

# Gaussian Process Models of Potential Energy Surfaces with Boundary Optimisation

Jack Broad,<sup>1, a)</sup> Simon Preston,<sup>1</sup> Richard J. Wheatley,<sup>2, b)</sup> and Richard S. Graham<sup>1</sup>

<sup>1</sup>*School of Mathematical Sciences, University of Nottingham, Nottingham NG7 2RD, United Kingdom*

<sup>2</sup>*School of Chemistry, University of Nottingham, Nottingham NG7 2RD, United Kingdom*

(Dated: 27 September 2021)

A strategy is outlined to reduce the number of training points required to model intermolecular potentials using Gaussian processes, without reducing accuracy. An asymptotic function is used at long range and the cross-over distance between this model and the Gaussian process is learnt from the training data. Results are presented for different implementations of this procedure, known as boundary optimisation, across the following dimer systems: CO-Ne, HF-Ne, HF-Na<sup>+</sup>, CO<sub>2</sub>-Ne and (CO<sub>2</sub>)<sub>2</sub>. The technique reduces the number of training points, at fixed accuracy, by up to  $\sim 49\%$ , compared to our previous work based on a sequential learning technique. The approach is readily transferable to other statistical methods of prediction or modelling problems.

## I. INTRODUCTION

In any molecular simulation, approximations of the potential energy surfaces (PESs) that describe the relevant interactions are a pre-requisite. Traditionally, these approximations have been made using empirical potentials (force-fields)<sup>1,2</sup>. However, such force-fields have closed functional forms, which limit their capacity to capture the complicated topography of the PES. Furthermore, they are laborious to produce and may fail to capture accurately even the fitted data. As the accuracy of the simulation depends on the potential employed, much work has been devoted to developing force-fields that provide approximations of PESs with quantum-mechanical accuracy.

Such work includes methods to generate *ab initio* force fields<sup>3</sup> and attempts to ‘learn’ the potential via a machine learning algorithm<sup>4,5</sup>. This article concerns the latter approach, which has also been applied in other fields in chemistry and materials science<sup>6–8</sup>. This approach proceeds by training a statistical technique on a relatively small set of data from *ab initio* calculations on the PES of interest, known as the training set. Many such techniques have been employed to predict the energy in these algorithms, including neural networks<sup>9–13</sup>, moment tensors<sup>14–16</sup> and Gaussian processes<sup>17–34</sup> (GPs).

Herein, GPs are employed as the statistical technique. An existing example of a force-field that uses GPs is FFLUX<sup>35</sup>, which has been used to approximate the energies of weakly bound complexes<sup>36</sup> and water clusters<sup>37</sup>, among other applications<sup>38–40</sup>. Furthermore, in the field of materials science, GP models that invoke a smooth overlap of atomic positions (SOAP) kernel<sup>41</sup> have seen successful applications to many systems<sup>29–31,42–45</sup>.

GP models have also produced promising results in applications to intermolecular interactions<sup>17–19</sup>. Initially<sup>17</sup>, the train-

ing sets for these models were constructed with Latin hypercube sampling<sup>46–48</sup> though sequential design strategies<sup>49,50</sup>, which achieve a prescribed accuracy with a smaller training set, have been shown to outperform such methods<sup>19</sup>. Regardless of training set design, a fixed cross-over distance,  $R_{\text{cross}}$ , was imposed prior to training. This parameter defines a boundary beyond which a simple, long-range asymptotic function takes over prediction from the GP<sup>17,19</sup>. A similar approach is used in materials science, in which the contribution by one atom to the neighbour density of another is assumed to be zero if the two are separated by a distance in excess of  $R_{\text{cross}}$ <sup>31,44</sup>. Cross-over distances have been applied alongside other statistical methods of prediction too<sup>5,51</sup>, with neural network<sup>9–12</sup> and moment tensor<sup>15</sup> models of PESs employing fixed, pre-determined cross-over distances.

Use of a cross-over distance limits the portion of the PES approximated by the statistical method to the region around the potential well. Consequently, fewer training points are required to develop an accurate model. Previously these distances were all fixed and specified *a priori*, however, they can be learnt from the training data. Herein the work of Uteva *et al.*<sup>17,19</sup> is extended into a sequential design strategy in which  $R_{\text{cross}}$  varies as a function of the number of training points and the species of the interacting atoms. This produces models of the same accuracy with fewer training points than the same design strategy with a fixed  $R_{\text{cross}}$ . The process by which  $R_{\text{cross}}$  is optimised is referred to here as boundary optimisation.

## A. GP Regression

All GP models herein make predictions via GP regression, with detailed descriptions of GP regression theory available elsewhere<sup>52–54</sup>. This section briefly introduces the key concepts in the context of intermolecular potentials. For any intermolecular interaction the PES can be thought of as a multivariate function  $f(\mathbf{x})$ ,  $\mathbf{x} \in \mathbb{R}^Z$ , where  $\mathbf{x}$  is a vector of inputs and  $Z$  is the number of elements in  $\mathbf{x}$ . Here the inputs are inverse interatomic separations, though promising results have also been obtained with Morse variables<sup>24,55</sup>. Only pair intermolecular interactions are considered, although GP models have been

<sup>a)</sup>Electronic mail: jack.broad@nottingham.ac.uk (author to whom correspondence should be addressed)

<sup>b)</sup>Electronic mail: richard.wheatley@nottingham.ac.uk

applied successfully to non-additive interactions<sup>17,19,40</sup> and the methodology outlined here can be extended to such cases straightforwardly.

When the outputs  $f(\mathbf{x}_i) = Y_i$  are available at  $N$  values of  $i$ , the set  $\{\mathbf{x}_i, Y_i\}_{i=1}^N$  forms the training set. Letting  $\mathbf{Y}$  be a vector of the observed energies from this set, GP regression can be used to approximate the value of  $f(\mathbf{x})$  at a new point  $\mathbf{x}_*$  as

$$f(\mathbf{x}_*) = \mathbf{K}_*^T \mathbf{K}^{-1} \mathbf{Y}, \quad (1)$$

where  $\mathbf{K}_*$  is a vector of the covariances between  $\mathbf{x}_*$  and all  $\mathbf{x}_i$ ,

$$\mathbf{K}_* = [k(\mathbf{x}_*, \mathbf{x}_1) \quad k(\mathbf{x}_*, \mathbf{x}_2) \quad \cdots \quad k(\mathbf{x}_*, \mathbf{x}_N)], \quad (2)$$

and  $\mathbf{K}$  is the positive-definite covariance matrix,

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) + \sigma_n^2 & k(\mathbf{x}_1, \mathbf{x}_2) & \cdots & k(\mathbf{x}_1, \mathbf{x}_N) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) + \sigma_n^2 & \cdots & k(\mathbf{x}_2, \mathbf{x}_N) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_N, \mathbf{x}_1) & k(\mathbf{x}_N, \mathbf{x}_2) & \cdots & k(\mathbf{x}_N, \mathbf{x}_N) + \sigma_n^2 \end{bmatrix}.$$

Here  $\sigma_n^2$  is the Gaussian noise variance, which accounts for the noise in the training data.

All entries in  $\mathbf{K}$  are found by evaluating a covariance function,  $k(\mathbf{x}, \mathbf{x}')$ , where  $\mathbf{x}$  and  $\mathbf{x}'$  are configurations from the training set. The same holds for  $\mathbf{K}_*$ , only with one training configuration replaced by the new point  $\mathbf{x}_*$ . A common example is the squared exponential covariance function, for which

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \prod_{d=1}^D \exp\left(-\frac{(x_d - x'_d)^2}{2l_d^2}\right). \quad (3)$$

Here  $D$  is the total number of interatomic distances which comprise unique atomic pairs,  $\sigma_f^2$  is the signal variance,  $x_d$  and  $x'_d$  are the inverse separations in  $\mathbf{x}$  and  $\mathbf{x}'$  which contain the  $d$ th atomic pair, and  $l_d$  is the lengthscale for the interaction between this pair of atoms. The hyperparameters of the covariance function are  $\sigma_n^2$ ,  $\sigma_f^2$  and the set of lengthscales,  $l_d$ .

This covariance function can be modified to account for the fact that two configurations can have equivalent energies due to symmetry. The result is the symmetric squared exponential covariance function,  $k_{\text{sym}}(\mathbf{x}, \mathbf{x}')$ , investigated by Uteva *et al.*<sup>17</sup>, which is used here. The set of all permutations of the interatomic distances in  $\mathbf{x}$  under which the energy surface is unchanged is denoted as  $P$ , while  $p$  is a single permutation within this set. Assuming that  $l_d$  is invariant under interchange of  $\mathbf{x}_d$  and  $\mathbf{x}'_d$ ,

$$k_{\text{sym}}(\mathbf{x}, \mathbf{x}') = \sum_{p \in P} k(p\mathbf{x}, \mathbf{x}'), \quad (4)$$

where  $k(\mathbf{x}, \mathbf{x}')$  is the squared exponential covariance function from equation 3. GPs also employ a mean function, which is taken here to be zero everywhere.

To specify the hyperparameters of  $k_{\text{sym}}(\mathbf{x}, \mathbf{x}')$  that produce a model that achieves the best fit to the training data, the log of the marginal likelihood function,  $\log(\mathcal{L})$ , is maximised<sup>52,54</sup>

$$\log(\mathcal{L}) = -\frac{1}{2} \mathbf{Y}^T \mathbf{K}^{-1} \mathbf{Y} - \frac{1}{2} \log|\mathbf{K}| - \frac{N}{2} \log(2\pi). \quad (5)$$

Equation 5 shows that this optimisation entails inversion of the  $N \times N$  covariance matrix. This incurs a sizable computational cost, which translates to scaling of order  $\mathcal{O}(N^3)$  for hyperparameter optimisation<sup>56</sup>. Consequently, though they achieve higher predictive accuracies than other statistical methods when modelling PESs<sup>51,57,58</sup>, GPs require greater computational effort to train<sup>51</sup>. For prediction, meanwhile,  $\mathbf{K}^{-1} \mathbf{Y}$  in equation 1 needs to be calculated once only, and as a result the cost scales linearly with  $N$ . Even so, molecular simulations that employ a GP model for which  $N$  is large are more computationally intensive than those that use a traditional force-field.

These issues have led to attempts to minimise the training set size required by a GP model to achieve a given error. For two training strategies that achieve the same error, that which does so with fewer training points is more computationally efficient. Attempts to develop more computationally efficient training strategies have involved active learning or sequential design methods<sup>19,29,59</sup>, composite kernels<sup>24</sup> and new sampling methods<sup>27,60</sup>.

However, no attempt was made in these previous works to increase training efficiency by learning the optimal cross-over distance from the data used in training. Our boundary optimisation approach has the potential to increase the gains in efficiency from these prior methods further still with little associated increase in computational time, provided a viable long-range approximation is available. Such an improvement is shown here, building upon the sequential design method of Uteva *et al.*<sup>19</sup>.

## B. Training Set Design

It was found by Uteva *et al.*<sup>17</sup> that the predictive performance of a GP for a PES is enhanced by using inverse interatomic separations rather than non-inverse separations to describe intermolecular configurations. This is because  $k_{\text{sym}}(\mathbf{x}, \mathbf{x}')$  is a stationary kernel (*i.e.* it relies only on the distance between the two configurations being compared), meaning it assumes that the rate of change of the output (*i.e.* the interaction energy) with the input (the interatomic distances) is constant. This is not the case for a PES, where the energy varies rapidly with separation in the short-range repulsive wall but barely at all in the long-range asymptotic region. The  $r \rightarrow r^{-1}$  conversion addresses this issue: for low  $r$  the rate of change in  $r^{-1}$  is faster than in  $r$ , meaning the change in the input more closely matches the change in the output under the former; meanwhile, for large  $r$ , the rate of change in  $r^{-1}$  is far smaller, as is that of the energy.

Previously, Latin hypercube (LHC) sampling, which is a space-filling design, has been employed to build data sets when modelling PESs with GPs<sup>17-19,61</sup>. The approach of Uteva *et al.*<sup>17</sup> entailed designing LHCs in inverse separations over a range of angles that specified the symmetry-distinct region<sup>17</sup>. The algorithm generates a large number of candidate LHCs and finds the minimum separation in each. The LHC with the largest such separation is selected under what is referred to here as the ‘maximin’ criterion. Quantum chemical

calculations are undertaken on the selected LHC only. This LHC is then subject to a high energy cut-off,  $E_{\text{cut}}$ , and any configuration for which the interaction energy exceeds  $E_{\text{cut}}$  is discarded from the LHC<sup>17</sup>. In addition, a geometric constraint of 8.5 Å was placed on all LHCs to ensure no configurations with minimum separations above this threshold were included<sup>17</sup>. It was observed that at separations greater than the geometric constraint the GP predictions tended towards a small, non-zero constant<sup>17</sup>. As these predictions should tend towards zero, a long-range function derived from the multipole expansion<sup>62</sup> was introduced to approximate energies at large separations.

In a later work, Uteva *et al.*<sup>19</sup> presented active and sequential learning approaches, in which the model of the PES is progressively refined by adding new points to the GP. In this approach, training and validation of the GP models involves three data sets: a training, a reference and a test set. The training set is used in the GP regression. The reference set provides a pool of configurations from which new training points can be selected, while the test set determines the GP's accuracy against an independent data set.

One sequential design strategy presented by Uteva *et al.*<sup>19</sup> is the highest error search. This approach selects new training points based on the configuration in the reference set with the largest GP prediction error. We adopt this approach herein. However, Uteva *et al.*<sup>19</sup> applied this strategy only at separations below the fixed, pre-determined cross-over distance<sup>19</sup>, meaning no new training data could be added at separations above  $R_{\text{cross}}$ .

The highest error search is best described as a sequential design rather than an active learning<sup>63,64</sup> method. This is because active learning methods are a subset of sequential design methods that compute the output corresponding to an input only at the point it is added to the training set. This means, in the context of modelling PESs, an active learning method would use a reference set comprising configurations for which the energies were not calculated. Active learning strategies have, however, been successfully used to develop GP<sup>19,29,59</sup>, moment tensor<sup>15</sup> and neural network<sup>13</sup> models of PESs, and the methodology used here for boundary optimisation could be altered for use with the two set method of Uteva *et al.*<sup>19</sup>, which is an active learning method.

When using a sequential design method, meanwhile, the energies in the reference and test sets are pre-computed using a relatively computationally inexpensive *ab initio* technique. Once a model with the requisite predictive accuracy is obtained, the energies in the minimal training set can be recalculated using a more accurate, and costly, *ab initio* method before use in applications. This process is known as transfer learning<sup>65,66</sup> and it allows a GP model of a PES to be produced with the accuracy of a high-level *ab initio* technique with relatively few computationally expensive calculations. For example, MP2<sup>67</sup> energies have been upgraded to CCSD(T)<sup>68</sup> energies for calculation of the CO<sub>2</sub>-CO second virial coefficient<sup>17</sup>.

TABLE I: The co-ordinates for the reference and test LHCs for each system.  $N_{\text{ref}}$  and  $N_{\text{test}}$  are the number of points in the reference and test sets respectively after application of the high-energy cut-off, while the maximum number of training points for models of each potential are given in the text. Also shown is the minimum energy across the reference and test sets,  $E_{\text{min}}$ , in Hartrees ( $E_h$ ), though no attempt was made to approximate the global minimum energy for any potential.

| System                          | Coordinate       | Range                        | $N_{\text{ref}}$ | $N_{\text{test}}$ | $E_{\text{min}}$          |
|---------------------------------|------------------|------------------------------|------------------|-------------------|---------------------------|
| CO-Ne                           | $r^{-1}$         | 0.01 to 0.67 Å <sup>-1</sup> | 1914             | 5718              | -1.502 x 10 <sup>-4</sup> |
|                                 | $\cos(\theta)$   | -1 to 1                      |                  |                   |                           |
| HF-Ne                           | $r^{-1}$         | 0.01 to 0.67 Å <sup>-1</sup> | 2148             | 6468              | -2.633 x 10 <sup>-4</sup> |
|                                 | $\cos(\theta)$   | -1 to 1                      |                  |                   |                           |
| HF-Na <sup>+</sup>              | $r^{-1}$         | 0.01 to 0.67 Å <sup>-1</sup> | 2760             | 8416              | -2.518 x 10 <sup>-2</sup> |
|                                 | $\cos(\theta)$   | -1 to 1                      |                  |                   |                           |
| CO <sub>2</sub> -Ne             | $r^{-1}$         | 0.01 to 0.67 Å <sup>-1</sup> | 5057             | 5072              | -2.895 x 10 <sup>-3</sup> |
|                                 | $\cos(\theta)$   | 0 to 1                       |                  |                   |                           |
| (CO <sub>2</sub> ) <sub>2</sub> | $r^{-1}$         | 0.01 to 0.67 Å <sup>-1</sup> | 5810             | 5837              | -1.975 x 10 <sup>-3</sup> |
|                                 | $\cos(\theta_1)$ | 0 to 1                       |                  |                   |                           |
|                                 | $\cos(\theta_2)$ | 0 to 1                       |                  |                   |                           |
|                                 | $\phi$           | 0 to 180°                    |                  |                   |                           |

## II. METHODOLOGY

### A. Overview

In previous applications of both LHC learning<sup>17</sup> and sequential design methods<sup>19</sup> a fixed value of the cross-over distance,  $R_{\text{cross}}$ , was determined *a priori*. Rather than fix  $R_{\text{cross}}$  prior to training, however, it is possible to determine its optimal value from the reference data. This is boundary optimisation and is possible because a sequential design method permits  $R_{\text{cross}}$  to be varied each time the GP is updated. Consequently, the size of the region over which GP regression is used for prediction will grow with the predictive accuracy of the GP.

Boundary optimisation may lead to increased accuracy for the following reason. When the number of training points,  $N_{\text{TP}}$ , is low the predictive accuracy of the long-range function at the outer edge of the potential well will exceed that of the GP. Consequently, it is anticipated that allowing the long-range function to approximate the energies of configurations in this region at low  $N_{\text{TP}}$  will increase the predictive accuracy of the overall model and facilitate more efficient model development.

## B. Data Set Generation

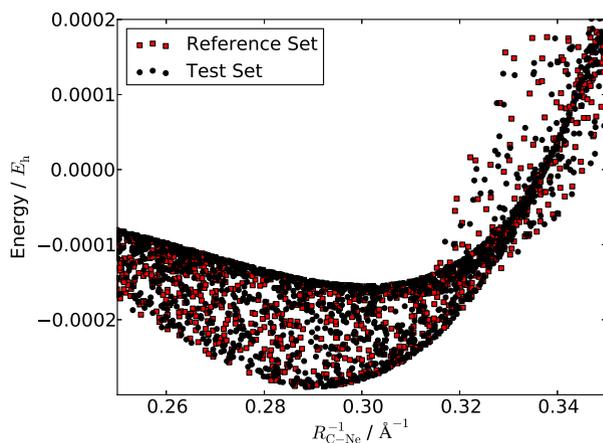


FIG. 1: The calculated energies in the reference set (red) and the test set (black) of the CO<sub>2</sub>-Ne potential for separations between 2.85 Å and 4 Å. Both sets contain ~ 5000 configurations (see table I).

Both the reference and test sets for each chemical system were generated under the LHC design strategy discussed in IB. All energies were calculated in Molpro<sup>69</sup> using MP2 with an aug-cc-pVTZ basis set and the counterpoise correction with the only exception being HF-Na<sup>+</sup>, which used an aug-cc-pwCVTZ basis set instead. The specifications of the LHCs for each system are given in table I. Therein  $r$  is the distance between the bond centres (not centres of mass),  $\theta$  is the angle between  $r$  and the bond axis of the molecule,  $\theta_1$  and  $\theta_2$  are the angles between  $r$  and the bond axis of the first and second CO<sub>2</sub> molecules respectively, and  $\phi$  is the torsional angle between the two CO<sub>2</sub> molecules. The molecules were kept rigid, with  $r_{\text{CO}} = 1.1283$  Å for CO,  $r_{\text{CO}} = 1.1632$  Å for CO<sub>2</sub> and  $r_{\text{HF}} = 0.9170$  Å. However, boundary optimisation could be applied to non-rigid systems straightforwardly, with the only additional requirement being generalisation of the long-range function to non-rigid molecules. A larger geometric constraint of 100 Å (instead of 8.5 Å<sup>17</sup>) was employed to probe the long range behaviour of the system. The maximum value of  $N_{\text{TP}}$  was 100 for all systems apart from (CO<sub>2</sub>)<sub>2</sub>, which used 300 training points at most.

A high energy cut-off,  $E_{\text{cut}}$ , was applied to the reference and test sets to remove configurations with interaction energies which exceeded its value.  $E_{\text{cut}} = 0.005 E_h$  for all systems apart from HF-Na<sup>+</sup>, for which  $E_{\text{cut}} = 0.05 E_h$  because the charge-dipole interaction increases the well-depth. ( $1 E_h \approx 27.211$  eV  $\approx 2625.5$  kJ mol<sup>-1</sup>.)

For systems with  $N_{\text{test}} \gg N_{\text{ref}}$  the independence of the test set is self-evident. However, for systems where  $N_{\text{test}} \approx N_{\text{ref}}$  it is also true that the test set is independent. This follows because, although the reference and test sets for each system were designed using the same ‘maximin’ strategy, the stochastic nature of the LHC algorithm means that separate LHCs contain completely independent sets of configurations. Fur-

TABLE II: The properties included in the multipolar long-range functions of each system, as well as the *ab initio* methods used in their calculation. All calculations were carried out using an aug-cc-pVQZ basis set apart from those for the dispersion coefficients, which used an aug-cc-pVTZ basis set.

| System             | Dipole                | Quadrupole | Polarizability | Dispersion |
|--------------------|-----------------------|------------|----------------|------------|
| CO-Ne              | ✓                     | X          | ✓              | ✓          |
| HF-Ne              | ✓                     | X          | ✓              | ✓          |
| HF-Na <sup>+</sup> | ✓                     | ✓          | ✓              | ✓          |
| Level of Theory    | MRCI <sup>71,72</sup> | MRCI       | MP2            | CCSD       |

thermore, the ‘maximin’ criterion is based on just one separation in the whole data set, meaning that two LHCs with similar maximin will still be dissimilar. This is demonstrated in figure 1, which shows the energies against the inverse C-Ne separation for the reference and test sets used in training models of the CO<sub>2</sub>-Ne potential.

## C. Long-range asymptotic functions

For the long-range energy model, multipole series were employed for all systems apart from HF-Na<sup>+</sup>. The contributions included in the multipolar long-range functions are shown in table II for each system apart from CO<sub>2</sub>-Ne, for which the function is already described in previous work<sup>17</sup>, and (CO<sub>2</sub>)<sub>2</sub>. The latter was developed prior to the other multipolar long-range functions and uses atomistic charge, dipole, quadrupole and polarizability contributions from Hartree-Fock<sup>70</sup> theory, which were scaled to give the known total molecular properties. For HF-Na<sup>+</sup> a fitted long-range function was used, which was derived by fitting a sum of two power laws between two points where the energy was predicted by GP regression. More information on the motivation for and derivation of this function is found in appendix A.

## D. Classification of Phase Space using a Boundary

When modelling PESs using GPs it is often necessary to classify configurations as suitable for prediction via the GP or via a long-range asymptotic function. In the work of Uteva *et al.*<sup>17,19</sup> the classifier formed a boundary from the superposition of atom-centred spheres defined by a single constant,  $R_{\text{cross}}$ . Specifically, if any interatomic distance was less than  $R_{\text{cross}}$  the GP was used. As  $R_{\text{cross}}$  was fixed, this classifier is referred to here as  $C_{\text{fixed}}$ . Denoting the region in which GP regression was used for prediction as  $A_{\text{GP}}$  and the region which employed the long-range function as  $A_{\text{LR}}$ , under  $C_{\text{fixed}}$  these regions were

$$A_{\text{GP}} = \{\mathbf{r} : \min(\mathbf{r}) \leq R_{\text{cross}}\} \quad (6)$$

and

$$A_{\text{LR}} = \{\mathbf{r} : \min(\mathbf{r}) > R_{\text{cross}}\}, \quad (7)$$

where  $\mathbf{r}$  is a set of intermolecular atom-atom distances, and  $\min(\mathbf{r})$  is the smallest separation in  $\mathbf{r}$ .

