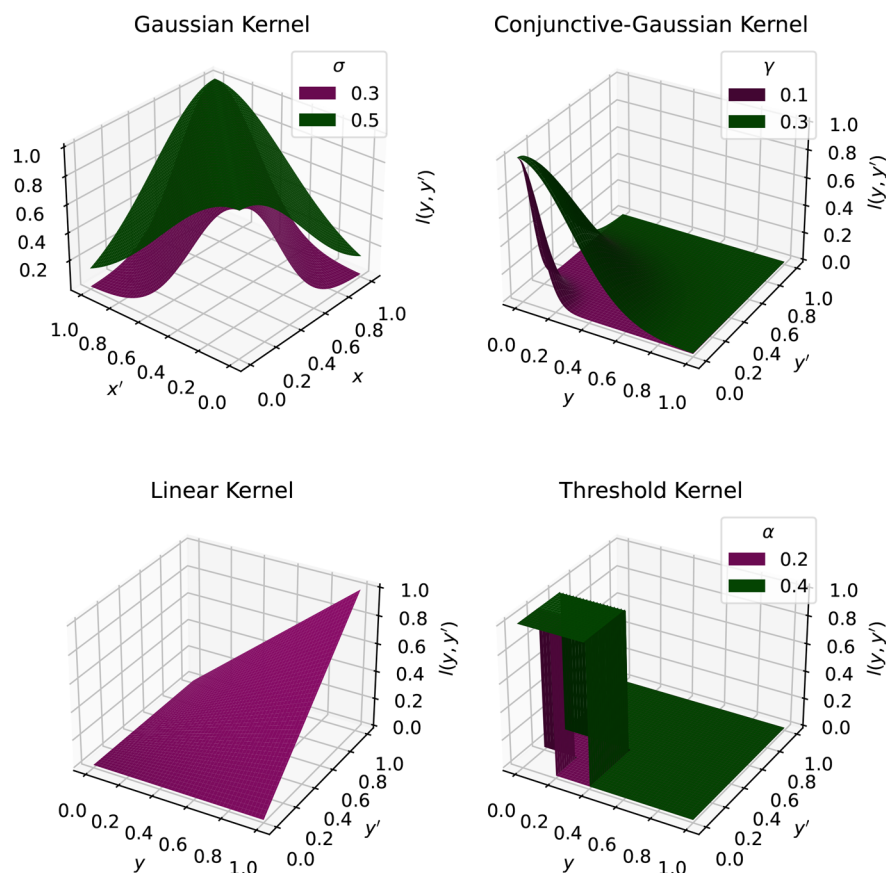


Figure 1. Plots of the tested kernel functions.

computationally infeasible. Similarly, a feature map  $\psi: \mathcal{Y} \rightarrow \mathcal{F}_y$  exists, such that  $l(y, y') := \langle \psi(y), \psi(y') \rangle_{\mathcal{F}_y}$ .

Now consider random variables  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y}$  with a joint distribution  $\mathbb{P}_{XY}$ . Variance-based sensitivity methods characterize the dependence of  $Y$  on  $X$  with  $\text{cov}[X, Y] = \mathbb{E}_{XY}[XY] - \mathbb{E}_X[X]\mathbb{E}_Y[Y]$ , which only picks up linear correlations. In principle, one may overcome this limitation by computing the covariance in feature space instead:  $\text{cov}[\phi(X), \psi(Y)] = \mathbb{E}_{XY}[\phi(X)\psi(Y)] - \mathbb{E}_X[\phi(X)]\mathbb{E}_Y[\psi(Y)]$ . However, the latter covariance is impractical because it directly operates in feature space and because the covariance is an element of a new Hilbert space  $\mathcal{F}_x \otimes \mathcal{F}_y$ , which troubles its direct interpretation. Gretton et al.<sup>25</sup> showed how to overcome both issues by rewriting the squared norm of the latter covariance,

$$\text{HSIC}(\mathcal{F}_x, \mathcal{F}_y, \mathbb{P}_{xy}) := \|\text{cov}[\phi(X), \psi(Y)]\|_{\mathcal{F}_x \otimes \mathcal{F}_y}^2 \quad (9)$$

in terms of kernels:

$$\begin{aligned} \text{HSIC}(\mathcal{F}_x, \mathcal{F}_y, \mathbb{P}_{xy}) &= \mathbb{E}_{XX'YY'}[k(X, X')l(Y, Y')] \\ &+ \mathbb{E}_{XX'}[k(X, X')]\mathbb{E}_{YY'}[l(Y, Y')] \\ &- 2\mathbb{E}_{XY}[k(X, X')]\mathbb{E}_Y[l(Y, Y')] \end{aligned} \quad (10)$$

This definition of HSIC yields a positive real number which increases as  $X$  and  $Y$  are more correlated. When  $k$  and  $l$  are linear kernels, HSIC reduces to  $|\text{cov}[X, Y]|^2$ , and there is no added value compared to variance-based sensitivity criteria. In case of nonlinear kernels, the HSIC also detects nonlinear correlations

between  $X$  and  $Y$ . Ideally,  $k$  and  $l$  are *characteristic kernels*,<sup>17</sup> for which  $\text{HSIC} = 0$  if and only if  $X$  and  $Y$  are independent ( $\mathbb{P}_{XY} = \mathbb{P}_X\mathbb{P}_Y$ ). Of the common kernels, the Gaussian kernel is characteristic, while the linear and polynomial kernels are not. It is possible to use noncharacteristic kernels for HSIC, but then one loses the guarantee that all dependence will be captured.<sup>28</sup>

It is usually not possible to obtain  $\mathbb{P}_{xy}$  explicitly. However, the HSIC can be estimated if we sample the distribution. We define a set of samples as

$$Z := (\mathbf{X}, \mathbf{y}) \quad (11)$$

where  $\mathbf{X} \in \mathbb{R}^{n \times d}$  is a matrix of  $n$  rows of uniformly randomly sampled parameter vectors  $\{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n\}$  and we let  $\mathbf{y} \in \mathbb{R}^n = \{\text{SSE}(\mathbf{X}_1), \text{SSE}(\mathbf{X}_2), \dots, \text{SSE}(\mathbf{X}_n)\}$  be the corresponding set of loss function values.

To measure the sensitivity of a parameter with index  $\omega$ , we take the associated column with  $n$  sample values,  $\mathbf{X}_{\omega}$ , and the loss values,

$$Z_{\omega} := (\mathbf{X}_{\omega}, \mathbf{y}) \quad (12)$$

and we make use of the *unbiased* estimator of the HSIC from Song et al.:<sup>28</sup>

$$\text{HSIC}(\mathcal{F}_x, \mathcal{F}_y, Z_{\omega}) \approx \frac{1}{n(n-3)}[\rho_1 + \rho_2 - \rho_3] \quad (13)$$

where

$$\rho_1 := \text{tr}(\tilde{\mathbf{K}}\tilde{\mathbf{L}})$$

$$\rho_2 := \frac{\mathbf{1}^T \tilde{\mathbf{K}} \mathbf{1} \tilde{\mathbf{L}}^T \mathbf{1}}{(n-1)(n-2)}$$

$$\rho_3 := \frac{2}{n-2} \mathbf{1}^T \tilde{\mathbf{K}} \tilde{\mathbf{L}} \mathbf{1}$$

in which

$$\tilde{K}_{ij} := (1 - \delta_{ij})k(x_i, x_j) \quad \forall x_i, x_j \in \mathbf{X}_\omega$$

$$\tilde{L}_{ij} := (1 - \delta_{ij})l(y_i, y_j) \quad \forall y_i, y_j \in \mathbf{y}$$

HSIC is a convenient metric with which to rank the influence of parameters on an output. However, to make interpretation even easier, we define the sensitivity of a parameter as the normalized HSIC value:<sup>17</sup>

$$s_\omega = \frac{\text{HSIC}(\mathcal{F}_x, \mathcal{F}_y, Z_\omega)}{\sum_{i=1}^d \text{HSIC}(\mathcal{F}_x, \mathcal{F}_y, Z_i)} \quad (14)$$

By construction, sensitivity values are bounded ( $0 \leq s_\omega \leq 1$ ) and  $\sum s_\omega = 1$ . This makes sensitivities a more intuitive value to interpret and understand than HSIC.

The advantage of the *biased* HSIC estimator<sup>28</sup> is that it is positive by construction. The *unbiased* estimator used here, although more accurate, runs the risk of producing small negative values for low-sensitivity items. In the context of the sensitivity analysis, where a qualitative consideration of the results is arguably more important than a quantitative one, we simply clip any negative HSIC values to zero before calculating the sensitivities (eq 14).

