On the capacity and superposition of minima in neural network loss function landscapes

IOP Publishing Mach. Learn.: Sci. Technol. 3 (2022) 025004 https://doi.org/10.1088/2632-2153/ac64e6

MACHINE LEARNING
Science and Technology

OPEN ACCESS

RECEIVED
10 November 2021

REVISED
22 March 2022

ACCEPTED FOR PUBLICATION
6 April 2022

PUBLISHED
20 April 2022

PAPER

On the capacity and superposition of minima in neural network loss function landscapes

Maximilian P Niroomand*, John W R Morgan, Conor T Cafolla and David J Wales*

Department of Chemistry, University of Cambridge, Cambridge, United Kingdom
* Authors to whom any correspondence should be addressed.

E-mail: mpn26@cam.ac.uk **and** dw34@cam.ac.uk

Keywords: ensemble learning, interpretability, loss function landscape, theoretical chemistry

Abstract

Minima of the loss function landscape (LFL) of a neural network are locally optimal sets of weights that extract and process information from the input data to make outcome predictions. In underparameterised networks, the capacity of the weights may be insufficient to fit all the relevant information. We demonstrate that different local minima specialise in certain aspects of the learning problem, and process the input information differently. This effect can be exploited using a meta-network in which the predictive power from multiple minima of the LFL is combined to produce a better classifier. With this approach, we can increase the area under the receiver operating characteristic curve by around 20% for a complex learning problem. We propose a theoretical basis for combining minima and show how a meta-network can be trained to select the representative that is used for classification of a specific data item. Finally, we present an analysis of symmetry-equivalent solutions to machine learning problems, which provides a systematic means to improve the efficiency of this approach.

1. Introduction

Deep learning with neural networks is a high-dimensional, non-convex optimisation problem for a loss function landscape (LFL). The coordinates of a *minimum* in the LFL are a *set of weights* for the machine learning model and a locally optimal *solution* to the learning problem, and these terms will therefore be used interchangeably throughout. It follows that the coordinates of the global minimum of the LFL are the weights that produce the lowest possible value of the loss function for the training data. The aim of machine learning is usually for the model to find a set of weights that fit the training data, but also generalise well to unseen testing data. Our approach extends this view. Instead of looking at just one minimum of the LFL, we are interested in the expressive power of multiple minima. To analyse how different minima extract and process information from the input data, we survey numerous low-lying minima of the LFL. Here, we employ tools from the energy landscape approach (Wales 2003) to gain new insight into machine learning LFLs (Ballard *et al* 2017). We note that the role of local minima is somewhat different in ML landscapes compared to molecular systems. While in a molecular energy landscape only minima provide valid configurations for a stable molecule, this restriction does not apply to LFLs for machine learning. In fact, some low-lying non-minima will have a smaller loss value and higher classification accuracy than a high-lying minimum. Here, we are interested in developing a better understanding of the capacity of diverse minima of the LFL, and we show that by combining the expressive power of different minima, we can build a better classifier. The compact form of this predictor provides a balance between accuracy and efficiency, which will be useful in applications where evaluation is a computational bottleneck.

1.1. Background

Machine learning models are structurally limited in the amount of data they can fit: their capacity is finite. The most commonly known measure of capacity is perhaps the Vapnik-Chervonenkis (VC) dimension (Vapnik and Chervonenkis 1971, Vapnik *et al* 1994). The higher the VC dimension, the more complex are the

data that can be fitted. More rigorously, VC dimension is defined as the largest cardinality of a set of data points that the NN can shatter (for our purpose, shatter means classify correctly). Thus, the weights of an underparameterised model (i.e. fewer parameters than training data points) may be incapable of fitting the entire test data set, but instead fit just parts of it.

The approach we employ to study the expressive power of combinations of individual minima is a variation of ensemble learning, where the results of multiple different predictors are combined to improve the overall accuracy of an approximation problem (Dong *et al* 2020). The idea of combining multiple sources of information, specifically the output predictions of multiple classifiers, has been considered for over two decades (Breiman 1996, Hashem 1997, Jin and Lu 2009). Two of the most important questions in ensemble learning are: which classifiers to consider, and how to combine the individual predictions (Wang 2008). For a detailed review see Kuncheva (2014).

1.2. Motivation

In this contribution, we are interested in quantitatively and systematically characterising a cornerstone of ensemble learning, namely classifier diversity. Logically, ensemble learning works if different classifiers extract different information from the input data or process it differently (Melville and Mooney 2005, Zaidi *et al* 2020). In the present work, the classifiers in question correspond to local minima of a reference neural network. We aim to visualise diversity of minima from the corresponding LFL and show how to select a few of them to produce a compact, yet more accurate, classification. We will show that different minima of the LFL successfully classify distinct subsets of the entire input dataset. Hence different local minima specialise in distinct parts of the test dataset, which we believe has not been shown before.

In summary, our main contributions are:

- Showing that different local minima specialise in distinct subsets of the input;
- MLSUP, a proof-of-concept method that exploits minima diversity to improve classification results for complex problems;
- An interpretation of the limitation of single-minimum models and visualisation of the differences between minima;
- Novel insights into the symmetry properties of minima in neural network LFLs.

2. Superposition of machine learning solutions: MLSUP

We will show that different local minima of a reference neural network extract different information from the input data and that combining just a few examples can improve classification significantly. To study this effect, we employ a modified stacking approach where multiple minima from the same classifier are combined, rather than multiple classifiers. We do not obtain these minima by different random initialisation but rather from exploring the LFL. This approach provides insight into the functional landscape and a deeper understanding of LFL minima. To answer the second important question in ensemble learning design, we employ a second neural network to select one of the local minima for a given input data item. This idea is related to previous theory, Jordan and Jacobs (1994) where a gating network chooses which classifier to apply to some problem (McGill and Perona 2017, Shazeer *et al* 2017). We call our method MLSUP, to denote a superposition of machine learning solutions (local minima of the LFL). We describe MLSUP in a four step process (figure 1).

The first step involves characterising local minima $\mathcal{M}$ by exploring the LFL during training. Next, we choose a subset of minima $\mathcal{M}' \subseteq \mathcal{M}$ and evaluate each $m \in \mathcal{M}'$ for every training datapoint, which reveals how well each of them can classify specific data items (step 2 in figure 1). A detailed discussion of how a few minima are selected for combination is included below. The superposition of chosen minima is done by training a second, meta-network (classifier 2, i.e. step 3 in figure 1) to learn which of the $m \in \mathcal{M}'$ minima is best suited to classify a specific input datapoint. Thus, the second network learns to apply different minima to classify different types of input data, as shown in Step 4 of figure 1. A pseudocode version of MLSUP is provided in the appendix.