In boundary optimisation,  $R_{\text{cross}}$  varies according to the GP accuracy. Thus  $R_{\text{cross}}$  is not constant but a model parameter, which is learnt from the reference data. As this classifier is still parameterised by a single value,  $R_{\text{cross}}$ , it is referred to here as  $C_{\text{single}}$ .

More elaborate classifiers are possible by using more detailed parametric forms to define the boundary region. A simple way of defining a more complex classifier is for the value of  $R_{\text{cross}}$  to depend on the atom types that comprise the interatomic distance. The resulting classifier is referred to here as  $C_{\text{multi}}$ . For a system of molecules with  $D$  interatomic pairs of chemically different atoms, using  $C_{\text{multi}}$  requires the vector of cross-over distances  $\mathbf{R}_{\text{cross}} = (R_1, \dots, R_d, \dots, R_D)$ . This defines a multiple-parameter boundary region, as follows,

$$C_{\text{multi}}(\mathbf{r}) = \begin{cases} A_{\text{GP}}, & \text{if } \min_d(\mathbf{r}) \leq R_d \text{ for any } d \\ A_{\text{LR}}, & \text{if } \min_d(\mathbf{r}) > R_d \text{ for all } d, \end{cases} \quad (8)$$

where  $\min_d(\mathbf{r})$  is the minimum separation in  $\mathbf{r}$  that involves an atomic pair of type  $d$ .

The optimum value of the classifier parameters are determined by minimising the error between the model and the reference set (*i.e.* by minimising the training error), meaning the sizes of  $A_{\text{GP}}$  and  $A_{\text{LR}}$  vary with the GP. The sum of squared errors,  $\text{SSE}_{\text{tot}}$ , over the two regions is

$$\text{SSE}_{\text{tot}} = \text{SSE}_{\text{GP}} + \text{SSE}_{\text{LR}}, \quad (9)$$

where

$$\text{SSE}_{\text{method}} = \sum_{i=1}^{N_{\text{method}}} (\hat{Y}_i - Y_i)^2. \quad (10)$$

Here, "method" denotes either GP or LR,  $N_{\text{method}}$  the number of points in  $A_{\text{method}}$ ,  $\hat{Y}_i$  the prediction of the energy for the  $i$ th configuration from the desired method and  $Y_i$  the calculated energy of the same configuration. The RMSE against the test set,  $\text{RMSE}_{\text{test}}$ , is given by

$$\text{RMSE}_{\text{test}} = \left( \frac{\text{SSE}_{\text{test}}}{N_{\text{test}}} \right)^{\frac{1}{2}}, \quad (11)$$

where  $\text{SSE}_{\text{test}}$  is  $\text{SSE}_{\text{tot}}$  over the test set and  $N_{\text{test}}$  is the number of configurations in the test set.  $\text{RMSE}_{\text{test}}$  is therefore a function of  $R_{\text{cross}}$  with discrete steps, as the RMSE changes only when a change in  $R_{\text{cross}}$  causes a configuration in the test set to be re-classified. Both  $C_{\text{single}}$  and  $C_{\text{multi}}$  are simple, parametric classifiers that pre-impose a mathematical form on the classification. Hence, neither is expected to be optimal with respect to the RMSE against the reference or test sets. That is, a more complicated boundary than that described by these classifiers will likely produce a lower RMSE against a given data set. However, it is shown later that an artificial 'ideal' classifier produces only very marginal improvements over  $C_{\text{multi}}$ , suggesting this classifier provides a very good balance of simplicity and accuracy.

## E. Overall algorithm

Using the data and calculation methods above, the algorithm generates a GP model sequentially as follows: train the GP to the current training set; select the classifier parameters by minimising the RMSE against the reference set; move a new point from the reference to the training set based on the largest error. Below are further details of each step, along with how the choice of classifier and placement strategy affects the algorithm.

### 1. GP training

All GPs described herein were trained using the GPy<sup>73</sup> package in Python 2.7. Optimisation of the hyperparameters was carried out by maximising  $\log(\mathcal{L})$  using 20 independent restarts whenever a configuration was added to the training set. Moreover, a gamma distribution with an expectation of one and a variance of two was used as a prior on all hyperparameters for all systems. This was to weakly penalise large hyperparameter values, given that the expected values are typically of order 0.1 or below.

### 2. Direct Search Algorithm

When using  $C_{\text{single}}$ ,  $R_{\text{cross}}$  can be optimised via a direct search that exploits how the RMSE varies as a piecewise constant function of  $R_{\text{cross}}$ . Because of this feature, all possible values of  $\text{SSE}_{\text{tot}}$  (and hence the RMSE) can be computed readily against the reference set. A direct search of these values is guaranteed to find the global minimum. Full details of the direct search are given in algorithm 1.

---

#### Algorithm 1 Direct Search Algorithm

---

- 1: Compute the squared errors for the GP and long-range function and  $\min(\mathbf{r})$  at each configuration in the reference set.
  - 2: Order the configurations from smallest to largest in terms of  $\min(\mathbf{r})$ .
  - 3: Approximate  $\text{SSE}_{\text{tot}}$  initially from the squared errors of the long-range function alone.
  - 4: Iterate through the  $\min(\mathbf{r})$  from the lowest to the highest:
    - a: set  $R_{\text{cross}} = \min(\mathbf{r})$ , which moves a single configuration from  $A_{\text{LR}}$  to  $A_{\text{GP}}$ ;
    - b: update  $\text{SSE}_{\text{tot}}$  by deducting the squared long range error of the moved point from  $\text{SSE}_{\text{LR}}$  and adding its squared GP error to  $\text{SSE}_{\text{GP}}$ ;
    - c: store the new  $\text{SSE}_{\text{tot}}$ .
  - 5: Select whichever value of  $R_{\text{cross}}$  corresponds to the smallest value of  $\text{SSE}_{\text{tot}}$ .
- 

The direct search algorithm is computationally cheap for the following reasons. Each instance of the direct search requires only a single set of predictions from each of the long range function and GP regression. As the long range function is fixed throughout the sequential design process, these predictions need only be calculated once at the start. The GP pre-

dictions need to be re-calculated whenever the GP is updated, which occurs once per step of the sequential design. However, these predictions are already required by a sequential design step when choosing the new training point. Furthermore, calculation of all possible squared error values (step 4) is cheap because calculating each value in order requires only a simple update of  $SSE_{LR}$  and  $SSE_{GP}$ . Consequently, the direct search is fast compared to the other steps of the sequential design algorithm and can therefore be undertaken at each design step with negligible additional computational effort.

### 3. Orthogonal direct search algorithm

An adaption of the direct search algorithm is required for  $C_{\text{multi}}$ , as this requires a multidimensional optimisation of multiple classifier parameters. This is called the orthogonal direct search as it optimises one element of  $\mathbf{R}_{\text{cross}}$  while keeping the others fixed. The single element that varies is optimised using the direct search algorithm, as this is guaranteed to return the best minimum along that 1D slice of  $\mathbf{R}_{\text{cross}}$ . Although orthogonal optimisations can be time-consuming, the speed of the one-dimensional direct search means that repeating it multiple times for all cross-over distances is feasible.

---

#### Algorithm 2 Orthogonal Direct Search Algorithm

---

- 1: Choose limits  $NI_{\text{max}}$  and  $NR_{\text{max}}$  on the number of iterations and restarts respectively (see the text for details).
  - 2: Assemble the array of minimum distances  $\mathbf{M}$ , as follows:
    - a: For every configuration, collect all interatomic distances which comprise the same atoms types into group  $d$  and find  $\min_d(\mathbf{r})$  for each  $d$ .
    - b: Arrange the lists of  $\min_d(\mathbf{r})$  values in an  $N \times D$  array,  $\mathbf{M}$ , where  $N$  is the number of configurations in the data set and  $D$  is the number of unique atomic pairs in  $\mathbf{r}$ .
    - c: For each column in  $\mathbf{M}$ , order the values from smallest to largest to produce an ordered list in each column.
  - 3: For each restart, select a row from  $\mathbf{M}$  at random to be the initial guess at  $\mathbf{R}_{\text{cross}}$ .
  - 4: For each  $d$  in turn, fix all cross-over distances apart from  $R_d$  and find the optimal value of  $R_d$  using a direct search.
  - 5: Repeat step 4 until  $NI = NI_{\text{max}}$  or the elements of  $\mathbf{R}_{\text{cross}}$  remain unchanged; save this  $\mathbf{R}_{\text{cross}}$  and its corresponding  $SSE_{\text{tot}}$ .
  - 6: Repeat steps 3-5 until  $NR = NR_{\text{max}}$ .
  - 7: Select the  $\mathbf{R}_{\text{cross}}$  that corresponds to the lowest value of  $SSE_{\text{tot}}$ .
- 

The orthogonal direct search proceeds via algorithm 2. As in the one-dimensional direct search, the square errors of the long-range and GP predictions at each configuration are pre-computed for this algorithm. The orthogonal search algorithm is not guaranteed to find the global minimum, as local minima may exist in the RMSE landscape. Hence,  $NR_{\text{max}}$  restarts are performed with randomly selected starting points. Each restart involves  $NI_{\text{max}}$  optimisations of each  $R_d$ . Here  $NI_{\text{max}} = 15$  and  $NR_{\text{max}} = 5$  for all systems. If the values in  $\mathbf{R}_{\text{cross}}$  converge such that further one-dimensional searches in any orthogonal direction do not change its elements prior to  $NI = NI_{\text{max}}$ ,  $\mathbf{R}_{\text{cross}}$  is saved and the next restart undertaken. In fact, it was rare that  $NI$  reached  $NI_{\text{max}}$  for any of the systems ex-

plored here. Moreover, despite the low  $NR_{\text{max}}$  employed, the same minimum was usually found across multiple restarts.

The direct search offers a method that is at once very fast, designed for the discrete-stepped surface and perfect in a single dimension. This latter property means it skips over any local minima in a given orthogonal direction. As such, optimal cross-over distances are obtained quickly and reproducibly under this method. Typically, about 80 % of random restarts return the same set of cross-over distances for a given system and a given training set.

### 4. Training Point Placement Methods

Once the classifier parameters are optimized, the PES model is specified and the next training point determined from the highest error method. The point with the greatest error can be selected either from  $A_{GP}$  alone or from the union of  $A_{GP}$  and  $A_{LR}$ . Using  $A_{GP}$  alone is referred to as the constrained placement method, which proceeds via algorithm 3. Steps 2-5 of this algorithm comprise a stage of training, with the  $RMSE_{\text{test}}$  also calculated at each such stage.

---

#### Algorithm 3 Constrained Placement Algorithm

---

- 1: Select the configuration in the reference set with the highest energy; add this configuration to the training set and remove it from the reference set.
  - 2: Retrain the GP to the updated training set.
  - 3: Determine the boundary that minimises the RMSE against the reference set (see subsection IIE 2).
  - 4: Find the configuration in  $A_{GP}$  for which the GP error is highest, where  $A_{GP}$  is defined by the boundary from the previous step.
  - 5: Add this configuration to the training set and remove it from the reference set.
  - 6: Repeat steps 2-5 until the desired RMSE or number of training points is reached.
- 

Choosing new points from the union of  $A_{GP}$  and  $A_{LR}$  is named the open placement method. This proceeded identically to the constrained placement method (algorithm 3) except the highest error point was from either  $A_{GP}$  or  $A_{LR}$ . In each region this point is found using the corresponding method of prediction (*i.e.* the error of the long-range function was used in  $A_{LR}$ ).

Models were trained under both point placement methods because each held potential advantages over the other. The constrained placement method ensures all training points are added in the GP region and so are of immediate use in prediction. Meanwhile, the open placement method is capable of immediately placing points in regions of the PES where the long-range function performed poorly, potentially transferring these regions to  $A_{GP}$  more rapidly than under the constrained placement approach.

### 5. Closest Model Training Strategy

In the present work, all training strategies used are combinations of a classifier and a point placement method. A further strategy, which is intended for comparison only, is the closest model training strategy. This employs the open placement method to select training points and classifies a configuration using  $C_{\text{optimal}}$ , where

$$C_{\text{optimal}}(\mathbf{r}) = \begin{cases} A_{\text{GP}}, & \text{if } SE_{\text{GP}} \leq SE_{\text{LR}} \\ A_{\text{LR}}, & \text{otherwise} \end{cases} \quad (12)$$

Here,  $SE_{\text{method}}$  is the squared error in the prediction from "method" at  $\mathbf{r}$ .  $C_{\text{optimal}}$  is so named because it classifies a configuration based on whether GP regression or the long-range function best approximate its energy.

Equation 12 shows that  $C_{\text{optimal}}$  employs no boundary and so requires prior knowledge of the energy of a configuration in order to classify it. Consequently, models obtained from the closest model strategy are unsuitable for prediction. However, this method represents an 'ideal' classifier that is guaranteed to find the optimal  $RMSE_{\text{test}}$  for a given GP model. Hence  $C_{\text{optimal}}$  is useful for estimating how inaccuracies in the parametric classifiers affect the training efficiency. If models from  $C_{\text{optimal}}$  significantly outperform the other classifiers, this suggests the other classifiers are too simple to properly approximate the true boundary.

There are circumstances where  $C_{\text{optimal}}$  may not give an optimal RMSE against the test set. There may be short-range hypersurfaces where the interaction energy is predicted near-exactly by the long-range function due to chance. Points in the reference set that are near these hypersurfaces will be classified by  $C_{\text{optimal}}$  as part of  $A_{\text{LR}}$  instead of  $A_{\text{GP}}$ . If the configuration density of the reference set around such a hypersurface is insufficiently high, no training points will be added in its vicinity. Consequently, points in the test set near to the hypersurface will be inadequately approximated by either method of prediction; the GP will perform poorly due to a lack of nearby training points and the long-range function will be inaccurate for short-range points that are not extremely close to the hypersurface. This problem was avoided by using dense reference sets and by not using the constrained placement strategy with  $C_{\text{optimal}}$ .

### 6. Combining Classifiers and Placement Strategies

The sequential design method, or training strategy, requires a choice of classifier and placement strategy. The combinations of these examined in this work are given in table III. Methods that do not involve  $C_{\text{optimal}}$  can make predictions and so are suitable for applications. The method involving  $C_{\text{optimal}}$  can only classify points if the true energy is already known, and so is only useful to estimate the loss in performance due to inaccuracies in the parametric classifiers. The fixed boundary method corresponds to the method of Uteva *et al.*<sup>19</sup> and is included to allow comparison with this prior method, on which the new methods build. The method of Uteva *et al.*<sup>19</sup>

TABLE III: The classifier and training point placement method for all training strategies examined here.

| Training Strategy  | Classifier           | Point Placement Method |
|--------------------|----------------------|------------------------|
| Single-Constrained | $C_{\text{single}}$  | Constrained Placement  |
| Multi-Constrained  | $C_{\text{multi}}$   | Constrained Placement  |
| Single-Open        | $C_{\text{single}}$  | Open Placement         |
| Multi-Open         | $C_{\text{multi}}$   | Open Placement         |
| Closest Model      | $C_{\text{optimal}}$ | Open Placement         |
| Fixed Boundary     | $C_{\text{fixed}}$   | Constrained Placement  |

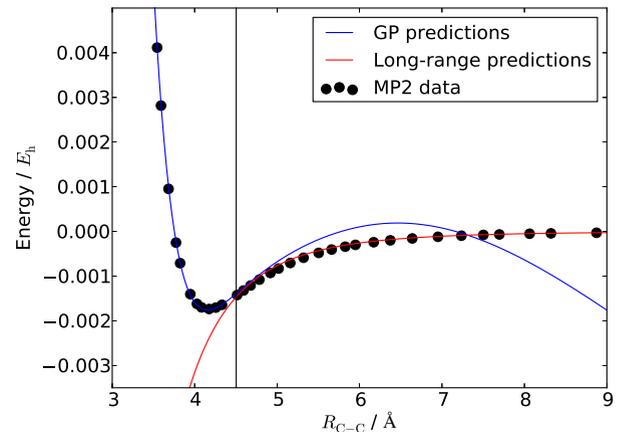


FIG. 2: The predictions of the GP (blue) and the long-range function (red) for a slice through the PES of the  $(\text{CO}_2)_2$  potential in which the two molecules are in a T-shaped configuration. The cross-over distance for this model is shown by the black line at 4.505 Å. Data points (black circles) are independent of the GP training data.

was previously shown to significantly reduce the number of training points compared to LHC design. It is demonstrated below that these new boundary optimisation methods improve further the already efficient methods of Uteva *et al.*<sup>19</sup>.

## III. RESULTS AND DISCUSSION

Comparisons of the performances of the different training strategies are made using the HF-Ne, HF- $\text{Na}^+$ , CO-Ne,  $\text{CO}_2$ -Ne and  $(\text{CO}_2)_2$  potentials. These were selected as they provide a range of interaction types and well depths to test the robustness of the new training strategies.

The number of training points is  $N_{\text{TP}}$ , which in each system was less than 10 % of the number of configurations in the corresponding reference set (see table I and the related text). Consequently, the training sets for the models discussed are candidates for transfer learning because their small size makes, for example, CCSD(T) calculations possible for the whole training set. Boundary optimisation is anticipated to improve training efficiency as, for small  $N_{\text{TP}}$ , the long-range function will outperform the GP at the outer edge of the potential well. This is illustrated in figure 2 for the  $(\text{CO}_2)_2$  system,

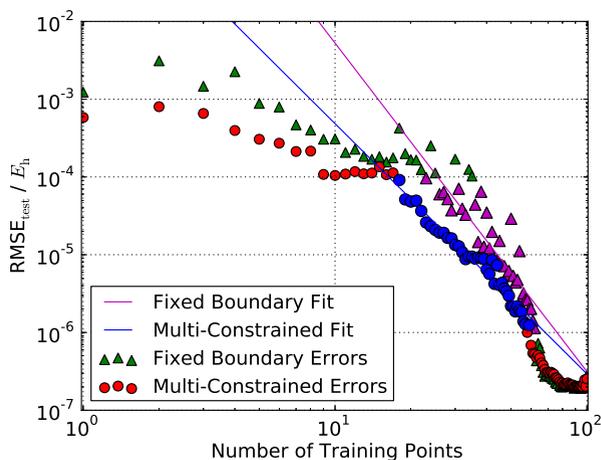


FIG. 3: Plot of  $\text{RMSE}_{\text{test}} / E_h$  against  $N_{\text{TP}}$  on a  $\log_{10}$  scale for models of the HF-Ne PES trained via the multi-constrained (circles) and fixed boundary (triangles) training strategies. The blue circles and pink triangles are points which were included in the fitting of the lines shown.

using a GP model trained under the single-constrained strategy up to  $N_{\text{TP}} = 21$ .

Figure 2 shows that the predictive accuracy of the long-range function exceeds that of GP regression for configurations with C-C separations above the cross-over distance. Despite this, in the fixed boundary method these configurations would be predicted with the GP. This not only reduces overall accuracy, but also means new training points must be placed at long range to address this, when allowing the long-range function to predict these energies would give adequate accuracy.

The impact of the above on training efficiency compared to fixed boundary training is shown for HF-Ne in figure 3. For  $C_{\text{multi}}$  the RMSE falls faster with  $N_{\text{TP}}$  compared to  $C_{\text{fixed}}$ , indicating improved training efficiency. This improvement is most pronounced when  $N_{\text{TP}}$  is low and so  $\text{RMSE}_{\text{test}}$  is high. This is expected because increasing the size of the training set increases the size of the GP region, meaning  $R_{\text{cross}}$  in the boundary-optimised strategies will approach the fixed value of 8.5 Å. Consequently, the difference between the fixed boundary and boundary-optimised models will close as  $N_{\text{TP}}$  increases. Equivalent plots to figure 3 for all training strategies for all systems are found in the supplementary material. The RMSE data are somewhat noisy. Possible causes of this noise are minor variations in the hyperparameters upon re-training, the discrete nature and stochastic design of the reference set and the low values of  $N_{\text{TP}}$  meaning that addition of a single point has a non-smooth effect on the RMSE. The new methods herein have considerably lower noise than the prior fixed boundary method.

To compare the efficiency gain over fixed boundary training across all new strategies from subsections IIE 5 and IIE 6, a metric,  $E$ , is used. This compares the  $N_{\text{TP}}$  required by two different strategies to achieve a given  $\text{RMSE}_{\text{test}}$ . When com-

TABLE IV: The equation of the lines of best fit for  $\text{RMSE}_{\text{test}}$  of models from all training strategies for CO-Ne as power laws in the number of training points,  $N_{\text{TP}}$ . Also shown are the  $R^2$  values of each fit on the data. All fits were over points in the range  $1 \times 10^{-6} E_h \leq \text{RMSE}_{\text{test}} \leq 1 \times 10^{-4} E_h$ .

| Training Strategy            | Line of best fit to $\text{RMSE}_{\text{test}}$ | $R^2$ |
|------------------------------|---|-------|
| Single-Constrained Placement | $4.694N_{\text{TP}}^{-3.842}$                   | 0.929 |
| Multi-Constrained Placement  | $1.551N_{\text{TP}}^{-3.647}$                   | 0.944 |
| Single-Open Placement        | $2.901N_{\text{TP}}^{-3.718}$                   | 0.921 |
| Multi-Open Placement         | $3.113N_{\text{TP}}^{-3.830}$                   | 0.958 |
| Closest Model                | $0.865N_{\text{TP}}^{-3.490}$                   | 0.944 |
| Fixed Boundary               | $546.0N_{\text{TP}}^{-5.069}$                   | 0.914 |

paring a new training strategy with fixed boundary training,

$$E(\text{RMSE}_{\text{test}}) = \frac{N_{\text{TP}}(\text{RMSE}_{\text{test}})}{N_{\text{TP, fixed}}(\text{RMSE}_{\text{test}})} \times 100\%, \quad (13)$$

where  $N_{\text{TP}}$  and  $N_{\text{TP, fixed}}$  are the numbers of training points required by the new strategy and fixed boundary training, respectively. These quantities and  $E$  are shown as functions of  $\text{RMSE}_{\text{test}}$  as they vary with its value.

For a given  $\text{RMSE}_{\text{test}}$ , the values of  $N_{\text{TP}}$  and  $N_{\text{TP, fixed}}$  are determined from least squares fits of  $\log_{10}(\text{RMSE}_{\text{test}})$  versus  $\log_{10}(N_{\text{TP}})$ , with examples of such fits shown in figure 3 for the HF-Ne potential. Fits are made in the region where the RMSE decays as a power law of  $N_{\text{TP}}$ . For all systems this corresponds to  $1 \times 10^{-6} E_h \leq \text{RMSE}_{\text{test}} \leq 1 \times 10^{-4} E_h$ , apart from HF- $\text{Na}^+$  where the region is  $1 \times 10^{-5} E_h \leq \text{RMSE}_{\text{test}} \leq 1 \times 10^{-3} E_h$  due to the larger high energy cut-off for this potential. The fits provide continuous lines to interpolate to any  $\text{RMSE}_{\text{test}}$  within the above range. This enables a comparison of  $N_{\text{TP}}$  between training strategies at fixed  $\text{RMSE}_{\text{test}}$  that accounts for the somewhat noisy data. Moreover, this is the range of errors in which the models are useful for applications and the decrease in  $\log(\text{RMSE}_{\text{test}})$  with  $\log(N_{\text{TP}})$  is linear.