For a more detailed description of HSIC, the interested reader is directed to Gretton et al.<sup>25</sup> for the original introduction and to De Lozzo and Marrel,<sup>23</sup> Da Veiga,<sup>24</sup> Song et al.,<sup>28</sup> and Spagnol et al.<sup>17</sup> for more recent works.

Note that it is also possible to consider a more fine-grained sensitivity approach and let  $\mathbf{y}$  be individual residual values (i.e.,  $y_i = b_{\text{ref},i} - b(\mathbf{x}, \mathbf{G}_i)$ ). This would allow one to determine sensitivities for particular training set items, providing further information to assist in training set design. This represents an extension of the current work, and it will not be dealt with here but is left as an avenue for further study.

**2.4. Kernels.** Figure 1 illustrates the kernel distributions used in this work, using different parameter values where appropriate.

We select the characteristic Gaussian kernel for the parameters kernel:

$$k(x, x') := \exp\left[-\frac{(x - x')^2}{2\sigma^2}\right] \quad (15)$$

This captures the (dis)similarity between parameters and allows the distribution to take an arbitrary shape. For the sake of clarity,  $\sigma$  as it is used in eq 15 is unrelated to the same symbol in eq 7.

We test several candidates for the loss kernel. We first consider the noncharacteristic threshold kernel introduced by Spagnol et al.<sup>17</sup>

$$l(y, y') := h(y)h(y') \quad (16)$$

where

$$h(y) := \begin{cases} 1 & y < q_\alpha \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

in which  $q_\alpha$  is the value corresponding to the  $\alpha$  quantile of the data. The above kernel is constructed to focus the weight of the distribution on the smallest numerical values of  $\mathbf{y}$  since these represent good loss function values that we would like the most information about. Spagnol et al.<sup>17</sup> showed that thresholding was extremely important in identifying sensitivities near minima and not allowing the calculation to be swamped by very large function values.

We then consider a characteristic and continuous approximation of the threshold kernel, which we call the conjunctive-Gaussian (CG) kernel:

$$l(y, y') := \exp\left[-\frac{(y^2 + y'^2)}{2\gamma^2}\right] \quad (18)$$

where lower values of the  $\gamma$  parameter focus the distribution more heavily on the best values.

For comparison, we also test a simple noncharacteristic linear kernel,

$$l(y, y') := yy' \quad (19)$$

and the Gaussian kernel introduced in eq 15. Other kernels are possible, but we leave their consideration for later work.

To stabilize the numerics and handle order-of-magnitude concerns, we transform and scale the loss values before applying the loss kernel to them:

$$\tilde{y}(y) = \frac{\ln y - \min(\ln \mathbf{y})}{\max(\ln \mathbf{y}) - \min(\ln \mathbf{y})} \quad (20)$$

### 3. SOFTWARE

Training data were calculated using the Amsterdam Modeling Suite (AMS) 2022.

The sensitivity calculation has been implemented in the Python programming language and integrated into a development version of ParAMS 2023,<sup>14</sup> a reparametrization toolkit which comes bundled with AMS. It is also used to handle the force field optimizations through the GloMPO optimization management software.<sup>4</sup>

### 4. THE REPARAMETRIZATION PROBLEM

Our main goal in this work is to demonstrate the method and utility of our new sensitivity approach. For this reason, the training set we have designed is smaller than what would typically be required to create a production-quality parametrization. Designing a large training set would require an entire publication in its own right and distract heavily from the focus of our article. A smaller training set allows us to keep the discussions more focused and clear.

In any case, training set design is always an iterative process where more items are added as validation tests on the new force fields demonstrate deficiencies in the original set.<sup>50</sup> This work can be considered a first iteration upon which later work can build.

**4.1. Initial Force Field.** We aim to create a new ReaxFF force field which correctly models the adsorption of H<sub>2</sub>S on ZnS, an adsorbent which has received some attention for its favorable electronic properties.<sup>51–55</sup> H<sub>2</sub>S is a common but toxic gas, and its adsorption behavior on ZnS has been investigated for both gas detection<sup>55</sup> and gas removal purposes.<sup>53</sup> As far as we are aware, however, no specially designed Zn/S/H ReaxFF force field exists to model this behavior.

Table 1 details the nomenclature we will use to refer to the various parameter blocks which compose the ReaxFF force field file.<sup>56</sup> For a discussion of the parameter blocks, see Section 2.1

**Table 1. Aliases Used to Refer to Different Parameter Blocks in the ReaxFF Force Field File Format**

Alias	Description	Number of parameters
GEN	General parameters	41
ATM:W	Parameters for atoms of type W	32 per atom type
BND:W.X	Parameters for W–X bonds	16 per W.X pair
OFD:W.X	Off-diagonal definitions for atom pairs	6 per W.X pair
ANG:W.X.Y	Parameters for angles formed by W–X–Y atoms	7 per angle group
TOR:W.X.Y.Z	Parameters for torsions formed by W–X–Y–Z atoms	7 per angle group
HBD:W.H.X	Hydrogen-bonding parameters for atoms W and X	4 per W.X pair

Our initial force field is an amalgamation of parameter values from already published force fields:

1. The ATM:S, BND:S.S, and ANG:S.S.S parameter blocks are filled with values from the Li/S force field of Islam et al.<sup>57</sup>
2. ATM:H, BND:S.H, and ANG:H.S.H values are taken from the C/H/O/S force field of Müller and Hartke.<sup>50</sup>
3. GEN, ATM:Zn, BND:H.H, BND:Zn.H, BND:Zn.Zn, BND:Zn.S, ANG:Zn.Zn.S, ANG:Zn.S.Zn, ANG:Zn.S.S, ANG:S.Zn.S, and HBD:S.H.S blocks are filled with parameters published in the Zn/O/H force field of Raymand et al.<sup>54</sup> This force field does not contain any sulfur compounds, so the corresponding oxygen-related parameter blocks are used. For example, BND:Zn.O parameters are filled into the BND:Zn.S block of the initial force field.

The force field contains 279 parameters. From this we use some common intuition based on the contents of the training set and descriptions of the parameters to select 53 for initial optimization. The selection is “greedy” in that a parameter is selected for optimization if:

1. it could conceivably affect the behavior of a training set item;

2. the training set contains items related to the parameter; and
3. the parameter is appropriate for optimization.

For example, the atomic masses of the elements are technically parameters, but changing them would not be appropriate. Similarly,  $\pi$  bonds are not present in the training set, and thus, their related parameters are not chosen for optimization. Finally, we do not optimize any atomic or general block parameters. These are generally known to be more “expert”-level parameters that are not suitable for preliminary-stage optimization.<sup>14</sup>

The 53 parameters selected for optimization are listed in Table 2. The default parameter ranges supplied by ParAMS<sup>14</sup> were used and can be found in the Supporting Information (SI).

**4.2. Training Data.** The training set, calculated with BAND using PBEsol and DZ or DZP basis set, consists of 16 jobs, from which 471 individual training points are extracted.

AMS BAND calculates charges via the Hirshfeld, Voronoi deformation, Mulliken, and CMS methods. We have included Hirshfeld charges in our training set. Although these charges are not without problems,<sup>59</sup> we believe they are the best available. The accuracy of these charges is also not overly important since we do not activate any charge-related parameters in our initial parameter selection (see Section 4.1), and the few charges we include in the training set are mainly used as a sanity check on the results.

The training set consists of:

1. energy–volume scans, charges, and enthalpies of formation for the cubic zincblende, cubic rocksalt, and hexagonal wurtzite structures of ZnS;
2. the optimized geometry, charges, H–S bond scan, and H–S–H angle scan for H<sub>2</sub>S;
3. the energy of formation, bond lengths, charges, and angles of a periodic zincblende surface;
4. the forces of a distorted zincblende surface;
5. the charges, bond length, and adsorption energy of H<sub>2</sub>S adsorbed on zincblende.

Details of the training set are included in SI Section F.

Default ParAMS values are used for  $\sigma$ . Individual energies, angles, distances, charges, torsions, and force groups are equally weighted in the initial training set. The full training set, job collection, and initial parameter interface are available in the SI.