3. Model

We consider a classification problem for C classes with a single hidden layer, as we are specifically interested in underparameterised networks. For some data $\mathcal{D} = (\mathbf{X}, \mathbf{c})$, the inputs are denoted $\mathbf{X} = \{\mathbf{x}^1, \ldots, \mathbf{x}^N\}$, where N is the number of data points in the training or testing set, which are denoted as $\mathbf{X}_{train}$ and $\mathbf{X}_{test}$ respectively. The correct label for some data point d is defined as c^d. We use tanh as the nonlinear activation function to the hidden layer, since it has continuous derivatives, which we require for optimisation. Outputs at node y_i

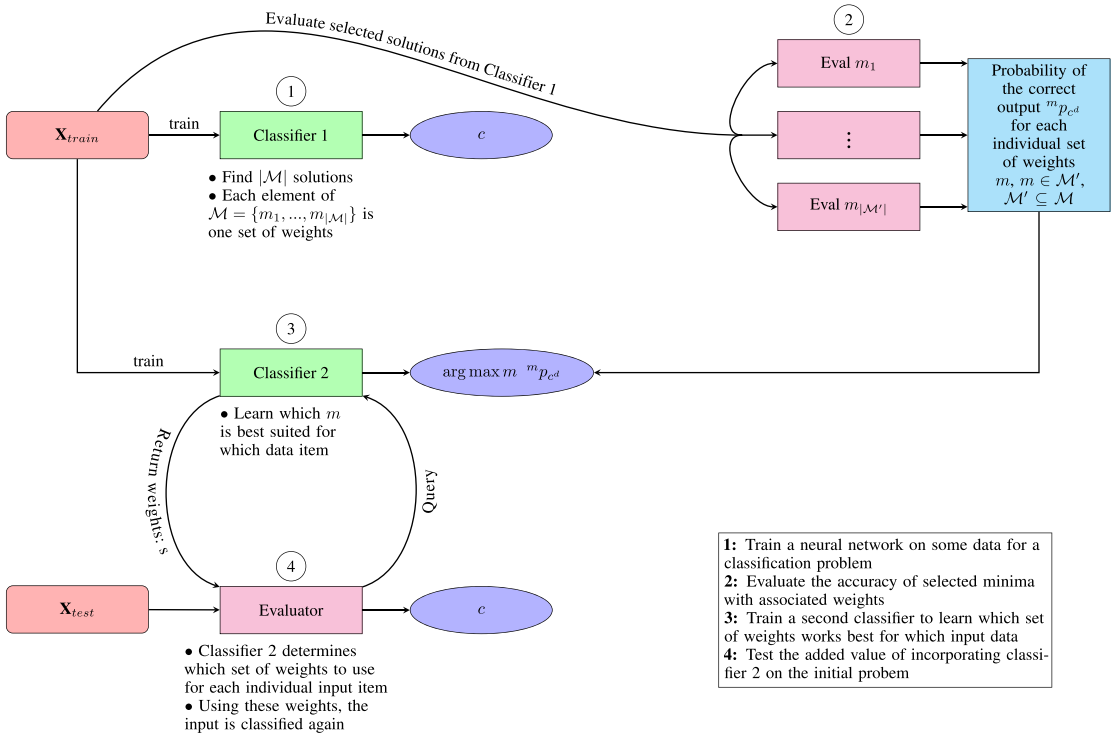

Figure 1. Diagrammatic representation of the MLSUP learning procedure. Ellipsoid boxes contain the true outputs, which are compared to the predicted output for training/evaluation purposes.

for input data item $\mathbf{x}$ are converted to softmax probabilities $p_i(\mathbf{W};\mathbf{x}) = \exp(y_i)/\exp(\sum_j e^{y_j})$, where $\mathbf{W}$ denotes the vector containing all weights. During training, we minimise a loss function $L(\mathbf{W};\mathbf{X})$ with respect to these weights. We use a cross-entropy loss function:

$$L(\mathbf{W};\mathbf{X}) = -\frac{1}{N}\sum_{d=1}^{N} \ln p_{c^d}(\mathbf{W};\mathbf{X}) + \lambda \mathbf{W}^2, \tag{1}$$

where c^d is the correct class for some data item $\mathbf{x}^d$. A L2 regularisation term $\lambda\mathbf{W}^2$ is added to eliminate zero Hessian eigenvalues, which by Noether's theorem arise as a consequence of continuous symmetries in the loss function when additively shifting all the output bias weights (Ballard *et al* 2017). We find $\lambda = 10^{-5}$ to be appropriate for the tests considered; our conclusions are largely insensitive to λ. This setup is used for the neural networks in steps 1, 2 and 4 of figure 1.

3.1. Defining the meta-network loss function

For step 3 of figure 1, the loss function is different from equation (1). This section describes Classifier 2 from figure 1, which is distinct from the one used for steps 1, 2 and 4. For Classifier 2, we are not interested in learning c^d, i.e. the correct output class, but rather the best local minimum to classify some input data item d, which is defined by the highest corresponding probability:

$$b^d = \arg\max m \ {}^{m}p_{c^d} \ \forall\, m \in \mathcal{M}', \tag{2}$$

for data item d, where m is one set of weights. This formulation changes the loss function to:

$$L(\widetilde{\mathbf{W}};\mathbf{X}) = -\frac{1}{N}\sum_{d=1}^{N} \ln p_{b^d}(\widetilde{\mathbf{W}};\mathbf{X}) + \lambda \widetilde{\mathbf{W}}^2, \tag{3}$$

with $\widetilde{\mathbf{W}}$ representing the weights for network 2. We evaluate the classification predictions of our model using the area under the receiver operating characteristic curve (ROC-AUC) (Fawcett 2006). The change in loss function also impacts the way we calculate the AUC for Step 4 in figure 1. In the usual case, the AUC is given as:

$$\mathrm{AUC} = \int_0^1 T(P)dF(P), \tag{4}$$

with the true positive rate for outcome number one, $T(P)$, and the false positive rate, $F(P)$:

$$T(P) = \frac{\sum_d^{N_{\text{data}}} \delta(c^d - 1)\Theta(p_1 - P)}{\sum_d^{N_{\text{data}}} \delta(c^d - 1)} \qquad F(P) = \frac{\sum_d^{N_{\text{data}}} [1 - \delta(c^d - 1)]\Theta(p_1 - P)}{\sum_d^{N_{\text{data}}} 1 - \delta(c^d - 1)}, \tag{5}$$

where $\delta(c^d - 1)$ is the Dirac delta function, and $\Theta(p_1 - P)$ the Heaviside step function, defined as:

$$\delta(c^d - 1) = \begin{cases} 1 & \text{if } c^d = 1, \\ 0 & \text{if } c^d \neq 1 \end{cases} \qquad \Theta(p_1 - P) = \begin{cases} 1 & \text{if } p_1 \geqslant P, \\ 0 & \text{if } p_1 < P \end{cases} \tag{6}$$

However, we now have $|M'|$ possibilities. Thus, we must evaluate the AUC using the minimum b^d that is chosen by Classifier 2 from figure 1:

$$T(P) = \frac{\sum_d^{N_{\text{data}}} \delta(c^d - 1)\Theta(p_1^{b^d} - P)}{\sum_d^{N_{\text{data}}} \delta(c^d - 1)} \qquad F(P) = \frac{\sum_d^{N_{\text{data}}} [1 - \delta(c^d - 1)]\Theta(p_1^{b^d} - P)}{\sum_d^{N_{\text{data}}} 1 - \delta(c^d - 1)}, \tag{7}$$

Unless described otherwise, Classifier 2 is a single hidden-layer network with 3 hidden nodes.

