

Predicting the Mpemba effect using machine learning

Felipe Amorim,^{*} Joey Wisely,[†] Nathan Buckley,[‡] Christiana DiNardo,[§] and Daniel Sadasivan^{||}
Ave Maria University, Ave Maria, Florida 34142, USA

 (Received 4 October 2022; revised 18 April 2023; accepted 1 August 2023; published 24 August 2023)

The Mpemba effect can be studied with Markovian dynamics in a nonequilibrium thermodynamics framework. The Markovian Mpemba effect can be observed in a variety of systems including the Ising model. We demonstrate that the Markovian Mpemba effect can be predicted in the Ising model with several machine learning methods: the decision tree algorithm, neural networks, linear regression, and nonlinear regression with the least absolute shrinkage and selection operator (LASSO) method. The positive and negative accuracy of these methods are compared. Additionally, we find that machine learning methods can be used to accurately extrapolate to data outside the range in which they were trained. Neural networks can even predict the existence of the Mpemba effect when they are trained only on data in which the Mpemba effect does not occur. This indicates that information about which coefficients result in the Mpemba effect is contained in coefficients where the results does not occur. Furthermore, neural networks can predict that the Mpemba effect does not occur for positive J , corresponding to the ferromagnetic Ising model even when they are only trained on negative J , corresponding to the antiferromagnetic Ising model. All of these results demonstrate that the Mpemba effect can be predicted in complex, computationally expensive systems, without explicit calculations of the eigenvectors.

DOI: [10.1103/PhysRevE.108.024137](https://doi.org/10.1103/PhysRevE.108.024137)

I. INTRODUCTION

A. Mpemba effect

The Mpemba effect is best known for the claim that hot water can freeze faster than cool water in the same environment. This was named after Erasto Mpemba, who most famously brought it to the attention of the scientific community [1] although it had been debated for thousands of years [2–4]. However, the Mpemba effect applies to a large number of situations beyond freezing water. We quote the definition given in Ref. [5], “The Mpemba effect is a counterintuitive relaxation phenomenon, where a system prepared at a hot temperature cools down faster than an identical system initiated at a cold temperature when both are quenched to an even colder bath.” Several recent theoretical explanations of the effect have been demonstrated. One shows that the Mpemba effect can be predicted from statistical mechanics [6–10]. Another explanation in Refs. [5,11] makes use of the framework developed in Ref. [11], which considers the Markovian dynamics as a cooling process in the framework of nonequilibrium thermodynamics. We refer to this process as the Markovian Mpemba effect (MME). The MME is further studied in Refs. [12,13]. Reference [11] studies the MME to an antiferromagnetic nearest-neighbor interacting Ising spin chain, generalizing the Mpemba effect to other systems. The one-dimensional Ising spin chain will be the main model used in this paper, enabling us to use machine learning methods to predict the effect, as described in Sec. II.

The Mpemba effect traditionally describes systems relaxing towards an equilibrium temperature that is lower than their initial temperature. However, a reverse effect occurs when two systems relax towards equilibrium at a higher temperature. Sometimes the former is referred to as the Mpemba effect and the latter is referred to as the inverse Mpemba effect, but both cases meet the definition of the Mpemba effect used in this paper, namely, that the system that is initially farther from equilibrium becomes closer after a finite time. Unless otherwise specified, we refer to the inverse Markovian Mpemba effect simply as the Mpemba effect.

This paper is organized as follows. Section IB discusses recent works with results that can be related to our results. Section II describes the formalism (laid out in Refs. [5,11]) used to determine whether the Mpemba effect occurs in the Ising model. Section III describes the data generated for machine learning and the methods used to generate this data. Section IV summarizes the various machine learning methods that are employed. The main results of this work are presented and discussed in Sec. V. We conclude in Sec. VI.

B. Applications and impact

The Mpemba effect has been observed in a number of systems. In addition to water [14,15], it has been predicted in magnetoresistance alloys [16], numerically predicted in colloids [17], analytically and numerically predicted in gasses [18,19], predicted in particles bounded by anharmonic potentials [20], observed in computational simulations and experiments in gas hydrates [21], and in computational simulations of carbon nanotube resonators [22]. It has been demonstrated to have application for faster heating with pre-cooling [23], in which the effect happens for temperatures dependent on the maximum work that can be extracted from the system [24], and where the Mpemba effect can lead to

^{*}FelipeAugusto.deAmor@my.avemaria.edu

[†]Peter.Wisely@my.avemaria.edu

[‡]Nathan.Buckley@my.avemaria.edu

[§]Christiana.George@my.avemaria.edu

^{||}Daniel.Sadasivan@avemaria.edu

quantum heat engines with greater power output and stability [25]. Furthermore, the Mpemba effect is closely linked to the Kovacs memory effect, by which systems out of equilibrium cannot merely be described by macroscopic thermodynamic variables. Knowledge of the Mpemba effect can be applied when studying memory effects in Refs. [26–30]. Notably, it has been demonstrated that the phenomenon of eigenvalue crossing, closely linked to the Mpemba effect, can be understood as a first order phase transition [31].

These works demonstrate the potential applications of the effect and motivate a better statistical understanding. The work presented in this paper uses machine learning methods to predict the Mpemba effect (ME) in the Ising model. The comparison of methods done in this analysis could be applied to the prediction of the ME in other systems.

Machine learning has previously been applied to the Ising model in a number of ways other than the Mpemba effect: to extract the region between phases using auto-encoding neural networks [32], to predict probability distributions with neural networks [33], and to compare various classification methods such as random forests, decision trees, k nearest neighbors, and artificial neural networks for the prediction of magnetization [34].

These works find several interesting results that can be compared to the work in this paper. First, Ref. [33] found that only the number of nodes in the first layer of the deep neural network could improve the accuracy of the prediction. Second, Ref. [34] found that the decision tree method was the most accurate predictor. Even though these results predict different things, it is interesting to see whether these results (namely, that the decision tree algorithm is the most accurate predictor and that deep learning improves prediction) also apply to predicting the Mpemba effect in the Ising model.

II. THE MARKOVIAN MPEMBA EFFECT FORMALISM

This work uses the formalism developed in Refs. [5,11] for the Mpemba effect in the Ising model. We summarize the key equations in the following.

The one-dimensional Ising model consists of a chain of N spins, s_i with values of 1 or -1 . The energy for a given set of spins is defined to be

$$E = -J \sum_{k=0}^N s_k s_{k+1} - h \sum_{k=1}^N s_k, \quad (1)$$

where J and h are parameters, respectively, giving the field of interaction between neighboring spins and the strength of the external field. There are multiple ways to handle the end points, s_1 and s_N , of the Ising chain. One possibility is to consider both end points to be connected to fixed spins as shown in Fig. 1 or Fig. 3 of Ref. [11]. This choice of endpoints means that for odd N , the ground state energy level is nondegenerate, whereas for even N , the ground state energy is degenerate. This can be seen in Fig. 1.

Another possibility is to have the two endpoints connected such that the Ising chain forms a ring. Figure 1 of Ref. [35] shows such a ring.