**Table 2. Initial Naive Parameter Selection (Equation Numbers and Parameter Names Refer to Notation Used in Reference 58)**

Block	Name	Equation	Description	Atoms
ANG	-p_hb2	18	Hydrogen bond/bond order	H.S.H
	-p_hb3	18	Hydrogen bond parameter	H.S.H
	p_hb1	2	Hydrogen bond energy	H.S.H
	p_val1	13a	Valence angle parameter	H.S.H, S.S.Zn, S.Zn.S, S.Zn.Zn, Zn.S.Zn
	p_val2	13a	Valence angle parameter	H.S.H, S.S.Zn, S.Zn.S, S.Zn.Zn, Zn.S.Zn
	p_val4	13b	Valence angle parameter	H.S.H, S.S.Zn, S.Zn.S, S.Zn.Zn, Zn.S.Zn
	p_val7	13c	Undercoordination	H.S.H, S.S.Zn, S.Zn.S, S.Zn.Zn, Zn.S.Zn
	Theta_0,0	13g	180° – (equilibrium angle)	H.S.H, S.S.Zn, S.Zn.S, S.Zn.Zn, Zn.S.Zn
HBD	r_hb^0	18	Hydrogen bond equilibrium distance	S.H.S
BND	D_e^sigma	6, 11a	$\sigma$ -Bond dissociation energy	H.S, S.S, S.Zn, Zn.Zn
	p_be1	6	Bond energy parameter	H.S, S.S, S.Zn, Zn.Zn
	p_be2	6	Bond energy parameter	H.S, S.S, S.Zn, Zn.Zn
	p_bo1	2	$\sigma$ -Bond order	H.S, S.S, S.Zn, Zn.Zn
	p_bo2	2	$\sigma$ -Bond order	H.S, S.S, S.Zn, Zn.Zn
	p_ovun1	11a	Overcoordination penalty	H.S, S.S, S.Zn, Zn.Zn




Figure 2. Optimizer trajectories for the 16 CMA-ES reparametrizations of the initial force field using the naive parameter selection.

## 5. RESULTS AND DISCUSSION

Our workflow proceeds as follows:

1. **reparametrization of the initial force field** based on the naive parameter selection (done purely for comparative purposes);
2. **running of the sensitivity calculation** to determine a parameter ordering (entirely independent of the previous step);
3. **rerunning the reparametrization** using 10, 20, 33, and 43 of the least and most sensitive parameters, as determined by the sensitivity analysis;
4. **running validation tests** using the initial parametrization and some of the best parametrizations found.

Our aim is to show that:

1. the HSIC sensitivity method correctly identifies the most sensitive parameters;
2. reparametrizations within a lower-dimensional space of very sensitive parameters can find force fields of similar quality in a shorter amount of time;
3. reparametrizations in higher dimension run the risk of overfitting.

**5.1. Initial Optimizations.** An initial set of 16 reparametrizations were performed using the covariance matrix adaptation–evolutionary strategy (CMA-ES),<sup>60</sup> which has been shown to work well on ReaxFF reparametrizations.<sup>4,13</sup> All of the optimizers were started from the initial force field values described in Section 4.1 with a wide initial sampling distribution ( $\sigma_0 = 0.5$ ). All parameters were scaled between 0 and 1 according to their bounds, which are automatically suggested by ParAMS.

Optimizers were stopped if:

1. within the last 2000 function evaluations their lowest value had not improved and their explored values were very similar; or
2. any of the default internal CMA-ES convergence criteria were triggered.

The optimizers were run in parallel, and the entire optimization was stopped after using a cumulative total of 300 000 function evaluations. These conditions are all quite strict and provided all of the optimizers a very long time to find the best possible minimum.

Figure 2 presents the optimizer trajectories for these reparametrizations, showing the original loss of 20 662 being

improved to between 410 and 308. These results will be discussed in more detail in Section 5.3 when comparing the results to other optimizations. Detailed results for every optimizer can be found in SI Section A.

**5.2. Sensitivity Analysis.** A total of 10 000 uniformly randomly generated samples of the original optimization space were collected for the sensitivity analysis. The sampling procedure took a total of 142 min on a 64 core node (AMD EPYC 7513 @ 2.6 GHz). This sampling took place once, and this same set of samples was used in all of the sensitivity calculations.

The distributions of sampled loss values ( $y$ ) and their transformed values ( $\tilde{y}$ ) are shown in Figure 3; the lowest value is 2651. Approximately half of the samples are actually worse than

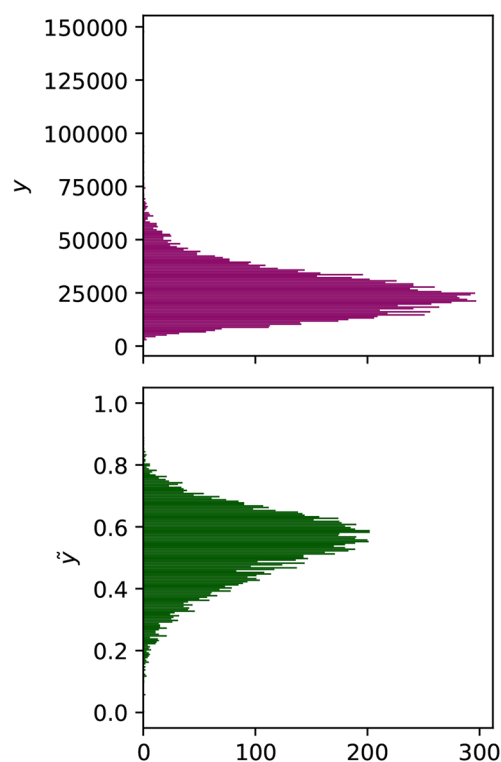


Figure 3. Distributions of (top) loss values and (bottom) loss values scaled according to eq 20, gathered during the sampling procedure and used during the sensitivity analyses.

the initial loss value, and the set clearly does not include the good minima found through the original optimizations in Section 5.1. We will demonstrate that having such minima in the sample set is not a requirement for having a good estimation of parameter sensitivity.

Instead of using all of the samples at once in a single HSIC calculation, we repeated it 10 times using a bootstrapping method. Each bootstrap used 2000 random subsamples from the original 10 000 sample set. Using fewer samples per calculation significantly speeds up the calculation, and the spread of the results gives us an indication of the error. Unless otherwise mentioned, the sensitivity values discussed below refer to the average value across the 10 repeats. Algorithm 1 provides a pseudocode of the calculation where we have abandoned notational precision for brevity and clarity.

**Algorithm 1** Sensitivity calculation workflow.

**Require:** SSE,  $k$ ,  $l$

- 1: randomly generate  $\mathbf{X} \in \mathbb{R}^{10000 \times 53}$
- 2:  $\mathbf{y} \in \mathbb{R}^{10000} \leftarrow \text{SSE}(\mathbf{X})$  ▷ Evaluate loss
- 3:  $\tilde{\mathbf{y}} \leftarrow \tilde{y}(\mathbf{y})$  ▷ Scale loss values
- 4: **for**  $i = 1, \dots, 10$  **do** ▷ Repeat the calculation
- 5:  $\mathbf{X}^{\text{sub}} \leftarrow$  2000 random rows of  $\mathbf{X}$
- 6:  $\tilde{\mathbf{y}}^{\text{sub}} \leftarrow$  corresponding 2000 loss values
- 7:  $\tilde{\mathbf{L}} \leftarrow l(\tilde{\mathbf{y}}^{\text{sub}})$
- 8: **for**  $\omega = 1, \dots, 53$  **do** ▷ Every parameter
- 9:  $\tilde{\mathbf{K}} \leftarrow k(\mathbf{X}_{\omega}^{\text{sub}})$
- 10:  $\mathbf{H}_{i\omega} \leftarrow \text{HSIC}(\tilde{\mathbf{K}}, \tilde{\mathbf{L}})$
- 11:  $\mathbf{h} \leftarrow \text{MEAN}(\mathbf{H}_1, \dots, \mathbf{H}_{10})$  ▷ Average repeats
- 12:  $s_j \leftarrow \frac{h_j}{\sum_i h_i}$  ▷ Calculate sensitivities


The following kernels were applied to the loss values: CG kernel ( $\gamma = 0.3$ ), Gaussian kernel ( $\sigma = 0.3$ ), linear kernel, and threshold kernel ( $\alpha = 0.4$ ). In all cases a Gaussian kernel ( $\sigma = 0.3$ ) was applied to parameter values. Since all values are scaled between 0 and 1, a value of  $\sigma = 0.3$  approximates the median distance between all samples in a uniform distribution between these values. This has repeatedly been reported as an appropriate value.<sup>17,61</sup>

The average time needed to perform the sensitivity calculation (regardless of loss kernel and including the repeats) was approximately 60 s. The optimizations require several days to complete, and thus, the inclusion of these sampling and sensitivity steps in a reparametrization workflow would not be the limiting step and, as will be shown below, comes with several advantages.

We note that we also repeated the calculations using a sample pool of 30 000 samples and 10 000 subsamples per calculation. These results produced tighter distributions between the repeats at the cost of more time, but the average sensitivity values were essentially the same. In other words, they were not worth the effort and are not presented below.