3.2. Optimisation routine

To survey the LFL we employ methods from the energy landscape approach, which has been widely used to study molecular and condensed matter systems in the physical sciences (Wales 2003). Specifically, global optimisation is performed using the basin-hopping method (Li and Scheraga 1987, Wales and Doye 1997) with a customised quasi-Newton L-BFGS (Nocedal 1980) optimiser (appendix). Candidates for transition states are obtained using a doubly-nudged (Trygubenko and Wales 2004a, 2004b) elastic band (Henkelman *et al* 2000, Henkelman and Jónsson 2000) approach and accurately refined by hybrid-eigenvector following (Munro and Wales 1999, Zeng *et al* 2014). The transition states are identified solely for visualisation of the global organisation of the landscape using the OPTIM program. Transition states are clearly distinguished from local minima by the Hessian eigenvalue structure, and only local minima are explored in the global optimisation procedure by the GMIN program (see appendix). Transition states are not used in training or testing procedures. They are used to understand how the local minima obtained in training are organised in the LFL. All the routines employed in LFL exploration are implemented in the GMIN (Wales 2003a), OPTIM (Wales 2003b) and PATHSAMPLE (Wales 2003c) programs.

4. Minima selection

The problem analogous to classifier selection in ensemble learning for MLSUP is minima selection. Here our motivation is to identify a small number of minima that significantly enhance the accuracy of the predictions, to produce a compact representation suitable for use in large-scale simulations that are typical in physical science applications. Hence we need to identify complementary minima that can be combined to improve the classification. We compare two methods for minima selection: one based on the theory of a thermodynamic analogue for the LFL, the other on more abstract machine learning concepts. We will also explain why simply picking minima that are 'far away' in Euclidean distance can be ineffective.

4.1. Euclidean distance

The Euclidean distance between two weight vectors is simply $|\mathbf{w}_1 - \mathbf{w}_2| = \sqrt{(\mathbf{w}_1 - \mathbf{w}_2)^2}$. However, minima 'far away' from each other in Euclidean space may not extract different information from the training data. In fact, there are symmetry-related minima, potentially distant in Euclidean space, which for the same input data will return exactly the same loss value. We refer to these solutions as permutational invariants or isomers, as for molecular systems. This sort of problem was considered in Dinh *et al* (2017), where redundant minima are shown to exist for non-negative homogeneous nonlinear functions, such as the rectifier function $\phi_{rect}(x) = \max(x, 0)$. However, these results do not apply here because tanh is not homogeneous. Instead, let $\mathcal{G}$ denote the permutation group for some set of weights w, i.e. all permutations allowed such that for the same input, they return the same output. By shuffling hidden nodes (Brea *et al* 2019), and the fact that tanh is an odd function, we can show that for each minimum, there exist:

$$|\mathcal{G}| = \prod_{l=1}^{H} (n_l! \times 2^{n_l}) \quad \text{or} \quad |\mathcal{G}| = \prod_{l=1}^{H-1} (n_l! \times 2^{n_l}) \times n_H!, \tag{8}$$

degenerate solutions for each minimum, depending on whether the number of hidden layers H is odd or even, respectively, where n_l is the number of nodes in hidden layer l. Because tanh is an odd function, we can

change the sign of the weights before and after each hidden node and each combination of any number of nodes for an odd number of hidden layers. This result increases the per-layer number of degenerate minima from $n!$ to $n! \times 2^n$, i.e. for each node swap, any nodes in each element of the power set of the nodes $\mathbb{P}(n)$ could be multiplied by a factor of -1 with no effect on the output. Thus, randomly picking some 'far away' minima might just return permutationally invariant solutions, which does not lead to any performance improvement in MLSUP. For further interest, we have included in the appendix an example of four such permutational invariants and a pairwise distance matrix of permutational isomers of one minimum. In conclusion, a large Euclidean distance for minima selection is neither a necessary, nor a sufficient selection condition.

4.2. Misclassification space

Instead of computing the distance between two sets of weights in Euclidean space, we can compute a metric in misclassification space. The rationale behind this approach is that minima with a high pairwise misclassification rate will be complementary, and therefore good candidates for MLSUP. We denote a vector containing the predicted output class for each data item $i \in X$ for some minimum m as $\mathbf{c}_i^{(m)}$, and define the set $\Xi = \{n \mid \mathbf{c}_n^{(1)} \neq \mathbf{c}_n^{(2)}\}$, which contains the indices of all elements that are classified differently by the two minima. To obtain a misclassification ratio Λ from this set, we simply represent the cardinality of Ξ as a fraction of the number of data points $\Lambda = |\Xi|/N$, i.e. the number of misclassified items divided by the total number of items. If we compute Λ for all $m \in \mathcal{M}$, where $\mathcal{M}$ again denotes the set of all known minima, we obtain a symmetric matrix of pairwise differences in classification space. We can then train MLSUP on selected minima that are distant in classification space. Below, we will abbreviate the Maximum Misclassification Distance to the Best Minimum as MMDBM. A more theoretical basis for weighting classifiers according to correlation of errors is given in Masegosa *et al* (2020).

4.3. Heat capacity (C_V)

Building on results for molecular systems, we can compute a theoretical analogue of the heat capacity (C_V) (Wales 2017). The heat capacity reports upon the changing occupation probabilities of local minima as the temperature changes, which depends on both the relative loss function value for the minimum, and a local density of states (the analogue of entropy) in weight space. By computing the partial derivatives of the occupation probabilities with respect to temperature, we can identify the minima with the largest rate of change of probability, which theory shows make the largest contribution to the heat capacity (Wales 2017). Peaks in the heat capacity correspond to changes in occupation between local minima with different loss function values and entropy. Thus, in physical systems, the minima with the largest positive and negative partial derivative around a peak generally have qualitatively different properties. Peaks in C_V therefore tell us where to find local minima with complementary properties (Wales 2017), which may provide good candidates for combinations in a new classifier. This reasoning was the original motivation for the MLSUP procedure, since the heat capacity analysis is computationally efficient and physically insightful.

