



OPEN

## Parsimonious neural networks learn interpretable physical laws

Saaketh Desai & Alejandro Strachan

Machine learning is playing an increasing role in the physical sciences and significant progress has been made towards embedding domain knowledge into models. Less explored is its use to discover interpretable physical laws from data. We propose parsimonious neural networks (PNNs) that combine neural networks with evolutionary optimization to find models that balance accuracy with parsimony. The power and versatility of the approach is demonstrated by developing models for classical mechanics and to predict the melting temperature of materials from fundamental properties. In the first example, the resulting PNNs are easily interpretable as Newton's second law, expressed as a non-trivial time integrator that exhibits time-reversibility and conserves energy, where the parsimony is critical to extract underlying symmetries from the data. In the second case, the PNNs not only find the celebrated Lindemann melting law, but also new relationships that outperform it in the pareto sense of parsimony vs. accuracy.

Machine learning (ML) can provide predictive models in applications where data is plentiful and the underlying governing laws are unknown<sup>1–3</sup>. These approaches are also playing an increasing role in the physical sciences where data is generally limited but underlying laws (sometimes approximate) exist<sup>4–9</sup>. For example, ML-based constitutive models are being used in electronic structure calculations<sup>10</sup> and molecular dynamics (MD) simulations<sup>11–13</sup>. One of the major drawbacks of the use of ML in the physical sciences is that models often do not *learn* the underlying physics of the system at hand, such as constraints or symmetries, limiting their ability to generalize. In addition, most ML models lack interpretability. That is, ML approaches generally neither learn physics nor can they explain their predictions. In many fields, these limitations are compensated by copious amounts of data, but this is often not possible in areas such as materials science where acquiring data is expensive and time consuming. To tackle this challenge, progress has been made towards using knowledge (even partial) of underlying physics to improve the accuracy of models and/or reduce the amount of data required during training<sup>14–16</sup>. Less explored is the use of ML for scientific discovery, i.e. extracting physical laws from observational data, see Refs.<sup>17–19</sup> for recent progress. In this letter, we combine neural networks (NNs) with stochastic optimization to find models that balance accuracy and parsimony and apply them to learn, solely from observational data, the dynamics of a particle under a highly non-linear potential, and expressions to predict the melting temperature of materials in terms of fundamental properties. Our hypothesis is that the requirement of parsimony will result in the discovery of the physical laws underlying the problem and result in interpretability and improved generalizability. We find that the resulting descriptions are indeed interpretable and provide insight into the system of interest. In the case of particle dynamics, the learned models satisfy non-trivial underlying symmetries embedded in the data which increases the applicability the parsimonious neural networks (PNNs) over generic NN models. Stochastic optimization has been previously used in conjunction with backpropagation to improve robustness or minimize overfitting in models<sup>20–26</sup>, this work extends these ideas to finding parsimonious models from data to learn physics.

The power of physics-based ML is well documented and remains an active area of research. Neural networks have been used to both parametrize and solve differential equations such as Navier Stokes<sup>14,15</sup> and Hamilton's equations of motion<sup>27</sup>. Recurrent architectures have also shown promise in predicting the time evolution of systems<sup>28,29</sup>. These examples focus on using prior knowledge of the underlying physics to guide the model, often as numerical constraints, or by using the underlying physics to numerically solve equations with variables predicted by the ML algorithms. In contrast, we are interested in learning physics, including the associated numerical solutions, directly from data, without prior knowledge. Pioneering work along these lines used symbolic regression methods, enhanced by matching partial derivatives to identify invariants<sup>17</sup>, or using dimensionality reduction and other symmetry-identifying methods to aid equation discovery<sup>30</sup>. These approaches also consider the tradeoff between parsimony and accuracy to develop simple models that describe the data well. On the other hand, neural networks such as time-lagged autoencoders have also proven useful at extracting laws that govern