The scaling of the sensitivity calculation is  $O(rdn^2)$ , where  $r$  is the number of repeats,  $d$  is the number of parameters, and  $n$  is the number of subsamples used per calculation. Increasing the size of the matrices has the largest impact on the calculation time, so being able to get robust parameter orderings from repeated calculations using small subsamples is very advantageous.


Figure 4 compares the sensitivity values determined by each calculation for each parameter. Despite the large differences between loss kernels, the ordering of parameters is very robust for the most sensitive parameters. Significant deviations only appear for the least sensitive parameters, which are all close to




**Figure 4.** Average sensitivity values (over 10 runs each) determined for each parameter using different kernels applied to the loss values. Parameters are ordered by the results of the CG kernel, which is used in the later reparametrizations. Black lines group the 10, 20, 33, and 43 most sensitive parameters.

one another numerically (note the logarithmic axis). This demonstrates that the HSIC approach we have taken is robust and not overly dependent on the choice of loss kernel or hyperparameters.

The loss kernel which struggles the most with differentiating sensitivities is the threshold kernel. We believe that this is because the kernel is discontinuous and provides very little information to the calculation because values are toggled to 0 or 1. Our investigation of this kernel was motivated by the work of Spagnol et al.<sup>67</sup> They demonstrated the necessity of this kernel to focus the calculation on small loss values (near the minima of interest) and negate the effects of large order-of-magnitude differences. As discussed in Section 2.4, our implementation always takes the logarithm of the loss values and then scales the result; it appears to address the order-of-magnitude problem.



**Figure 5.** Grouped sensitivities as determined by the HSIC calculation using the CG kernel, grouped by (left) parameter group and (right) parameter name.



**Figure 6.** Running best value seen by any of the 16 parallel optimizers for each of the optimization configurations.

Random sampling will also generally not contain very good minima (true in this case, as discussed above), making the second advantage of the threshold kernel less important. The continuity provided by the other kernels seems to make converging the sensitivity values for the parameters easier.

We have also explored the effect of using different kernel parameters in SI Section B. This has a minor effect on the sensitivity results, provided that a reasonable kernel parameter value is selected so that the kernel values are not all very close to 0 or 1.

Given the similarities in orderings obtained by the various loss kernels, we will continue discussing only the results obtained when using the CG kernel ( $\gamma = 0.3$ ). This is a somewhat arbitrary choice, but it can be argued that it has the best theoretical foundation by emphasizing the sensitivities for low minima.

Figure 5 shows the sensitivities obtained grouped by (a) parameter group and (b) parameter name. For example, “S.Zn” refers to all of the bond parameters associated with sulfur–zinc bonds, and “p\_bo2” refers to all of the p\_bo2 bond parameters regardless of atoms involved. p\_bo2 appears in the calculation of  $\sigma$ -bond orders and is an exponential term. D\_e^sigma is a linear parameter which is multiplied by the  $\sigma$ -bond orders to determine the bond energy contribution to the overall potential (see eq 6).

The identification of these  $\sigma$ -bond parameters, particularly for sulfur–zinc and zinc–zinc bonds, is quite reasonable given the composition of the training set. The degree to which the sensitivity is dominated by only these parameters, however, may be surprising. Nevertheless, the results seem intuitive.

**5.3. Reparametrizations with Sensitive Parameters Only.** Based on the sensitivities determined above, we ran reconfigured reparametrizations using the most sensitive 10, 20,

33, and 43 parameters and the least sensitive 10, 20, 33, and 43 parameters. The most sensitive groupings are shown in Figure 4. We made these selections so that the set of 10 most sensitive parameters and the set of 43 least sensitive parameters are complementary; together they account for the original 53 parameters. The other combinations were handled similarly. Each of the reparametrizations was conducted in the same way as the original described in Section 5.1.

In a more practical setting we do not advocate blindly activating only the most sensitive  $n$  parameters; we do so here only to avoid introducing human decision-making into the results. In practice, one should use the sensitivities as a guide to better understand how the training data affect the loss function and use some human intuition when selecting parameters.

For the sake of clarity, we introduce the following nomenclature to refer to the different optimizations:

O53 reparametrizations using the original 53 active parameters;

M# reparametrizations using the most sensitive # parameters;

L# reparametrizations using the least sensitive # parameters.

Figure 6 shows the running best loss value seen by any of the 16 parallel optimizers for each of the optimization configurations (SI Section A contains the results for individual optimizers). Table 3 shows the difference between the lowest loss value and initial loss value as a fraction of the initial loss value, in other words, the percentage of the initial loss which was “removed” during the optimization. Unsurprisingly, O53 produces the lowest loss because it had the most degrees of freedom. Any reduction in the number of active parameters can be expected to worsen achievable loss values. Indeed, all of the other optimizations find worse minima than the original. However,

**Table 3. Fractions of the Initial Loss Value Removed during Optimization**

	Most Sensitive (%)	Least Sensitive (%)
10	94.2	3.61
20	97.7	10.9
33	98.2	29.4
43	98.4	87.4
53	98.5	

the difference between using the most sensitive and least sensitive parameters is marked.

In all cases, using the most sensitive parameters allows us to remove >93% of the original error. In fact, some M33 and M43 optimizers find better minima than some O53 optimizers. In contrast, using the same number of least sensitive parameters produces very poor optimizations which are unable to meaningfully reduce the loss value. Especially noteworthy is that M10 is able to locate better minima than its complement L43. Similarly, L10 is only able to reduce the loss by 4% compared to 94% achieved by M10.

It is also important to note that M10 and M20 are able to converge long before the 300 000 function evaluation limit and use approximately one-third and two-thirds of the time used by O53. This represents significant time savings, as O53 took approximately 2 days to complete, more than the time required to run the sensitivity analysis and M10 or M20.

These results show that our proposed HSIC sensitivity method can quickly identify the most sensitive parameters for reparametrization and that parameter selection guided by it can produce good minima in a shorter time.

**5.4. Force Field Comparison and Validation.** In order to verify the quality of the new force fields and demonstrate the risks of overfitting, we perform several validation tests using the initial force field and the force fields with lowest error produced

by O53, M10, and M20 across the 16 optimizations each one performed. We consider:

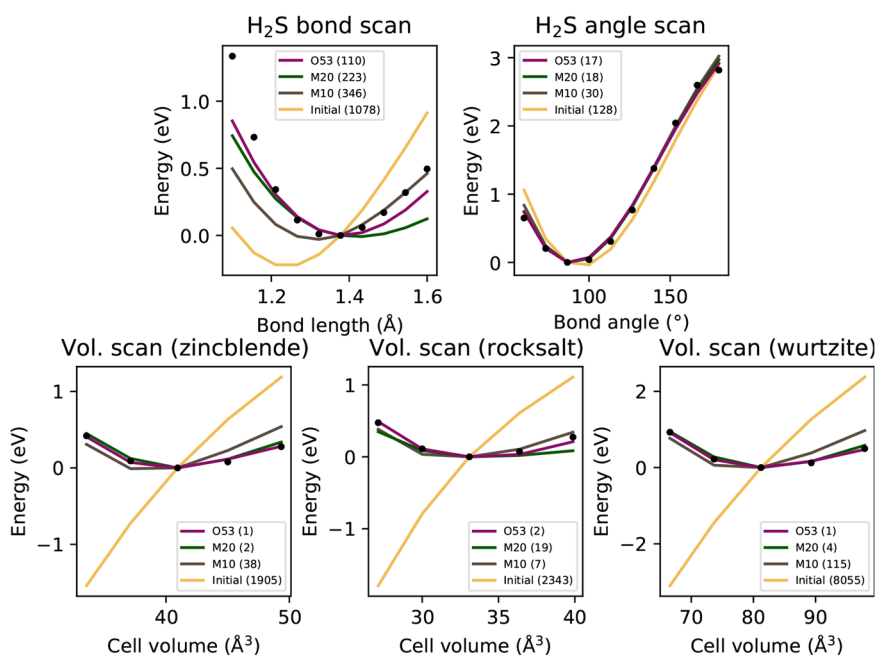
1. the error on the training set items;
2. an MD simulation of H<sub>2</sub>S adsorption on zincblende and wurtzite slabs;
3. an MD simulation of zincblende bulk crystal;
4. the adsorption energy of a H<sub>2</sub>S molecule on a (110) zincblende surface; and
5. the surface energies of (110) zincblende and (11 $\bar{2}$ 0) wurtzite.

We compare the results to reference DFT calculations as well as literature values. For calculations involving crystal surfaces, we use the (110) face of zincblende and the (1120) face of wurtzite ZnS, as they have been found to be the most energetically favorable.<sup>52,62</sup> Slabs were always constructed from preoptimized lattice parameters calculated with BAND using the PBEsol/DZP level of theory.