4.4. Data

To test MLSUP, we chose to study the spiral data problem (Lang and Witbrock 1988). An example can be seen in figure 4. This problem is considered to be relatively difficult because of its high degree of non-linear separability. We further increase the difficulty by adding a small uniform noise term to each datapoint. Our models are trained on 3200 data points and tested on 800. To show that the MLSUP approach is generalisable, we report results for other datasets in the appendix.

5. Results

In this section, we report general results for MLSUP and specifically the impact on the AUC for different minima selection methods and different numbers of minima combined in MLSUP. We benchmark the MLSUP AUC in four ways. First, we compare it to the best individual AUC of the minima that we choose for MLSUP. Second, we compare it to the best AUC of the non-MLSUP system (classifier 1 in figure 1) from which the MLSUP minima are selected. Third, in order to evaluate MLSUP, we compare its AUC with the best possible AUC achievable from combining the selected n minima. To compute this theoretical maximum AUC (TMA), for each datapoint we select the maximum Softmax probability of the correct class amongst all minima considered. The TMA gives an upper bound for the accuracy of MLSUP. Thus, to evaluate the quality of Classifier 2 alone, an AUC achieved with MLSUP should be compared against the TMA. Lastly, we compare it with a simple majority vote (with ties considered as correctly classified).

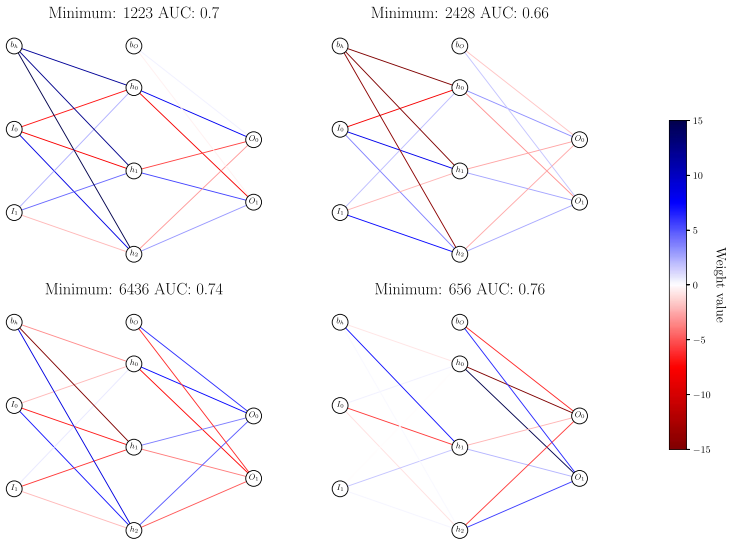

Figure 2. Four sets of weights, shown for a small network. All four minima have very different sets of weights, as evidenced by the colours. However, their AUC values for the test data are similar.

Table 1. AUC improvements for different MLSUP settings.

		AUC			
Selection method	# minima	Best individual	MLSUP	TMA	Majority vote
C_V peak 1	2	0.70	0.88	0.98	0.76
C_V peak 2	2	0.70	0.86	0.97	0.73
MMDBM	2	0.76	0.91	0.97	0.78
MMDBM & C_V peak 1	4	0.76	0.93	0.99	0.86
MMDBM & C_V peak 2	4	0.76	0.94	0.99	0.85
Max Euclidean distance	2	0.65	0.67	0.68	0.66

5.1. MLSUP effect

A substantial increase in the AUC can be achieved with MLSUP. For the original network (Classifier 1 in figure 1), we found 7842 minima and a maximum AUC value of 0.76. Using MLSUP with two of these minima (6436 and 656), we obtain an AUC of 0.91. By combining four minima, the AUC increases to 0.94 (table 1). The weights for these four minima are shown in figure 2. There is no apparent relationship between these four sets, yet they are all minima of the original LFL and have similar AUCs. We also find that MLSUP outperforms a majority vote and include results for other benchmark datasets, higher-dimensional problems and MLSUP hyperparameters in the appendix.

5.2. Minima selection

Table 1 clearly shows why a minima selection routine is necessary. Choosing the two minima with the largest Euclidean distance, we obtain an AUC of 0.67, substantially lower than for either of the systematic minima selection methods. Furthermore, the Pearson correlation coefficient between distance in Euclidean and misclassification space is only 0.06, which demonstrates that Euclidean distance alone is not useful for minima selection. Additional details can be found in the appendix.

5.2.1. Misclassification space

The distance between minima in misclassification space is the best minima selection method for an MLSUP meta-network based on two minima. By combining the best individual AUC minimum (number 656) with the minimum furthest away from it in misclassification space (number 6436), we obtain an AUC value of 0.91 (MMDBM). We note that the individual AUC values of both these minima were reasonably high already, at 0.74 and 0.76. The best possible AUC that MLSUP could have achieved, if for each data item the better minimum had been chosen for this system (TMA) is 0.97 (table 1).

5.2.2. Heat capacity

The heat capacity (C_V) curve for the LFL has two peaks, a small one at low T and a larger peak at $T \approx 3 \times 10^{-3}$ (figure 3). Such a C_V curve in a molecular system would indicate a solid-solid phase transition

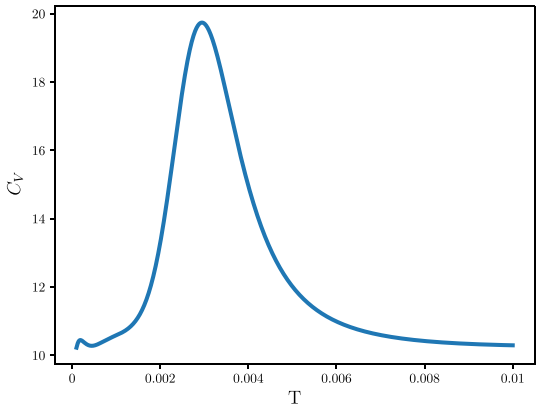

Figure 3. Heat capacity (C_V) curve for ML landscape. As the temperature is increased, the occupation probabilities change and more minima become accessible. C_V is calculated using an analogue of normal mode analysis for each minimum (Ballard *et al* 2017).

at low temperature (first peak), and a solid-liquid transition (melting) at higher temperature (second peak) (Wales 2003). In the neural network case, these peaks can be understood in terms of transitions between two energetically/entropically different sets of minima. By combining the minima that make the greatest contributions to C_V at the larger peak (figure 3) in MLSUP, we improve the AUC to 0.86 (table 1). For the minima that produce the smaller peak, the improvement is similar, with an AUC value of 0.88 (table 1). The TMA of MLSUP is 0.98 for the first and 0.97 for the second peak.