A probability distribution $\bar{p}(t)$ that represents the probability of finding the system in one of the 2^N microstates

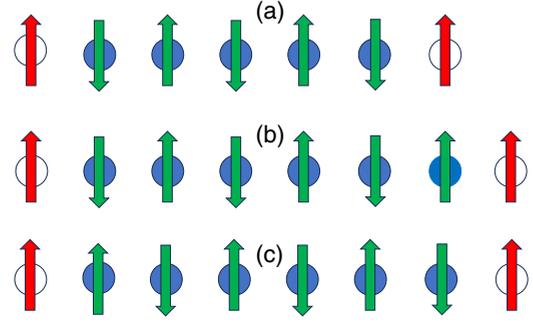


FIG. 1. Ising chains of green arrows in blue shaded circles with fixed endpoints shown as red arrows in unshaded circles. For the antiferromagnetic Ising model, there is a nondegenerate ground state at zero temperature where each spin is antialigned with its neighbors. This is shown for $N = 5$ in the upper row (a). However, for even $N = 6$, it is not possible for each spin to be antialigned, as shown in lower rows (b) and (c).

is defined, for the case of a finite number of states, by the equation

$$\frac{d\bar{p}(t)}{dt} = R(T_b)\bar{p}(t), \quad (2)$$

where T_b is the bath temperature. Matrix multiplication is assumed. One approach is to use heat bath dynamics (employed in Refs. [5,11]). In this case, the transition rate matrix from a state j to state i , $R_{ij}(T_b)$, is defined as

$$R_{ij} = \begin{cases} \Gamma e^{-\frac{1}{2}\beta_b(E_i - E_j)} & \text{if } i \text{ and } j \text{ differ by one spin flip} \\ 0 & \text{if } i \text{ and } j \text{ differ by } >1 \text{ spin flip} \\ -\sum_{k \neq j} R_{kj}, & i = j. \end{cases} \quad (3)$$

While this rate matrix can provide an exact description of the system's evolution over time starting from any initial probability distribution, the computational costs for larger N are prohibitive. The method of coarse graining can reduce the computational cost. In this method, microstates of the same energy are grouped together. This method is employed in Refs. [35–37]. If coarse graining is used, the rate matrix is given by Ref. [35] (in the limit where the system is weakly coupled to the external environment) as

$$R_{ij}^{\text{weak}} = \frac{G_{ij}}{\Omega_j} \frac{1}{1 + e^{\beta_b(E_i - E_j)}}, \quad (4)$$

where Ω_j is the number of microstates with energy E_j and G_{ij} is the number of transitions between microstates of energy E_i and E_j . The coarse-graining approach can considerably reduce the size of the rate matrix. Instead of 2^N states, there are now $2 + (N/2)^2$ states. This allows calculation for larger N . This work makes the calculation for both methods. The former method with fixed endpoints is used in Secs. V A–V C, and the latter method with the endpoints forming a ring is used in Sec. V D.

At a given temperature, a system will have an equilibrium state,

$$\bar{\pi}(T) = \frac{e^{-E_i/k_B T}}{\sum_i e^{-E_i/k_B T}}. \quad (5)$$

Because the eigenvalues form a complete basis, any state can be expressed as

$$\bar{p} = \bar{\pi}(T) + \sum_{i>1} a_i \bar{v}_i, \quad (6)$$

where T is the temperature, v_i are the eigenvectors of the rate transition matrix, and a_i are the coefficients corresponding to the vectors v_i . $\bar{\pi}$ is the eigenvector corresponding to an eigenvalue of 0. This eigenvector can be determined by the Boltzmann probability distribution.

The general solution to Eq. (2) is given by

$$\bar{p}(T; t) = e^{Rt} \bar{\pi}(T) = \bar{\pi}(T_b) + \sum_{i>1} a_i(T) e^{\lambda_i t} v_i, \quad (7)$$

where λ_i are the eigenvalues of the transition rate matrix and T_b is the bath temperature.

The distance from the equilibrium function, $D[\bar{p}(t); T_b]$, is given by

$$D[\bar{p}(t); T_b] = \sum_i \left(\frac{E_i(p_i - \pi_i^b)}{T_b} + p_i \ln p_i - \pi_i^b \ln \pi_i^b \right). \quad (8)$$

Multiple possible distance functions can be used. We use the distance function employed in Ref. [11]; however, others are possible such as the Kullback-Leibler divergence used in Ref. [31]. The Mpemba effect occurs for any two probability distributions, \bar{p}_H and \bar{p}_C , if $D[\bar{p}_H(0); T_b] > D[\bar{p}_C(0); T_b]$ but $D[\bar{p}_H(t'); T_b] < D[\bar{p}_C(t'); T_b]$ for some time t' . Reference [11] shows that this will occur if $|a_2^H| < |a_2^C|$. The H in \bar{p}_H refers to hot and the C in \bar{p}_C refers to cold. However, in the case of the inverse Mpemba effect in which the initial probability relaxes to a hotter bath temperature, the probability vector corresponding to the hotter initial system will have a smaller distance function than the vector corresponding to the colder system.

Additionally, Ref. [5] defines the strong Mpemba effect. This occurs when a_2 in Eq. (7) is 0. This means that for large t , the time evolution is governed by a_3 . A system with an initial \bar{p} such that $a_2 = 0$ will approach equilibrium exponentially slower than a state with $a_2 \neq 0$. In contrast, cases where the Mpemba effect occurs when the hotter system has initial \bar{p} with $a_2 \neq 0$ is known as the weak Mpemba effect. The Mpemba effect can be further divided into the direct Mpemba effect in which the bath temperature is lower than the temperature of the initial states and the inverse Mpemba effect in which the bath temperature is lower. In this work, we primarily classify the weak inverse Mpemba effect; however, we discuss classification of the strong inverse Mpemba effect in Sec. VC.

III. DATA GENERATION

The data used for the various machine learning methods are generated in the following way. The matrix R_{ij} , given in Eq. (3), is calculated for a given J , h , T_b , and N . Sparse matrix methods are used to improve the speed of the computation. The Arnoldi method [38] is used to calculate the second greatest (least negative) eigenvalue and the corresponding eigenvector. These correspond to λ_2 and v_2 from Eq. (7) and

can be used to calculate a_2 using the method described in Sec. II in Ref. [5].

When we use smaller N , we use values between 5 and 15. We also make predictions for data with larger N in Sec. VD. The values of h are between -10 and 10 , the values of J are between -10 and 0 , and the values of T are between 1 and 30 .

For the data set with smaller N , we work only with odd-numbered spins. As discussed, the ground state with fixed endpoints for even N is degenerate, whereas it is nondegenerate for odd N . We find that the difference between the dynamics in the case of odd and even spins is large enough that none of the Machine learning (ML) methods were able to accurately predict even N in the case of small N with fixed endpoints. We thus exclude even-valued N . This should be understood as a distinct limitation of the method. It should be noted that Monte Carlo methods for numerically determining the Mpemba effect used in Ref. [23] do not suffer from this limitation. However, for larger N in the weak coupling limit using Glauber dynamics, we are able to make accurate predictions for even N , even if the data are only trained on odd N .

For each data set, we work with a bath temperature of $T_b = 30$. The bath temperature is a physical property of the system and could be varied for these methods. However, we work with a single bath temperature so that the temperature of the initially hotter system, T_h , and the temperature of the initially colder system, T_c , are selected from the same range. It is worth noting that as far as the Mpemba effect is concerned, our results that do not vary β_B but do vary J can be mapped to results that vary β_B but not J . This is because the time evolution of a probability distribution [with an initial value given by Eq. (5)] depends on the rate matrix given by Eq. (3) which depends on the energy given in Eq. (1). The following mapping would leave $\pi(\bar{T})$ and R_{ij} unchanged:

$$J \rightarrow \alpha J, \quad h \rightarrow \alpha h, \quad T_c \rightarrow \alpha T_c, \quad T_h \rightarrow \alpha T_h, \quad \beta_b \rightarrow \beta_b / \alpha. \quad (9)$$

If α is chosen to be $\alpha = 1/|J|$, then $J = -1$ for each data point and the new bath temperature becomes $\beta_b |J|$. Thus, such a mapping would allow results in this paper to be related to other results which vary β_b but not J .

IV. MACHINE LEARNING METHODS

We use four different machine learning methods to predict the occurrence of the Mpemba effect in an Ising chain. All of the employed methods have certain free parameters called fit parameters and a loss function which measures the distance of the prediction from the data. The methods use various algorithms to fit the data or match the model to the data.