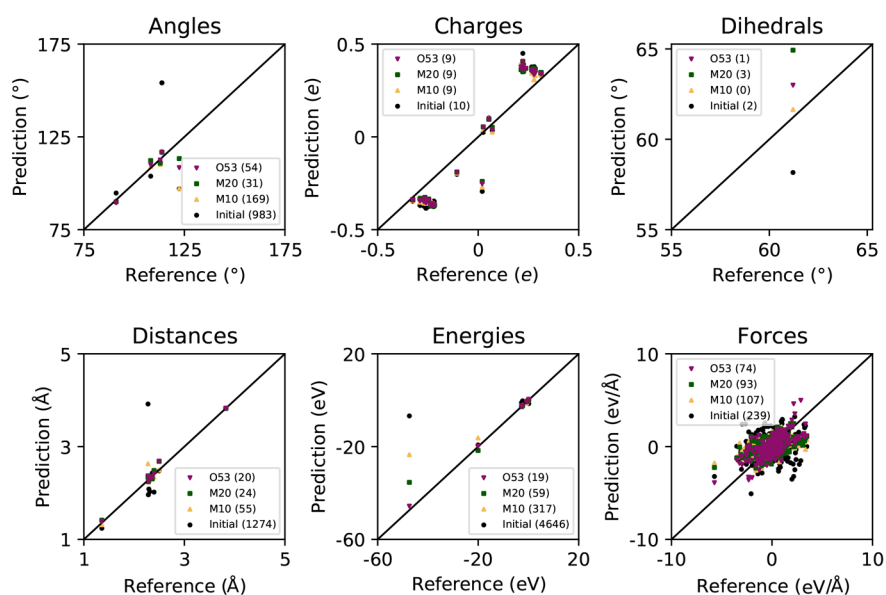
**5.4.1. Training Set Errors.** We start our analysis by decomposing the overall loss values for each of the force fields into their individual contributions. Figures 7 and 8 compare force field predictions to reference values for each of the training set items.

Figure 7 shows training set items which can be grouped together into PES scans along some coordinate. With the exception of the H<sub>2</sub>S angle scan, the original force field performs extremely poorly for these items, and in all cases the optimized force fields produce significantly improved results. The O53 force field is the most accurate for all the scans, but M20 is generally quite competitive and has similar errors. The one exception to this is the H–S bond scan in the H<sub>2</sub>S molecule, which M20 did not replicate well.


Figure 8 shows the remaining training set items where we compare reference values to predictions made by the selected force fields. Although angles and forces were improved, significant errors remained in most cases, and a large fraction



**Figure 7.** Training set items which form part of various PES scans. Reference values are shown as black dots, and loss values for each force field are shown in parentheses in the legends.



**Figure 8.** Comparison of reference to predicted values for each of the training set items and force fields, grouped by item type. Loss values for a group and force field are included in parentheses in the figure legends.




**Figure 9.** One surface in the final frames of 75 ps MD simulations of 100  $\text{H}_2\text{S}$  gas molecules above and below a 10 nm thick wurtzite slab with  $(11\bar{2}0)$  surfaces using the (a) initial, (b) O53, (c) M10, and (d) M20 force fields. Sulfur atoms in yellow, zinc in silver, hydrogen in green.

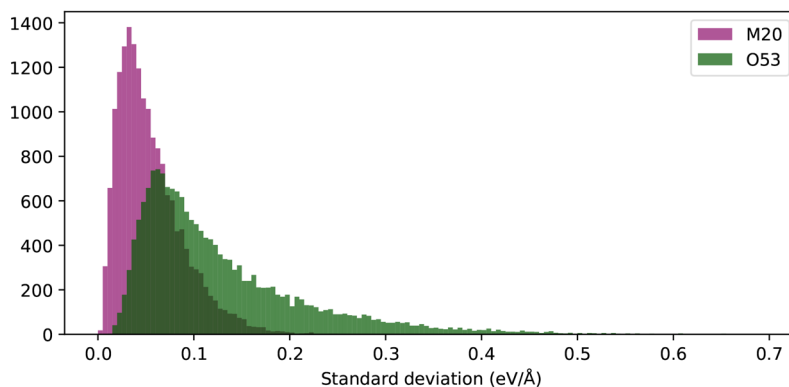
of the original forces' error remains in the final fields. Energies, which form a large part of the overall loss value, actually appear generally well-fitted, and most of the error comes from a single item: the energy difference between a distorted and regular (110) zincblende surface. We can also see that charges were effectively unimproved during the reparametrization. This is because charge-related parameters were not activated in our parameter selection. Interatomic distances were the most improved class of training set items. This can be explained by the fact that the original sulfur-related parameters were actually

trained for oxygen, and thus, the initial force field bond lengths were too short to accommodate the larger sulfur atoms and had to be lengthened.

In most cases the M10 errors are substantially worse than the M20 ones. On the other hand, M20 errors are often very similar to (or in the case of angles) better than the O53 errors. In comparison to O53, M20 performs worst for forces and energies. However, as mentioned above, the bulk of the energy error is concentrated in one item, and no force field appears to predict forces very well. From these observations, there does not appear



**Figure 10.** One surface in the final frames of 75 ps MD simulations of 100  $\text{H}_2\text{S}$  gas molecules above and below a 10 nm thick zincblende slab with (110) surfaces using the (a) initial, (b) O53, (c) M10, and (d) M20 force fields. Sulfur atoms in yellow, zinc in silver, hydrogen in green.



**Figure 11.** Standard deviation on the forces felt by all atoms in all directions at all time steps in bulk simulations of zincblende crystal using committee simulations of the 16 M20 and O53 force fields.


to be a strong signal that O53 is a significantly better force field than M20.

**5.4.2. MD Simulations. Adsorption to Crystal Surfaces.** Our first validation test is an MD simulation of the adsorption of  $\text{H}_2\text{S}$  molecules on ZnS slabs. We are guided by literature in constructing a realistic simulation scenario. Zhang et al.<sup>62</sup> showed computationally and experimentally that cubic ZnS with a particle size smaller than  $\sim 7$  nm is not stable at room temperature and can easily convert to the hexagonal polymorph. Experimentally, Dloczik et al.<sup>51</sup> were able to produce stable ZnS columns with wall thicknesses between 10 and 30 nm where both cubic and wurtzite phases were detected. Qi et al.<sup>55</sup> later

studied the adsorption of  $\text{H}_2\text{S}$  on ZnS surfaces at room temperature.

Using these sources, we create MD simulations for both morphologies using slabs  $\sim 10$  nm thick. The temperature is maintained at 298 K, and 100  $\text{H}_2\text{S}$  molecules are randomly positioned above and below the slabs. The simulations are conducted for 300 000 time steps of 0.25 fs each.

Figures 9 and 10 show the final surface geometries for the wurtzite and zincblende simulations, respectively. Unsurprisingly, the initial force field performs poorly, as the slabs immediately collapse. The M10 and M20 simulations, however, perform better. The slabs maintain their crystalline shapes and adsorb the  $\text{H}_2\text{S}$  molecules to the slab surfaces.



**Figure 12.** (a)  $\text{H}_2\text{S}$  adsorption energy, (b) wurtzite (11 $\bar{2}$ 0) surface energy, and (c) zincblende (110) surface energy versus SSE for the best force fields found by each of the 16 optimizers run in the O53 and M20 reparametrizations. Reference DFT results are shown with horizontal yellow lines.

The O53 simulations do not behave in the same way. These slabs struggle to maintain their surfaces. In the wurtzite case, the surface layers have become almost amorphous, and although the bulk crystal is relatively well-maintained, the extent of the deformation extends several layers deep. The zincblende surface is more ordered, but it is not as rigidly maintained as M10 or M20, and it appears that some  $\text{H}_2\text{S}$  molecules are being chemically absorbed into the structure.

As discussed above, literature suggests that these slabs should be stable rigid solids under these conditions, making the behavior of the O53 force field surprising. We will discuss the reasons for this behavior in Section 5.4.3.

**Bulk Crystal.** The above simulation provides a qualitative indication that the lower loss value of the O53 force field does not guarantee better performance. However, we would like to demonstrate quantitatively that M20 parametrizations produce more robust results.

To do this, we ran a second MD simulation of a bulk zincblende crystal (64 atoms total) at 500 K for 10 ps using a time step of 0.5 fs. Two simulations were run: one used the best force fields produced by each of the 16 M20 optimizers, and the other used the best force fields produced by each of the 16 O53 optimizers. In each simulation the force fields were used as a committee. In other words, each force field was applied to the same crystal geometries at each time step. The forces of these results were then averaged to update the atomic positions for the next time step.