6. Discussion

Combining different minima in one classifier can substantially improve the ROC-AUC for a complex problem like spiral data. This result strongly suggests that different minima extract different information from the input and/or process it in a different way. Our results shed new light on the LFLs of complex problems. Even the global minimum is limited in the amount of information it can extract and process. To maximise the efficiency of MLSUP, we therefore aim to identify complementary high-AUC minima for which the intersection of correctly classified datapoints is small. This selection translates to a large distance in misclassification space. Such minima are ideal candidates for MLSUP.

6.1. Computational resources

Training MLSUP for four minima is fast and hence does not constitute a bottleneck. All tasks were run on six-core dual Xeon X5650 (2.6 GHz) nodes with 24Gb RAM per node. Using just one core, MLSUP is trained within a few hours. The initial analytical landscape exploration requires more subtle considerations. Due to the nature of the LFL (theoretically infinitely large, yet implicitly bounded by L2-regularisation), more minima can generally be found. For reference, finding 7842 minima took us around 92 hours. The computing time required for landscape exploration depends on system size, number of data points, and various parameters in the optimisation routine, but also the structure of the LFL, which is specific to a given problem. Landscape exploration can be parallelised, so there is scope to speed up the calculations significantly in the future.

6.1.1. Applications for MLSUP

We emphasise that our main objective in the present contribution is to improve our understanding of the LFL and investigate whether the MLSUP approach might be useful. In particular, we have investigated the underlying assumptions concerning the information extracted by different minima for underparameterised networks. For larger, overparameterised networks, a single minimum may suffice to achieve near-perfect AUC. However, evaluating a complex classifier could become a computational bottleneck in large-scale simulations of physical systems. Hence, a compact MLSUP representation will be useful in such situations. For example, we have an application where predicted values of specific molecular properties are required at every step in a molecular simulation (Dedmon *et al* 2005, Lindorff-Larsen *et al* 2005, Vendruscolo and Dobson 2005, Clore and Schwieters 2006, Cavalli *et al* 2007). The properties depend on the atomic coordinates, which vary throughout the simulation, and we need analytical derivatives of the properties (outputs) with respect to the coordinates (inputs). A compact but computationally efficient prediction

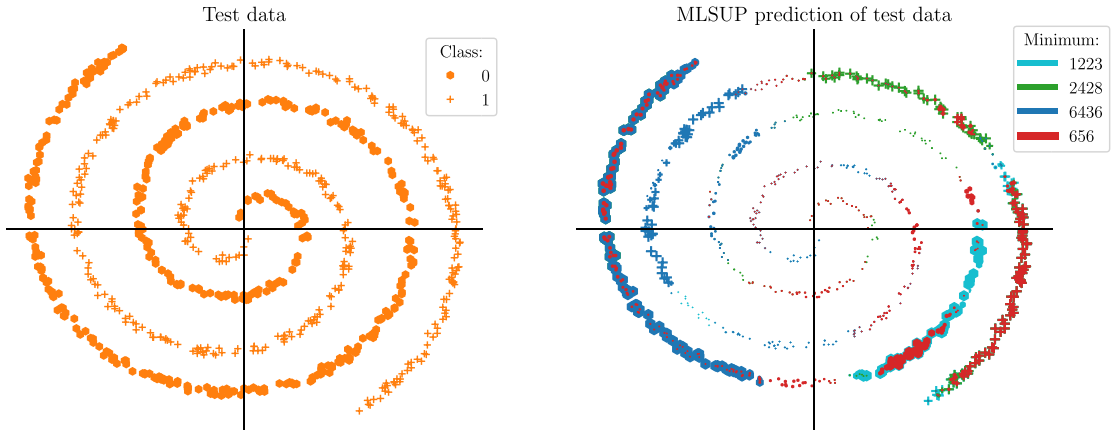

Figure 4. Contribution of each of four selected minima (right) to the MLSUP prediction of binary test data (left). Marker shape indicates the label, colour of the respective minimum, and size the Softmax probability that the respective minimum obtains for the correct class of this datapoint (scaled quadratically). Only points that are classified correctly by the respective minimum are included.

engine, such as MLSUP, is needed to prevent these computations from becoming a bottleneck in the simulations.

6.1.2. Minima selection

Selecting the correct minima for MLSUP is a key problem. Even when combining just two out of $|\mathcal{M}|$ minima, there are $\binom{|\mathcal{M}|}{2}$ possible combinations. While a choice guided by Euclidean distance may be appealing, this approach is ineffective due to low misclassification distance or permutationally equivalent solutions. Hence, we have analysed two other methods, one based on a theoretical analogue of the heat capacity, and the other on misclassification space. Combining minima of high misclassification distance achieves the best results. However, evaluating the misclassification space is relatively expensive and scales as $\mathcal{O}(n^2)$, where at each step, all data points must be evaluated for one of n minima. In contrast, the analysis of C_V scales linearly as $\mathcal{O}(n)$. Additionally, a high misclassification rate between two individually poor minima is unlikely to substantially improve results, as the confidence in each minimum would be low, and a superposition would not necessarily improve this situation. We have not observed any such issues with the C_V approach. The superior results of C_V as opposed to random initialisation are striking, and demonstrate how a methodology developed for molecular energy landscapes can provide new insight into machine learning LFLs.

6.1.3. Interpretation of differences between minima

Figure 4 provides an explanation for the substantial improvements that are achieved in MLSUP. While figure 2 shows that the individual minima are very different from each other, it does not show how these differences are propagated to classification decisions. Figure 4 shows that individual minima are good at predicting specific ranges of the input data, and only in combination do they manage to classify points all around the spiral. Note that the Softmax probability cannot generally be read as true probability because it lacks uncertainty quantification (Sensoy *et al* 2018). However, it can be used here as a visualisation proxy due to the well-behaved nature of the synthetic data. We see that the number of large points per minimum correlates with the AUC. Minima 6436 and 656 have AUCs of 0.74 and 0.76, respectively. In the right panel of figure 4, large parts of the MLSUP predictions are red or blue, corresponding to these two minima. In contrast, minimum 2428, which only seems to be able to accurately predict datapoints from Class 1 in the upper right quadrant, has a lower AUC (0.66). Another feature clearly illustrated in figure 4 is the input-sign dependence of the minima. In particular, minima 1223 and 2428 only have high confidence for datapoints that lie in one of the four segments. While 2428 is good at classifying Class 1 datapoints with positive sign for both inputs, minimum 1223 works best for classifying Class 0 datapoints with positive x and negative y. Thus, if MLSUP receives an input data point where both coordinates are positive (and perhaps at least one of them large), it is likely to classify it using the weights from minimum 2428.