A. Decision tree algorithm (DT)

For the decision tree, the classification of data is divided into two classes: for each set of initial conditions, the system can either undergo the Inverse Markovian Mpemba Effect (IMME) or not undergo it. We use a data set in the format of

$$\{N, J, h, T_c, T_h\} = \begin{cases} 1 & \text{if ME occurs} \\ 0 & \text{if ME does not occur.} \end{cases} \quad (10)$$

The decision tree is trained using the classification and regression tree (CART) algorithm [39]. Just like all ML methods, this method minimizes a loss function. However, the loss function for this method is different than the loss function of the other used methods. Therefore, we can compare the accuracy of this method to the accuracy of other methods, but cannot compare the loss function of this method to the other employed methods.

B. Neural network (NN)

Neural networks can be trained to predict the parameters $|a_2|$ when given a set of Mpemba parameters. Prediction of a_2 is sufficient to determine whether the Mpemba effect occurs. The network is trained on four input parameters,

$$\{N, J, h, T\} = |a_2|. \tag{11}$$

Because all models rely on fit parameters, multiplying $|a_2|$ by a constant factor will not affect the accuracy. If the fit parameters of each model are multiplied by the same factor, the fitting procedure will be mathematically the same as if no factor had been applied to either. However, certain algorithms automatically select natural sized starting values. To ensure that the automatic starting values are optimized, we multiply $|a_2|$ by a factor of 10 000. This leads to the exact same local minimum, but allows for more efficient training. This adjustment prevents the need to change the starting parameters of each algorithm.

We work with a network consisting of three hidden layers with 50, 20, and 5 nodes, respectively. Our explanation for this architecture is given in Appendix B. Each layer as well as the output layer uses the Scaled Exponential Linear Unit (SELU) activation function. The loss function is the mean squared logarithmic error (MSLE) [40]. This is given with the formula

$$L(y, \hat{y}) = \frac{1}{n} \sum_{i=0}^n [\log(y_i + 1) - \log(\hat{y}_i + 1)]^2, \tag{12}$$

where y_i are the actual values, n is the number of data points, and \hat{y}_i are the predicted values.

The neural network can also be used to directly predict the Mpemba effect rather than predicting $|a_2|$. This is done by training the network with data given by Eq. (10). The network remains the same, except that the output layer has a sigmoid activation function and a binary cross-entropy loss function is used. The case where the Mpemba effect is predicted directly is referred to as NN2.

C. Linear regression (LR)

Linear regression is used to predict the $|a_2|$ coefficients with the data. We once again minimize the MSLE given in Eq. (12). The predicted values \hat{y}_i are given by

$$\hat{y}_i = \sum_{j=0}^4 c_j x_{ij},$$

$$x_{i0} = 1; \quad x_{i1} = N_i; \quad x_{i2} = J_i; \quad x_{i3} = h_i; \quad x_{i4} = T_i. \tag{13}$$

The fit parameters are c_j .

TABLE I. Upper row: The linear regression fit parameters. The parameter in column i corresponds to the fit parameters c_i in Eq. (13). The first column is labeled c_0 . The other columns are labeled with the parameters that c_i is multiplied by. Lower row: The correlation coefficients between each parameter and a_2 .

	c_0	N	J	h	T
Regression coefficients	10.191	84.359	-78.725	-9.889	-36.903
Correlation coefficients		-0.005	-0.370	-0.027	-0.601

The MSLE is chosen as the loss function rather than the sum of least squares because it more accurately captures the behavior of data with very different orders of magnitude.

Linear regression cannot be used to determine whether the Mpemba effect occurs because, for the effect to occur, $|a_2|$ must increase as T increases. If $|a_2|$ is given by a monotonically decreasing linear function of T , the T that is farther from equilibrium will always correspond to a larger value of $|a_2|$.

Nevertheless, linear regression is useful for two reasons. First, the MSLE can be compared to other methods, giving information about their accuracy. Second, the values of the fit parameters c_j can be used as a crude approximation of the importance of each category of data for predicting $|a_2|$.

The fit parameters are given in Table I. We also quote the correlation coefficients which correspond to an even simpler model. Note that because the linear regression in the fit was performed by minimizing the MSLE rather than the least squares, the regression coefficients are not the same as multivariate correlation coefficients.

D. Nonlinear regression (NLR) with the LASSO method

Nonlinear models are able to have more fit parameters than linear models. We use a general expansion of the form

$$\hat{y}_i = \sum_{k_1=0}^M \dots \sum_{k_4=0}^M \delta_{M, \sum_j k_j} c_{\underbrace{1 \dots 1}_{k_1 \text{ times}} \dots \underbrace{4 \dots 4}_{k_4 \text{ times}}} \prod_{j=1}^4 P_{k_j}(x_{ij}), \tag{14}$$

where $P_{k_j}(x_{ij})$ are the Legendre polynomials mapped to the range of each data set, M is the order up to which we consider, $\delta_{M, \sum_j k_j}$ is the Kronecker delta function, and $c_{\underbrace{1 \dots 1}_{k_1 \text{ times}} \dots \underbrace{4 \dots 4}_{k_4 \text{ times}}}$ are

the fit parameters. In the case when $M = 1$, this expansion reduces to linear regression. The error that is minimized is once again the MSLE given in Eq. (12). In our notation, coefficients such as c_{1144} and c_{4411} are multiplied by identical Legendre polynomials. Each coefficient corresponds to an identical term as each other coefficient with subscripts that are a permutation of its subscript. We include only one of each set of identical coefficients.

The Legendre polynomials were chosen rather than simple polynomials because of their orthonormality. This allows for a continuous function to be fit with fewer free parameters than would be necessary with nonorthogonal functions [41].

For large M , the number of fit parameters leads to a prohibitive computational cost of the minimization. Not all fit parameters at a given order are equally effective at minimization. The LASSO method [42–45] is employed in order to determine which fit parameters are most effective and which

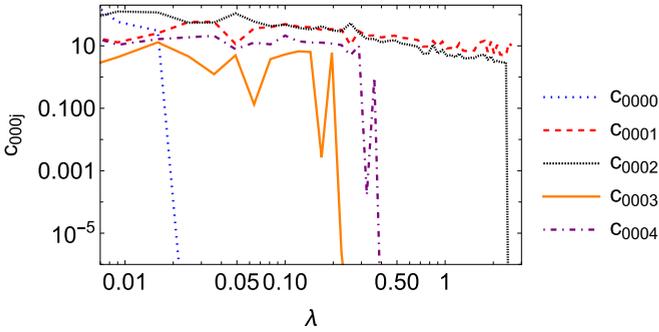


FIG. 2. The best-fit values of parameters c_{000j} for j , from 1 to 4, defined in Eq. (14). The parameters are obtained by minimizing the loss function given in Eq. (15) for various λ . As λ increases, c_{000j} will become small. The order in which the parameters approach zero gives the order of importance of each parameter in minimizing the MSLE. Note that the parameter c_{0001} does not descend to zero in this plot. If a larger range of λ were chosen, it would.

can be neglected to reduce computational cost. This is done in the following way. We randomly select 1000 data points and use them to calculate a modified loss function given by

$$L(y, \hat{y}) = \frac{1}{n} \sum_{i=0}^n [\log(y_i + 1) - \log(\hat{y} + 1)]^2 + \lambda \sum |c_{ijkl}|. \quad (15)$$

The term λ is a penalty term that can be varied. The data points are used to find the fit parameters that lead to local minima for various values of λ . When $\lambda = 0$, all fit parameters are nonzero. As λ becomes large, all fit parameters approach zero. A number of fits are performed with all fit parameters up to the fourth order with a gradually increasing λ . The fit parameters are ranked according to the value of λ necessary to reach an absolute value less than 10^{-6} . This list is ordered in terms of importance for the minimization of the MSLE. A visualization of the ordering can be found in Fig. 2, which shows the values of the second order parameters as λ is increased. The lower the value of λ at which the parameters are close to zero, the less important they are.