Figure 11 shows the standard deviation of the forces across the 16 force fields. We are not interested in the results at particular times, for particular atoms, or in particular directions because the forces always average out to zero. Thus, the histograms simply show all of the standard deviations regardless of the above distinctions. The 16 M20 force fields show significantly more agreement than the O53 ones. This suggests that the lower-dimensional optimization is much more robust and produces force fields with a tighter distribution of results. The higher-dimensional optimization has many more minima in which optimizers might be trapped. Since these extra minima are created by the presence of insensitive parameters, they are more likely to represent overfittings than genuine attractors driven by important parameters. Interestingly, the favorable spread of results seen here for M20 could not be assumed from the spread in loss values among its 16 force fields, which was 98 compared to 33 for O53.

**5.4.3. Adsorption and Surface Energies.** Our more quantitative validation tests are:

1. the calculation of the adsorption energy ( $E_{\text{ads}}$ ) of a  $\text{H}_2\text{S}$  molecule on the (110) surface of zincblende; and
2. surface energy ( $E_{\text{surf}}$ ) calculations for the (110) surface of zincblende and (11 $\bar{2}$ 0) surface of wurtzite.


Details of these calculations are included in SI Sections C and D, respectively.

In this case we use the best force fields from every optimizer in the O53 and M20 sets to analyze the distribution of results. We do not consider all 16 M10 results, as they perform poorly. The results are shown in Figure 12 as a function of the force fields' associated SSE. Detailed results for individual optimizers can be found in SI Section E. In these plots, we have also included DFT, initial force field, and best M10 force field results for comparison. The initial and M10 results are not shown with their corresponding loss values, as they are too large and would make the plots illegible. In the discussions that follow, we will continue to use “best” to refer to the force field with the lowest loss value (SSE); it does not mean the force field with the most accurate validation result.

**Adsorption Energy.** We begin our comments by considering the adsorption energy results. None of the force fields are very accurate in this regard in comparison to our DFT reference. However, the best M20 force field is more accurate than that from O53. Interestingly, the initial force field produces a better prediction than M10 and is only slightly worse than the best O53 force field. More broadly, however, all of the M20 force fields predict adsorption energies in a narrow range from  $-0.89$  to  $-0.81$  eV, which is much more reproducible than the O53 results, which range from  $-1.14$  to  $-0.43$  eV.

This variability might be explained by the odd behavior seen in several O53 geometry optimizations. Figure 13 shows the final frames the adsorbed molecule geometry optimizations using DFT and the best O53 force field. The  $\text{H}_2\text{S}$  molecule is initially placed close to the zinc atom highlighted in blue and settles into this position using DFT. However, using the O53 force field, the  $\text{H}_2\text{S}$  molecule moves away from this atom toward a row of sulfur atoms which are, in turn, repelled by it. The geometry optimization is converged, which means that this desorbed position with deformed surface is more energetically favorable than an adsorbed position on zinc—which is incorrect. It is possible that the strange surface effects seen in Section 5.4.2 may be associated with this.

**Surface Energy.** If we consider the surface energy results, the distributions of the O53 force fields are consistently larger than those of the M20 results. On average, the O53 results are similarly accurate to M20 for the wurtzite surface. For the



**Figure 13.** Final frame of a geometry optimization of a  $\text{H}_2\text{S}$  molecule adsorbed on a (110) zincblende surface using (a) DFT and (b) the best O53 force field. The molecule is originally positioned near the zinc atom shown in blue.

zincblende surface, one of the M20 clusters is actually quite accurate.

Overall, none of the force fields performed generally well across all of the validation tests, but they are improvements to the initial force field. If more accuracy is desired, then more attention needs to be given to the composition of the training set. However, that is not the focus of this article.

**5.4.4. Correlations with Loss Values.** The most interesting conclusions from these results come from comparing them to the force fields' associated SSEs. All 16 O53 optimizers found SSE values tightly clustered between 308 and 410, which are much lower and less variable than the SSEs of the M20 optimizers, which are clustered around two values between 485 and 684. If one only considers the loss values, then one might expect that predictions made by the O53 force fields will be similarly more precise than those of M20; however, we see that this is not the case. O53 validation results are significantly more variable than those of M20 and, on average, seem to be only slightly more accurate.

Most crucially, the high variability of the O53 validations is not correlated with the SSE. In other words, a better prediction of the training set items is not correlated with a more accurate prediction of the validation tests. We believe that these facts and the unusual behavior of the O53 MD simulations in Section 5.4.2 can be attributed to overfitting. Interestingly, the M20 results seem to be *inversely* correlated with the loss values, i.e., a higher loss leads to a more accurate validation result. This might suggest that M20 also sees some overfitting effects. However, there is not enough data to verify this because all optimizers fall near one of only two SSE values.

Overfitting is a common occurrence during ReaxFF reparametrization<sup>9,13,50</sup> and becomes more likely as more

parameters are allowed to change. These validation tests highlight the dangers of activating too many dimensions during a reparametrization, as more degrees of freedom create opportunities for overfitting. We present two dangers which occur when too many parameters are used.

First, it is important to emphasize the difference between a parameter's *sensitivity* to some loss function (which is a function of a user-generated training set) and its true *importance* in the potential function. If a training set does not include enough items to capture the importance of a parameter, then it will register as insensitive. If an important but insensitive parameter is activated during reparametrization, then the optimizer will change it quasi-randomly. This can result in force fields with very low loss values that perform very poorly in production runs; sensitive parameters have been set correctly, but other important ones have been changed incorrectly.

A second danger is the fact that parameters can have a compensatory effect. If many parameters are active, then it is possible that multiple parameter sets can satisfy the training data, but the potential will perform poorly during simulation. By using too many parameters, the optimizers have an opportunity to move insensitive ones to reduce the loss rather than identifying the "correct" values for a small number of appropriate parameters.

In our example, the 53-dimension reparametrizations seem to have activated too many dimensions. Although they produced the lowest overall loss values, they did not achieve the lowest loss for every class of items, nor did they always make the most accurate or precise predictions in our validation tests. In the MD simulations we see odd behavior which, along with the precision problems discussed previously, we believe to be evidence of overfitting.

Conversely, the reparametrizations using only the 10 most sensitive parameters seem to not have had enough degrees of freedom. These force fields produce significant training set errors compared to O53 and M20 and perform poorly in the validation tests.

Using 20 of the most sensitive parameters produces good-quality force fields which perform better than the O53 force fields in some tests, despite having marginally worse training set losses. The results of the validation tests were also consistently more precise than those of O53 and of comparable accuracy. The MD simulation aligned with expectations, and the force fields were produced in a shorter time than that required for the original optimization.

## 6. OUTLOOK

We believe that there are several areas available which warrant further study. First, the training set used here could be used as a starting point to produce a better-quality force field for Zn/S/H. This would require a significant expansion of the training set and involve all of the well-known and thorny issues associated with training set design.<sup>50</sup>

Second, a parameter's sensitivity is a strong function of the range in which it is allowed to vary. In this work, we have not needed to give this much consideration since the ParAMS package provides recommended ranges. However, it is not certain that these are always appropriate. A formalization of appropriate ranges or a robust mechanism to determine them would go a long way to ensuring that sensitivities are being appropriately determined.

Third, we have identified sampling as the limiting step of the sensitivity procedure. It may be interesting to investigate the

extraction of samples from optimizer trajectories. However, since the HSIC requires uniformly distributed random samples, one would need to extract samples carefully from the trajectories. One possibility is the Kennard–Stone algorithm,<sup>63</sup> but this is slow and sequential. Nevertheless, “closing the loop” and allowing users to extract sensitivities from optimization results seems like an appealing prospect.

## 7. CONCLUSIONS

As an introductory work, we have demonstrated that an HSIC sensitivity analysis applied to a ReaxFF reparametrization can successfully identify the most sensitive parameters. We have further shown that using only the most sensitive parameters during optimization leads to faster convergence and a reduced chance of overfitting. Even qualitatively, the use of such a sensitivity analysis can provide valuable insights for the user into the composition of the training set. Overall, we believe that the HSIC sensitivity analysis is a useful, robust, and easy to use tool which has the potential to greatly aid in the reparametrization of ReaxFF force fields.