Combining the above results with figure 2, we see that minimum 6436 has strongly positive weights going into the upper output node O_0 and strongly negative weights into O_1, while the opposite holds for minimum 656. Indeed, figure 4 confirms that 6436 is very good at classifying data belonging to output class 0, while 656

is very good at classifying data belonging to output class 1. We therefore conclude that different minima truly 'specialise' in parts of the input data, thus providing the foundations for the MLSUP approach. Visualising the results we provides an intuitive way to confirm that different minima process the same information in a different way. This observation raises various questions around initialisation methods and interpretability of ML models. Instead of considering model interpretability, it may be more appropriate to discuss the interpretability of different minima, as various solutions may extract, and hence focus on, completely different information from the input data. The overparameterisation of neural networks (Belkin *et al* 2019) may explain why different minima achieve such good results individually. However, to explain neural network decision making, understanding that different minima process the input differently seems a crucial first step.

7. Conclusions

Different sets of weights associated with different minima of the LFL extract different information from the input data, or process it differently. By training a meta-network to decide which minimum should be applied for any specific input data point, we can substantially improve classification accuracy. These results explain why different solutions to the same learning problem, obtained with different initialisation methods, have varying degrees of accuracy. Furthermore, we show that if the capacity of the network is limited, individual minima specialise in parts of the network. In future work, it will be interesting to consider these results in the light of mathematical measures of network capacity, such as VC-dimensions.

In addition to this insight, we also provide an analysis of redundant minima which, although they might appear to be different, return the same loss value for some input. Such symmetries make minima selection more challenging. Simple measures, such as the Euclidean distance in weight space, are insufficient due to low correlation with misclassification distance. To solve this problem, we can exploit thermodynamic analogues for the LFL by analogy to molecular systems. In particular, we find that defining a heat capacity analogue for the LFL provides an efficient and direct approach for locating complementary local minima. Computing the heat capacity, and identifying the minima that make the largest contributions to the peaks, is a promising method for constructing a compact but relatively accurate classifier in the MLSUP approach. This representation may prove useful for applications where evaluation of the classifier is a potential bottleneck, as for simulations in the physical sciences, where predictions of molecular properties are required at every step (Dedmon *et al* 2005, Lindorff-Larsen *et al* 2005, Vendruscolo and Dobson 2005, Clore and Schwieters 2006, Cavalli *et al* 2007). The analogy to physical systems provides further motivation for pursing the energy landscape view of machine learning LFLs.

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

M N acknowledges funding from Downing College, Cambridge and specifically the Voellm-Hruska PhD studentship. D J W gratefully acknowledges funding from the EPSRC and an International Chair at the Interdisciplinary Institute for Artificial Intelligence at 3iA Cote d'Azur, supported by the French government, reference ANR-19-P3IA-0002, which has provided interactions that furthered the present research project. We thank Prof. Benedict Leimkuhler for suggesting the spiral dataset. We declare no competing interests.

Appendix

Further experiments

Table 2 includes some additional results that reinforce our conclusions. Importantly, we also want to test MLSUP on other datasets to provide evidence that the performance increase is not specific to the spiral dataset. We first considered the well known Iris flower dataset and, to make it more difficult, mislabelled 45% of the training samples. The second additional test for MLSUP was for the checkerboard dataset, another synthetic, nonlinear problem. Similar to Iris, we achieve significant improvements using MLSUP. When the number of minima is increased substantially (e.g. to 20) the AUC only increases slightly. Training MLSUP for 20 minima takes around 10 times as long as for four minima using the same setup as above. These results again suggest that MLSUP can be effective for a small number of local minima, with the potential to provide

Table 2. AUC improvements by MLSUP.

Description	# minima	AUC		
		Best individual	MLSUP	TMA
Large hidden layer ($n = 100$)	2	0.99	0.99	1
Small training set $N = 400$	2	0.62	0.78	0.86
5 MLSUP minima	5	0.76	0.94	0.99
20 MLSUP minima	20	0.76	0.97	1
100 randomly chosen (mean)	2	0.74 (0.67)	0.7–0.89 (0.78)	0.98
100 randomly chosen (mean)	4	0.76 (0.67)	0.78–0.93 (0.84)	0.99
Mislabelled Iris flower dataset	2	0.85	0.96	0.99
Checkerboard	2	0.68	0.84	0.96

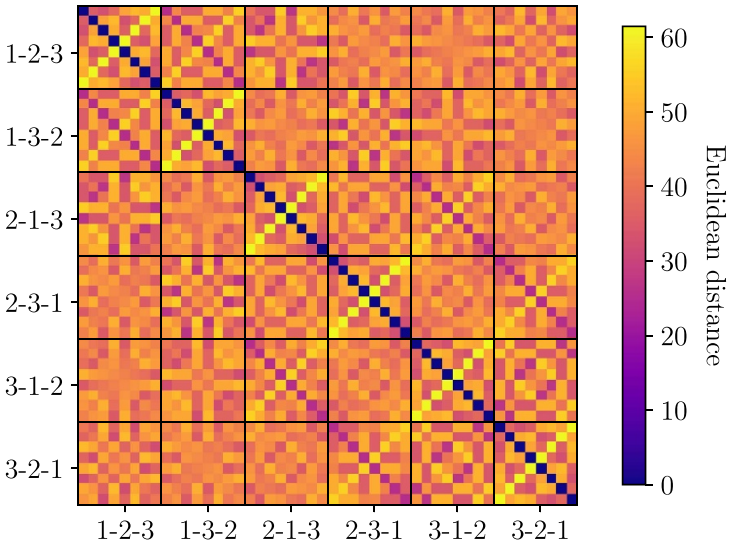

Figure 5. Euclidean distance between 48 permutationally equivalent local minima of the global minimum for a particular LFL.

useful accuracy at a low computational cost. This balance of accuracy and efficiency will be valuable in applications where reasonable predictions are required without a significant overhead.

Permutational invariants

Examples

Here we provide additional insight into the issue of permutationally equivalent sets of weights.

Figure 5 shows how distant permutationally equivalent local minima can be in Euclidean space. The system contains a single hidden layer with $n = 3$ nodes. Shuffling the nodes leads to $n! = 6$ permutations and within each permutation, the weights going in and out of any combination of the hidden nodes can be multiplied by -1. Thus, for each shuffle-permutation there are an additional $2^n = 8$ isomers, producing 48 in total. In misclassification space, the distance between all these 48 minima would be 0, as in figure 6, which shows how far apart these equivalent solutions can appear to be.

Comparison of Euclidean and misclassification space

This section presents further evidence that the Euclidean distance alone is insufficient for minima selection. Figure 7 shows that only a very weak correlation exits between Euclidean and misclassification distance. Large misclassification distance can be observed roughly as often for points of large Euclidean distance as it is observed for points close together in the parameter space defined by the weights. Given the good results of misclassification distance for minima selection, the low correlation coefficient indicates that Euclidean distance alone is not useful.