The parameters are then validated by calculating the MSLE with new data. First the data are validated with a fit involving all parameters. Next, the same validation is performed while excluding the parameter determined to be least important. After that, the least important two parameters are excluded. This process continues until all parameters have been excluded. This process generates a list of validation MSLE with parameters excluded in order of importance, from least important to most important. The validation MSLE is plotted in Fig. 3.

The plot shows a minimum when 22 parameters are excluded. Thus, 22 parameters are excluded for the nonlinear regression. This allows for quicker minimization.

When fitting the data, the following procedure is performed. First, we fit the parameters c_{000i} to the data from i , from 0 to 4, using the Newton method. The other fit parameters are kept constant at 0. Next we fit the parameters c_{000ij} with i and j from 0 to 4. The initial values of c_{000i} are the best-fit values from the previous minimization. The other fit parameters are given initial values of zero. All parameters that

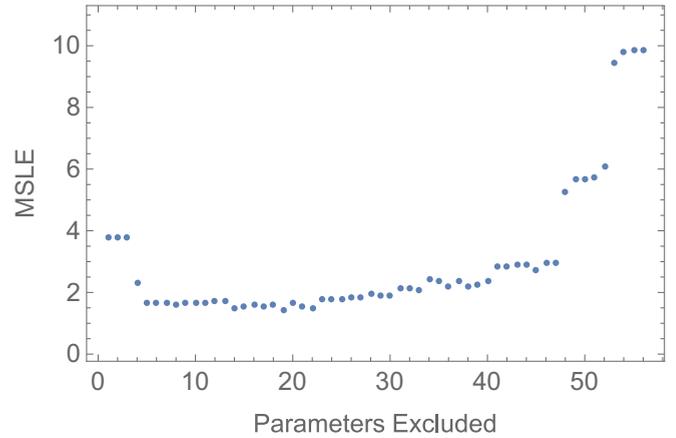


FIG. 3. The validation MSLE varying the number of excluded parameters. The data were fit to 1000 training points. The parameters are excluded in order of impact, as described in the text. The minimum corresponds to a number of excluded parameters that neither overfit nor underfit the 1000 data points.

are not used in the fit are once again set to zero. The same process is repeated for c_{0ijk} , and finally for c_{ijkl} .

This method of finding minimum parameters is employed in order to find a stable local minimum. If all parameters were fitted at once without appropriate starting values, small changes in the fit, such as changing the number of data points or increasing the penalty term, would result in drastically different local minima. However, it causes the fit to rely most heavily on lower-order parameters.

The best-fit parameters are given in Table VII in the Appendix. Perhaps because of the fit procedure, far more higher-order parameters than lower-order parameters are found by the least absolute shrinkage and selection operator (LASSO) to be unnecessary. The four fourth-order parameters that are not excluded are c_{1114} , c_{3334} , and c_{4441} . All three of these parameters involve the temperature T . This could indicate that the correlations between the temperature and other input variables are important. This emphasis on the importance of T in determining a_2 is in agreement with the linear correlation coefficients given in Table I; however, as described in the next section, nonlinear regression is not found to be the most accurate method to be employed. Thus, any speculations about the meaning of fourth-order parameters is highly uncertain.

V. RESULTS AND DISCUSSION

A. Comparison of methods for small N

In this section, as well as Secs. VB and VC, we discuss the methods used to predict the Mpemba effect for small N (equal to or less than 15). We use heat-bath dynamics with the rate matrix given in Eq. (3) with endpoints connected to fixed up spins. Because both endpoints are connected to fixed spins, there is a considerable difference between the case when N is odd and N is even. For odd N , in the lowest-energy state, each dipole will have the opposite spin of the two dipoles next to it. This is not possible for even N . We were unable to obtain good predictions of the Mpemba effect for even N in this section. This can be contrasted with the case of larger N in a ring, in

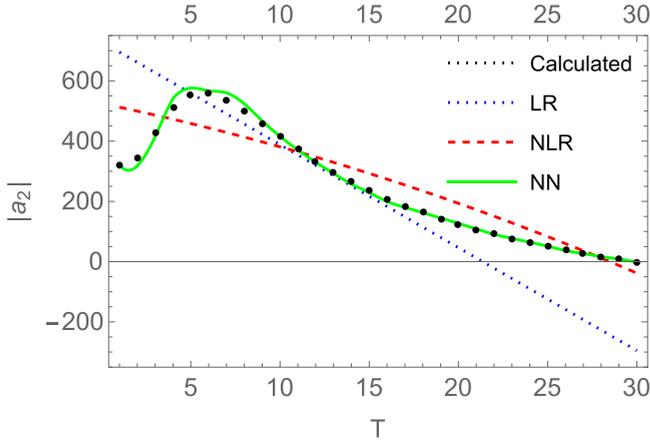


FIG. 4. An example of $|a_2|$ as a function of temperature. The black dots show the actual calculations compared with the predictions for various methods (colored curves). This occurs for $N = 7$, $J = -0.965$, and $h = -7.01$. The NN gives the most accurate prediction, followed by NLR and then LR. Note that the models are not fit to this data, but to a large set of data that includes various N , J , h , and T .

which we were able to obtain good predictions for even N , even when the even values were not included in the training data.

An example of the various methods compared with the data is shown in Fig. 4. Only the methods that predict $|a_2|$ directly are shown. The black dots give the actual computed data. At very low temperatures, they increase slightly with temperature, indicating a very weak Mpemba effect. After that, they smoothly decrease to equilibrium.

In order to compare the effectiveness of the various methods, several metrics are used. All of these metrics distinguish between training data, which are used to fix the free parameters or nodes of the model, and validation data, which are not used to determine the free parameters but are used to assess the accuracy.

The validation MSLE can be used to compare the accuracy of the neural network (NN), linear regression (LR), and non-linear regression (NLR). The decision tree (DT) and neural network trained directly on the Mpemba effect (NN2) cannot be compared because they minimize a different loss function.

The positive and negative validation accuracies are also computed. These are calculated by predicting whether the Mpemba effect occurs for randomly selected J , h , N , and two temperatures. The positive accuracy is the percentage of cases where the method predicts the Mpemba effect in which the effect actually occurs. The negative accuracy is the percentage of times the method correctly predicts that the effect does not occur. As discussed in Sec. IV C, the accuracy is meaningless for LR.

For each method, the accuracy increases and the error decreases as the amount of data increases; however, it eventually reaches a point at which additional data do not improve the predictions. Table II gives the calculations for the four methods for various numbers of data points.

The negative accuracy listed in the table for all four methods is much larger than the positive accuracy. This can be

TABLE II. The error, accuracy, and computational cost for several machine learning methods for varying numbers of training data points. The error is calculated with the MSLE on validation data. The positive (negative) accuracy is the fraction of correct results when the method predicts that the effect occurs (does not occur). The asterisk (*) is placed next to numbers which do not beat the baseline RRC described in the text. The training time gives the time spent training the data and the computation time gives the time for the machine learning method to validate a data set or make a prediction. Training time is not necessarily linear and thus we give the time for various sizes of training data sets. Computation time is linear so times for multiple data sets are not necessary.