The fitted equations for each training strategy for the CO-Ne system are given in table IV as power laws in  $N_{\text{TP}}$ . Equivalent tables for all other systems are found in the supplementary material. As the data are noisy, these tables also show the  $R^2$  value of each fit. In the case of CO-Ne, these evidence the high quality of the fits. In fact, no fit from any training strategy for any system achieves an  $R^2$  value lower than 0.8907 (from the fixed boundary training strategy on the HF- $\text{Na}^+$  system). Furthermore, figure 3 and the corresponding plots in the supplementary material show that the RMSE data follow a straight line (in a log-log plot). This implies that the  $R^2$  arises from scatter in the data rather than unsuitability of the fitting function.

Re-arranging the equations in table IV provides expressions for  $N_{\text{TP}}$  in terms of  $\text{RMSE}_{\text{test}}$  for the CO-Ne potential. These give  $E$  as a function of  $\text{RMSE}_{\text{test}}$ , via equation 13, for all new training strategies herein. This was done for all other potentials as well, with plots of  $E$  against  $\text{RMSE}_{\text{test}}$  for all training strategies for each system given in figure 4.

As observed earlier, the training efficiency gains in figure 4

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.  
PLEASE CITE THIS ARTICLE AS DOI:10.1063/5.0063534

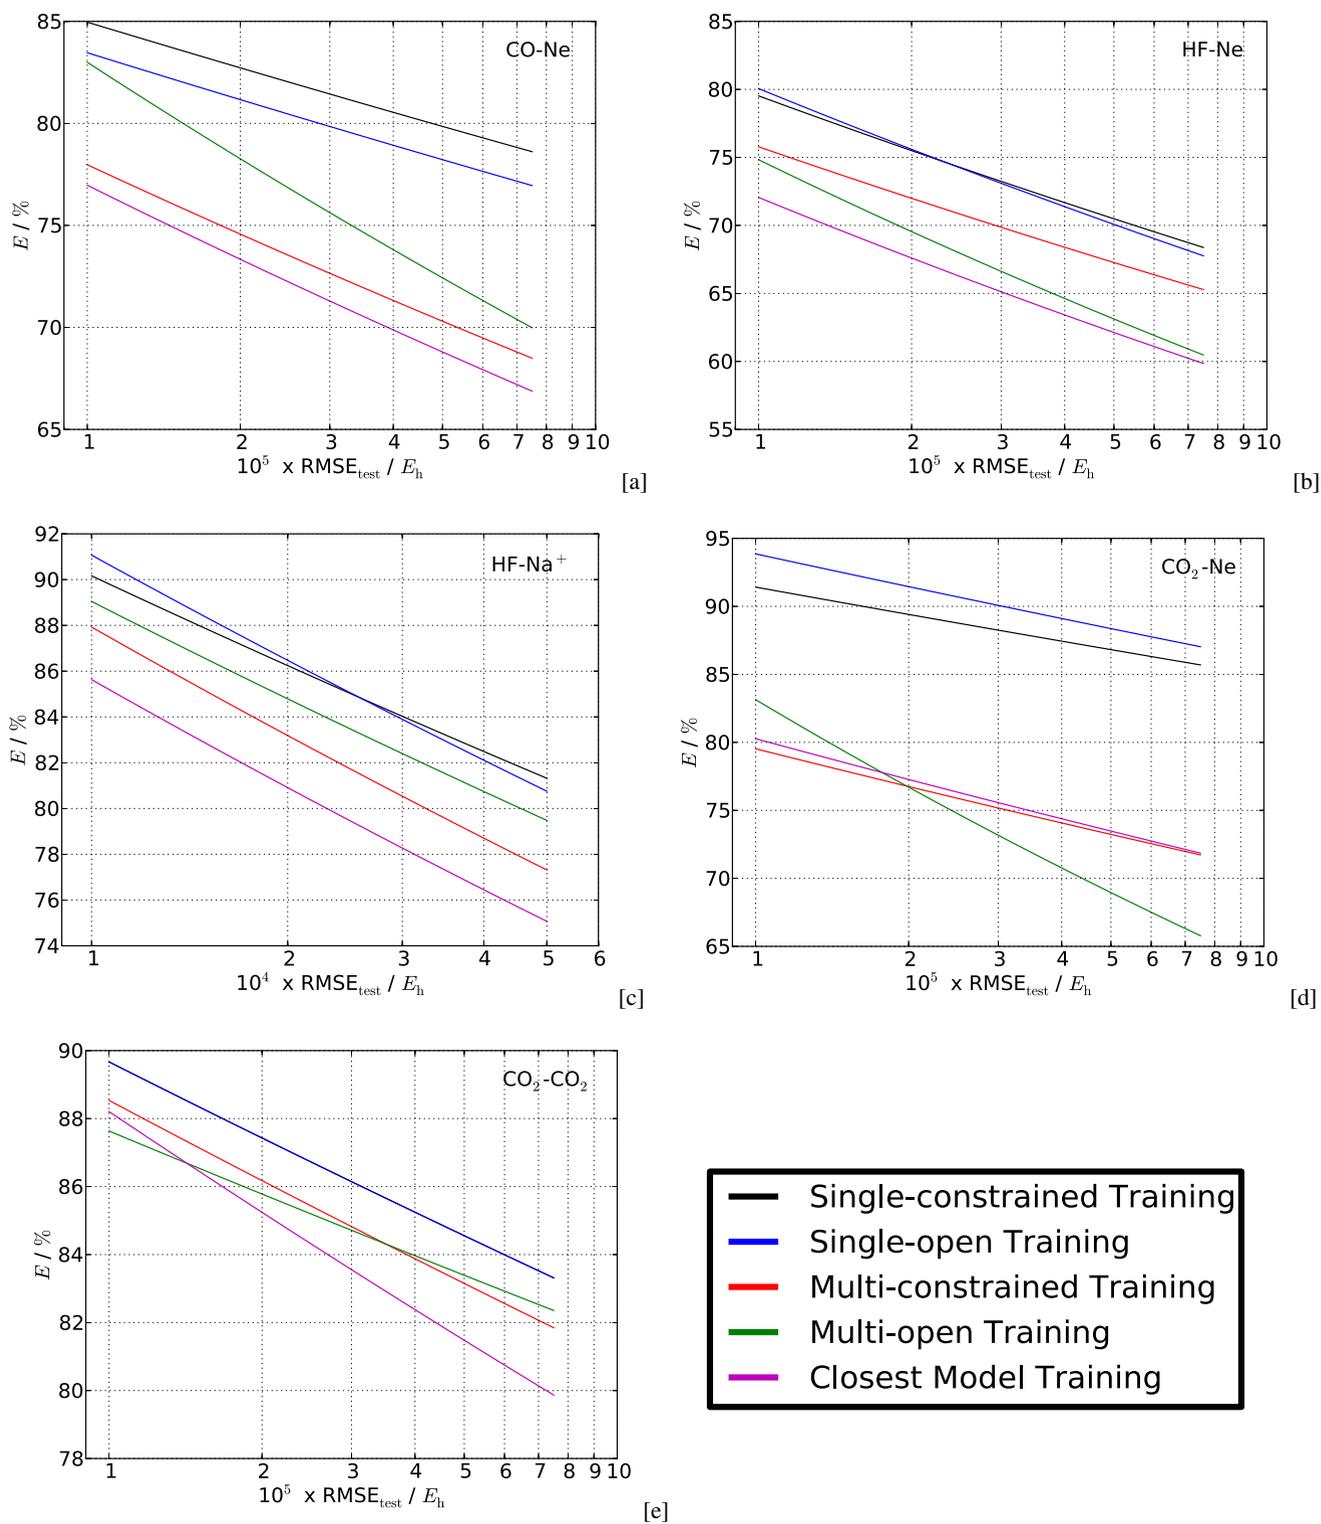


FIG. 4: Plots of  $E$  against  $\text{RMSE}_{\text{test}}$  for all potentials are shown in parts (a) to (e). The potential referred to in each frame is shown in the upper right corner.

are more pronounced at high  $\text{RMSE}_{\text{test}}$  for all training strategies across all systems. This suggests that boundary optimisation is most effective when the training set is small, making this technique ideal for applications where a computationally cheap but less accurate potential is required. Nevertheless, significant reductions in the number of training points are also obtained in the RMSE range where PESs become useful for first principles predictions. For example, a GP potential with an RMSE of  $3 \times 10^{-4}$  eV per atom ( $1.1 \times 10^{-5} E_h$  per atom) was employed in a recent simulation of the thermal properties of  $\beta\text{-Ga}_2\text{O}_3$ <sup>31</sup>. Furthermore, Uteva *et al.* successfully determined the  $\text{CO}_2\text{-CO}$  second virial coefficient using a GP PES with an RMSE of  $2.4 \times 10^{-5} E_h$ <sup>17</sup>. In this RMSE range the boundary optimisation methods typically reduce the required number of training points by 15-33% (see figure 4).

Generally, the closest model strategy generates the largest efficiency gain, while the smallest improvement comes from strategies involving  $C_{\text{single}}$ . Efficiency gains only slightly below those from the closest model strategy are achieved by the strategies that use  $C_{\text{multi}}$ . This hierarchy of improvement implies that the choice of classifier, rather than point placement strategy, is most important because strategies with the same classifier perform more similarly than those with the same point placement method. The closest model strategy is included only to illustrate the total training efficiency gain possible from an ‘ideal’ classifier, and it is encouraging that the best boundary optimisation methods are close to this ideal case. Indeed, the difference in  $E$  between this ideal method and the closest-performing  $C_{\text{multi}}$  strategy never exceeds  $\sim 3\%$  for any system. This implies that  $C_{\text{multi}}$  captures the true shape of the boundary region for the systems explored sufficiently. Thus introducing a more detailed classifier would not be worth the increased cost of evaluating the PES.

$C_{\text{multi}}$  always outperforms  $C_{\text{single}}$ . However, for the  $\text{CO}_2\text{-Ne}$  potential (figure 4d),  $C_{\text{multi}}$  also outperforms the closest model strategy. This is because the long-range function is nearly exact for a group of short-range configurations in the reference set. This is shown in figure 5 by the thin ‘peninsula’ of points that are best estimated by the long range function (in red) which encroaches deep into the GP region (in blue). This ‘peninsula’ exists because the long-range function is of higher energy than the MP2 data in some regions of the PES and of lower energy in others, meaning there must be some hypersurface in between where the two are equal. Prediction for test configurations close to the ‘peninsula’ is problematic under  $C_{\text{optimal}}$  unless the reference set is very dense in this region. This is because the exact predictions of the long-range function on the ‘peninsula’ mean no training points are added there, leading the closest model strategy to perform relatively poorly for the  $\text{CO}_2\text{-Ne}$  potential.

The total gain in training efficiency achieved by boundary optimisation varies somewhat between systems. While the best-performing training strategy for HF-Ne improved training efficiency by between 25-39% (*i.e.*  $E = 61\text{-}75\%$ ), for  $(\text{CO}_2)_2$  the gain was only 12-18% ( $E = 82\text{-}88\%$ ). The more limited gains for  $(\text{CO}_2)_2$  may be because the *a priori* choice of  $R_{\text{cross}} = 8.5 \text{ \AA}$ , required for the fixed boundary method, is reasonable for this system. Nevertheless, even for the  $(\text{CO}_2)_2$  po-

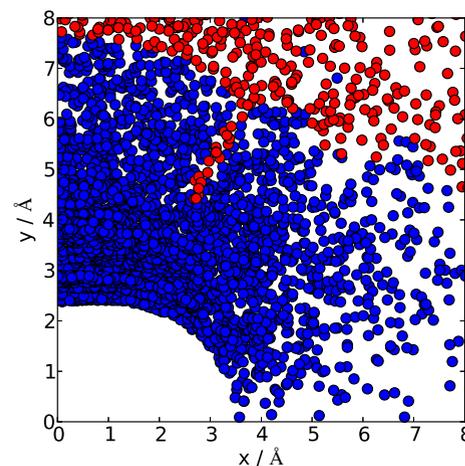


FIG. 5: Plot of the  $x$  and  $y$  coordinates of the Ne in the  $\text{CO}_2\text{-Ne}$  system for configurations in the reference set. The  $\text{CO}_2$  molecule is aligned along the  $x$ -axis with the C at the origin. Points are classified as GP points (blue) or asymptotic points (red) by  $C_{\text{optimal}}$  using a model from the closest model strategy trained up to  $N_{\text{TP}} = 100$ .

tential, use of multi-constrained training confers an efficiency gain of  $\sim 18\%$  over fixed boundary training. This means that for all of the potentials explored, use of a training strategy that employs  $C_{\text{multi}}$  confers a useful improvement over fixed boundary training.

Boundary optimisation improves the training efficiency due to more effective placement of training points. This is illustrated in figure 6, which shows, for HF-Ne, the differences in training point placement for three training methods: fixed boundary, single-constrained and closest model. While all place most points at separations below  $3 \text{ \AA}$ , indicating that the repulsive wall is the hardest region to model, the placement of points at larger separations diverges between methods. For the first 20 training points few configurations are placed beyond  $3 \text{ \AA}$  for the single-constrained and closest model strategies. Thereafter, the distance at which training points are added slowly increases, even for the closest model strategy, which does not employ a boundary directly. In fact, single-constrained training adds training points beyond  $8 \text{ \AA}$  only after  $\sim 90$  training points have been placed and the closest model strategy does not add any training points above  $7 \text{ \AA}$  at all.

In contrast, the fixed boundary training method adds its fifth training point at a minimum separation above  $5 \text{ \AA}$  and its eighth at  $8.5 \text{ \AA}$ . This demonstrates that a fixed boundary strategy switches between placing training points in the repulsive wall and at the boundary from the onset of training. This difference is because fixed boundary training requires the GP to predict energies at separations up to  $8.5 \text{ \AA}$  from the start of training. Consequently, training points must be added at separations near the  $8.5 \text{ \AA}$  boundary from the onset, even though the energies at these separations are very small and generally well-approximated by the long-range function. This contrasts with the boundary-optimised and closest model strategies, which allow the long-range function to approximate con-

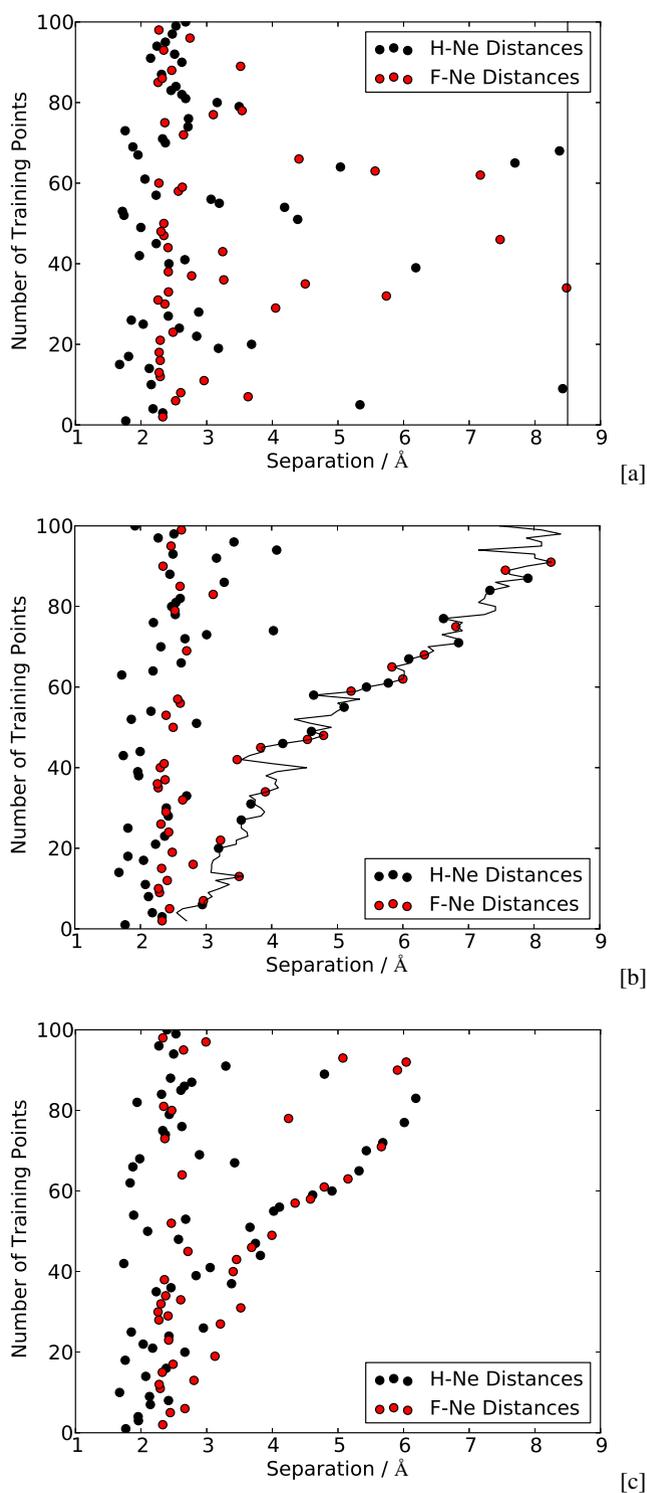


FIG. 6: Plots showing training point placement for the first 100 training points for models of the HF-Ne PES trained using the fixed boundary (a), single-constrained placement (b) and closest model (c) strategies. These points are coloured based on the shortest interatomic distance in the configuration, with only this shortest distance shown for each. Boundary values are represented by black lines where applicable.

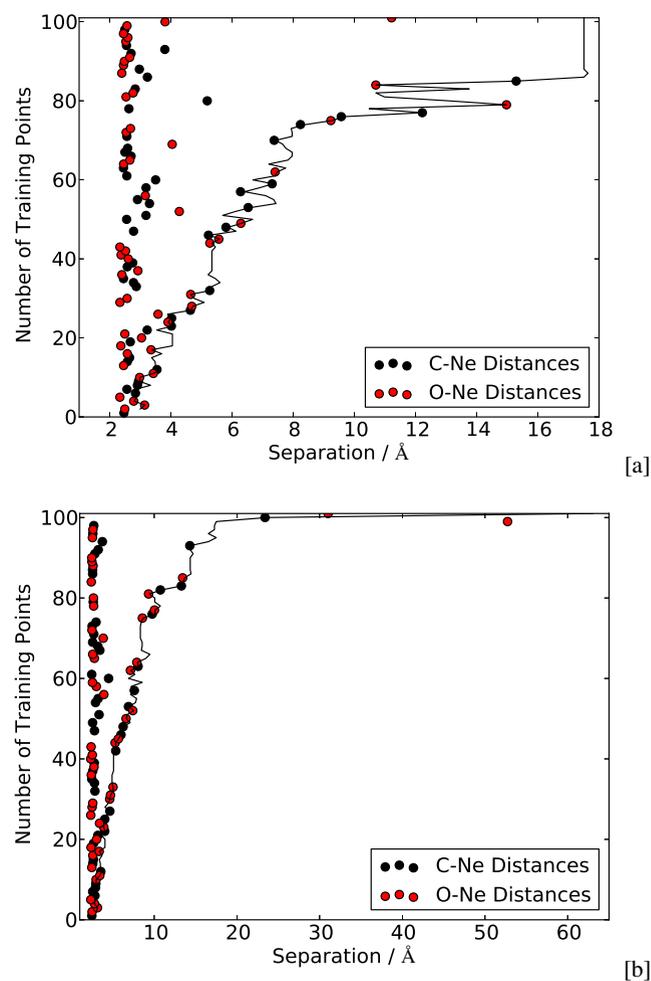


FIG. 7: Plots showing training point placement for the first 100 training points for models of the CO-Ne PES trained using the single-constrained placement (a) and single-open placement (b) strategies. These points are coloured based on the shortest interatomic distance in the configuration, with only this shortest distance shown for each. Boundary values are represented by black lines.

figures at separations around 8.5 Å when the number of training points is low. This produces more efficient model development under boundary-optimised and closest model training because point placement can be focused on the short-range region of the PES, where energy varies more rapidly with configuration.

Training point plots for the CO-Ne potential, given in figure 7, show that the single-open strategy extends  $A_{GP}$  much further than single-constrained training. This suggests that the capacity to place points in  $A_{LR}$  facilitates faster expansion of the boundary for this system. However, these plots also indicate that this discrepancy is most noticeable when the number of training points is large; specifically, rapid expansion of the boundary under single-open training was instigated by placement of a single training point at long range, which facilitated

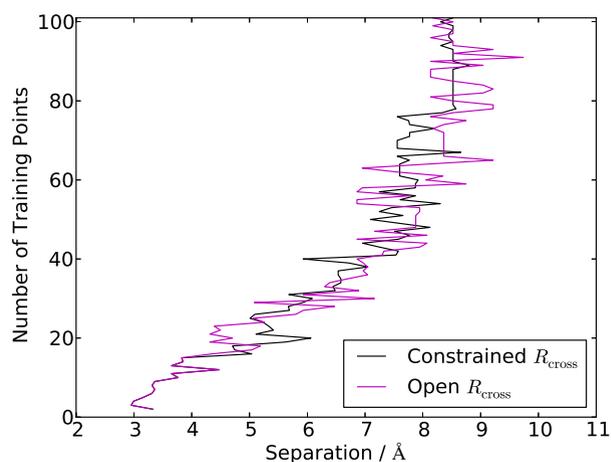
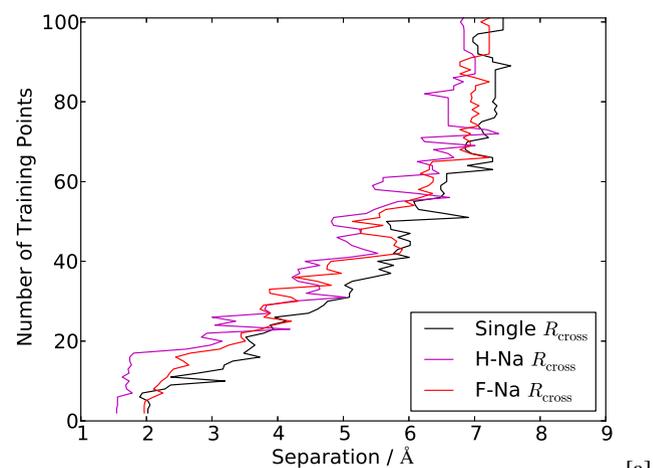


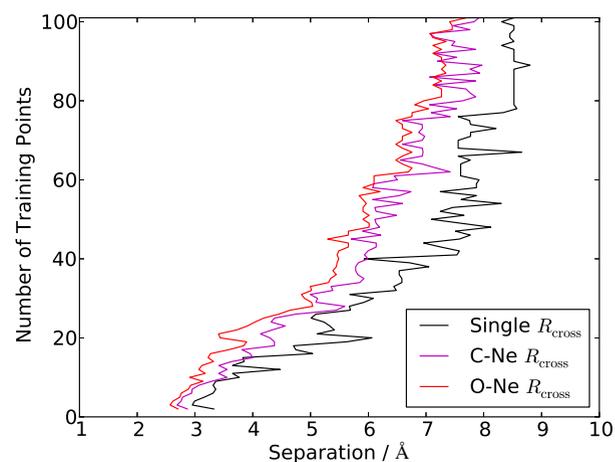
FIG. 8: Plot showing the values of  $R_{\text{cross}}$  achieved for the first 100 training points by single-open and single-constrained training for models of the  $\text{CO}_2\text{-Ne}$  potential.

the increase in  $R_{\text{cross}}$  to 63.0 Å. Prior to this, the cross-over value from single-open training was quite similar to that from single-constrained training (both were  $\sim 17$  Å). Hence the two training strategies differ significantly only when the predictions from GP regression are already highly accurate, as at this stage the predictions of the long-range function are poor enough by comparison to merit placement of a training point at long-range. The other systems also show close similarity between the cross-over distance for the single-constrained and single-open strategies, with equivalent plots to figure 8 found in the supplementary material. Also, Figure 8 shows that for the  $\text{CO}_2\text{-Ne}$  potential the values of  $R_{\text{cross}}$  from the two strategies are similar throughout training. Such a trend suggests that the choice of point placement strategy does not significantly change the value of  $R_{\text{cross}}$ , implying once more that the choice of classifier is of greater import.