MLSUP algorithm

The MLSUP algorithm 1 can be represented in pseudocode as:

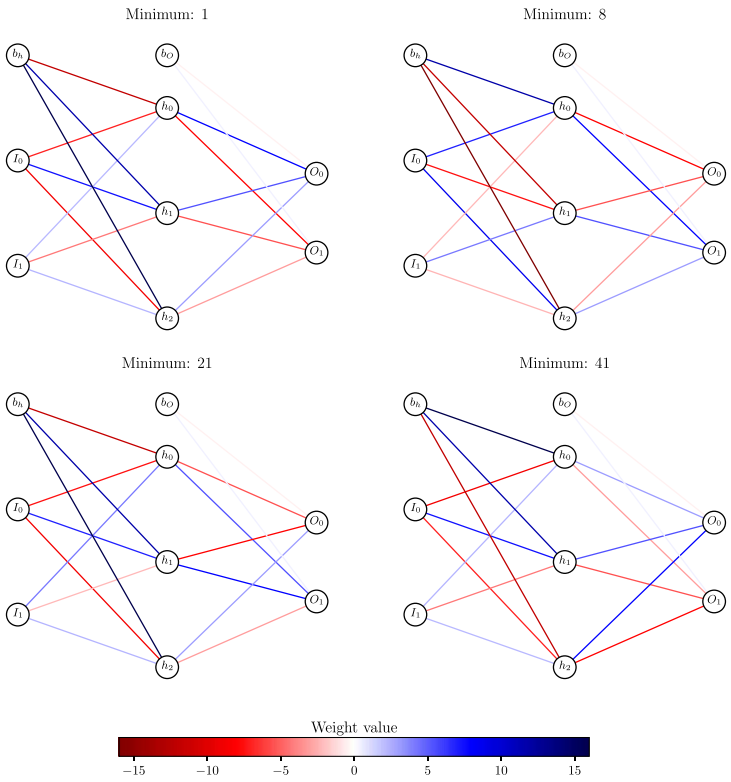

Figure 6. An example of four very different looking minima that are actually permutationally equivalent, i.e. redundant solutions. For any input, these four networks will all return exactly the same loss value.

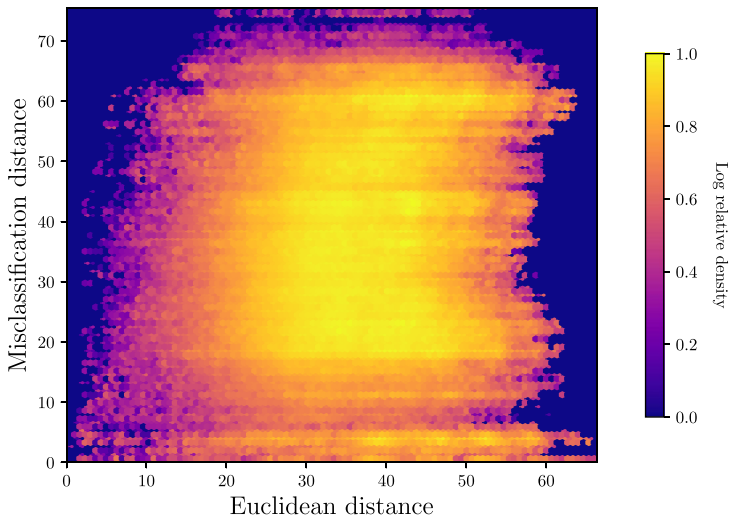

Figure 7. Euclidean vs Misclassification distance for each of the $\binom{n}{2}$ combinations of n minima of the spiral data LFL. The Pearson correlation coefficient for Euclidean distance vs misclassification is 0.06.

Algorithm 1. Train MLSUP.

Require: $\mathcal{D} = (\mathbf{X}, \mathbf{c})$
Sample LFL, find $|\mathcal{M}|$ minima
Pick some $|\mathcal{M}'|$ minima candidates
Evaluate each minimum $m \in |\mathcal{M}'|$ on training data
Train second nn to learn $\arg\max m \ {}^{m}p_{c^d}$
for $d \in \mathcal{D}$ **do**
Evaluate first network, query second network for which minimum to use
end for

Choosing m' minima by pairwise misclassification distance

One may wish to compute MLSUP with m' minima, chosen to have the maximum misclassification distance from another. This is a combinatorial optimisation problem, analogous to having m cities with pairwise distances and trying to find the m' cities that are furthest from each other. This problem can be solved by quadratic programming with convex relaxation. Let $\mathbf{D}$ be an upper triangular matrix with element $\mathbf{D}_{ij}$ denoting the misclassification distance from point i to point j, where $i<j$. The square matrix $\mathbf{D}$ has 0 diagonal. Denote x as a m-dimensional binary vector, which contains the value 1 at the position of a chosen minimum and 0 elsewhere. The objective is then to maximise $x^\top \mathbf{D}x$. In practice, the values of x will be relaxed so they need not be discrete, and the optimisation procedure becomes:

$$\max_{y} y^\top \mathbf{D}y \quad \text{subject to} \quad 0 \leqslant y_i \leqslant 1 \; \forall i, \; 1 \cdot y = m'. \tag{9}$$

The final task is then to discretise y back to x such that $y^\top \mathbf{D}y$ is very close to $x^\top \mathbf{D}x$. We have achieved this mapping by simply setting the m' largest elements of y to 1 and all others to 0, but other methods of approximation may work equally well or even better.

Basin-hopping global optimisation

The first step in surveying the LFL for a neural network is to locate the global minimum. We perform global optimisation using the basin-hopping method, which uses a Metropolis criterion to accept or reject proposed moves between minima (Li and Scheraga 1987, Wales and Doye 1997). For local minimisation we use a modified quasi-Newton L-BFGS (Limited-memory Broyden-Fletcher-Goldfarb-Shanno) algorithm (Nocedal 1980). After a new local minimum is found via L-BFGS optimisation, a basin-hopping step to that minimum is accepted if the energy, here loss value, is lower than the current minimum. If the loss value of the new minimum is higher, it is accepted with probability:

$$P \propto \exp\left(-\frac{\Delta\widetilde{E}}{k_{\mathrm{B}}T}\right), \tag{10}$$

where $\Delta\widetilde{E}$ is the difference in loss value between the current minimum and the new minimum, k_{B} the Boltzmann constant and T a fictitious temperature. Intuitively, if the energy difference between the old and new minima is large, the move is less likely to be accepted. Local minima are defined as stationary points with an RMS gradient below 10^{-10} and no negative Hessian eigenvalues. Each basin-hopping run naturally produces a sampling of low-lying minima as it progresses, in addition to the global minimum.