## ■ ASSOCIATED CONTENT

### SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.2c01320>.

Full optimization results, calculation methodologies for adsorption and surface energies, full adsorption and surface energy results, and training set composition (PDF)

YAML files (human-readable and used by ParAMS<sup>14</sup>): training set geometries and job calculation settings (`job_collection.yaml`), task settings for the input geometries (`engine_collection.yaml`); reference values, weights, and  $\sigma$  values for training data (`training_set.yaml`); original ReaxFF parameter values and ranges (OS3 parameters set as active) (`initial_parameters.yaml`) (ZIP)

## ■ AUTHOR INFORMATION

### Corresponding Author

Toon Verstraelen – Center for Molecular Modeling, Ghent University, B-9052 Ghent, Belgium; [orcid.org/0000-0001-9288-5608](https://orcid.org/0000-0001-9288-5608); Email: [toon.verstraelen@ugent.be](mailto:toon.verstraelen@ugent.be)

### Authors

Michael Freitas Gustavo – Center for Molecular Modeling, Ghent University, B-9052 Ghent, Belgium; Software for Chemistry and Materials BV, 1081 HV Amsterdam, The Netherlands

Matti Hellström – Software for Chemistry and Materials BV, 1081 HV Amsterdam, The Netherlands; [orcid.org/0000-0003-3053-5658](https://orcid.org/0000-0003-3053-5658)

Complete contact information is available at: <https://pubs.acs.org/doi/10.1021/acs.jctc.2c01320>

### Author Contributions

M.F.G. designed and implemented the sensitivity test, ran the test work, and wrote this article. M.H. designed the training set and guided some of the validation tests. T.V. provided supervision, revisions, guidance, and advice.

### Notes

The authors declare no competing financial interest.

## ■ ACKNOWLEDGMENTS

The authors thank the Flemish Supercomputing Centre (VSC), funded by Ghent University, FWO, and the Flemish Government, for use of their computational resources. Funding for the project was provided by the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement 814143. T.V. and M.F.G. were also supported by the Research Board of Ghent University (BOF) under Award 01IT2322.

## ■ REFERENCES

- (1) van Duin, A. C. T.; Dasgupta, S.; Lorant, F.; Goddard, W. A. ReaxFF: a reactive force field for hydrocarbons. *J. Phys. Chem. A* **2001**, *105*, 9396–9409.
- (2) Senftle, T. P.; Hong, S.; Islam, M. M.; Kylasa, S. B.; Zheng, Y.; Shin, Y. K.; Junkermeier, C.; Engel-Herbert, R.; Janik, M. J.; Aktulga, H. M.; Verstraelen, T.; Grama, A.; van Duin, A. C. T. The ReaxFF reactive force-field: development, applications and future directions. *npj Comput. Mater.* **2016**, *2*, 15011.
- (3) Komissarov, L.; Rüger, R.; Hellström, M.; Verstraelen, T. ParAMS: parameter optimization for atomistic and molecular simulations. *J. Chem. Inf. Model.* **2021**, *61*, 3737–3743.
- (4) Freitas Gustavo, M.; Verstraelen, T. GloMPO (globally managed parallel optimization): a tool for expensive, black-box optimizations, application to ReaxFF reparameterizations. *J. Cheminf.* **2022**, *14*, 7.
- (5) Chenoweth, K.; van Duin, A. C.; Goddard, W. A. ReaxFF reactive force field for molecular dynamics simulations of hydrocarbon oxidation. *J. Phys. Chem. A* **2008**, *112*, 1040–1053.
- (6) Barcaro, G.; Monti, S.; Sementa, L.; Carravetta, V. Parametrization of a reactive force field (ReaxFF) for molecular dynamics simulations of Si nanoparticles. *J. Chem. Theory Comput.* **2017**, *13*, 3854–3861.
- (7) Bae, G.-T.; Aikens, C. M. Improved ReaxFF force field parameters for Au–S–C–H systems. *J. Phys. Chem. A* **2013**, *117*, 10438–10446.
- (8) Liu, Y.; Hu, J.; Hou, H.; Wang, B. Development and application of a ReaxFF reactive force field for molecular dynamics of perfluorinatedketones thermal decomposition. *Chem. Phys.* **2020**, *538*, 110888.
- (9) Hubin, P. O.; Jacquemin, D.; Leherste, L.; Vercauteren, D. P. Parameterization of the ReaxFF reactive force field for a proline-catalyzed aldol reaction. *J. Comput. Chem.* **2016**, *37*, 2564–2572.
- (10) Hu, X.; Schuster, J.; Schulz, S. E. Multiparameter and parallel optimization of ReaxFF reactive force field for modeling the atomic layer deposition of copper. *J. Phys. Chem. C* **2017**, *121*, 28077–28089.
- (11) LaBrosse, M. R.; Johnson, J. K.; van Duin, A. C. Development of a transferable reactive force field for cobalt. *J. Phys. Chem. A* **2010**, *114*, 5855–5861.
- (12) Larsson, H. R.; van Duin, A. C.; Hartke, B. Global optimization of parameters in the reactive force field ReaxFF for SiOH. *J. Comput. Chem.* **2013**, *34*, 2178–2189.
- (13) Shchygol, G.; Yakovlev, A.; Trnka, T.; van Duin, A. C.; Verstraelen, T. ReaxFF parameter optimization with Monte–Carlo and evolutionary algorithms: guidelines and insights. *J. Chem. Theory Comput.* **2019**, *15*, 6799–6812.
- (14) *Parameterization Tools for AMS*, 2022.204.r105828; Software for Chemistry and Materials, 2022; <https://www.scm.com/doc/params> (accessed: June 2022).
- (15) Saltelli, A.; Aleksankina, K.; Becker, W.; Fennell, P.; Ferretti, F.; Holst, N.; Li, S.; Wu, Q. Why so many published sensitivity analyses are false: a systematic review of sensitivity analysis practices. *Environ. Modell. Software* **2019**, *114*, 29–39.
- (16) Baroni, G.; Francke, T. An effective strategy for combining variance- and distribution-based global sensitivity analysis. *Environ. Modell. Software* **2020**, *134*, 104851.
- (17) Spagnol, A.; Le Riche, R.; Da Veiga, S. Global sensitivity analysis for optimization with variable selection. *SIAM/ASA J. Uncertainty Quantif.* **2019**, *7*, 417–443.
- (18) Borgonovo, E. A new uncertainty importance measure. *Reliab. Eng. Syst. Saf.* **2007**, *92*, 771–784.