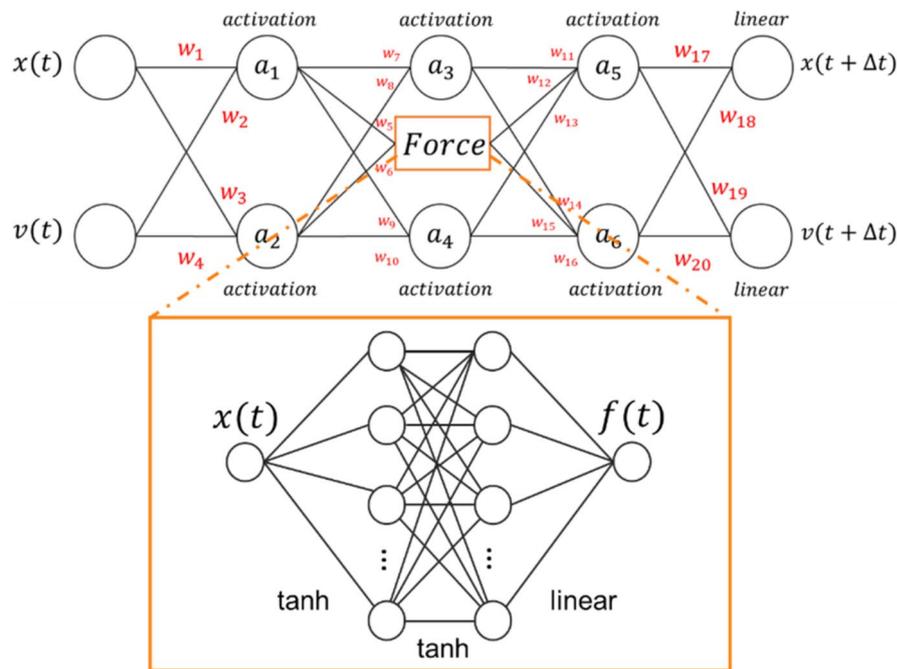
School of Materials Engineering and Birck Nanotechnology Center, Purdue University, West Lafayette, IN 47907, USA. email: strachan@purdue.edu

the time evolution of systems from data<sup>31</sup>, where the encoder networks attempt to learn features relevant to the problem. Advances here have considered networks with custom activation functions whose weights decide the functional form of the equation<sup>32,33</sup>. Lastly, other approaches to learning physics from data have focused on discovering partial differential equations directly from data, either using a library of candidate derivatives coupled with linear regression<sup>18,19</sup>, or using neural networks coupled with genetic algorithms to identify differential equations from an incomplete library<sup>34</sup>. We build on and extend these ideas to propose PNNs, models designed to balance parsimony with accuracy in describing the training data. The PNN approach allows complex compositions of functions via the use of neural networks, while balancing for parsimony using genetic algorithms. As will be shown with two examples, our approach is quite versatile and applicable to situations where an underlying differential equation may not exist. We first apply PNNs to learn the equations of motion that govern the Hamiltonian dynamics of a particle under a highly non-linear external potential with and without friction. Our hypothesis is that by requiring parsimony (e.g. minimizing adjustable parameters and favoring linear relationships between variables) the resulting model will not only be easily interpretable but also will be forced to tease out the symmetries of the problem. We find that the resulting PNN not only lends itself to interpretation (as Newton's laws) but also provides a significantly more accurate description of the dynamics of the particle when applied iteratively as compared to a flexible feed forward neural network. The resulting PNNs conserve energy and are time reversible, i.e., they learn non-trivial symmetries embedded in the data but not explicitly provided. This versatility and the generalizability of PNNs is demonstrated with a second, radically different, example: discovering models to predict the melting temperature of materials from atomic and crystal properties. By varying the relative importance of parsimony and accuracy in the genetic optimization, we discover a family of melting laws that include the celebrated Lindemann law<sup>35</sup>. Quite remarkably the Lindemann law, proposed in 1910, is near (but not on) the resulting pareto front.