Training points	DT	NN	NN2	LR	NLR
Error					
100		3.14		4.86	2.44
1000		0.57		4.08	1.65
10000		0.05		4.05	1.77
50000		0.03		4.14	1.66
200000		0.03		4.02	1.70
Positive accuracy (%)					
100	$1 \pm 14^*$	13 ± 9.5	67 ± 18		24 ± 21
1000	30 ± 6	18 ± 17	61 ± 14		15 ± 8.5
10000	46 ± 2.3	62 ± 8.5	79 ± 6.9		17 ± 7.7
50000	64 ± 1.4	77 ± 5.9	79 ± 4.6		20 ± 11
200000	66 ± 2.3	82 ± 7.2	91 ± 5.2		19 ± 00
Negative accuracy (%)					
100	98 ± 0.001	$92 \pm 4.5^*$	99 ± 0.6		$82 \pm 13^*$
1000	99 ± 0.5	99 ± 0.9	99 ± 0.2		$89 \pm 6.9^*$
10000	99 ± 0.02	99 ± 0.05	0.99 ± 0.09		$90 \pm 7.6^*$
50000	99 ± 0.02	99 ± 0.01	99 ± 0.03		$88 \pm 4.5^*$
200000	99 ± 0.01	99 ± 0.01	99 ± 0.05		$88 \pm 6.6^*$
Training time (seconds)					
100	0.58	0.4	0.8	0.2	10
1000	0.87	2.0	2.8	1.7	69
10000	2.08	9.4	14	38.2	365
50000	9.56	43	61	256	1624
Computation time (seconds)					
50000	1.41	1.52	1.48	0.20	1.71

explained by the fact that the Mpemba effect occurs in approximately 2% of cases. Methods are more likely to minimize a loss function by predicting that the effect does not occur. For reference, a method that always predicts that the effect does not occur regardless of training data would have a 0% positive accuracy and approximately 98% negative accuracy.

This work establishes a baseline accuracy for predicting the Mpemba effect with various methods which can be used for future comparison of machine learning predictions. To compare the results in this work, we consider baseline “classifiers” of binary outcomes used in Refs. [33,46,47] to compare to predictions in a field that was not previously classified. The first of these is the zero-rate classifier (ZRC), which classifies all cases as the more likely of the outcomes. In this case, the more likely outcome is for the effect not to occur. Thus, the ZRC would have a total accuracy of 98%. It is less helpful to use the ZRC to make a prediction of positive or negative accuracy. No matter what the data are, the ZRC will always

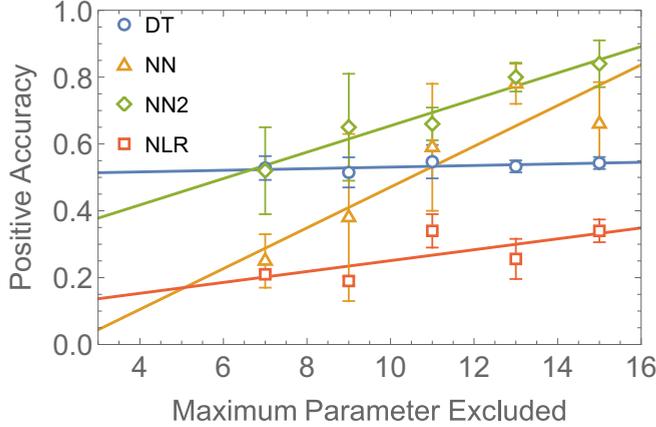


FIG. 5. A plot of the positive accuracy vs maximum value of N excluded. The full data are given in Table VI. Best-fit lines are shown to roughly estimate the relationship between the variables.

give one outcome a 0%. To establish a baseline for positive and negative accuracy, the random rate classifier (RRC) is used. The RRC classifies by randomly selecting one of the binaries with a likelihood proportional to the total fraction of positive or negative values. In this case, the RRC would guess that the effect does not occur with a likelihood of 98%. This corresponds to a total accuracy of $0.98^2 + 0.02^2 \approx 0.96$ and a positive and negative accuracy of 98% and 2%, respectively. This is the baseline that our results can be compared to.

We find that for large data sets, NN2 is most accurate, though NN is almost as accurate and has the added benefit of predicting the values of a_2 . For very small training sets, NLR is the most effective method at predicting $|a_2|$, although it fails to reach a high level of accuracy with larger data sets.

B. Extrapolation

In the previous section, the parameters for the validation data were different from those of the training data, but they were generated in the same range. Machine learning can also be used to extrapolate by validating the model with data in a different range. This is particularly beneficial if the models can make accurate predictions on systems that are more computationally expensive than the systems they are trained on. Out of the parameters N , J , h , and T , only N determines the computational complexity of the calculation. The time to compute $|a_2|$ does not depend on J , h , or T .

Figure 5 plots the positive accuracy of various methods for predicting the Mpemba effect in the case when $N = 15$, when it has only been trained on data with $N < 15$. The full results are given in Table VI in the Appendix. The lines shown in Fig. 5 give a rough approximation of how the accuracy depends on the maximum parameter that is excluded. Uncertainties are given by the standard deviation of five resamples.

When values of N close to 15 are included in the training data, the neural networks are once again the most accurate methods; however, the DT is relatively accurate at making predictions even when only $N = 5$ is included. In fact, changing which values of N are included has very little impact on the accuracy of the predictions, indicating that N has very little effect on the prediction with this algorithm. These results agree

TABLE III. The predictions for the Mpemba effect when the neural network is only trained on data where the Mpemba effect is not present. The neural network is trained on 20 000 data points.

	TME	ME
Positive accuracy	0.7764 ± 0.064	0.3035 ± 0.134
Negative accuracy	0.8487 ± 0.070	0.9848 ± 0.012

with the crude prediction from the correlation coefficients, given in Table I.

It is reasonable to assume that the methods that are best at long-range extrapolation in the ranges that we consider will also be most effective at long-range extrapolation outside the range we consider or for different systems than the Ising model. However, this has not been demonstrated for certain and merits further investigation.

Long-range extrapolation could be useful due to the high computational cost for a_2 for higher N . The number of states for a given N is 2^N , leading to a $2^N \times 2^N$ transition matrix. The computational cost of the eigenvalues for a $M \times M$ matrix is $O(M^2)$ [48]. Thus the computational cost for calculating a_2 is 2^{2N} . Thus, a calculation with $N = 5$ is quicker than a calculation with $N = 15$ by a factor of 2^{20} . Accurate prediction trained on much simpler systems could save considerable computational cost. It should be noted that ML methods are not the only numerical methods to avoid computational costs. The Monte Carlo methods used in Ref. [23] are used to make calculations that would otherwise be impossible within realistic time constraints. For this method to be of maximum use, the computational cost recorded in Table II would need to be lower than equivalent values for Monte Carlo methods.

Notably, the Mpemba effect can be predicted even if it is trained only with data in which it does not occur. Results for this extrapolation are given in Table III. The first prediction of the Mpemba effect (ME) is the same as in the previous sections. The second method of prediction, referred to as the total Mpemba effect (TME), measures the accuracy of predicting whether or not the Mpemba effect occurs for any temperature for a given J , h , and N . Put another way, neural networks can identify the nonmonotonic temperature dependence of the coefficient a_2 even when they are only trained on monotonically decreasing a_2 . This suggests that the Mpemba effect is encoded in the antiferromagnetic Ising model. This further confirms the results of Refs. [35,37].

The machine learning methods can also be extrapolated to J outside of the range the method was trained on. Specifically, our data set only includes negative J . No Mpemba effect has been found for positive J . To test whether the neural networks can identify this, we made predictions on over 10^6 data points for various N , and h . For each case, the algorithm correctly predicts that no Mpemba effect occurs for positive J . This result would not be as notable if the neural network was identifying that the effect never occurs for positive J because more negative J tend to cause the Mpemba effect. Figure 6 shows the regions where the Mpemba effect occurs on the $J - h$ plane. The Mpemba effect does occur for negative values of J close to zero and does not occur for very negative values of J . This means that the neural network's identification that the effect only occurs for negative J is a more subtle inference.