Similarities in the evolution of the boundary are seen when the cross-over distances achieved under  $C_{\text{multi}}$  and  $C_{\text{single}}$  are compared for a given point placement method. Figure 9 shows the results of such a comparison for the  $\text{HF-Na}^+$  and  $\text{CO}_2\text{-Ne}$  systems. The cross-over distances generally grow at a similar rate, but the value of  $R_{\text{cross}}$  is consistently larger than all values in  $\mathbf{R}_{\text{cross}}$ . This suggests that there are differences between  $C_{\text{single}}$  and  $C_{\text{multi}}$ , which manifest in both the training efficiency and boundary placement. The larger GP region under  $C_{\text{single}}$  compared with  $C_{\text{multi}}$  is explained by noting that both are approximations of an ‘ideal’ classifier. When  $N_{\text{TP}}$  is low (*i.e.* one or two), such a classifier will attempt to make  $A_{\text{GP}}$  as small as possible because the training set comprises configurations from the repulsive wall only. Consequently, given the same training set,  $C_{\text{single}}$  and  $C_{\text{multi}}$  also try to minimise the size of  $A_{\text{GP}}$ . As  $C_{\text{multi}}$  is more flexible its approximation of the ‘ideal’ classifier will be closer than that of  $C_{\text{single}}$ , meaning  $A_{\text{GP}}$  under  $C_{\text{single}}$  will be initially larger. Other than the repulsive wall, the largest errors tend to occur at separations around the boundary. Thus, due to its larger  $A_{\text{GP}}$ , a  $C_{\text{single}}$



[a]



[b]

FIG. 9: Plots showing the value of  $R_{\text{cross}}$  and entries in  $\mathbf{R}_{\text{cross}}$  for the first 100 training points for models of the  $\text{HF-Na}^+$  potential from single/multi-open training (a) and the  $\text{CO}_2\text{-Ne}$  potential from single/multi-constrained training (b).

strategy will place its first long-range training point at a larger separation than a  $C_{\text{multi}}$  strategy. This facilitates faster expansion of the boundary under  $C_{\text{single}}$  than  $C_{\text{multi}}$  regardless of which point placement method is used.

Figure 9a shows that  $R_{\text{H-Na}^+} < R_{\text{F-Na}^+}$  throughout most of training for the  $\text{HF-Na}^+$  potential, meaning that the interaction involving the larger atom (F) obeys a larger cross-over distance. Such ordering of the cross-over distances also applies to the  $\text{HF-Ne}$  system for most of training. Moreover, figure 9b shows that the ordering of the cross-over distances for the  $\text{CO}_2\text{-Ne}$  potential is consistent, with  $R_{\text{O-Ne}} < R_{\text{C-Ne}}$  throughout training. In fact, for the  $(\text{CO}_2)_2$  and  $\text{CO-Ne}$  potentials it holds generally throughout training that  $R_{\text{O-O}} < R_{\text{O-C}} < R_{\text{C-C}}$  and  $R_{\text{O-Ne}} < R_{\text{C-Ne}}$  respectively for both point placement methods. This implies that the ordering of the cross-over distances under  $C_{\text{multi}}$  is consistent between systems as well as being physically reasonable.

Additionally, figures 6-9 show that the cross-over distance increases with the size of the training set. This is expected

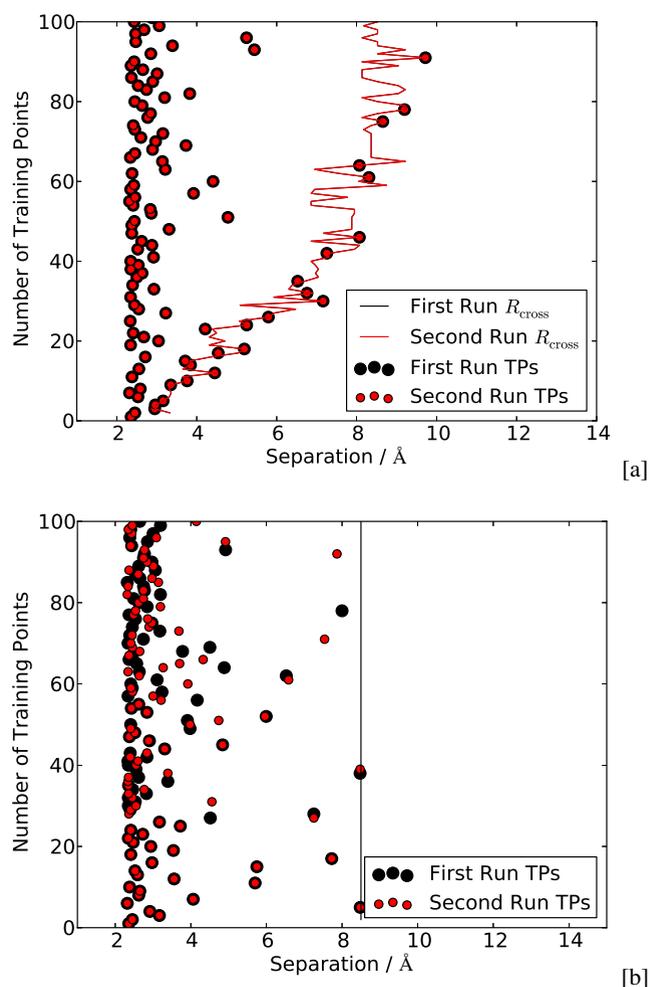


FIG. 10: Plots showing the first 100 training points and values of  $R_{\text{cross}}$  from two runs of single-open training (a) and fixed boundary training (b) for models of the  $\text{CO}_2\text{-Ne}$  potential.

because a larger training set means the GP has more information with which to infer the function describing the PES. The resultant more accurate GP enables larger cross-over distance(s). This is the case for all training strategies across all systems.

Boundary optimisation methods decrease the cost of using PES, by reducing the number of training points and the size of the GP region. Fewer training points means GP evaluations are cheaper, as the GP cost is proportional to  $N_{\text{TP}}$ . Additionally, the lower  $R_{\text{cross}}$  or  $\mathbf{R}_{\text{cross}}$  of the boundary-optimised model means that in any application the GP would be used less often, in favour of the much cheaper asymptotic function. Hence boundary optimisation produces models that are more efficient to implement in applications than fixed boundary training, with no reduction in overall accuracy.

Boundary optimisation also increases the reproducibility of training. That is, for two separate models from the same training strategy and reference set, the results will be more sim-

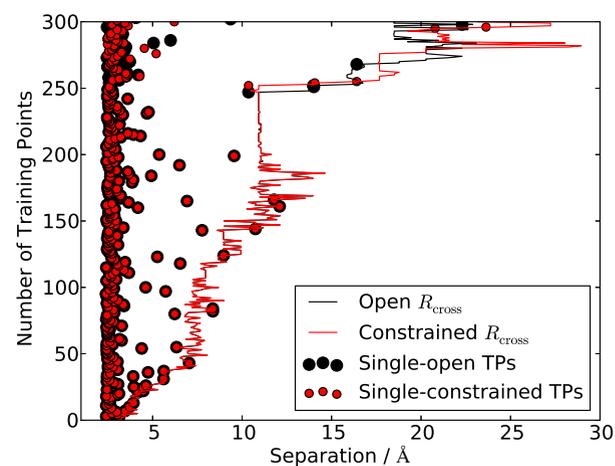


FIG. 11: The first 300 training points and values of  $R_{\text{cross}}$  from single-open and single-constrained training for the  $(\text{CO}_2)_2$  potential.

ilar for a boundary-optimised strategy compared to using a fixed boundary. For example, when modelling the  $\text{CO}_2\text{-Ne}$  potential with single-open training the results from two separate training runs were identical. This is shown in figure 10. This is noteworthy because the values of the hyperparameters selected when maximising  $\log(\mathcal{L})$  can vary slightly, even for the same training set. Such variations can alter the predictions of the GP, leading to different values for  $R_{\text{cross}}$  or  $\mathbf{R}_{\text{cross}}$  and the selection of different training points. Thus, that two separate runs of the same training strategy are totally identical is encouraging.

Furthermore, figure 10 shows that for fixed boundary training separate runs were identical only up to  $N_{\text{TP}} = 27$ . While the exact reproducibility in figure 10a is not present for the other potentials, there is generally significantly less difference between independent runs of the boundary-optimised training methods than for fixed boundary training. For example, models of the  $\text{HF-Ne}$  potential from single-constrained and single-open training are reproducible up to  $N_{\text{TP}} = 30$  and  $N_{\text{TP}} = 41$  respectively, compared with  $N_{\text{TP}} = 16$  for fixed boundary training. Equivalent plots to figure 10 for this system under these strategies are given in the supplementary material. This reproducibility increase compared to fixed boundary training does not transfer to the use of the  $\mathbf{C}_{\text{multi}}$  strategies, likely because a direct search is not as reproducible in multiple dimensions as in a single dimension. However, from the  $R^2$  values in table IV it can be seen that use of either  $\mathbf{C}_{\text{multi}}$  or  $\mathbf{C}_{\text{single}}$  reduces the scatter in the  $\text{RMSE}_{\text{test}}$  data relative to use of  $\mathbf{C}_{\text{fixed}}$ , as evidenced by the larger  $R^2$  values achieved when training with the former two. This trend is repeated across all other potentials examined here, which suggests that use of boundary optimisation leads to more stability in selection of the hyperparameters and hence more consistent predictions.

For the  $(\text{CO}_2)_2$  system, reproducibility is seen not just for repeat runs of identical training strategies but between the training strategies that use  $\mathbf{C}_{\text{single}}$ . This is illustrated in figure

11, which shows that the single-constrained and single-open training strategies choose identical training points until  $N_{\text{TP}} = 210$ . Such an observation explains why the two methods have identical  $E$  in figure 4e. This is because when  $N_{\text{TP}} = 210$  the  $\text{RMSE}_{\text{test}}$  values for single-open and single-constrained placement were  $2.817 \times 10^{-7} E_h$  and  $2.784 \times 10^{-7} E_h$  respectively. Consequently the error was too low at this  $N_{\text{TP}}$  to be included in the fit, meaning the two strategies were identical over the  $\text{RMSE}_{\text{test}}$  values used in fitting.

#### IV. CONCLUSION

It has been shown that boundary optimisation produces GP PESs of the same accuracy using fewer training points than fixing  $R_{\text{cross}}$  *a priori*. This improvement in efficiency is hierarchical, with a boundary defined by a single, variable cross-over distance offering a modest improvement and a boundary defined by multiple such distances facilitating a further gain.

The results presented imply that the classifier is more important to the training strategy than the point placement method. In the RMSE range that is suitable for first principles calculations ( $\sim 2 \times 10^{-5} E_h$ ) the boundary optimisation methods typically reduce the required number of training points by 15-33% relative to a training strategy that is already established as efficient<sup>19</sup>. Because of their reduced training set size, the resulting boundary-optimised PESs are strong candidates for transfer learning, in which the existing *ab initio* calculations are upgraded to a higher level of theory. Furthermore, as the size of the GP region increases with the size of the training set, only as needed, the resulting GPs are also less computationally intensive in applications than fixed boundary methods, as they employ the GP over a smaller region of phase space.

The classifier  $C_{\text{multi}}$ , which uses different cross-over distances for difference atomic pairs, performed almost as well as an 'ideal' classifier. Across all systems, the largest difference in performance between the closest model strategy, which uses an 'ideal' classifier, and nearest boundary-optimised strategy was  $\sim 3\%$ . In all cases the best-performing boundary-optimised strategy employed  $C_{\text{multi}}$ , implying that a classifier comprising of a spherical boundary with a unique radius on each unique atomic pair, captures the true boundary effectively. Further refinement of the classifier would not result in a sufficient reduction in training points to justify the extra classifier expense.

The cross-over distance(s) are learned from the reference data under boundary optimisation using a direct search. This is sufficiently fast to be used at every stage of training whether one or many cross-over distances are employed and in the multi-dimensional case returns cross-over distances in a physically reasonable order. Moreover, both direct search algorithms can be used easily in conjunction with another machine learning technique or on another chemical interaction. In fact, a direct search can be applied to any problem whereby a boundary is sought between a good method of approximation in one region of phase space (the long-range function here) and a machine learning technique in another. This means that

the methodology which underpins boundary optimisation is both fast and flexible, in addition to being effective in solving the problem of reducing the computational expense associated with training a GP model of an intermolecular PES.

Physical systems in which the behaviour crosses over from a simple asymptotic function to more complicated behaviour are common in many fields. A prominent example is the transition from ideal to non-ideal gas behaviour. As the boundary optimization techniques herein exploit this cross-over in behaviour, there are many potential applications of this technique to physical problems beyond intermolecular potentials.

#### V. SUPPLEMENTARY MATERIAL

See supplementary material for more plots of results from the explored systems.

#### VI. ACKNOWLEDGEMENTS

The computations of all reference and test sets described in this paper were performed using the University of Nottingham's Augusta HPC service, which provides a High Performance Computing service to the University's research community. See <https://www.nottingham.ac.uk/it-services/research/hpc/> for more details. The authors are also thankful for funding provided by the Leverhulme Trust Doctoral Scholarship through Modelling and Analytics for a Sustainable Society (MASS) at the University of Nottingham.

#### VII. DATA AVAILABILITY

All test and reference data sets used in this work are available in the Boundary Optimisation Data repository at <https://doi.org/10.6084/m9.figshare.16610713.v1>, reference number [74].

#### VIII. AUTHOR DECLARATION

The authors have no conflicts to disclose.

#### Appendix A: Derivation of an Empirical Long-range Function

$\text{HF-Na}^+$  was the only intermolecular potential examined above for which a multipolar long-range function was not used. Instead a fitted, empirical long-range function was employed. Although determining an optimal long-range function was not the goal of this work, one that was accurate enough that all configurations were not rapidly transferred to the GP region was needed for proof of concept. The multipolar long-range function was unsuitable for this because of the discrepancy between its predictions and the MP2 energies.

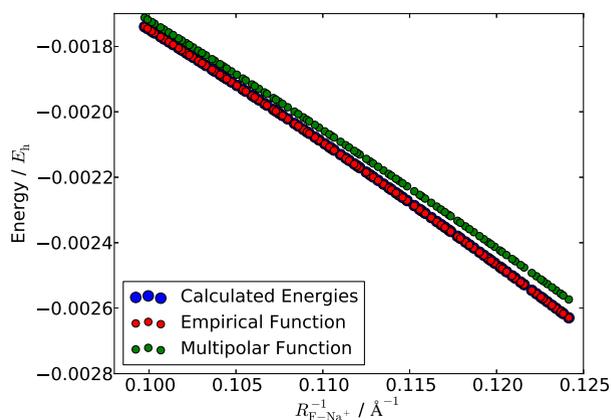


FIG. 12: Plots of the predicted energies of the long-range functions derived by the empirical method (red) and the multipole method (green) for the HF-Na<sup>+</sup> potential in a linear configuration with the F proximal to the Na<sup>+</sup> (*i.e.* with  $\cos(\theta) = 1$ ). Also shown for comparison are MP2 energies (blue) for this configuration.

Evidence of this is given in figure 12, which also highlights the superiority of the predictions of the empirical long-range function. The RMSEs of each method of prediction were separated by two orders of magnitude, with the empirical function achieving an RMSE of  $3.07 \times 10^{-7} E_h$  versus  $4.15 \times 10^{-5} E_h$  for the multipolar function against the MP2 data shown in the figure.

Figure 12 evidences that the multipolar long-range function captured the power law of the interaction energy for this system in this configuration, with the source of its error being that it was offset from the calculated energies. This offset was a product of the energies being calculated using MP2 while many of the properties used to derive the multipolar long-range function were calculated at higher levels of theory (see table II). In the other systems, the small magnitude of the long-range energies meant that this disagreement was negligible and the multipolar long-range function was usable. However, in the case of HF-Na<sup>+</sup>, long-range energies with larger magnitudes were commonplace due to the strong repulsive and attractive interactions between the H-F dipole and the Na<sup>+</sup> cation. This exacerbated the difference between the predictions of the multipolar long-range function and the MP2 energies.

To approximate accurately the long-range data without upgrading the reference and test data to a higher level of theory, a long-range function was produced by fitting directly to these data. This was the empirical long-range function. Taking  $r$  to be the distance between the centre of the H-F bond (not centre of mass) and the Na<sup>+</sup>, this function estimated the energy,  $E$ , as a sum of power laws,

$$E = Ar^{-2} + Br^{-3}. \quad (\text{A1})$$

In doing so, the empirical long-range function exploits that the dominant powers of  $r$  in the HF-Na<sup>+</sup> interaction are known to be -2 and -3 but assumes that the coefficients of these terms are unknown.

As the energy varies differently with  $r$  when  $\theta$  changes, a sum such as in equation A1 must be fitted for every configuration in a given data set for which  $\theta$  is unique. For a given  $\theta$  the coefficients in equation A1 can be found using simultaneous equations, which are set up by following algorithm 4. In all cases,  $r_{\min} = 8.5 \text{ \AA}$  and  $r_{\max} = 10.5 \text{ \AA}$ , though the power laws which resulted were accurate up to  $100 \text{ \AA}$ . Furthermore, both GPs in the algorithm were trained on  $\sim 300$  training points from a LHC design.

---

#### Algorithm 4 Coefficients for the Empirical Long-range Function

---

- 1: Train a GP,  $\text{GP}_{\min}$ , on a range of  $\theta$  values at separation  $r = r_{\min}$ ; do the same for another GP,  $\text{GP}_{\max}$ , at  $r = r_{\max}$ .
  - 2: For a given  $\theta$ , predict  $E_{\min}$  and  $E_{\max}$  using  $\text{GP}_{\min}$  and  $\text{GP}_{\max}$  respectively.
  - 3: Set up simultaneous equations of the form shown in equation A1: one with  $E = E_{\min}$  and  $r = r_{\min}$ , and another with  $E = E_{\max}$  and  $r = r_{\max}$ .
  - 4: Solve the equations from step three for the coefficients for the current  $\theta$ .
- 

By repeating steps 2-4 in algorithm 4, a sum of power laws was determined for every  $\theta$  value in the reference set and test sets. Fitting to the latter was possible as knowledge of the energies in the set over which fitting was undertaken was unnecessary. This was because the GP predictions were based on their respective training sets, and  $r$  and  $\theta$  were found from the inverse interatomic separations at each configuration alone.

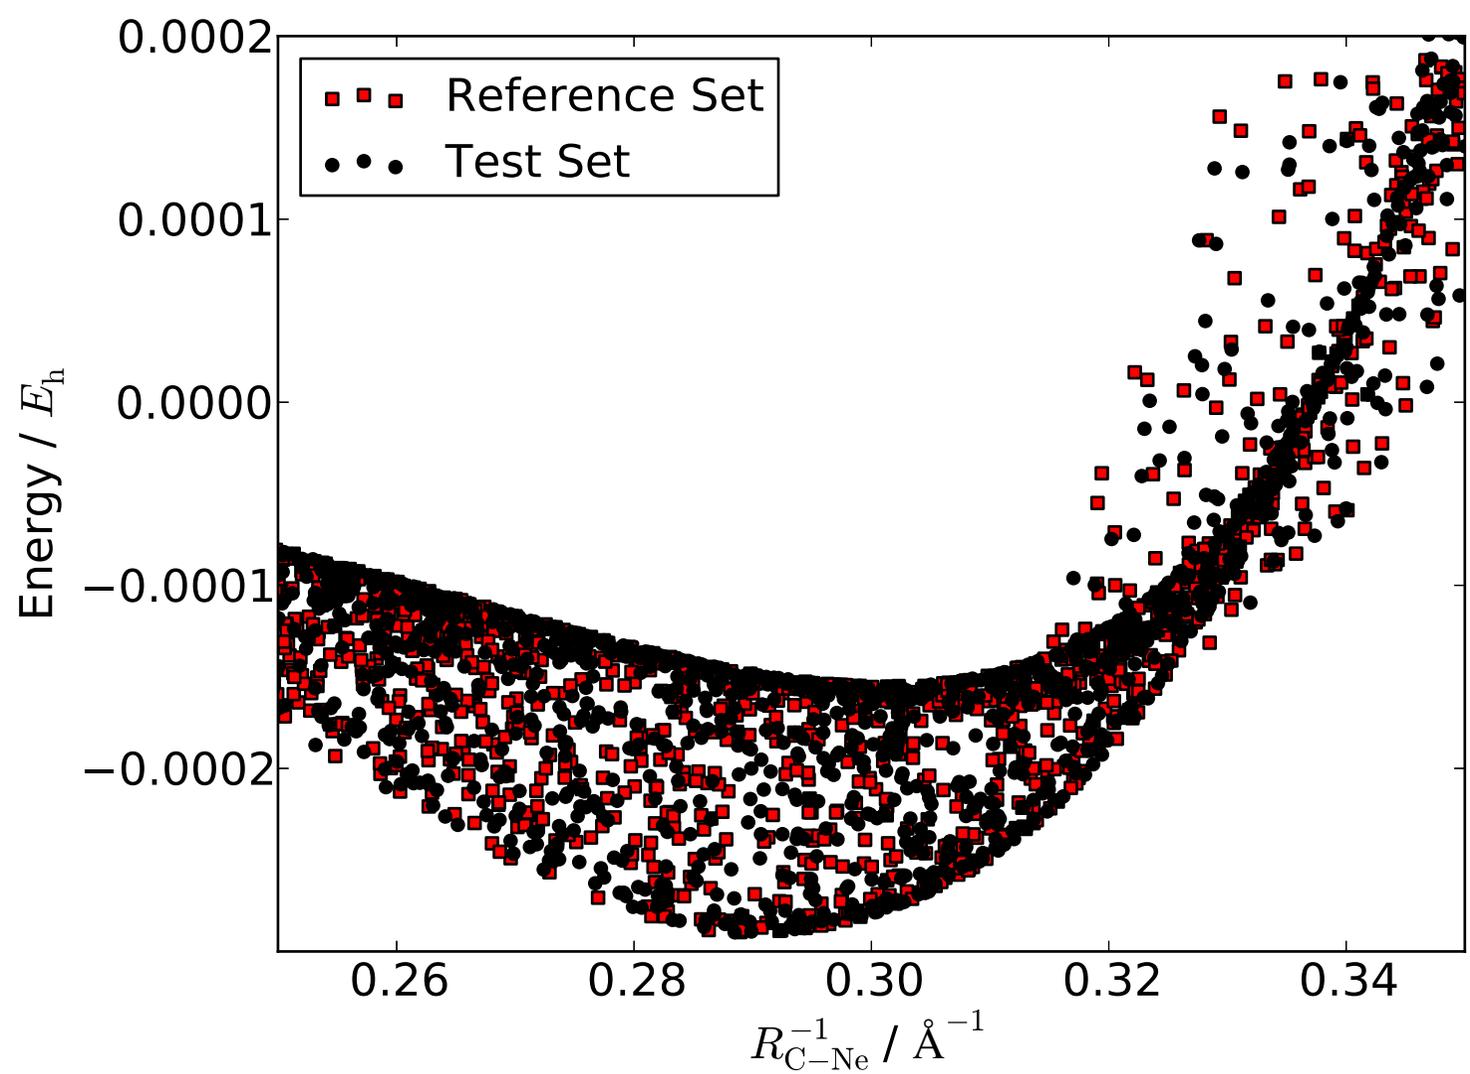
Deriving a long-range function from the predictions of two GPs could be problematic if a transfer learning approach were to be invoked as the data used to train these GPs would itself need to be upgraded. For the HF-Na<sup>+</sup> potential this would not be an issue because increasing the quality of the training data would increase the quality of the fit from the multipolar long-range function.