## Results

**Discovering integration schemes from data.** As a first example, we consider the dynamics of a particle under an external Lennard–Jones (LJ) potential with and without friction. In both cases the training data is obtained from accurate numerical trajectories with various total energies. The input and output data are positions and velocities at a given time and one timestep ahead, respectively (this timestep is ten times what was used to generate the underlying trajectories). The numerical data was divided into training and validation sets and an independent testing set was generated at a different energy, see [Methods](#) and section S1 of the Supplementary Material (SM). The input data in this example has been generated numerically for convenience but could have been obtained experimentally, as will be shown in the second example. Before describing the PNN model, we establish a baseline by training a standard feed forward neural network (FFNN) on our data for the case without friction. The details of this architecture can be found in section S2 of the SM and can be accessed for online interactive computing on nanoHUB<sup>36</sup>. We find the FFNN to be capable of matching the training/validation/test data well, with root mean squared errors (RMSEs) across all sets on the order of  $10^{-5}$  in LJ units for both positions and velocities (see Figure S1 in SM). However, the network has poor predictive power. Using it iteratively to find the temporal evolution of the particle results in significant drifts in total energy over time, and a lack of time reversibility. Reversibility is judged by applying the network sequentially 1,000 times, followed by time reversal (changing the sign of the particle's velocity) and applying the NN for a second set of 1,000 steps. We find that deeper architectures do not improve the RMSE, reversibility or energy conservation. Needless to say, these FFNNs are not interpretable. These results highlight the difficulty of the problem at hand. Hamilton's equations for classical mechanics represent a stiff set of differential equations and small errors in each timestep accumulate rapidly resulting in diverging trajectories. Prior attempts to address this challenge in the context of discovering differential equations explicitly trained models for multiple steps using a recurrent architecture<sup>37</sup>. The resulting models are interpretable and improve accuracy over the number of steps used in training the recurrent network but accumulate relatively high errors over multiple steps. In contrast, we are interested in solutions stable over timescales far greater than those typically accessed by current recurrent architectures, while also favoring the discovery of constants relevant to the physical problem. Finding such models is non-trivial and the development of algorithms to integrate equations of motion with good energy conservation and time reversibility has a rich history<sup>38–42</sup>. An example of such algorithms is the popular Verlet family of integrators<sup>38,39</sup> that are both reversible and symplectic<sup>43</sup>; their theoretical justification lies in Trotter's theorem<sup>44</sup>.

**Parsimonious neural networks.** Having established the shortcomings of the state-of-the-art neural networks, we introduce parsimonious neural networks (PNN). We begin with a generic neural network shown in Fig. 1 and use genetic algorithms to find models with controllable parsimony. In this first example, the neural network consists of three hidden layers and an output layer with two values, the position and velocity of the particle one timestep ahead of the inputs. Each hidden layer has two neurons, and the central layer includes an additional force sub-net, a network pre-trained to predict the force on the atom given its position. The use of a pre-trained force sub-net is motivated by the prior success of neural networks in predicting interatomic forces in a wide variety of materials significantly more complex than our example<sup>45–47</sup>. The architecture of the force sub-net was designed to be the simplest network that can predict the force with sufficient accuracy to result in accurate dynamics. Our architecture is similar to those used to predict interatomic forces, where atomic neighborhoods are encoded in descriptors, which are then mapped to the atomic forces via a one or two-layer shallow neural network<sup>12</sup>. In our one-dimensional case, the input is simply the atom coordinate. In addition, our focus is on learning classical dynamics and the use of a force sub-net only incorporates the physical insight that the force is an important quantity. As a second baseline, we trained a feed forward NN including a pre-trained force sub-net. This second network's performance is as poor as the previous baseline feed-forward network, see section S3



**Figure 1.** Neural network used as the starting point to find the parsimonious neural network as the network that explains the data in the simplest manner possible. The force sub-network is highlighted in orange and is fed into the neural network as a pre-trained model, whose weights are subsequently kept fixed throughout.