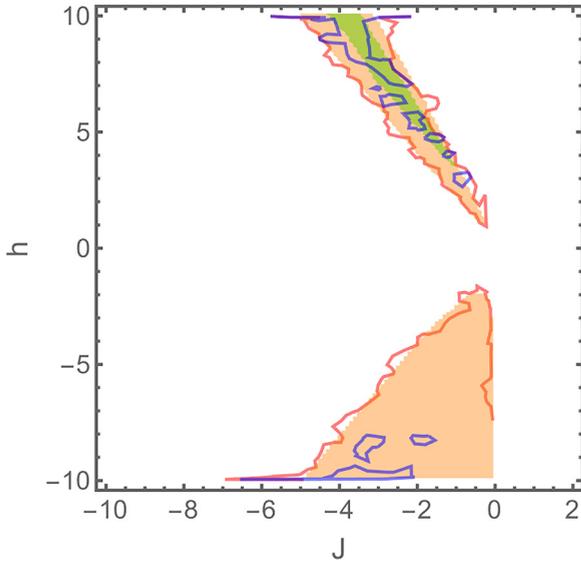


FIG. 6. The J - h plane. The shaded orange region (lighter-shaded region) shows the NN2 prediction for the region in which the Mpemba effect occurs. The red contour (lighter contour) shows the region generated with a contour from the data. The shaded green region (darker-shaded region) shows the predicted region for the strong Mpemba effect and the blue contour (darker contour) shows the contour generated from the data. We observe good agreement between the results and the prediction, especially for the ordinary Mpemba effect.

C. The strong Mpemba effect

In addition to the Mpemba effect, machine learning methods can be used to predict the strong Mpemba effect, defined in Ref. [5]. In this situation, the hot system cools exponentially faster than the cold system.

The strong Mpemba effect occurs when $a_2 = 0$ and $T \neq T_b$. The neural network is trained on three input variables N , J , and h . The output is 1 if the strong Mpemba effect occurs for any temperature; the output is 0 if the strong Mpemba effect does not occur. Results are given in Table IV.

The machine learning algorithm can be used to generate plots of the J - h plane. Figure 6 shows a plot of which values the Mpemba effect occurs at and which values the strong Mpemba effect occurs at. This is compared to the predictions made by the neural network. This diagram is inspired by the phase diagrams of Ref. [35] (Fig. 2), although it differs in that it plots the J - h plane rather than the T_b - h plane.

D. Predictions for large N using coarse graining

Much larger values of N can be computed for a given computational cost using the method of coarse graining. We generate a data set using the rate matrix given in Eq. (4). This

TABLE IV. The validation accuracy for the strong Mpemba effect. It is trained with 20 000 data points.

Positive accuracy	Negative accuracy
0.4633 ± 0.043	0.995 ± 0.0016

TABLE V. Validation accuracy and error of predictions by NN1 and NN2. Error is for NN1 only. Extrapolations are given above the double line. In this case, all values of N listed as excluded in addition to all values of N greater than the predicted value are excluded. Below the double line, predictions for validation data within the range of the training data are included.

Excluded	Predicted	NN1+	NN2+	NN1-	NN2-	Error
55, 57, 59	59	0.401	0.438	0.972	0.935	14.2
53	53	0.731	0.715	0.981	0.988	3.47
51, 53	53	0.530	0.627	0.932	0.988	5.91
49, 51, 53	53	0.475	0.517	0.986	0.989	8.64
50	50	0.515	0.985	0.974	0.838	3.31
	All	0.972	0.946	0.781	0.760	4.12
	53	0.991	0.986	0.669	0.799	5.69

data set includes a_2 for the weak-coupling limit for spins in a ring rather than with fixed end points.

To test the methods, we generate a data set with odd numbered N from 15 to 53, as well as sets with $N = 59$ and $N = 50$. This range is comparable to the size of the Ising chain in Ref. [35]; however, an exact coarse graining allowed a computation of an Ising chain with 1000 spins in Ref. [37]. Furthermore, the thermodynamic limit of large N has been studied through the phenomenon of eigenvalues crossing (closely linked to the Mpemba effect) [31].

We find that NN1 and NN2 can be used to make accurate predictions even outside the range that the model is trained on. Results are shown in Table V.

Notably, predictions can be made for $N = 50$. As noted, in the previous sections, predictions could not be made for even N ; however, with larger N that form a ring, predictions can be made for even N , even if only odd N are included in the training data. However, the positive accuracy of each of these methods noticeably decreases when more values of N are excluded. This suggests a limitation to long-range extrapolation for these Ising chains.

VI. CONCLUSION

The Mpemba effect has been shown to occur in many systems, beyond the freezing of water. These applications motivate better predictions of the effect. In order to understand this effect, this work applies statistical methods to the Mpemba effect in the Ising model.

For the case of 5–15 spins in an Ising chain, we demonstrate that a number of machine learning methods can be used. Neural networks are the most effective method when a large enough training data set is used. For very small data sets, nonlinear regression with the LASSO method may be more effective at predicting a_2 . These methods can predict the Mpemba effect in systems that are much more complex than those on which they were trained. The decision tree method may be the most effective method at making predictions far outside the range it was trained on. Additionally, the Mpemba effect can be predicted with neural networks when it is trained only on data where the Mpemba effect does not occur. This indicates that information about the Mpemba effect exists in situations in which it does not occur. Furthermore, we

TABLE VI. Extrapolation accuracy. The data is trained on 50 000 data points.

Excluded values of N	DT	NN	NN2	NLR
Positive accuracy				
15	0.5436 ± 0.018	0.6638 ± 0.125	0.8421 ± 0.0724	0.3435 ± 0.03047
13,15	0.5334 ± 0.0184	0.77960 ± 0.0633	0.8040 ± 0.04347	0.25629 ± 0.06267
11,13,15	0.5479 ± 0.0503	0.5922 ± 0.1994	0.6611 ± 0.04952	0.3467 ± 0.05206
9,11,13,15	0.5155 ± 0.04509	0.3826 ± 0.2561	0.6472 ± 0.1640	0.1909 ± 0.01386
7,9,11,13,15	0.5283 ± 0.0356	0.2741 ± 0.08402	0.5222 ± 0.1288	0.2184 ± 0.009012
Negative accuracy				
15	0.996207 ± 0.00245	0.99723 ± 0.000954173	0.997952 ± 0.00042289	0.718728 ± 0.0392236
13,15	0.99757 ± 0.001162	0.992953 ± 0.00388419	0.996802 ± 0.00108171	0.83444 ± 0.04182
11,13,15	0.998175 ± 0.000508	0.987025 ± 0.00648745	0.997013 ± 0.000912489	0.890572 ± 0.0353344
9,11,13,15	0.996978 ± 0.001237	0.958839 ± 0.0328791	0.996679 ± 0.00173007	0.837763 ± 0.034161
7,9,11,13,15	0.99767 ± 0.00033	0.960547 ± 0.0529037	0.992296 ± 0.01823	0.93159 ± 0.00735792

demonstrate that the strong Mpemba effect can also be predicted using machine learning methods. It should be noted that a limitation to these methods (in comparison to numerical methods in other works) is that they can only be successfully applied to chains with an odd number of spins.

In addition, we have studied the effect for a larger number of spins in the spin chain using coarse graining in the weak-coupling limit. In this method, neural networks can be used to make accurate predictions for the Mpemba effect. For these spin chains, accurate predictions can be made for an even number of spins, even if the network is only trained with data from odd-numbered spin chains. However, a limitation of predicting the Mpemba effect in these systems is that long-range extrapolation to much larger spin chains than the networks are trained on is not as accurate.

This work can be extended in several ways. First, future analysis could vary the bath temperature. Such an investigation could more optimally test the effect by selecting a single bath temperature and scanning for nonmonotonicity. Furthermore, these predictions were performed for the Mpemba effect in the Ising model; however, the relative accuracy of each method might hold for the Mpemba effect in other systems. Determination of the accuracy of various methods in other systems merits further investigation.