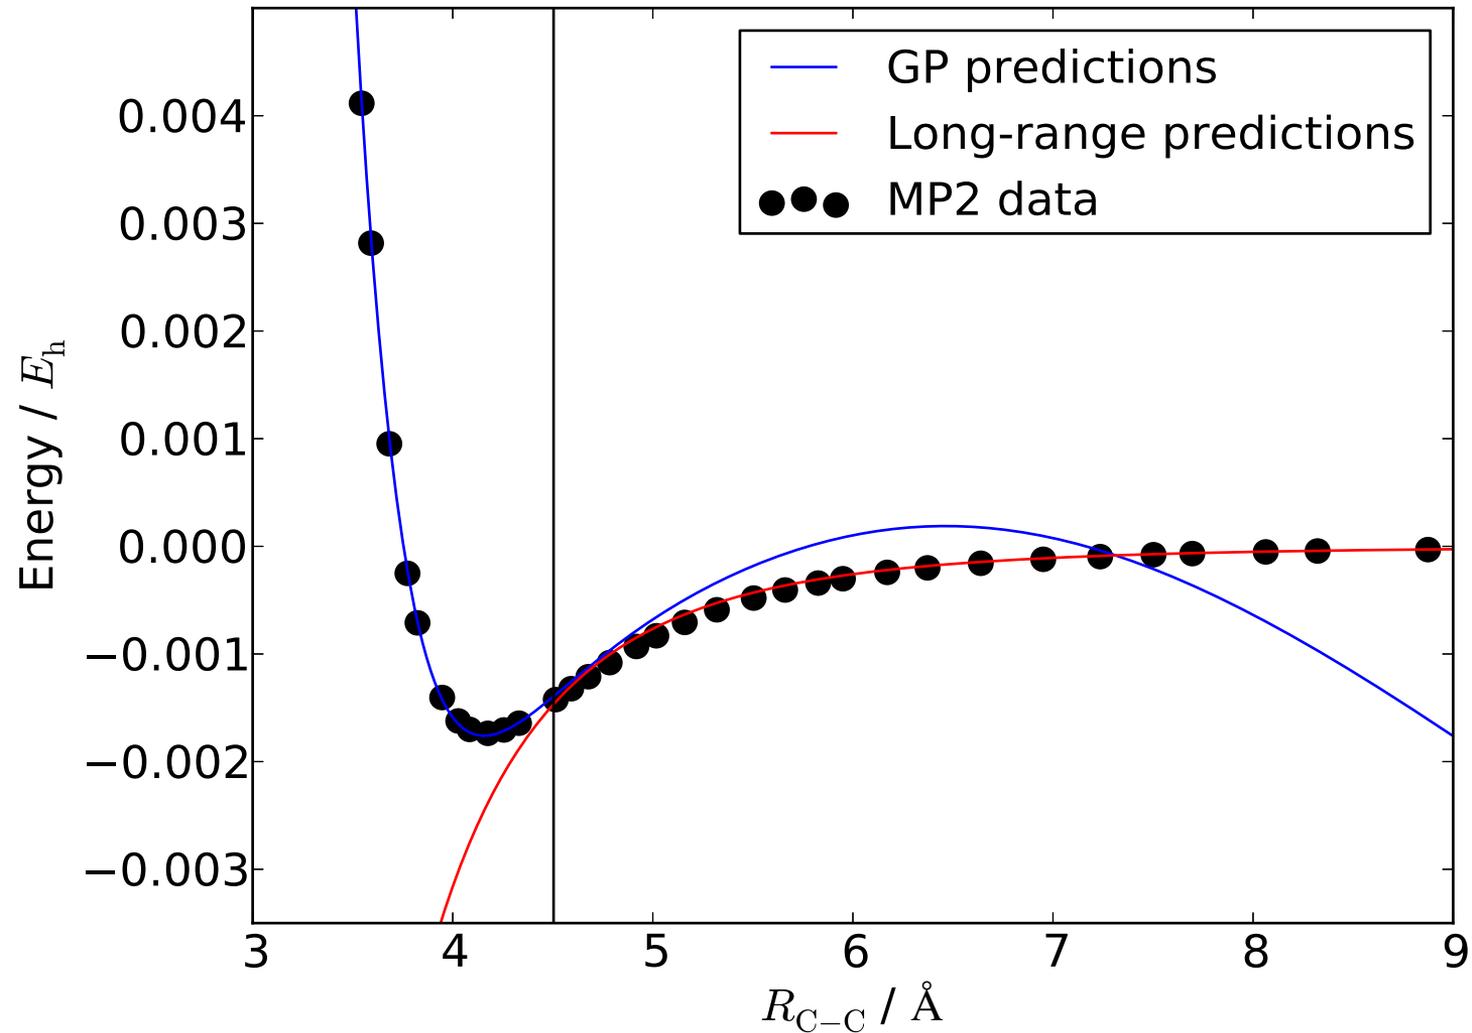
However, when an empirical function is the only option for modelling the long-range data, upgrading the data used in training  $\text{GP}_{\min}$  and  $\text{GP}_{\max}$  would be of considerable computational expense. This is because, currently, each comprise  $\sim 300$  configurations. Such an issue could be circumvented by using a sequential design strategy to build minimal training sets for these GPs, which could then be upgraded instead. Furthermore, the training sets used for  $\text{GP}_{\min}$  and  $\text{GP}_{\max}$  are independent of that used to train a GP on the wider PES. Letting the number of training points in the training sets of  $\text{GP}_{\min}$  and  $\text{GP}_{\max}$  be  $Y$  and that in the training set of the other GP be  $Z$ , the cost of short-range predictions in any simulation would scale linearly with  $Z$  and the cost of any long-range predictions would scale linearly with  $2Y$ . Given that the training of  $\text{GP}_{\min}$  and  $\text{GP}_{\max}$  take place at a fixed separations, it is likely that  $2Y < Z$  if all GPs were trained under a sequential design strategy. This means that the predictions of the empirical long-range function would not be the computational bottleneck and that such a function is suitable for use in simulations. Finally, the method is sufficiently flexible that the GP predictions can easily be replaced by those of another statistical method.

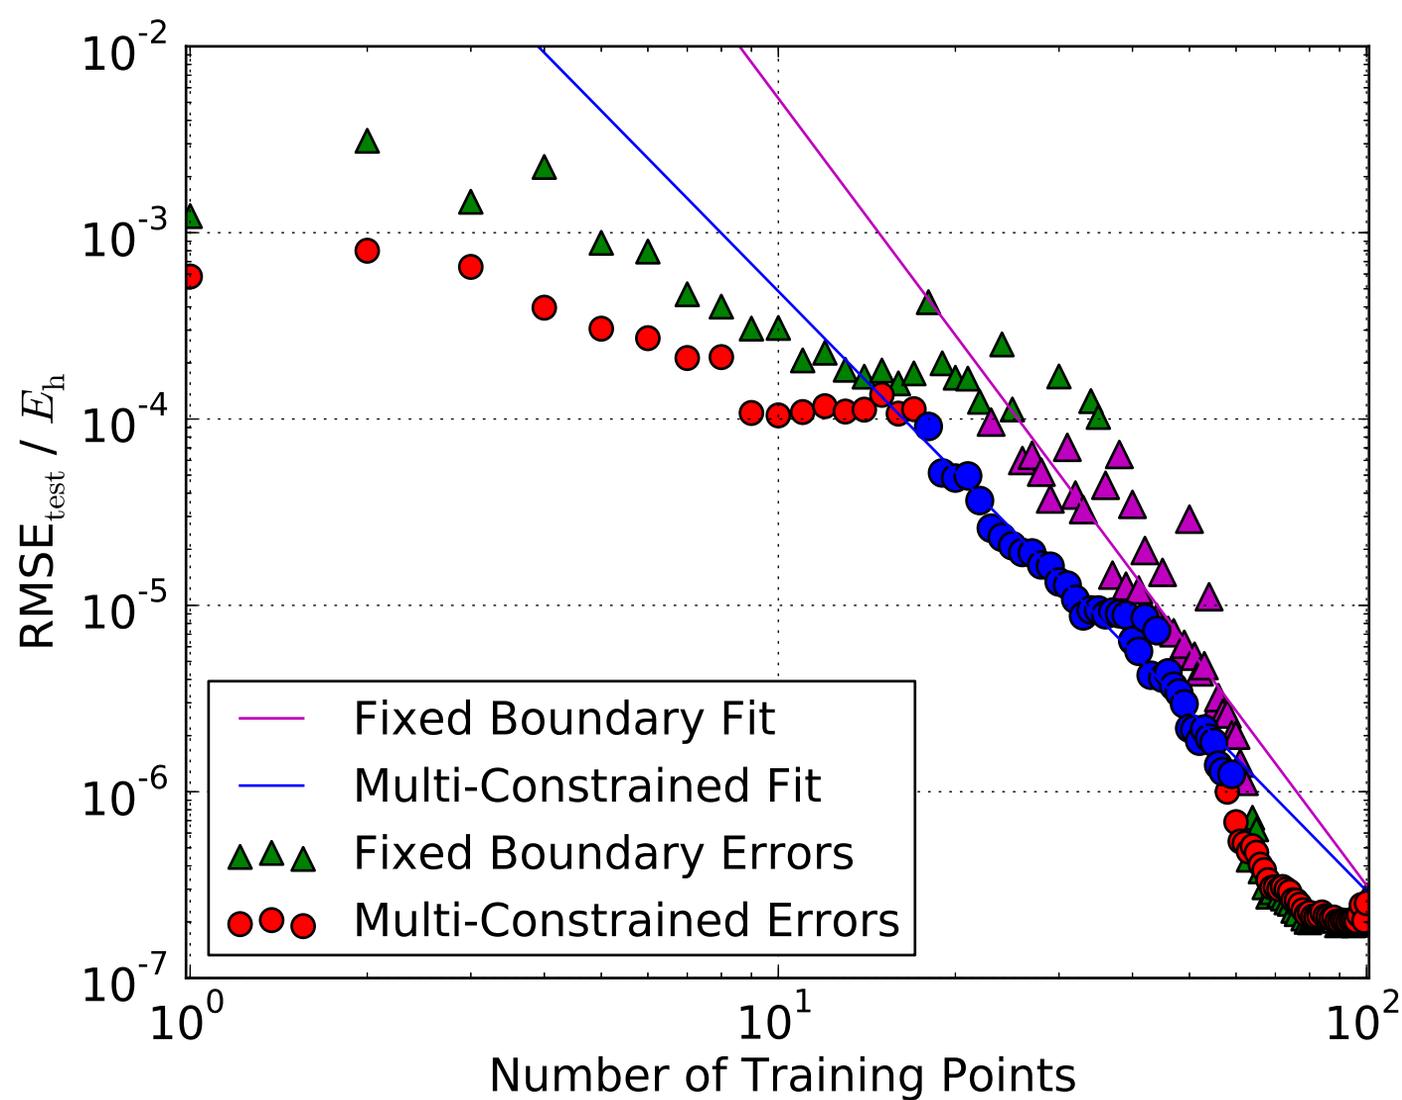
- <sup>1</sup>T. A. Halgren and W. Damm, "Polarizable force fields," *Curr. Opin. Struct. Biol.* **11**, 236–242 (2001).
- <sup>2</sup>A. Warshel, M. Kato, and A. V. Pislakov, "Polarizable force fields: history, test cases, and prospects," *Journal of Chemical Theory and Computation* **3**, 2034–2045 (2007).
- <sup>3</sup>P. Xu, E. B. Guidez, C. Bertoni, and M. S. Gordon, "Perspective: Ab initio force field methods derived from quantum mechanics," *J. Chem. Phys.* **148**, 090901 (2018).
- <sup>4</sup>J. Behler, "Perspective: Machine learning potentials for atomistic simulations," *J. Chem. Phys.* **145**, 170901 (2016).
- <sup>5</sup>T. Mueller, A. Hernandez, and C. Wang, "Machine learning for interatomic potential models," *J. Chem. Phys.* **152**, 050902 (2020).
- <sup>6</sup>B. Selvaratnam and R. T. Koodali, "Machine Learning in Experimental Materials Chemistry," *Catal. Today* (2020).
- <sup>7</sup>J. Noh, G. H. Gu, S. Kim, and Y. Jung, "Machine-enabled inverse design of inorganic solid materials: promises and challenges," *Chem. Sci.* **11**, 4871–4881 (2020).
- <sup>8</sup>P. O. Dral, "Quantum chemistry in the age of machine learning," *J. Phys. Chem. Lett.* **11**, 2336–2347 (2020).
- <sup>9</sup>J. Behler and M. Parrinello, "Generalized neural-network representation of high-dimensional potential-energy surfaces," *Phys. Rev. Lett.* **98**, 146401 (2007).
- <sup>10</sup>N. Artrith, T. Morawietz, and J. Behler, "High-dimensional neural-network potentials for multicomponent systems: Applications to zinc oxide," *Phys. Rev. B* **83**, 153101 (2011).
- <sup>11</sup>N. Artrith, B. Hiller, and J. Behler, "Neural network potentials for metals and oxides - First applications to copper clusters at zinc oxide," *Phys. Status Solidi B* **250**, 1191–1203 (2013).
- <sup>12</sup>J. Behler, "Representing potential energy surfaces by high-dimensional neural network potentials," *J. Phys. Condens. Matt.* **26**, 183001 (2014).
- <sup>13</sup>Q. Lin, Y. Zhang, B. Zhao, and B. Jiang, "Automatically growing global reactive neural network potential energy surfaces: a trajectory free active learning strategy," *J. Chem. Phys.* (2020).
- <sup>14</sup>A. V. Shapeev, "Moment tensor potentials: A class of systematically improvable interatomic potentials," *MMS* **14**, 1153–1173 (2016).
- <sup>15</sup>E. V. Podryabinkin and A. V. Shapeev, "Active learning of linearly parametrized interatomic potentials," *Comput. Mater. Sci.* **140**, 171–180 (2017).
- <sup>16</sup>I. I. Novoselov, A. V. Yanilkin, A. V. Shapeev, and E. V. Podryabinkin, "Moment tensor potentials as a promising tool to study diffusion processes," *Comput. Mater. Sci.* **164**, 46–56 (2019).
- <sup>17</sup>E. Uteva, R. J. Wheatley, R. D. Wilkinson, and R. S. Graham, "Interpolation of intermolecular potentials using Gaussian processes," *J. Chem. Phys.* **147**, 161706 (2017).
- <sup>18</sup>A. J. Cresswell, R. J. Wheatley, R. D. Wilkinson, and R. S. Graham, "Molecular simulation of the thermophysical properties and phase behaviour of impure CO<sub>2</sub> relevant to CCS," *Faraday Discussions* **192**, 415–436 (2016).
- <sup>19</sup>E. Uteva, R. J. Wheatley, R. D. Wilkinson, and R. S. Graham, "Active learning in Gaussian process interpolation of potential energy surfaces," *J. Chem. Phys.* **149**, 174114 (2018).
- <sup>20</sup>C. M. Handley, G. I. Hawe, D. B. Kell, and P. L. A. Popelier, "Optimal construction of a fast and accurate polarisable water potential based on multipole moments trained by machine learning," *Phys. Chem. Chem. Phys.* **11**, 6365–6376 (2009).
- <sup>21</sup>M. J. L. Mills and P. L. A. Popelier, "Intramolecular polarisable multipolar electrostatics from the machine learning method kriging," *Comput. Theor. Chem.* **975**, 42–51 (2011).
- <sup>22</sup>M. J. L. Mills and P. L. A. Popelier, "Polarisable multipolar electrostatics from the machine learning method Kriging: an application to alanine," *Theor. Chem. Acc.* **131**, 1137 (2012).
- <sup>23</sup>S. M. Kandathil, T. L. Fletcher, Y. Yuan, J. Knowles, and P. L. A. Popelier, "Accuracy and tractability of a Kriging model of intramolecular polarizable multipolar electrostatics and its application to histidine," *J. Comput. Chem.* **34**, 1850–1861 (2013).
- <sup>24</sup>J. Dai and R. V. Krems, "Interpolation and extrapolation of global potential energy surfaces for polyatomic systems by Gaussian processes with composite kernels," *J. Chem. Theory Comput.* **16**, 1386–1395 (2020).
- <sup>25</sup>H. Sugisawa, T. Ida, and R. V. Krems, "Gaussian process model of 51-dimensional potential energy surface for protonated imidazole dimer," *J. Chem. Phys.* **153**, 114101 (2020).
- <sup>26</sup>G. Schmitz, E. L. Klitting, and O. Christiansen, "A Gaussian process regression adaptive density guided approach for potential energy surface construction," *J. Chem. Phys.* **153**, 064105 (2020).
- <sup>27</sup>M. J. Burn and P. L. A. Popelier, "Creating Gaussian process regression models for molecular simulations using adaptive sampling," *J. Chem. Phys.* **153**, 054111 (2020).
- <sup>28</sup>M. A. Boussaidi, O. Ren, D. Voytsekhovskiy, and N. Manzhos, "Random Sampling High Dimensional Model Representation Gaussian Process Regression (RS-HDMR-GPR) for multivariate function representation: application to molecular potential energy surfaces," *J. Phys. Chem. A* **124**, 7598–7607 (2020).
- <sup>29</sup>G. Sivaraman, A. N. Krishnamoorthy, M. Baur, C. Holm, M. Stan, G. Csányi, C. Benmore, and Á. Vázquez-Mayagoitia, "Machine Learning Inter-Atomic Potentials Generation Driven by Active Learning: A Case Study for Amorphous and Liquid Hafnium dioxide," *arXiv preprint arXiv:1910.10254* (2019).
- <sup>30</sup>M. A. Caro, G. Csányi, T. Laurila, and V. L. Deringer, "Machine learning driven simulated deposition of carbon films," *Phys. Rev. B* **102**, 21 (2020).
- <sup>31</sup>Y. Liu, J. Yang, G. Xin, L. Liu, G. Csányi, and B. Cao, "Machine learning interatomic potential developed for molecular simulations on thermal properties of  $\beta$ -Ga<sub>2</sub>O<sub>3</sub>," *J. Chem. Phys.* **153**, 144501 (2020).
- <sup>32</sup>A. Glielmo, P. Sollich, and A. De Vita, "Accurate interatomic force fields via machine learning with covariant kernels," *Phys. Rev. B* **95**, 214302 (2017).
- <sup>33</sup>A. Glielmo, C. Zeni, and A. De Vita, "Efficient nonparametric n-body force fields from machine learning," *Phys. Rev. B* **97**, 184307 (2018).
- <sup>34</sup>C. Zeni, K. Rossi, A. Glielmo, Á. Fekete, N. Gaston, F. Baletto, and A. De Vita, "Building machine learning force fields for nanoclusters," *J. Chem. Phys.* **148**, 241739 (2018).
- <sup>35</sup>P. I. Maxwell, N. di Pasquale, S. Cardamone, and P. L. A. Popelier, "The prediction of topologically partitioned intra-atomic and inter-atomic energies by the machine learning method kriging," *Theor. Chem. Acc.* **135**, 195 (2016).
- <sup>36</sup>P. I. Maxwell and P. L. A. Popelier, "Accurate prediction of the energetics of weakly bound complexes using the machine learning method kriging," *Struct. Chem.* **28**, 1513–1523 (2017).
- <sup>37</sup>Z. E. Hughes, E. Ren, J. Thacker, B. Symons, A. F. Silva, and P. L. A. Popelier, "A FFLUX water model: Flexible, Polarizable and with a Multipolar Description of Electrostatics," *J. Comput. Chem.* **41**, 619–628 (2020).
- <sup>38</sup>T. L. Fletcher and P. L. A. Popelier, "FFLUX: Transferability of polarizable machine-learned electrostatics in peptide chains," *J. Comput. Chem.* **38**, 1005–1014 (2017).
- <sup>39</sup>J. C. R. Thacker, A. L. Wilson, Z. E. Hughes, M. J. Burn, P. I. Maxwell, and P. L. A. Popelier, "Towards the simulation of biomolecules: optimisation of peptide-capped glycine using FFLUX," *Mol. Simul.* **44**, 881–890 (2018).
- <sup>40</sup>A. Konovalov, B. Symons, and P. L. A. Popelier, "On the many-body nature of intramolecular forces in fflux and its implications," *J. Comput. Chem.* **42**, 107–116 (2020).
- <sup>41</sup>K. R. Bartók, A. P. and G. Csányi, "On representing chemical environments," *Phys. Rev. B* **87**, 184115 (2013).
- <sup>42</sup>F. C. Mocanu, K. Konstantinou, T. H. Lee, N. Bernstein, V. L. Deringer, G. Csányi, and S. R. Elliott, "Modeling the phase-change memory material, Ge<sub>2</sub>Sb<sub>2</sub>Te<sub>5</sub>, with a machine-learned interatomic potential," *J. Phys. Chem. B* **122**, 8998–9006 (2018).
- <sup>43</sup>D. Dragoni, T. D. Daff, G. Csányi, and N. Marzari, "Achieving DFT accuracy with a machine-learning interatomic potential: Thermomechanics and defects in BCC ferromagnetic iron," *Phys. Rev. Mater.* **2**, 013808 (2018).
- <sup>44</sup>A. P. Bartók, J. Kermode, N. Bernstein, and G. Csányi, "Machine learning a general-purpose interatomic potential for silicon," *Phys. Rev. X* **8**, 041048 (2018).
- <sup>45</sup>V. L. Deringer, M. A. Caro, and G. Csányi, "A general-purpose machine-learning force field for bulk and nanostructured phosphorus," *Nat. Commun.* **11**, 1–11 (2020).
- <sup>46</sup>M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics* **42**, 55–61 (2000).
- <sup>47</sup>M. Stein, "Large sample properties of simulations using latin hypercube sampling," *Technometrics* **29**, 143–151 (1987).
- <sup>48</sup>W.-L. Loh *et al.*, "On latin hypercube sampling," *Annals of statistics* **24**,