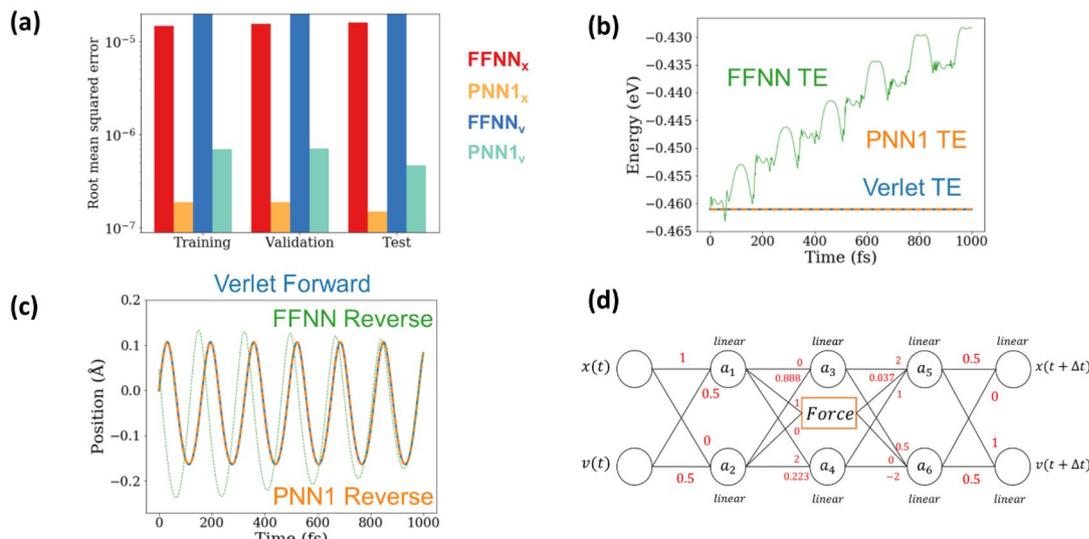
in the SM for details. This shows that adding the information about the force is not the key to the development of accurate models for classical mechanics, parsimony is.

The starting neural network provides a highly flexible mapping from input positions and velocities to output positions and velocities and PNNs seeks to balance simplicity and accuracy in reproducing the training data. This is an optimization problem in the space of functions spanned by the possible activations and weights of the network. We consider four possible activation functions in this example: linear, rectified linear unit (relu), hyperbolic tangent (tanh), and exponential linear unit (elu). The weights connecting the artificial neurons can be either *fixed* or *trainable*, with the fixed category allowing the following values: 0,  $\frac{1}{2}$ , 1, 2,  $\frac{\Delta t}{2}$ ,  $\Delta t$ , and  $2 \Delta t$ , with  $\Delta t$  the timestep separating the inputs and outputs. This is motivated by the fact that physical laws often involve integer or simple fractional coefficients and that the timestep represents important information. Our network has twenty weights (each with eight possible settings) and six activation functions to optimize, see Fig. 1 (top panel). A brute force approach to finding a PNN model would require training  $\sim 10^{21}$  neural networks, an impossible computational task even for the relatively small networks here. We thus resort to evolutionary optimization, using a genetic algorithm to discover models that balance accuracy and parsimony.

PNNs use an objective function that includes measures of the accuracy in the test set and parsimony. The latter term favors: i) linear activation functions over non-linear ones, and ii) non-trainable weights with simple values over optimizable weights. The objective function for the genetic optimization is defined as:

$$F = f_1(E_{test}) + p \left( \sum_{i=1}^{N_N} w_i^2 + \sum_{j=1}^{N_w} f_2(w_j) \right) \quad (1)$$

where  $E_{test}$  represents the mean squared error of the trained PNN on the testing set and  $f_1$  is a logarithmic function that converts the wide range errors into a scale comparable to the parsimony terms, see section S4 of the SM. The second term runs over the  $N_N$  neurons in the network and is designed to favor simple activation functions. The *linear*, *relu*, *tanh* and *elu* activation functions are assigned scores of  $w_i = 0, 1, 2$  and  $3$ , respectively. The third term runs over the network weights and biases ( $N_w$ ) and favors fixed, simple weights over trainable ones. A fixed weight value of 0 is assigned a score of 0, while other fixed weights are assigned the score 1, and a trainable weight is assigned a score of 2. The parsimony parameter  $p$  determines the relative importance of parsimony and accuracy. As will be shown with two examples, PNNs of interest will correspond to parameters  $p$  where both accuracy and parsimony affect model selection. We use the DEAP package for the evolutionary optimization<sup>48</sup> and Keras<sup>49</sup> to train individual networks, see [Methods](#). We note that our approach is similar in spirit to recent work combining genetic algorithms with neural networks to discover partial differential equations<sup>34</sup>, but PNNs are more versatile in terms of parsimony, composition of functions, and are applicable to situations where an underlying differential equation may not exist, as we will see in the second example discussed in this paper. We also note that evolutionary optimization is not the only way to achieve parsimony. For example, one could include hidden layers containing a library of possible activation functions and use sparsity to prune unnecessary activations. This has recently been used to discover simple kinematics equations<sup>33</sup>. An advantage of this approach over ours is simplicity and computational expedience since such networks can be trained using backpropagation