ACKNOWLEDGMENTS

This work was funded by an Ave Maria University undergraduate research grant by Michael and Lisa Schwartz.

TABLE VII. LASSO parameters for a fit with 500 000 data points. Fit parameters $c_{k_1 k_2 k_3 k_4}$ are given in Eq. (14). All parameters not present in this table are either excluded by the LASSO as described in the main text or unnecessary because they are multiplied by a term identical to that of one of the included parameters.

c_{0000}	12.0493	c_{0001}	56.5418	c_{0002}	-77.7456	c_{0003}	-3.88612	c_{0004}	4.93836
c_{0011}	11.9538	c_{0012}	-389.284	c_{0013}	-12.5517	c_{0014}	-6.88152	c_{0023}	4.51196
c_{0024}	0.37328	c_{0034}	0.780776	c_{0044}	67.0015	c_{0111}	-11.4937	c_{0112}	-77.7841
c_{0113}	-5.35089	c_{0114}	-1.60171	c_{0221}	-308.538	c_{0222}	-11.4937	c_{0224}	87.8123
c_{0331}	-801.778	c_{0332}	67.5452	c_{0333}	458.608	c_{0334}	183.366	c_{0441}	335.435
c_{0442}	-418.124	c_{0443}	37.4534	c_{0444}	-38.2862	c_{0123}	-0.00109161	c_{0124}	16.838
c_{0134}	0.431855	c_{0234}	0.00739725	c_{1114}	0.294404	c_{3334}	-16.9157	c_{4441}	-193.343

Thanks to Tomás Licheri for completing crucial tasks necessary for this research.

APPENDIX A: ADDITIONAL DATA

In this section, we give data for reference that is not used in the main analysis and discussion of the paper. Table VI gives the extrapolation accuracy of the methods. Positive accuracies of this table are plotted in Fig. 5 and the format is equivalent to Table II.

Table VII gives the best-fit parameters for nonlinear regression. Parameters not included in the table are either mathematically identical to an excluded parameter or shown by the LASSO method to have less impact. The fit was performed with 50 000 data points and corresponds to the fifth row of each section of Table II for NLR; however, NLR for any amount of data points greater than 1000 had similar minima. Note that the fit was done to data with $|a_2|$ multiplied by a factor of 10 000 as described in the text. To relate these results to data that have not been prepared in this way, one should divide each parameter by 10 000.

APPENDIX B: NEURAL NETWORK ARCHITECTURE TESTS

In this work, we employ a neural network with three hidden layers with a total of 75 nodes: 50 in the first layer, 20 in the second, and 5 in the third. In this Appendix, we explain our choice of architecture.

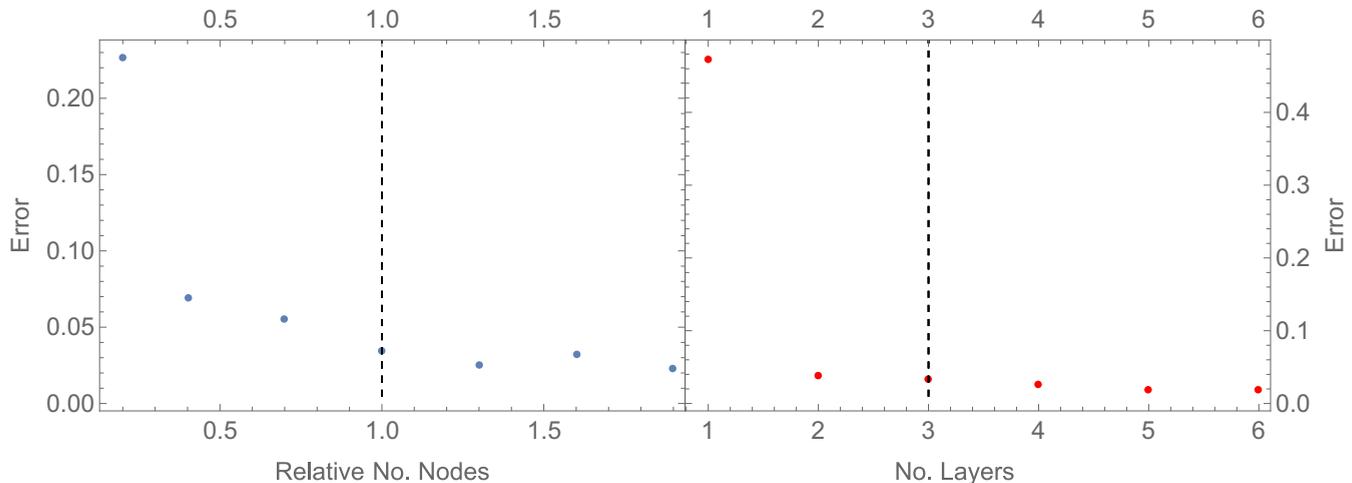


FIG. 7. Plots of the accuracy of the neural network (NN1) vs quantities of the architecture. The left plot varies the relative number of nodes. In this plot, there are three hidden layers of nodes. The ratios of each layer are held constant as the number of nodes is varied. The relative number is the number of nodes employed divided by our chosen number of hidden nodes (75). The right plot varies the number of layers. In this plot, the number of nodes is held constant. They are divided among the number of hidden layers on the horizontal axis. In both plots, the vertical dashed line indicates the values employed in this work.

There is no one unique best architecture for a neural network. However, if a neural network has too few nodes or too few layers, it will not be able to achieve a small error. If a neural network has too many nodes or too many layers, it can take prohibitively long to train or find weights that are at a local minimum that is nowhere near the global minimum. Our approach is to vary both the number of nodes and the number of layers until we find the smallest possible network that cannot significantly reduce the error by increasing the number of nodes or layers.

This variation is shown in Fig. 7. Our process was to vary the number of layers and number of nodes as described in the caption. We used the ADAM optimizer with a training data and a validation data of 50 000

data point. 200 epochs were employed with a batch size of 200.

Above the chosen values denoted with the dashed line, there is a slight decrease in the error but it is very small. For values below the dashed line, there is a very large increase in the error.

Our choices are not the only possible choices. Slightly different values could also be chosen and these might lead to results that are slightly better, but these results give evidence that a significant improvement in the accuracy will not occur. On the other hand, if the system had been much smaller either in nodes or layers, the error and hence the predictions would be significantly worse. This shows that the chosen architecture is reasonable.

-
- [1] E. B. Mpemba and D. G. Osborne, Cool? *Phys. Educ.* **4**, 172 (1969).
- [2] H. D. P. Lee, *Loeb Classical Library 397* (Harvard University Press, Cambridge, MA, 1952).
- [3] F. Bacon, *The Novum Organum of Sir Francis Bacon* (Thomas Lee London, England, 1677), pp. [3], 32.
- [4] H. C. Burridge and F. P. Linden, Questioning the Mpemba effect: Hot water does not cool more quickly than cold, *Sci. Rep.* **6**, 37665 (2016).
- [5] I. Klich, O. Raz, O. Hirschberg, and M. Vucelja, Mpemba Index and Anomalous Relaxation, *Phys. Rev. X* **9**, 021060 (2019).
- [6] A. Lasanta, F. Vega Reyes, A. Prados, and A. Santos, When the Hotter Cools More Quickly: Mpemba Effect in Granular Fluids, *Phys. Rev. Lett.* **119**, 148001 (2017).
- [7] M. Baity-Jesi, E. Calore, A. Cruz, L. A. Fernandez, J. M. Gil-Narvión, A. Gordillo-Guerrero, D. Iñiguez, A. Lasanta, A. Maiorano, E. Marinari, V. Martin-Mayor, J. Moreno-Gordo, A. M. Sodupe, D. Navarro, G. Parisi, S. Perez-Gaviro, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, S. F. Schifano, B. Seoane *et al.*, The Mpemba effect in spin glasses is a persistent memory effect, *Proc. Natl. Acad. Sci.* **116**, 15350 (2019).
- [8] A. Torrente, M. A. López-Castaño, A. Lasanta, F. V. Reyes, A. Prados, and A. Santos, Large Mpemba-like effect in a gas of inelastic rough hard spheres, *Phys. Rev. E* **99**, 060901(R) (2019).
- [9] A. Gijón, A. Lasanta, and E. R. Hernández, Paths towards equilibrium in molecular systems: The case of water, *Phys. Rev. E* **100**, 032103 (2019).
- [10] A. Megías and A. Santos, Mpemba-like effect protocol for granular gases of inelastic and rough hard disks, *Front. Phys.* **10**, 971671 (2022).
- [11] Z. Lu and O. Raz, Nonequilibrium thermodynamics of the Markovian Mpemba effect and its inverse, *Proc. Natl. Acad. Sci. USA* **114**, 5083 (2017).
- [12] D. Busiello, D. Gupta, and A. Maritan, Inducing and optimizing Markovian Mpemba effect with stochastic reset, *New J. Phys.* **23**, 103012 (2021).