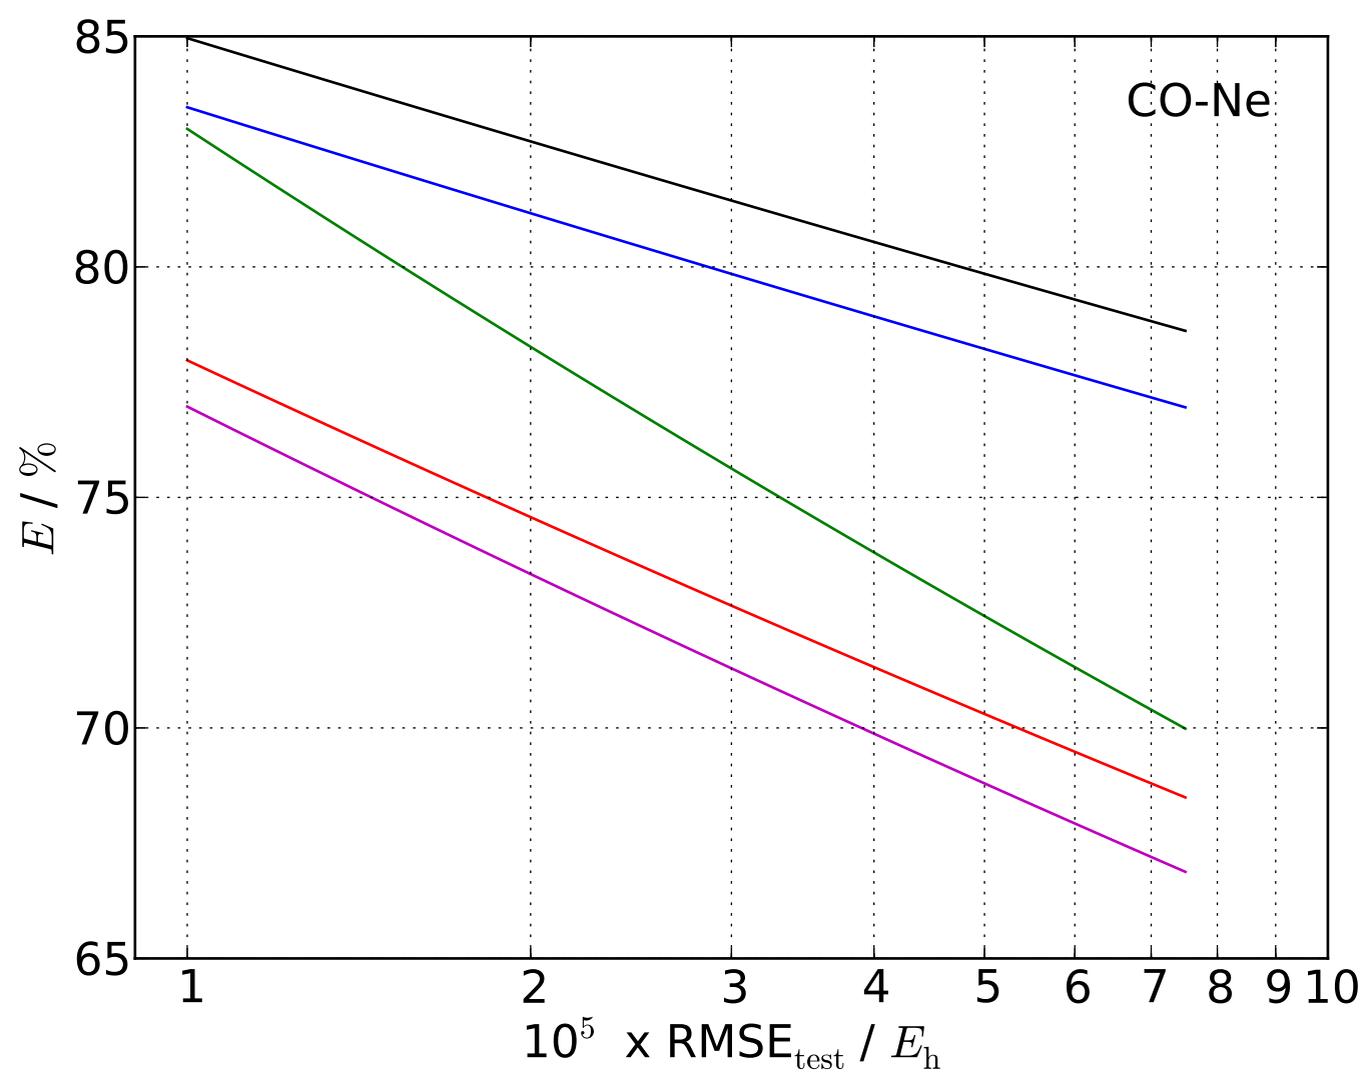
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.  
PLEASE CITE THIS ARTICLE AS DOI:10.1063/5.0063534

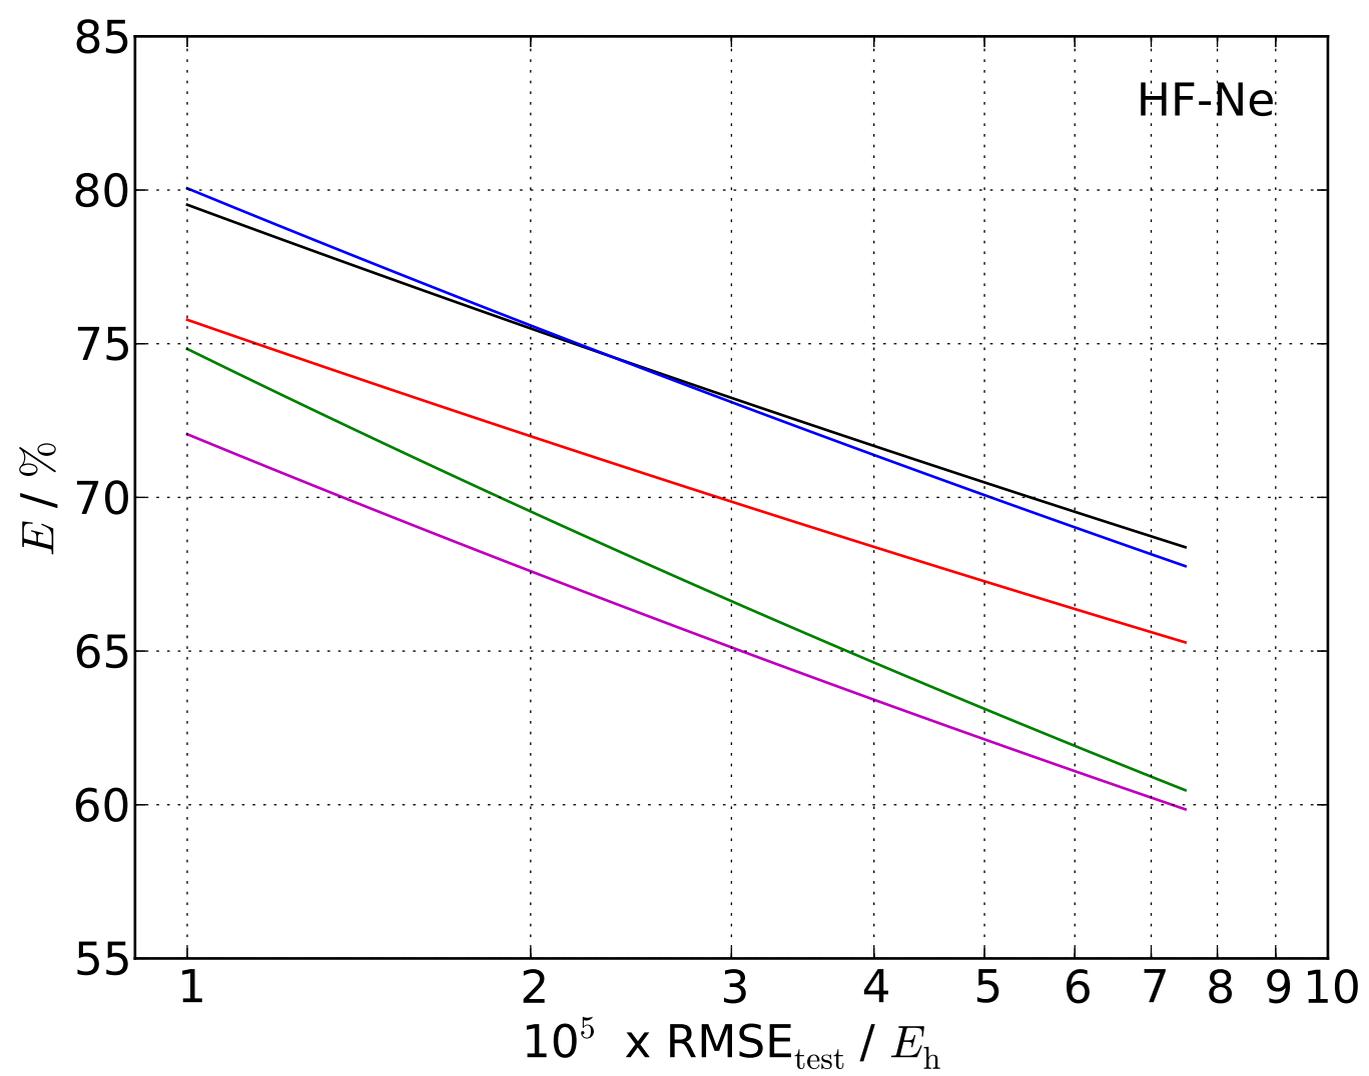
- 2058–2080 (1996).
- <sup>49</sup>H. Robbins, “Some aspects of the sequential design of experiments,” *Bulletin of the American Mathematical Society* **58**, 527–535 (1952).
- <sup>50</sup>H. Chernoff, “Sequential design of experiments,” *The Annals of Mathematical Statistics* **30**, 755–770 (1959).
- <sup>51</sup>Y. Zuo, C. Chen, X. Li, Z. Deng, Y. Chen, J. Behler, G. Csányi, A. V. Shapeev, A. P. Thompson, M. A. Wood, *et al.*, “Performance and Cost Assessment of Machine Learning Interatomic Potentials,” *J. Phys. Chem. A* **124**, 731–745 (2020).
- <sup>52</sup>D. J. MacKay, “Introduction to Gaussian processes,” *NATO ASI series F computer and systems sciences* **168**, 133–166 (1998).
- <sup>53</sup>M. N. Gibbs and D. J. MacKay, “Variational Gaussian process classifiers,” *IEEE Transactions on Neural Networks* **11**, 1458–1464 (2000).
- <sup>54</sup>C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006).
- <sup>55</sup>C. Qu, Q. Yu, B. L. Van Hoozen Jr, J. M. Bowman, and R. A. Vargas-Hernández, “Assessing Gaussian process regression and permutationally invariant polynomial approaches to represent high-dimensional potential energy surfaces,” *J. Chem. Theory Comput.* **14**, 3381–3396 (2018).
- <sup>56</sup>J. W. Ng and M. P. Deisenroth, “Hierarchical mixture-of-experts model for large-scale gaussian process regression,” *arXiv preprint arXiv:1412.3078* (2014).
- <sup>57</sup>T. T. Nguyen, E. Székely, G. Imbalzano, J. Behler, G. Csányi, M. Ceriotti, A. W. Götz, and F. Paesani, “Comparison of permutationally invariant polynomials, neural networks, and gaussian approximation potentials in representing water interactions through many-body expansions,” *J. Chem. Phys.* **148**, 241725 (2018).
- <sup>58</sup>A. Kamath, R. A. Vargas-Hernández, R. V. Krems, T. Carrington Jr, and S. Manzhos, “Neural networks vs Gaussian process regression for representing potential energy surfaces: A comparative study of fit quality and vibrational spectrum accuracy,” *J. Chem. Phys.* **148**, 241702 (2018).
- <sup>59</sup>R. V. Krems, “Bayesian machine learning for quantum molecular dynamics,” *Phys. Chem. Chem. Phys.* **21**, 13392–13410 (2019).
- <sup>60</sup>M. P. Metz and K. Szalewicz, “A statistically guided grid generation method and its application to intermolecular potential energy surfaces,” *J. Chem. Phys.* **152**, 134111 (2020).
- <sup>61</sup>J. Cui and R. V. Krems, “Efficient non-parametric fitting of potential energy surfaces for polyatomic molecules with Gaussian processes,” *J. Phys. B At. Mol. Opt.* **49**, 224001 (2016).
- <sup>62</sup>A. Stone, *The Theory of Intermolecular Forces* (OUP oxford, 2013).
- <sup>63</sup>L. E. Atlas, D. A. Cohn, and R. E. Ladner, “Training connectionist networks with queries and selective sampling,” in *Advances in neural information processing systems* (Citeseer, 1990) pp. 566–573.
- <sup>64</sup>D. J. MacKay, “Information-based objective functions for active data selection,” *Neural computation* **4**, 590–604 (1992).
- <sup>65</sup>L. Torrey and J. Shavlik, “Transfer learning,” in *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques* (IGI global, 2010) pp. 242–264.
- <sup>66</sup>S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering* **22**, 1345–1359 (2009).
- <sup>67</sup>C. Møller and M. S. Plesset, “Note on an approximation treatment for many-electron systems,” *Phys. Rev.* **46**, 618 (1934).
- <sup>68</sup>R. J. Bartlett and M. Musiał, “Coupled-cluster theory in quantum chemistry,” *RMP* **79**, 291 (2007).
- <sup>69</sup>H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, and M. Schütz, “Molpro: a general-purpose quantum chemistry program package,” *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **2**, 242–253 (2012).
- <sup>70</sup>V. Fock, “Näherungsmethode zur lösung des quantenmechanischen mehrkörperproblems,” *Zeitschrift für Physik* **61**, 126–148 (1930).
- <sup>71</sup>H.-J. Werner and P. J. Knowles, “An efficient internally contracted multiconfiguration–reference configuration interaction method,” *The Journal of chemical physics* **89**, 5803–5814 (1988).
- <sup>72</sup>P. J. Knowles and H.-J. Werner, “An efficient method for the evaluation of coupling coefficients in configuration interaction calculations,” *Chemical physics letters* **145**, 514–522 (1988).
- <sup>73</sup>GPY, “GPY: A gaussian process framework in python,” <http://github.com/SheffieldML/GPy> (since 2012).
- <sup>74</sup>J. Broad, R. J. Wheatley, and R. S. Graham, “Boundary Optimisation Data,” <https://doi.org/10.6084/m9.figshare.16610713.v1> (2021).