**Figure 2.** (a) PNN model 1 RMSEs on the training/validation/test sets compared to the feed forward network (b) We see that energy conservation between PNN1 and the verlet integrator is comparable (TE: total energy) (c) Forward and reverse trajectories generated by PNN1 show good reversibility (d) A visualization of PNN model 1 found by the genetic algorithm, attempting to predict positions and velocities one step ahead.

alone. However, the evolutionary approach used in PNNs offers significant advantages including a more efficient exploration of function space and avoiding local minima, flexibility in the definition of parsimony, and composition of functions via the neural network.

The PNNs resulting from a genetic optimization with  $p = 1$  reproduce the training, validation and testing data more accurately than the architecturally complex FFNNs. Figure 2(a) compares the RMSE for positions and velocities from the optimal PNN (denoted PNN1) to the FFNN. Remarkably, PNN1 also results in excellent long-term energy conservation and time reversibility, evaluated using the same procedure as before. Figures 2(b) and 2(c) compare the total energy and trajectories generated by PNN1, the FFNN, and the state-of-the-art velocity Verlet integrator. We see that PNN1 *learns* both time-reversibility and that total energy is a constant of motion. This is in stark contrast to the physics-agnostic FFNN and even naïve physics-based models like a first order Euler integration. A few of the top ranked PNNs perform similarly to PNN1 and they will be discussed in section S7 of the SM.

Having established that the PNNs *learn* the physics of the system and result in stable and accurate integrators, we now explore their interpretability in the hope of finding out how time-reversibility and energy conservation are achieved. In short: can the PNNs *teach* us what they learned? We find that the PNNs discover simple models, with many weights taking fixed values (including zero) and all activations functions taking the simplest possible alternative (linear functions). As an example, the parameters corresponding to PNN1 are shown in Fig. 2(d), and other PNNs with comparable (but higher) objective functions are shown in section S6 of the SM. This simplicity allows us to trivially obtain position and velocity update equations. The equations of motion learned by PNN1, rewritten in terms of relevant quantities such as timestep and mass, are:

$$x(t + \Delta t) = x(t) + 1.0001 v(t)\Delta t + 0.9997 \frac{1}{2} F \left( x(t) + v(t) \frac{\Delta t}{2} \right) \frac{\Delta t^2}{m} \tag{2}$$

$$v(t + \Delta t) = v(t) + 0.9997 F \left( x(t) + v(t) \frac{\Delta t}{2} \right) \frac{\Delta t}{m} \tag{3}$$

Inspecting Fig. 2(d) and Eqs. (2,3) we find that PNN1 achieves time-reversibility by evaluating the force at the midpoint between inputs and outputs, this central force evaluation is key to many advanced numerical methods. In fact, PNN1 represents the position Verlet algorithm<sup>39</sup> except that the NN training makes an error in the mass of approximately 3 in 10,000. This algorithm is both reversible and symplectic, i.e. it conserves volume in phase space. The small error in mass seems to originate from the small inaccuracies of the force sub-net to describe the Lennard -Jones potential, see section S6 of the SM.

Inspecting Fig. 2(d) and Eqs. (2-3) we find that PNN1 achieves time-reversibility by evaluating the force at the midpoint between inputs and outputs, this central force evaluation is key to many advanced numerical methods. In fact, PNN1 represents the position Verlet algorithm<sup>39</sup> except that the NN training makes a small error in the mass, of approximately 3 in 10,000. This algorithm is both reversible and symplectic, i.e. it conserves volume in phase space. The small error in mass actually seems to originate from the small inaccuracies of the force sub-net to describe the Lennard -Jones potential, see section S6 of the SM.