- (19) Saltelli, A.; Ratto, M.; Andres, T.; Campolongo, F.; Cariboni, J.; Gatelli, D.; Saisana, M.; Tarantola, S. *Global Sensitivity Analysis: The Primer*; John Wiley & Sons, 2008.
- (20) Herman, J. D.; Kollat, J. B.; Reed, P. M.; Wagener, T. Technical note: method of Morris effectively reduces the computational demands of global sensitivity analysis for distributed watershed models. *Hydrol. Earth Syst. Sci.* **2013**, *17*, 2893–2903.
- (21) Nossent, J.; Elsen, P.; Bauwens, W. Sobol' sensitivity analysis of a complex environmental model. *Environ. Modell. Software* **2011**, *26*, 1515–1525.
- (22) Campolongo, F.; Saltelli, A.; Cariboni, J. From screening to quantitative sensitivity analysis. A unified approach. *Comput. Phys. Commun.* **2011**, *182*, 978–988.
- (23) De Lozzo, M.; Marrel, A. New improvements in the use of dependence measures for sensitivity analysis and screening. *J. Stat. Comput. Simul.* **2016**, *86*, 3038–3058.
- (24) Da Veiga, S. Global sensitivity analysis with dependence measures. *J. Stat. Comput. Simul.* **2015**, *85*, 1283–1305.
- (25) Gretton, A.; Bousquet, O.; Smola, A.; Schölkopf, B. Measuring Statistical Dependence with Hilbert-Schmidt Norms. In *Algorithmic Learning Theory: Proceedings of the 16th International Conference, ALT 2005*, Singapore, October 2005; Jain, S., Simon, H. U., Tomita, E., Eds.; Springer: Berlin, 2005; pp 63–77.
- (26) Yamada, M.; Tang, J.; Lugo-Martinez, J.; Hodzic, E.; Shrestha, R.; Saha, A.; Ouyang, H.; Yin, D.; Mamitsuka, H.; Sahinalp, C.; Radivojac, P.; Menczer, F.; Chang, Y. Ultra high-dimensional nonlinear feature selection for big biological data. *IEEE Trans. Knowl. Data Eng.* **2018**, *30*, 1352–1365.
- (27) Yamada, M.; Jitkrittum, W.; Sigal, L.; Xing, E. P.; Sugiyama, M. High-dimensional feature selection by feature-wise kernelized lasso. *Neural Comput.* **2014**, *26*, 185–207.
- (28) Song, L.; Smola, A.; Gretton, A.; Bedo, J.; Borgwardt, K. Feature selection via dependence maximization. *J. Mach. Learn. Res.* **2012**, *13*, 1393–1434.
- (29) Climente-González, H.; Azencott, C.-A.; Kaski, S.; Yamada, M. Block HSI lasso: model-free biomarker detection for ultra-high dimensional data. *Bioinformatics* **2019**, *35*, i427–i435.
- (30) Feng, L.; Di, T.; Zhang, Y. HSI-based kernel independent component analysis for fault monitoring. *Chemom. Intell. Lab. Syst.* **2018**, *178*, 47–55.
- (31) Qian, J.; Lu, M.; Tian, F.; Liu, R. Study on sensor array optimization of medical electronic nose for wound infection detection. *IEEE Trans. Circuits Syst. II: Express Briefs* **2022**, *69*, 1867–1871.
- (32) Li, Y.; Pogodin, R.; Sutherland, D. J.; Gretton, A. Self-Supervised Learning with Kernel Dependence Maximization. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*; Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P. S., Vaughan, J. W., Eds.; Curran Associates, 2021; pp 15543–15556.
- (33) Wang, T.; Li, W. Kernel learning and optimization with Hilbert–Schmidt independence criterion. *Int. J. Mach. Learn. Cybern.* **2018**, *9*, 1707–1717.
- (34) Kaymak, M. C.; Rahnamoun, A.; O'Hearn, K. A.; van Duin, A. C. T.; Merz, K. M. J.; Aktulga, H. M. JAX-ReaxFF: a gradient-based framework for fast optimization of reactive force fields. *J. Chem. Theory Comput.* **2022**, *18*, 5181–5194.
- (35) Pahari, P.; Chaturvedi, S. Determination of best-fit potential parameters for a reactive force field using a genetic algorithm. *J. Mol. Model.* **2012**, *18*, 1049–1061.
- (36) Trnka, T.; Tvaroška, I.; Koča, J. Automated training of ReaxFF reactive force fields for energetics of enzymatic reactions. *J. Chem. Theory Comput.* **2018**, *14*, 291–302.
- (37) Jaramillo-Botero, A.; Naserifar, S.; Goddard, W. A. General multiobjective force field optimization framework, with application to reactive force fields for silicon carbide. *J. Chem. Theory Comput.* **2014**, *10*, 1426–1439.
- (38) Boothroyd, S.; Wang, L.-P.; Mobley, D. L.; Chodera, J. D.; Shirts, M. R. Open force field evaluator: an automated, efficient, and scalable framework for the estimation of physical properties from molecular simulation. *J. Chem. Theory Comput.* **2022**, *18*, 3566–3576.
- (39) Jónsson, E. O.; Rasti, S.; Galynska, M.; Meyer, J.; Jónsson, H. Transferable potential function for flexible H<sub>2</sub>O molecules based on the single-center multipole expansion. *J. Chem. Theory Comput.* **2022**, *18*, 7528–7543.
- (40) Mondal, A.; Young, J. M.; Barckholtz, T. A.; Kiss, G.; Koziol, L.; Panagiotopoulos, A. Z. Genetic algorithm driven force field parameterization for molten alkali-metal carbonate and hydroxide salts. *J. Chem. Theory Comput.* **2020**, *16*, 5736–5746.
- (41) Verstraelen, T.; Bultinck, P.; van Speybroeck, V.; Ayers, P. W.; van Neck, D.; Waroquier, M. The significance of parameters in charge equilibration models. *J. Chem. Theory Comput.* **2011**, *7*, 1750–1764.
- (42) He, Z.; Xiong, X.; Yang, B.; Li, H. Aerodynamic optimization of a high-speed train head shape using an advanced hybrid surrogate-based nonlinear model representation method. *Optim. Eng.* **2022**, *23*, 59–84.
- (43) Wang, T.; Wang, C.; Xu, Z.; Cui, C.; Wang, X.; Demitrack, Z.; Dai, Z.; Bagtzoglou, A.; Stuber, M. D.; Li, B. Precise control of water and wastewater treatment systems with non-ideal heterogeneous mixing models and high-fidelity sensing. *Chem. Eng. J.* **2022**, *430*, 132819.
- (44) Zhang, T.; Di, X.; Chen, G.; Zhu, L. Parameterization of a COMPASS force field for single layer blue phosphorene. *Nanotechnology* **2020**, *31*, 145702.
- (45) Gopalakrishnan, S.; Dash, S.; Maranas, C. K-FIT: an accelerated kinetic parameterization algorithm using steady-state fluxomic data. *Metab. Eng.* **2020**, *61*, 197–205.
- (46) Mohammed, I.; Gnedin, N. Y. Baryonic effects in cosmic shear tomography: PCA parameterization and the importance of extreme baryonic models. *Astrophys. J.* **2018**, *863*, 173.
- (47) Quinn, K. N.; Wilber, H.; Townsend, A.; Sethna, J. P. Chebyshev approximation and the global geometry of model predictions. *Phys. Rev. Lett.* **2019**, *122*, 158302.
- (48) Mortier, W. J.; Ghosh, S. K.; Shankar, S. Electronegativity-equalization method for the calculation of atomic charges in molecules. *J. Am. Chem. Soc.* **1986**, *108*, 4315–4320.
- (49) Verstraelen, T.; Ayers, P. W.; van Speybroeck, V.; Waroquier, M. ACKS2: atom-condensed Kohn–Sham DFT approximated to second order. *J. Chem. Phys.* **2013**, *138*, 074108.
- (50) Müller, J.; Hartke, B. ReaxFF reactive force field for disulfide mechanochemistry, fitted to multireference ab initio data. *J. Chem. Theory Comput.* **2016**, *12*, 3913–3925.
- (51) Dloczik, L.; Engelhardt, R.; Ernst, K.; Fiechter, S.; Sieber, I.; Könenkamp, R. Hexagonal nanotubes of ZnS by chemical conversion of monocrystalline ZnO columns. *Appl. Phys. Lett.* **2001**, *78*, 3687.
- (52) Hamad, S.; Cristol, S.; Catlow, C. R. A. Surface structures and crystal morphology of ZnS: computational study. *J. Phys. Chem. B* **2002**, *106*, 11002–11008.
- (53) Li, S.; Wei, X.; Zhu, S.; Gui, Y. First principles analysis of SO<sub>2</sub>, H<sub>2</sub>S adsorbed on Fe–ZnS surface. *Sens. Actuators, A* **2021**, *329*, 112827.
- (54) Raymand, D.; van Duin, A. C.; Spangberg, D.; Goddard, W. A.; Hermansson, K. Water adsorption on stepped ZnO surfaces from MD simulation. *Surf. Sci.* **2010**, *604*, 741–752.
- (55) Qi, G.; Zhang, L.; Yuan, Z. Improved H<sub>2</sub>S gas sensing properties of ZnO nanorods decorated by a several nm ZnS thin layer. *Phys. Chem. Chem. Phys.* **2014**, *16*, 13434–13439.
- (56) *Force Field Format Specification*; Software for Chemistry and Materials, 2022; [https://www.scm.com/doc/ReaxFF/ffield\\_descr.html](https://www.scm.com/doc/ReaxFF/ffield_descr.html) (accessed June 2022).
- (57) Islam, M. M.; Ostadhossain, A.; Borodin, O.; Yeates, A. T.; Tipton, W. W.; Hennig, R. G.; Kumar, N.; van Duin, A. C. T. ReaxFF molecular dynamics simulations on lithiated sulfur cathode materials. *Phys. Chem. Chem. Phys.* **2015**, *17*, 3383–3393.
- (58) *ReaxFF: SCM Developer Documentation*; Software for Chemistry and Materials, 2022; [https://www.scm.com/doc/ReaxFF/ffield\\_descr.html#equation-reference](https://www.scm.com/doc/ReaxFF/ffield_descr.html#equation-reference) (accessed June 2022).
- (59) Verstraelen, T.; Vandenbrande, S.; Heidar-Zadeh, F.; Vanduyfhuys, L.; van Speybroeck, V.; Waroquier, M.; Ayers, P. W. Minimal basis iterative stockholder: atoms in molecules for force-field development. *J. Chem. Theory Comput.* **2016**, *12*, 3894–3912.

- (60) Hansen, N.; Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evol. Comput.* **2001**, *9*, 159–195.
- (61) Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; Smola, A. A kernel two-sample test. *J. Mach. Learn. Res.* **2012**, *13*, 723–773.
- (62) Zhang, H.; Huang, F.; Gilbert, B.; Banfield, J. F. Molecular dynamics simulations, thermodynamic analysis, and experimental study of phase stability of zinc sulfide nanoparticles. *J. Phys. Chem. B* **2003**, *107*, 13051–13060.
- (63) Kennard, R. W.; Stone, L. A. Computer aided design of experiments. *Technometrics* **1969**, *11*, 137–148.