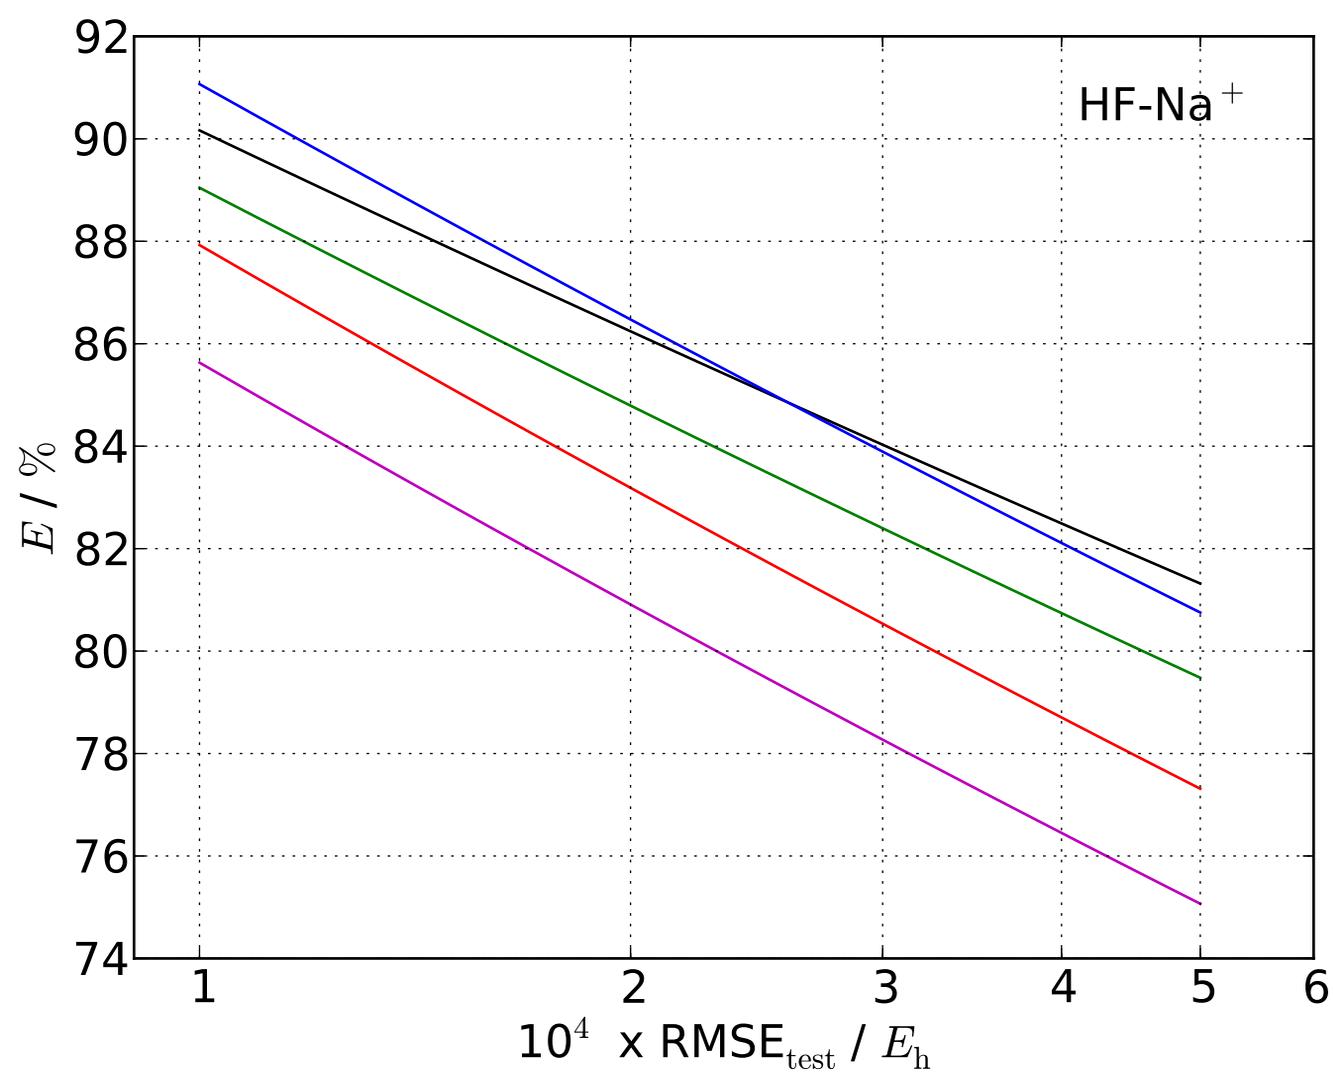


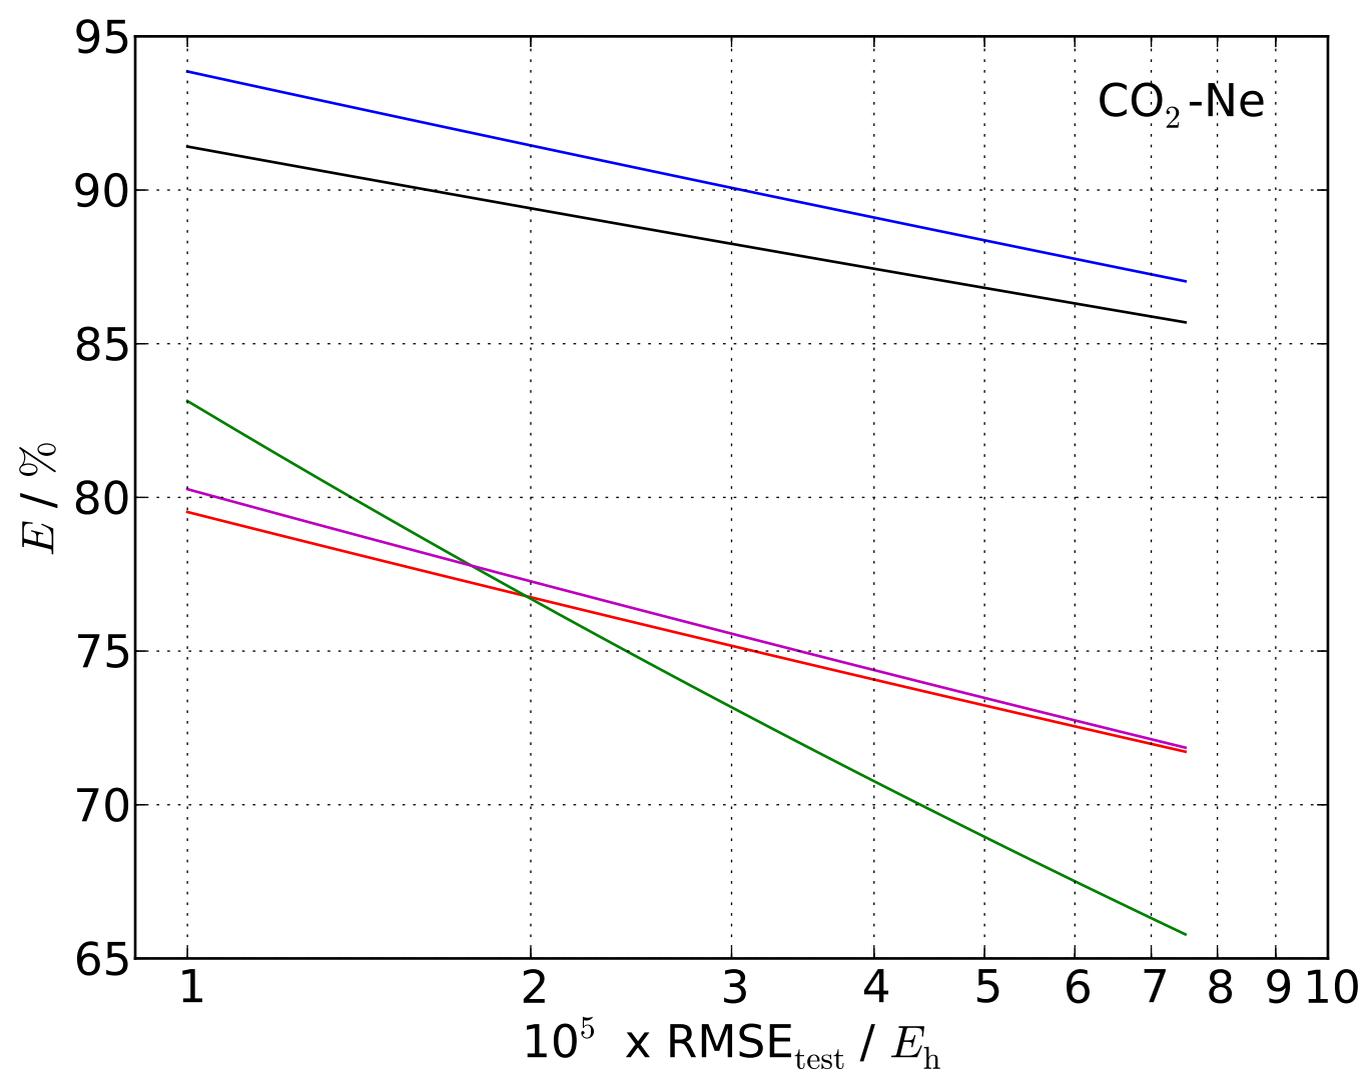


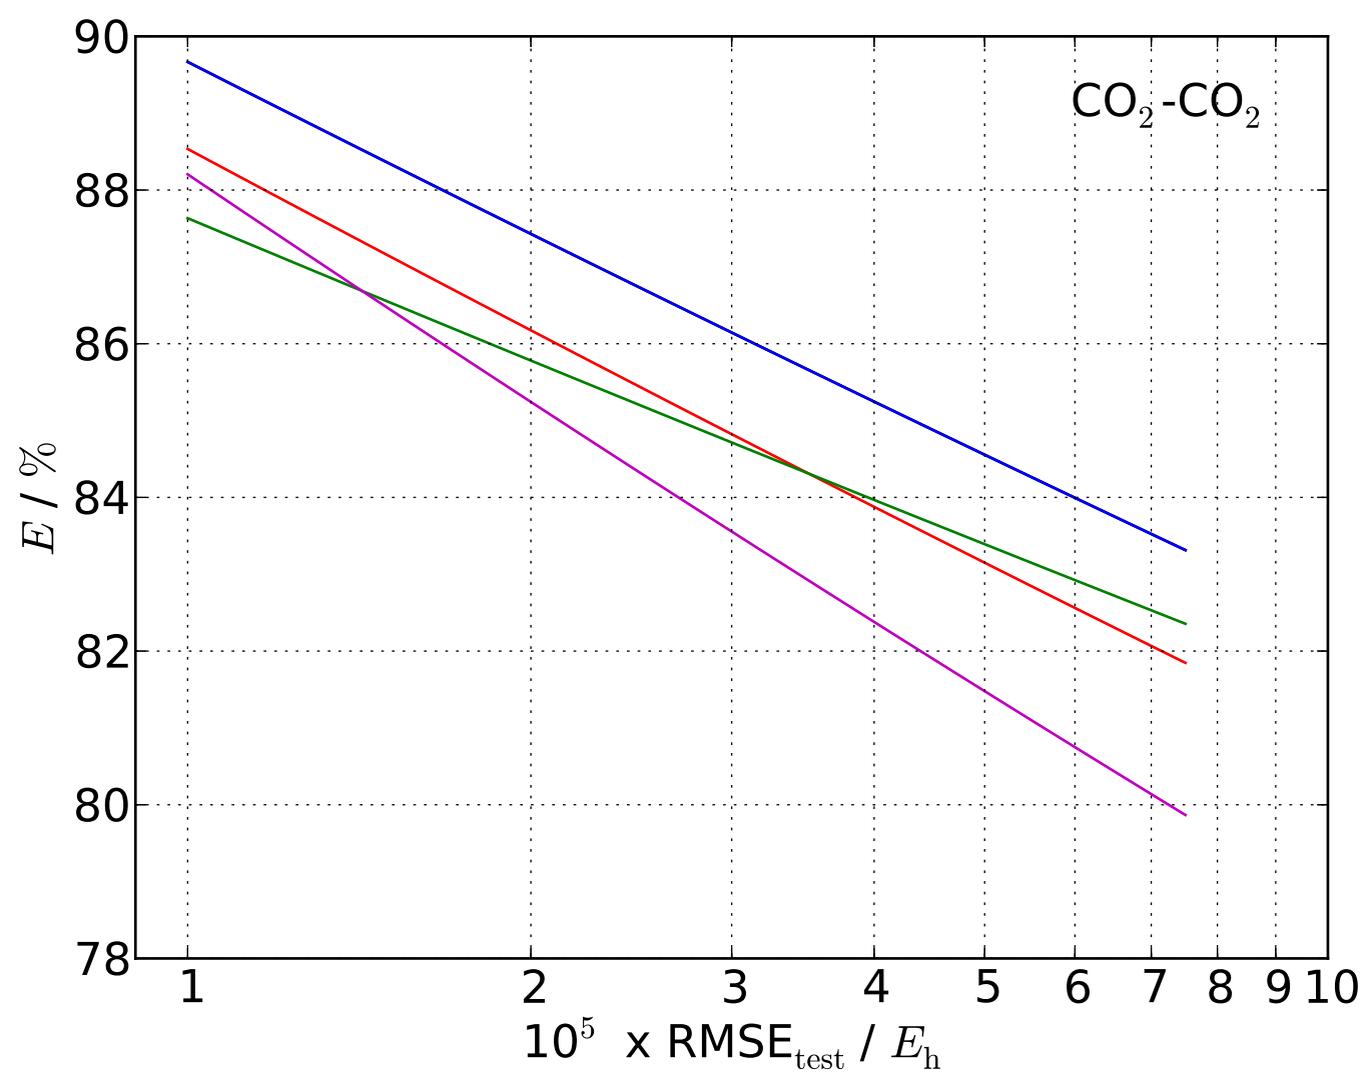




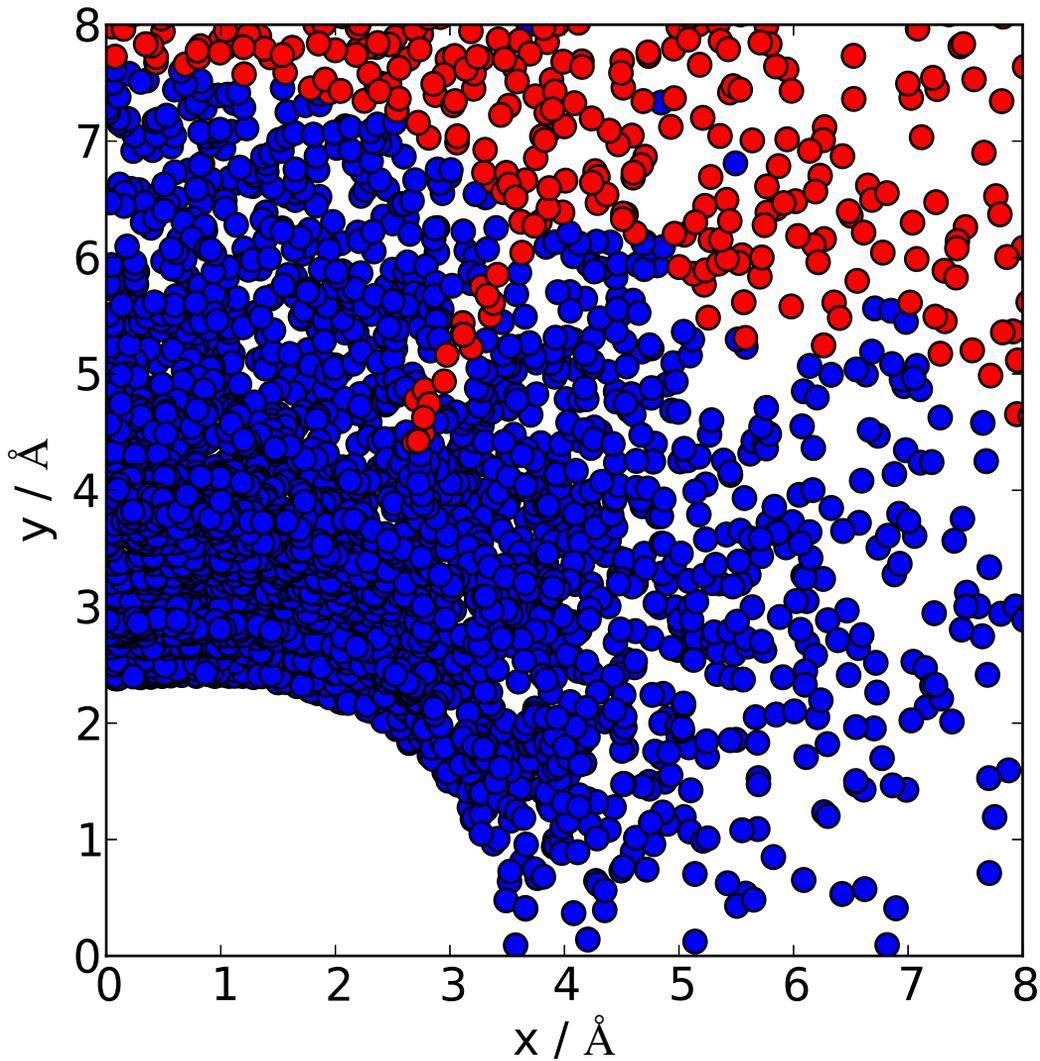


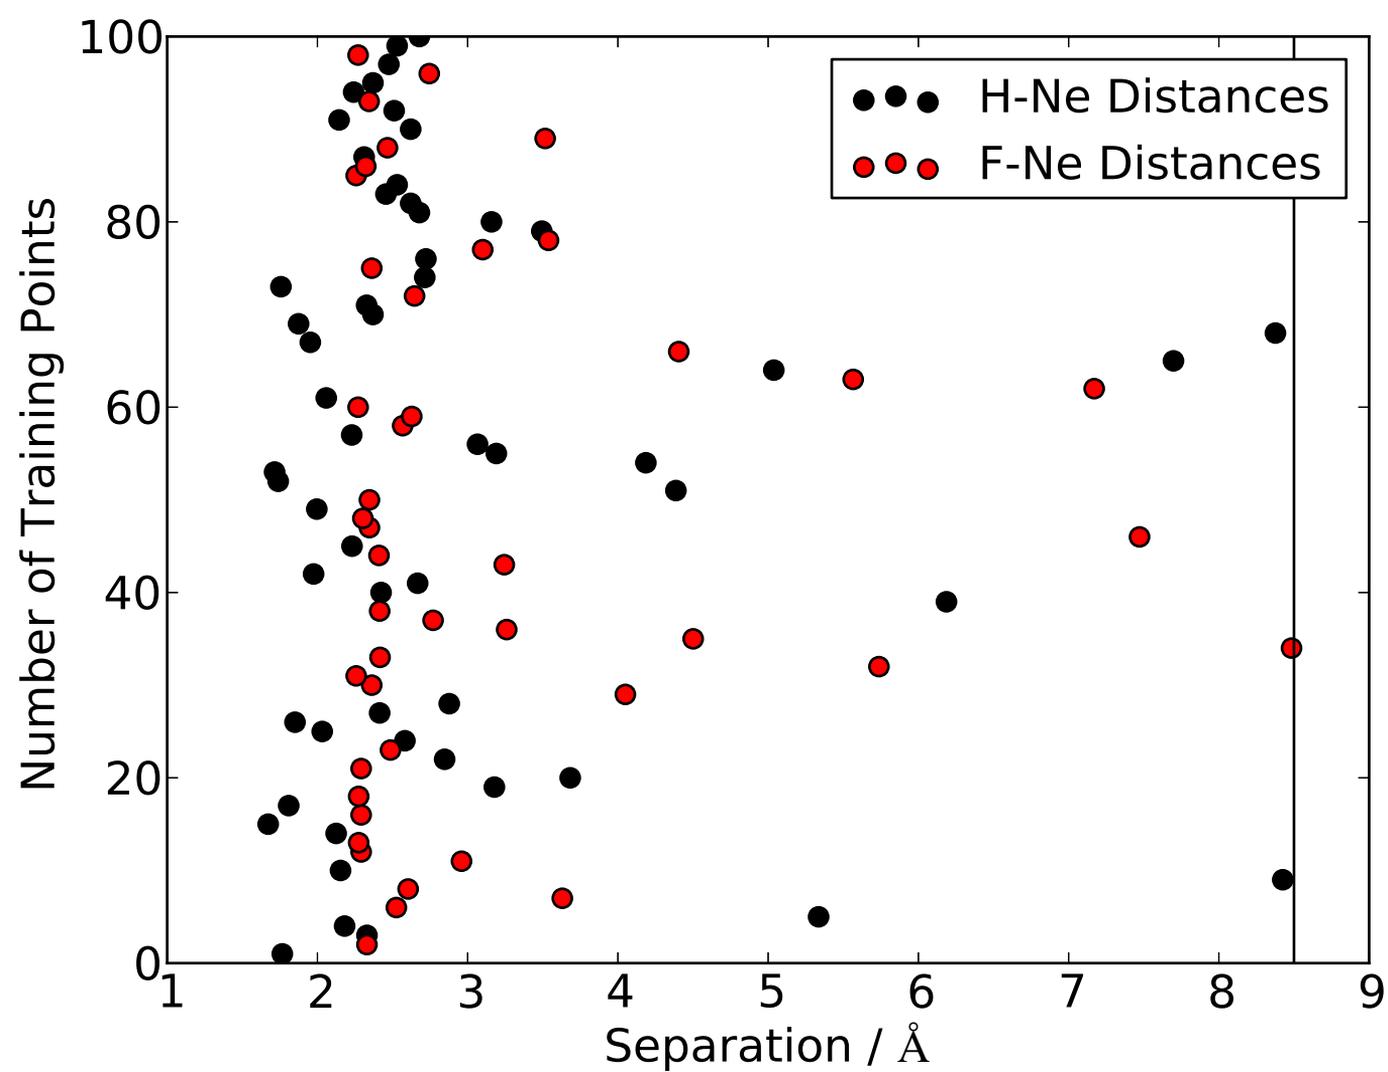




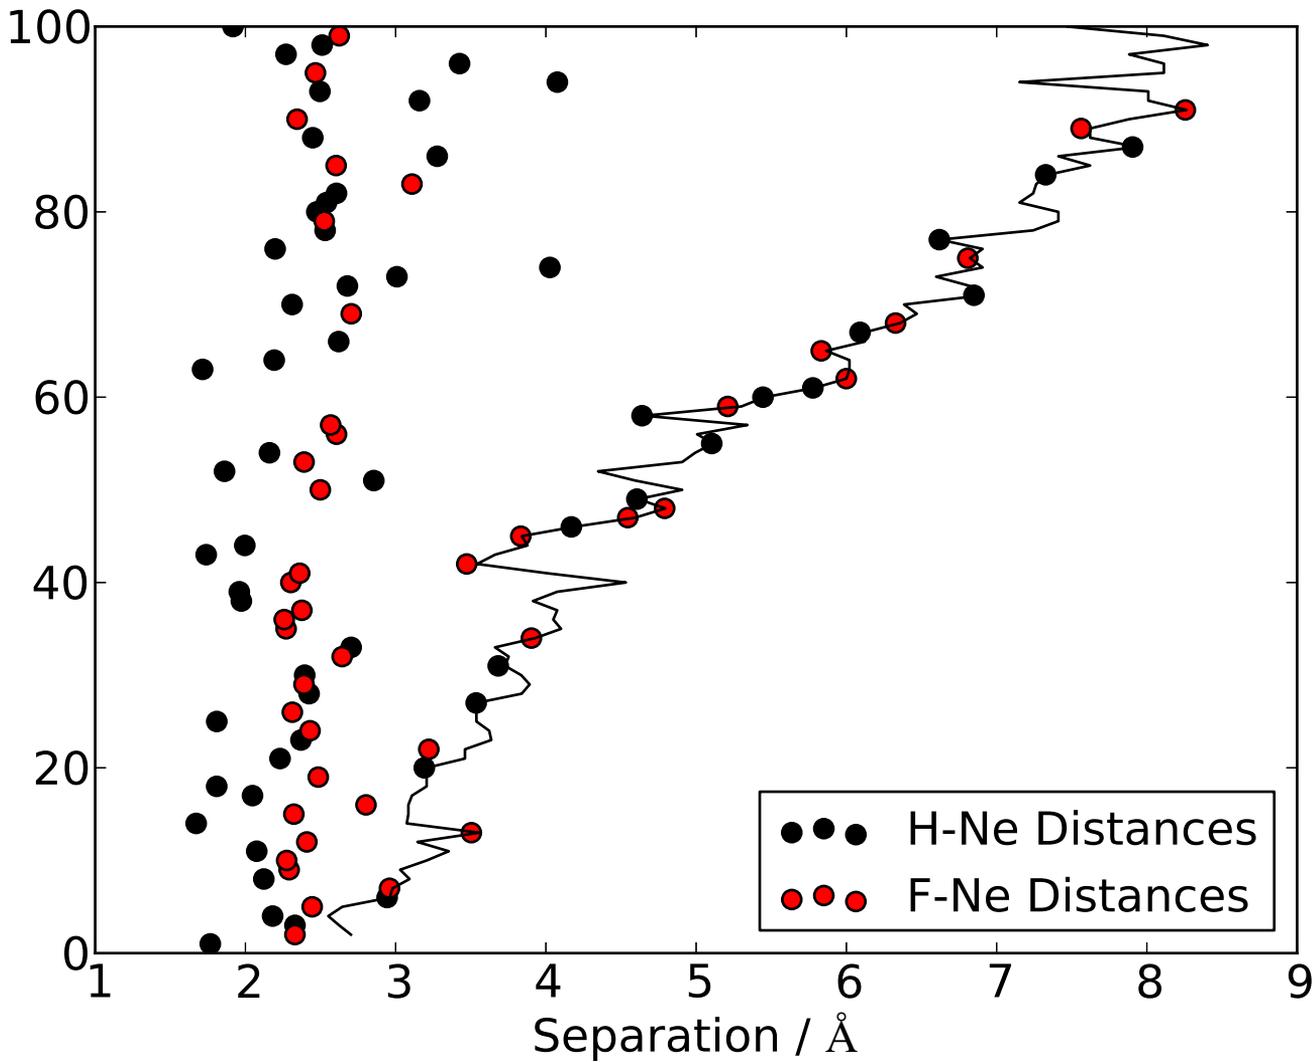


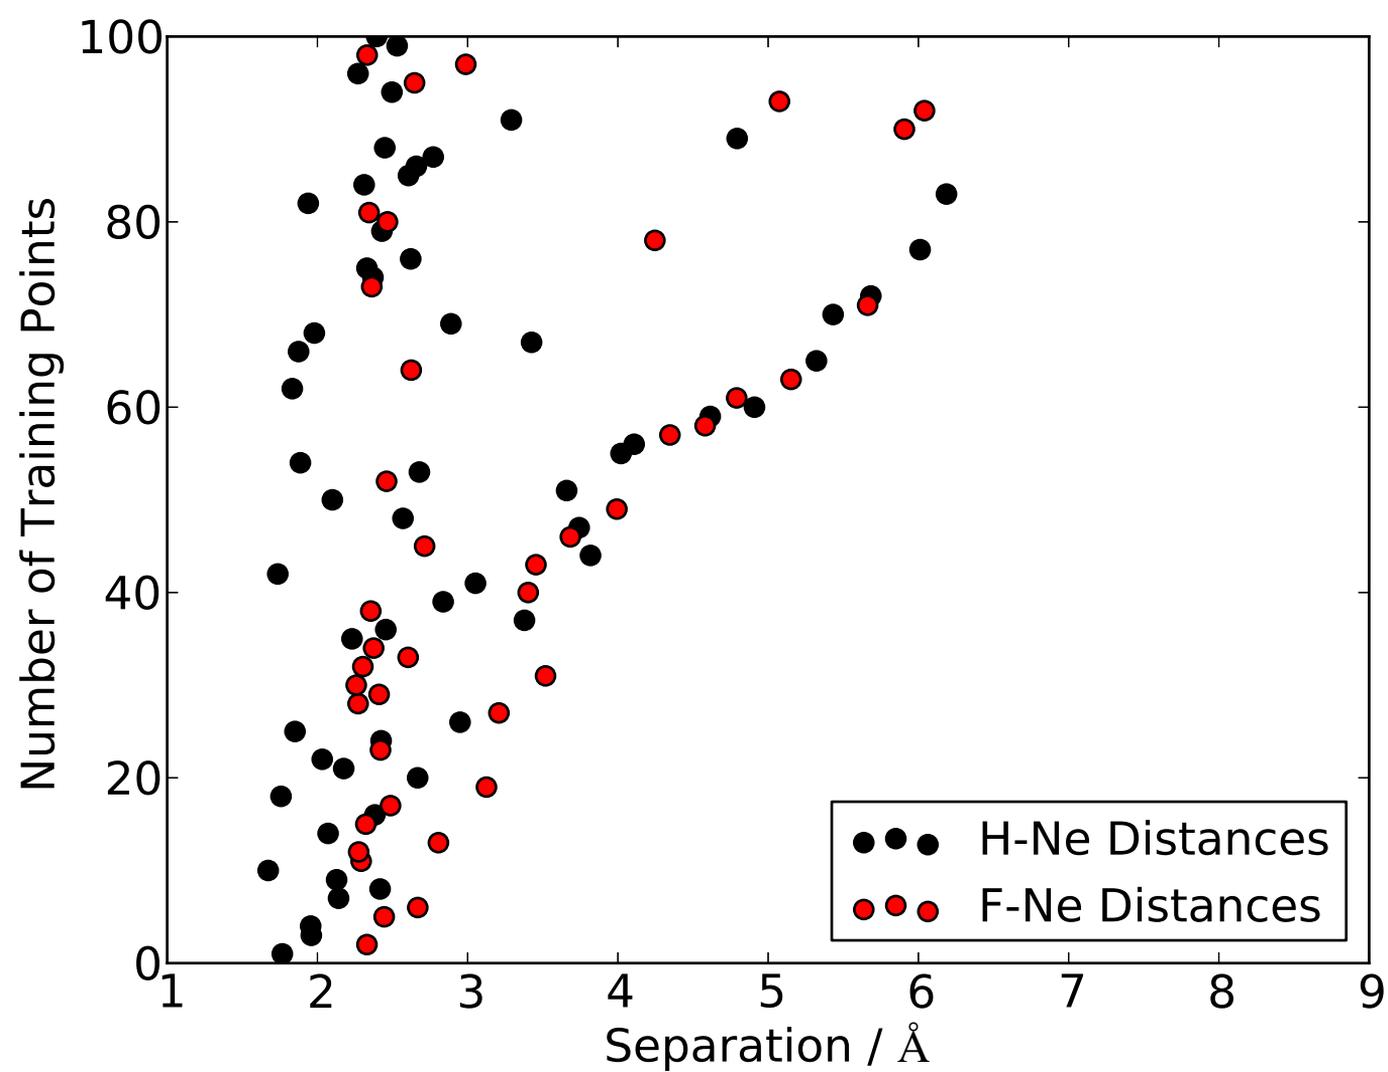
- Single-constrained Training
- Single-open Training
- Multi-constrained Training
- Multi-open Training
- Closest Model Training