ORCID iDs

Maximilian P Niroomand https://orcid.org/0000-0002-7189-0456
Conor T Cafolla https://orcid.org/0000-0003-2021-974X
David J Wales https://orcid.org/0000-0002-3555-6645

References

Ballard A J, Das R, Martiniani S, Mehta D, Sagun L, Stevenson J D and Wales D J 2017 Energy landscapes for machine learning *Phys. Chem. Chem. Phys.* **19** 12585–603

Belkin M, Hsu D, Siyuan M and Mandal S 2019 Reconciling modern machine-learning practice and the classical bias-variance trade-off *Proc. Natl Acad. Sci.* **116** 15849–54

Brea J, Simsek B, Illing B and Gerstner W 2019 Weight-space symmetry in deep networks gives rise to permutation saddles, connected by equal-loss valleys across the loss landscape (arXiv:1907.02911 [cs, stat])

Breiman L 1996 Bagging predictors *Mach. Learn.* **24** 123–40

Cavalli A, Salvatella X, Dobson C M and Vendruscolo M 2007 Protein structure determination from nmr chemical shifts *Proc. Natl Acad. Sci. USA* **104** 9615–20

Clore G M and Schwieters C D 2006 Concordance of residual dipolar couplings, backbone order parameters and crystallographic b-factors for a small α/β protein: a unified picture of high probability, fast atomic motions in proteins *J. Mol. Biol.* **355** 879–86

Dedmon M M, Lindorff-Larsen K, Christodoulou J, Vendruscolo M and Dobson C M 2005 Mapping long-range interactions in alpha-synuclein using spin-label nmr and ensemble molecular dynamics simulations *J. Am. Chem. Soc.* **127** 476–7

Dinh L, Pascanu R, Bengio S and Bengio Y 2017 Sharp minima can generalize for deep nets *Proc. 34th Int. Conf. on Machine Learning (Proc. Machine Learning Research* vol 70), ed D Precup and Y W Teh (PMLR) pp 1019–28

Dong X, Zhiwen Y, Cao W, Shi Y and Qianli M 2020 A survey on ensemble learning *Frontiers Comput. Sci.* **14** 241–58

Fawcett T 2006 An introduction to ROC analysis *Pattern Recognit. Lett.* **27** 861–74

Hashem S 1997 Optimal linear combinations of neural networks *Neural Netw.* **10** 599–614

Henkelman G and Jónsson H 2000 Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points *J. Chem. Phys.* **113** 9978–85

Henkelman G, Uberuaga B P and Jónsson H 2000 A climbing image nudged elastic band method for finding saddle points and minimum energy paths *J. Chem. Phys.* **113** 9901–4
Jin H and Lu Y 2009 The optimal linear combination of multiple predictors under the generalized linear models *Stat. Probab. Lett.* **79** 2321–7
Jordan M I and Jacobs R A 1994 Hierarchical mixtures of experts and the em algorithm *Neural Comput.* **6** 181–214
Kuncheva L I 2014 *Combining Pattern Classifiers: Methods and Algorithms* (Hoboken, NJ: Wiley)
Lang K J and Witbrock M J 1988 Learning to tell two spirals apart *Proc. 1988 Connectionist Models Summer School* (*San Mateo*) pp 52–9
Li Z and Scheraga H A 1987 Monte Carlo-minimization approach to the multiple-minima problem in protein folding *Proc. Natl Acad. Sci.* **84** 6611–15
Lindorff-Larsen K, Best R B, DePristo M A, Dobson C M and Vendruscolo M 2005 Simultaneous determination of protein structure and dynamics *Nature* **433** 128–32
Masegosa Aes R, Lorenzen S S, Igel C and Seldin Y 2020 Second order pac-bayesian bounds for the weighted majority vote *NeurIPS* (arXiv:2007.13532)
McGill M and Perona P 2017 Deciding how to decide: Dynamic routing in artificial neural networks *Int. Conf. on Machine Learning* (PMLR) pp 2363–72
Melville P and Mooney R J 2005 Creating diversity in ensembles using artificial data *Inf. Fusion* **6** 99–111
Munro L J and Wales D J 1999 Defect migration in crystalline silicon *Phys. Rev.* B **59** 3969–80
Nocedal J 1980 Updating quasi-Newton matrices with limited storage *Math. Comput.* **35** 773–82
Sensoy M, Kaplan L and Kandemir M 2018 Evidential deep learning to quantify classification uncertainty (arXiv:1806.01768)
Shazeer N, Mirhoseini A, Maziarz K, Davis A, Quoc L, Hinton G and Dean J 2017 Outrageously large neural networks: the sparsely-gated mixture-of-experts layer (arXiv:1701.06538)
Trygubenko S A and Wales D J 2004a Analysis of cooperativity and localization for atomic rearrangements *J. Chem. Phys.* **121** 6689–97
Trygubenko S A and Wales D J 2004b A doubly nudged elastic band method for finding transition states *J. Chem. Phys.* **120** 2082–94
Vapnik V N and Chervonenkis A Y 1971 On the uniform convergence of relative frequencies of events to their probabilities *Theor. Probab. Appl.* **16** 264–80
Vapnik V, Levin E and Le Cun Y 1994 Measuring the VC-dimension of a learning machine *Neural Comput.* **6** 851–76
Vendruscolo M, Dobson C M, McMillan P F and Clary D C 2005 Towards complete descriptions of the free–energy landscapes of proteins *Phil. Trans. R. Soc.* A **363** 433–52
Wales D J 2003 *Energy Landscapes (Cambridge Molecular Science)* (Cambridge, UK: Cambridge University Press)
Wales D J 2003a GMIN: a program for basin-hopping global optimisation, basin-sampling, and parallel tempering
Wales D J 2003b OPTIM: a program for geometry optimisation and pathway calculations
Wales D J 2003c PATHSAMPLE: a program for generating connected stationary point databases and extracting global kinetics
Wales D J 2017 Decoding heat capacity features from the energy landscape *Phys. Rev.* E **95** 6
Wales D J and Doye J P K 1997 Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms *J. Phys. Chem.* A **101** 5111–6
Wang W 2008 Some fundamental issues in ensemble methods *2008 IEEE Int. Conf. on Neural Networks (IEEE World Congress on Computational Intelligence)* (Hong Kong: IEEE) pp 2243–50
Zaidi S, Zela A, Elsken T, Holmes C, Hutter F and Teh Y W 2020 Neural ensemble search for performant and calibrated predictions (arXiv:2006.08573)
Zeng Y, Xiao P and Henkelman G 2014 Unification of algorithms for minimum mode optimization *J. Chem. Phys.* **140** 044115