- [13] G. Teza, R. Yaacoby, and O. Raz, Relaxation Shortcuts through Boundary Coupling, *Phys. Rev. Lett.* **131**, 017101 (2023).
- [14] Z. Tang, W. Huang, Y. Zhang, Y. Liu, and L. Zhao, Direct observation of the Mpemba effect with water: Probe the mysterious heat transfer, *InfoMat* **5**, e12352 (2022).
- [15] M. Jeng, The Mpemba effect: When can hot water freeze faster than cold? *Am. J. Phys.* **74**, 514 (2006).
- [16] P. Chaddah, S. Dash, K. Kumar, and A. Banerjee, Overtaking while approaching equilibrium, [arXiv:1011.3598](https://arxiv.org/abs/1011.3598).
- [17] F. J. Schwarzendahl and H. Löwen, Anomalous Cooling and Overcooling of Active Colloids, *Phys. Rev. Lett.* **129**, 138002 (2022).
- [18] A. Megías, A. Santos, and A. Prados, Thermal versus entropic Mpemba effect in molecular gases with nonlinear drag, *Phys. Rev. E* **105**, 054140 (2022).
- [19] A. Biswas, V. V. Prasad, and R. Rajesh, Mpemba effect in anisotropically driven inelastic Maxwell gases, *J. Stat. Phys.* **186**, 45 (2022).
- [20] J. Meibohm, D. Forastiere, T. Adeleke-Larodo, and K. Proesmans, Relaxation-speed crossover in anharmonic potentials, *Phys. Rev. E* **104**, L032105 (2021).
- [21] Y.-H. Ahn, H. Kang, D.-Y. Koh, and H. Lee, Experimental verifications of Mpemba-like behaviors of clathrate hydrates, *Korean J. Chem. Eng.* **33**, 1903 (2016).
- [22] P. A. Greaney, G. Lani, G. Cicero, and J. C. Grossman, Mpemba-like behavior in carbon nanotube resonators, *Metall. Mater. Trans. A* **42**, 3907 (2011).
- [23] A. Gal and O. Raz, Precooling Strategy Allows Exponentially Faster Heating, *Phys. Rev. Lett.* **124**, 060602 (2020).
- [24] R. Chétrite, A. Kumar, and J. Bechhoefer, The metastable Mpemba effect corresponds to a nonmonotonic temperature dependence of extractable work, *Front. Phys.* **9**, 654271 (2021).
- [25] J. Lin, K. Li, J. He, J. Ren, and J. Wang, Power statistics of Otto heat engines with the Mpemba effect, *Phys. Rev. E* **105**, 014104 (2022).
- [26] A. Militaru, A. Lasanta, M. Frimmer, L. Bonilla, L. Novotny, and R. Rica, Kovacs Memory Effect with an Optically Levitated Nanoparticle, *Phys. Rev. Lett.* **127**, 130603 (2021).
- [27] E. Mompó, M. Á. López-Castaño, A. Lasanta, F. Reyes, and A. Torrente, Memory effects in a gas of viscoelastic particles, *Phys. Fluids* **33**, 062005 (2021).
- [28] I. González-Adalid Pemartín, E. Mompó, A. Lasanta, V. Martín-Mayor, and J. Salas, Slow growth of magnetic domains helps fast evolution routes for out-of-equilibrium dynamics, *Phys. Rev. E* **104**, 044114 (2021).
- [29] A. Patrón, B. Sánchez-Rey, and A. Prados, Strong nonexponential relaxation and memory effects in a fluid with nonlinear drag, *Phys. Rev. E* **104**, 064127 (2021).
- [30] V. Uskokovic, ...And all the world a dream: Memory effects outlining the path to explaining the strange temperature-dependency of crystallization of water, a.k.a. the Mpemba effect, *Substantia* **4**, 59 (2020).
- [31] G. Teza, R. Yaacoby, and O. Raz, Eigenvalue Crossing as a Phase Transition in Relaxation Dynamics, *Phys. Rev. Lett.* **130**, 207103 (2023).
- [32] N. Walker, K.-M. Tam, and M. Jarrell, Deep learning on the 2-dimensional Ising model to extract the crossover region with a variational autoencoder, *Sci. Rep.* **10**, 13047 (2020).
- [33] A. Morningstar and R. Melko, Deep learning the Ising model near criticality, *J. Mach. Learn. Res.* **18**, 1 (2017).
- [34] N. Portman and I. Tamblin, Sampling algorithms for validation of supervised learning models for Ising-like systems, *J. Comput. Phys.* **350**, 871 (2017).
- [35] G. Teza and A. L. Stella, Exact Coarse Graining Preserves Entropy Production out of Equilibrium, *Phys. Rev. Lett.* **125**, 110601 (2020).
- [36] G. Teza, Out of equilibrium dynamics: From an entropy of the growth to the growth of entropy production (unpublished).
- [37] G. Teza, R. Yaacoby, and O. Raz, Far from equilibrium relaxation in the weak coupling limit (unpublished).
- [38] W. E. Arnoldi, The principle of minimized iterations in the solution of the matrix eigenvalue problem, *Q. Appl. Math.* **9**, 17 (1951).
- [39] L. Breiman, J. H. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*, CART (Wadsworth and Brooks, Monterey, CA, 1984).
- [40] T. O. Hodson, T. M. Over, and S. S. Foks, Mean squared error, deconstructed, *J. Adv. Model. Earth Syst.* **13**, e2021MS002681 (2021).
- [41] B. Higgins, Fitting data with Legendre polynomials (unpublished).
- [42] R. Tibshirani, Regression shrinkage and selection via the LASSO, *J. R. Stat. Soc. B* **58**, 267 (1996).
- [43] B. Guegan, J. Hardin, J. Stevens, and M. Williams, Model selection for amplitude analysis, *JINST* **10**, P09002 (2015).
- [44] D. Sadasivan, M. Mai, and M. Döring, S- and p-wave structure of $S = -1$ meson-baryon scattering in the resonance region, *Phys. Lett. B* **789**, 329 (2019).
- [45] J. Landay, M. Mai, M. Döring, H. Haberzettl, and K. Nakayama, Towards the Minimal Spectrum of Excited Baryons, *Phys. Rev. D* **99**, 016001 (2019).
- [46] J.-L. Xu, X. Lin, and A. A. Gowen, Combining machine learning with meta-analysis for predicting cytotoxicity of micro- and nanoplastics, *J. Hazard. Mater. Adv.* **8**, 100175 (2022).
- [47] I. Kraljevski, F. Duckhorn, C. Tschöpe, F. Schubert, and M. Wolff, Paper tissue softness rating by acoustic emission analysis, *Appl. Sci.* **13**, 1670 (2023).
- [48] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. (Cambridge University Press, New York, NY, 1992).