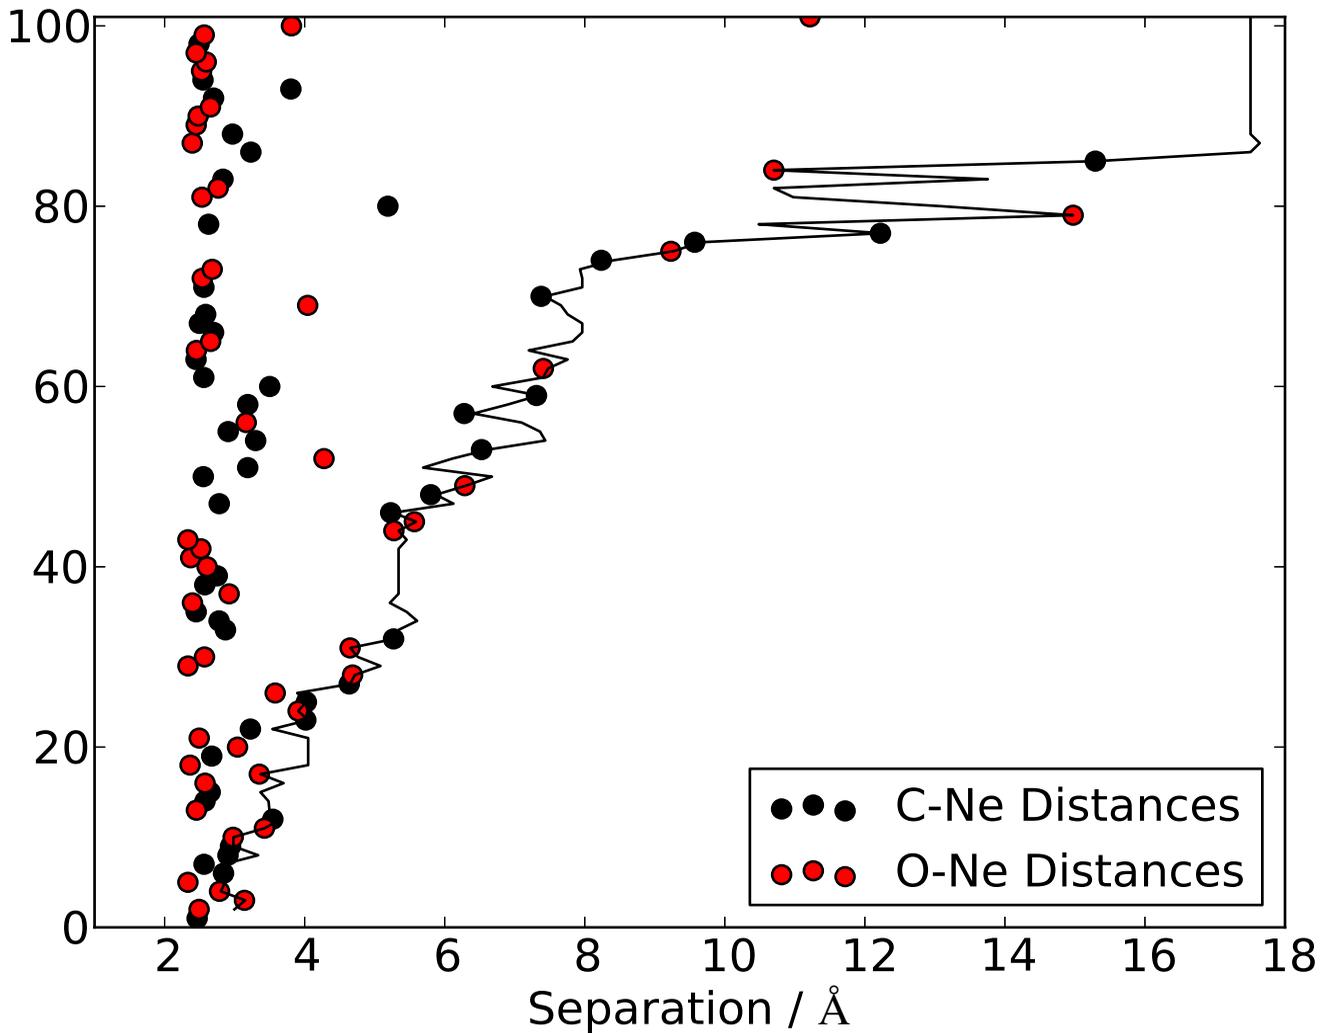


Number of Training Points

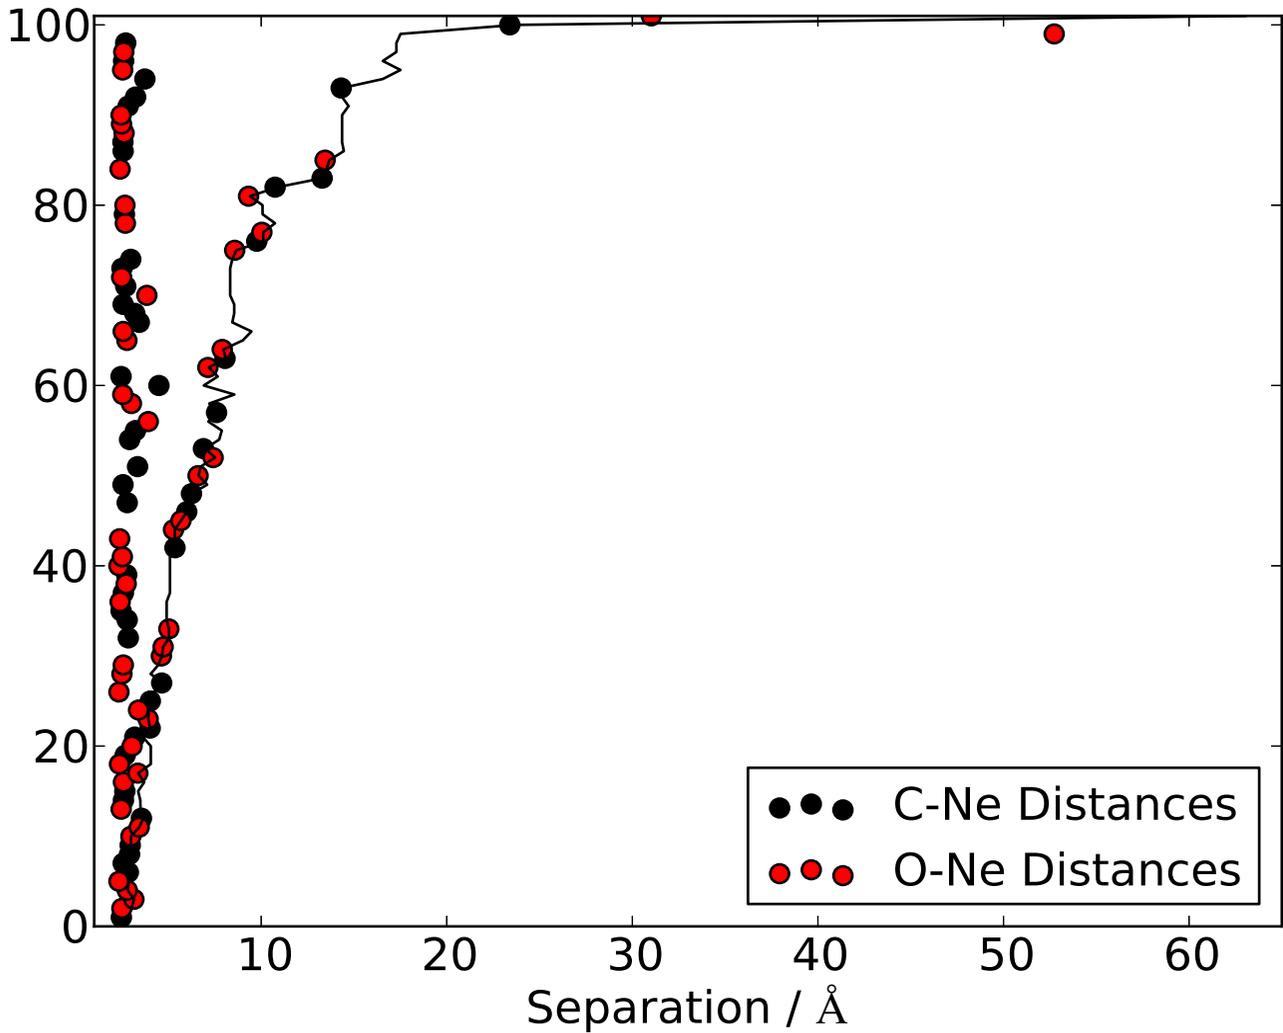


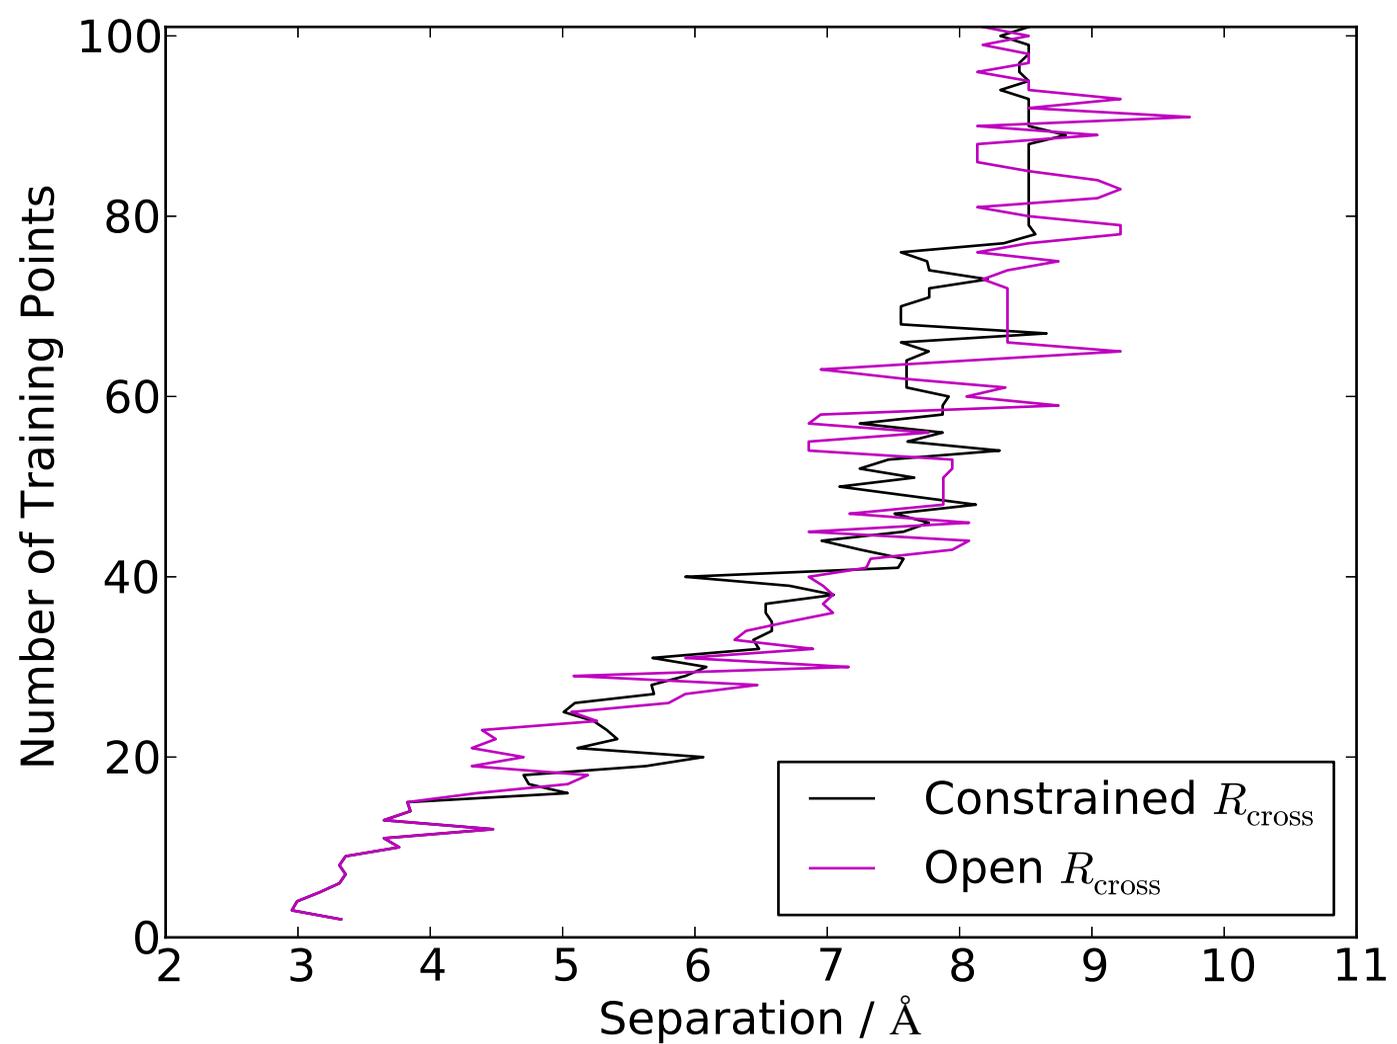


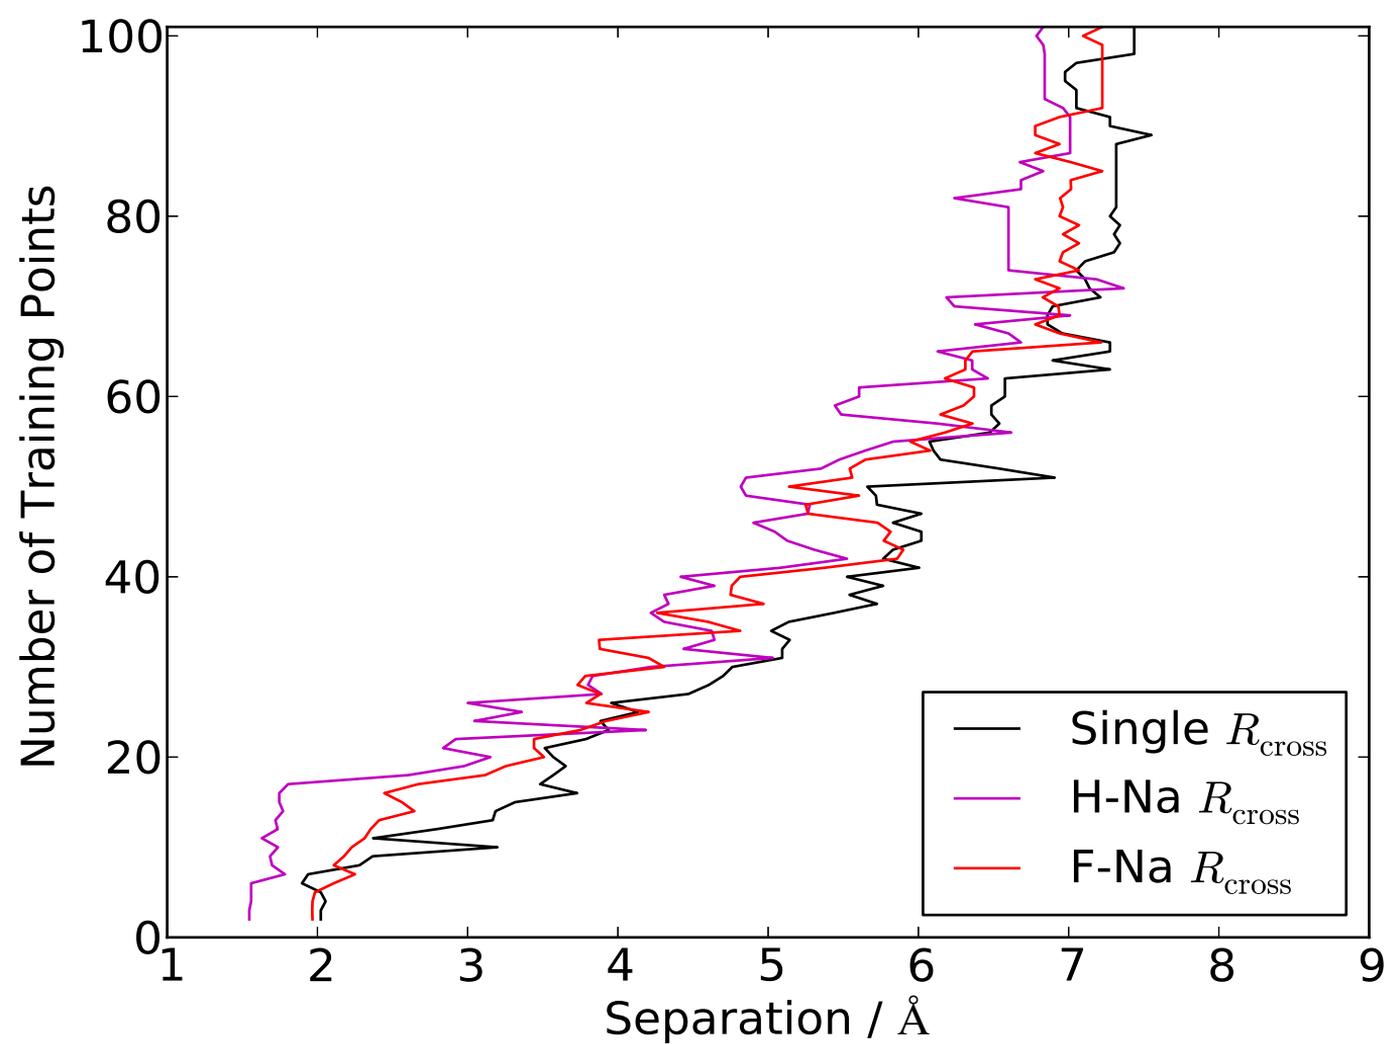
Number of Training Points

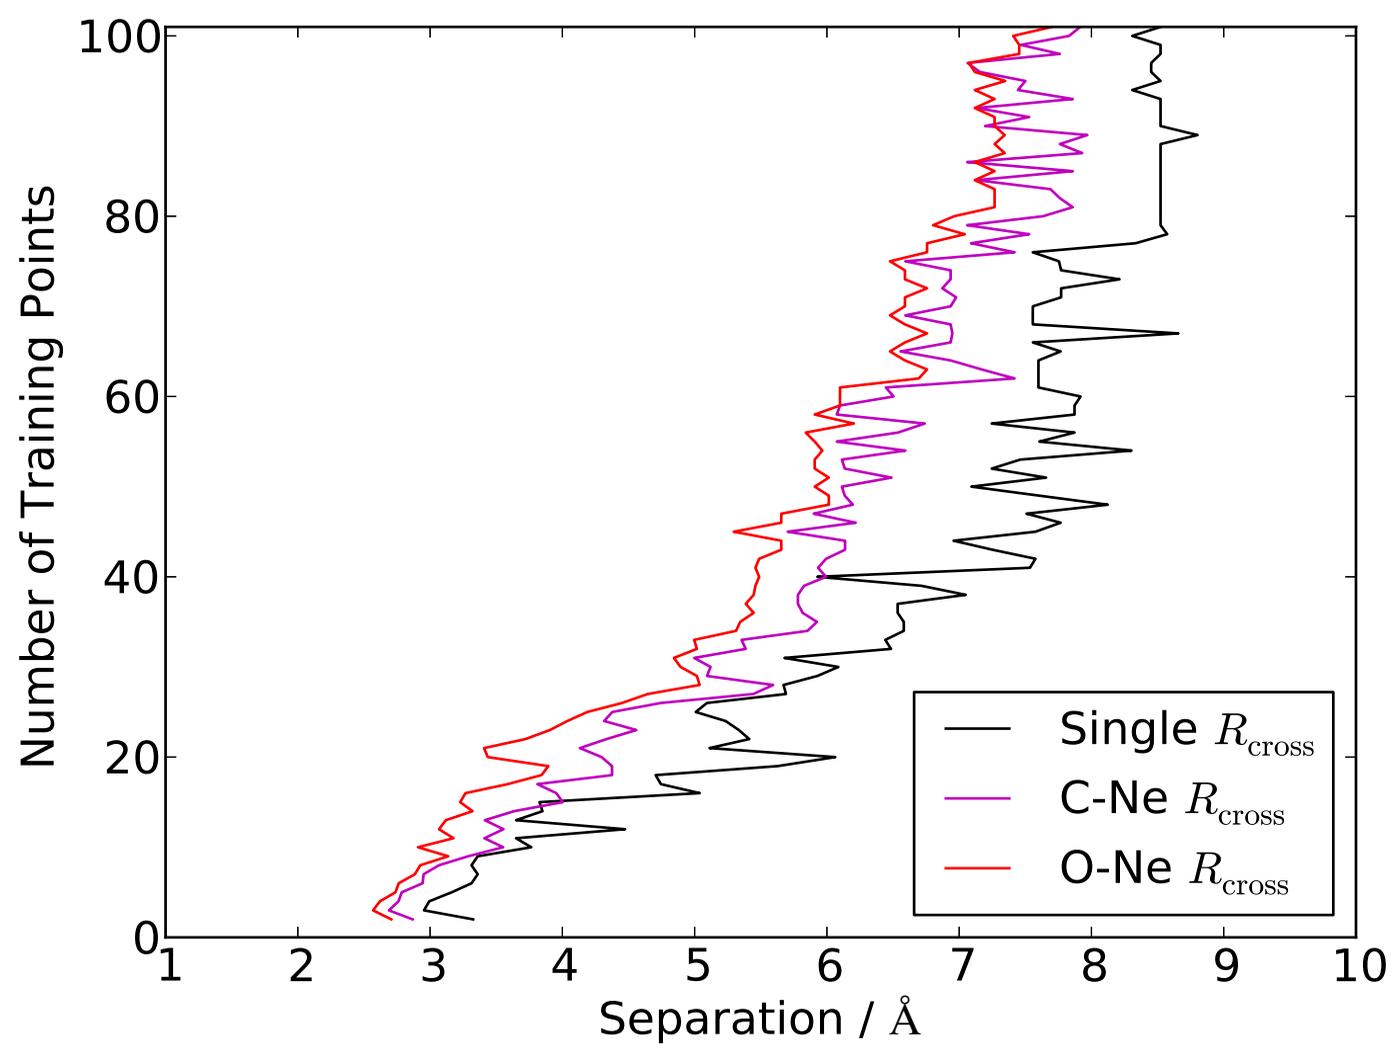


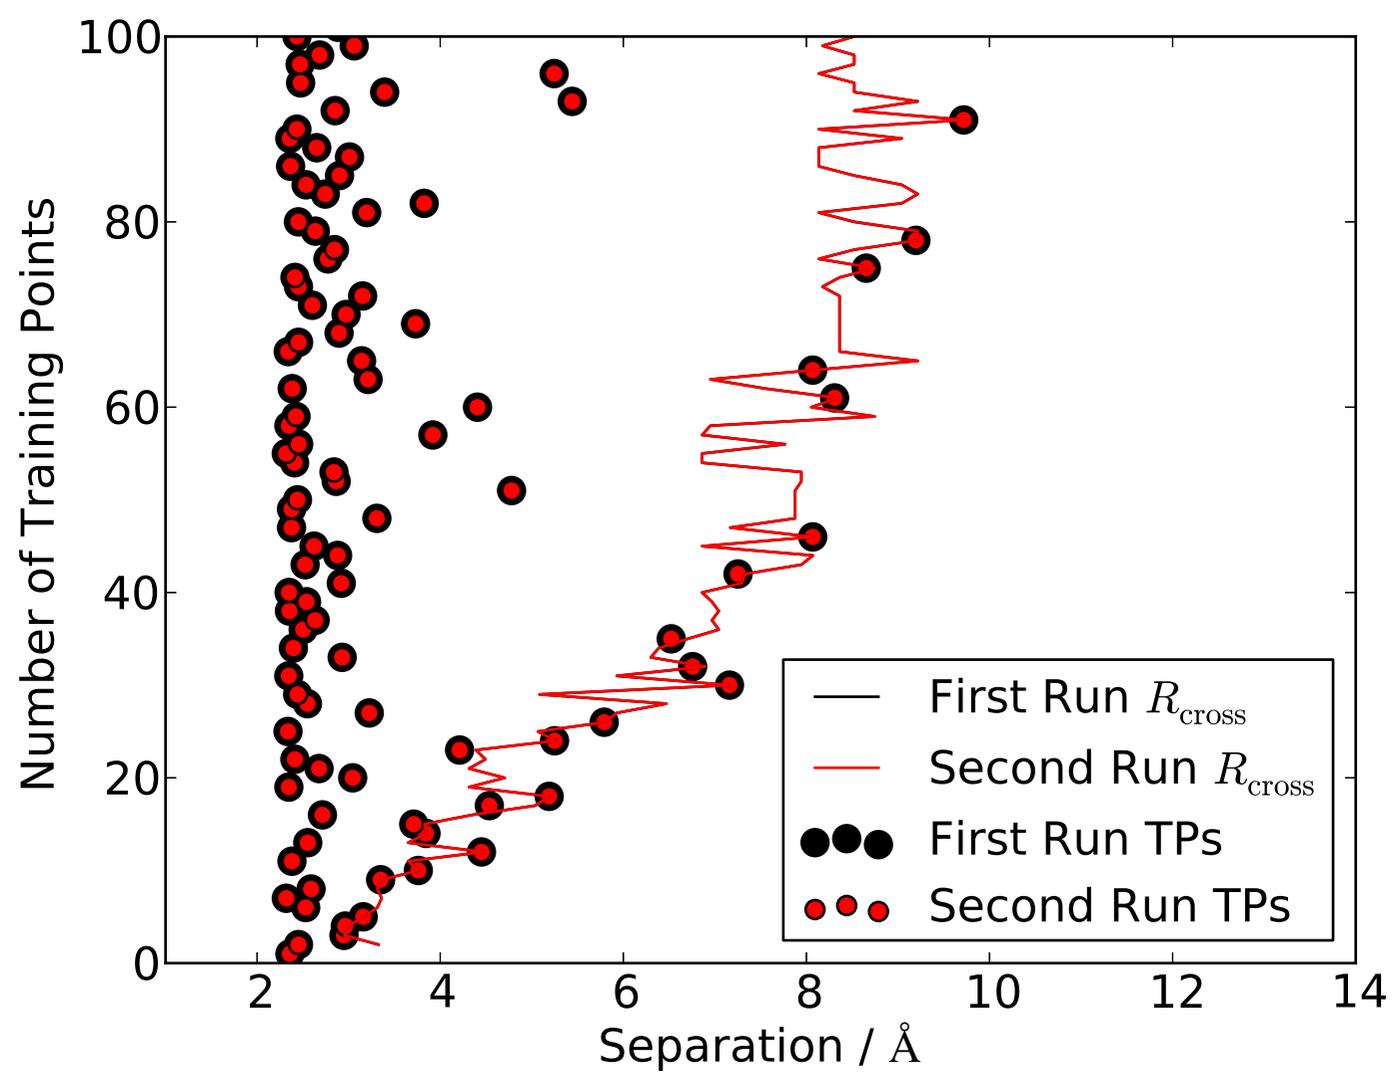
Number of Training Points











Number of Training Points