The genetic optimization provides an ensemble of models and inspecting slightly sub-optimal ones provides interesting insights. The SM provides the equations of motion predicted by PNN2 and 3. These are similarly

interpretable and, quite remarkably, they also learn to evaluate the force at the half-step. They represent a slightly inaccurate version of the position Verlet algorithm with minor energy drifts due to a slight asymmetry in effective mass in the position and velocity update equations. Finally, changing the parsimony parameter  $p$  in the objective function, allows us to generate a family of models with different tradeoffs between accuracy and parsimony; see Figure S8 in the SM. Interestingly, we find models that reproduce the training and testing data more accurately than PNN 1 and Verlet. However, these models are not time reversible and their energy conservation is worse than PNN1, see section S7 in the SM, stressing the importance of parsimony.

Along similar lines, we tested the ability of the PNNs to discover the physics governing a damped dynamical system, see [Methods](#). The equations learned by the top PNN, with  $\gamma$  the damping constant, are:

$$x(t + \Delta t) = x(t) + 1.0008 v(t) * \left( \Delta t - \frac{\gamma \Delta t^2}{2m} \right) + 0.9991 \frac{1}{2} F \left( x(t) + v(t) \frac{\Delta t}{2} \right) \frac{\Delta t^2}{m} \quad (4)$$

$$v(t + \Delta t) = \left( 1 - \frac{\gamma \Delta t}{m} \right) v(t) + 1.0002 F \left( x(t) + v(t) \frac{\Delta t}{2} \right) \frac{\Delta t}{m} \quad (5)$$

In this second example, PNNs learn classical mechanics, the fact that the frictional force is proportional to negative the velocity, and discover the same stable integrators based on the position Verlet method, all from the observational data.

We consider the emergence of Verlet style integrators from data remarkable. This family of integrators is the preferred choice for molecular dynamics simulations due to its stability. Unlike other algorithms such as the Runge–Kutta family or the first order Euler method, Verlet integrators are symplectic and time reversible<sup>50</sup>. This class of integrators has been long known, and proposed independently by several researchers over decades (see Ref.<sup>50</sup> for a review), but a detailed understanding of their properties and their justification from Trotter's theorem are relatively modern<sup>39</sup>. Importantly, we find more complex models that reproduce the data more accurately than PNN1 but do not exhibit time reversibility nor conserve energy. This shows that parsimony is critical to learn models that can provide insight into the physical system at hand and for generalizability. We stress that the equations of motion and an advanced integrator were obtained from observational data of the motion of a particle and the force–displacement relationship alone. We believe that, at the expense of computational cost, the force sub-net could be learned together with the integrators (effectively learning the acceleration) from large-enough dynamical datasets. This assertion is based on the observation that on fixing some of the network parameters that result in a Verlet integrator, the remaining parameters and the force subnet can be learned from the observational data used above, see section S7 of the SM.

**Melting temperature laws.** To demonstrate the versatility and generalizability of PNNs, we now apply them to discover melting laws from experimental data. Our goal is to predict the melting temperature of materials from fundamental atomic and crystal properties. To do this, we collected experimental melting temperatures for 218 materials (including oxides, metals, and other single elements crystals) as well as fundamental physical quantities including: bulk modulus  $K$ , shear modulus  $G$ , density  $\rho$ , a characteristic atomic distance  $a$  (the cube root of the volume per atom), and mean atomic mass  $m$ .

Before feeding this data to PNNs, we perform a standard dimensionality analysis to use dimensionless inputs and output. For convenience we first define an effective sound speed,  $v_m$ , from density and elasticity moduli, see section S8 of the SM. From these fundamental quantities, we define four independent quantities with the dimensions of temperature:

$$\theta_0 = \frac{\hbar v_m}{k_b a} \quad (6)$$

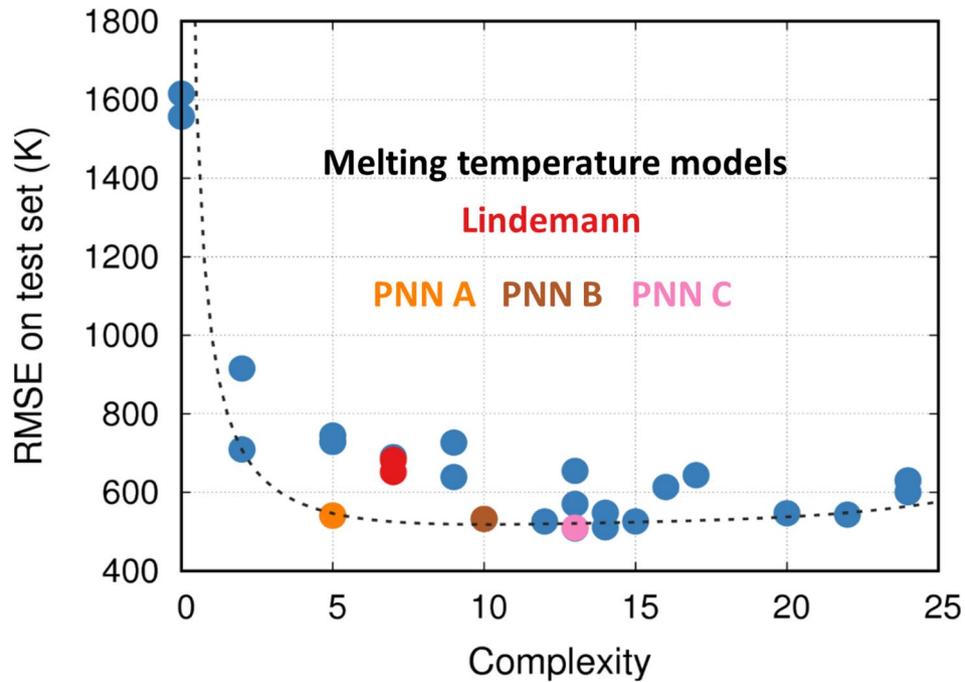
$$\theta_1 = \frac{\hbar^2}{m a^2 k_b} \quad (7)$$

$$\theta_2 = \frac{a^3 G}{k_b} \quad (8)$$

$$\theta_3 = \frac{a^3 K}{k_b} \quad (9)$$

where  $\hbar$  is Planck's constant and  $k_b$  is the Boltzmann's constant. All variables have physical meanings, for example,  $\theta_0$  is proportional to the Debye temperature. The inputs to the PNNs are the last three quantities normalized by  $\theta_0$  and the output melting temperature is also normalized by  $\theta_0$ . Additional details on the preprocessing steps as well as network architecture, including custom activations, can be found in the section S8 of the SM.

Armed with dimensionless inputs and outputs, we use PNNs to discover melting laws. Varying the parsimony parameter in the objective function, Eq. 1, results in a family of melting laws. These models are presented in Fig. 3 in terms of their accuracy with respect to the testing set and their complexity. The latter is defined as the sum of the second and third terms of the objective function Eq. 1, i.e., the sum of the activation function and weight



**Figure 3.** Melting laws discovered by PNNs. The red points show the celebrated Lindemann law, while the blue points show other models discovered. The black dotted line denotes the Pareto front of models, with some of the models performing better than the Lindemann law while also being simpler. Three models are highlighted and labeled.

terms. PNN models represent various tradeoffs between accuracy and parsimony from which we can define a Pareto front of optimal models (see dashed line).

The PNN approach finds several simple yet accurate expressions. The simplest non-trivial relationship is given by PNN A, it approximates the melting temperature as proportional to the Debye temperature:

$$T_m^{PNN A} = 21.8671 \theta_0 \quad (10)$$

This makes physical sense as the Debye temperature is related to the characteristic atomic vibrational frequencies and stiffer and stronger bonds tend to lead to higher Debye and melting temperatures. Next in complexity, PNN B adds a correction proportional to the shear modulus:

$$T_m^{PNN B} = 17.553 \theta_0 + 0.001985 \theta_2 \quad (11)$$

This is also physically sensible as shear stiffness is closely related to melting. This fact is captured by the classic Born instability criteria<sup>51</sup> that associates melting to loss of shear stiffness. Just above PNN B in complexity, PNN finds the celebrated Lindemann melting law<sup>35</sup>:

$$T_m^{lind} = C \frac{\theta_0^2}{\theta_1} = \frac{k_b}{9\hbar^2} f^2 a^2 m T_D^2 \quad (12)$$

Written in its classical form in the third term; here  $T_D$  is the Debye temperature of the material and  $f$  and  $C$  are empirical constants. Remarkably, this law, derived using physical intuition in 1910, is very close to, but not on, the optimal Pareto front in accuracy-complexity space. For completeness, we describe the model with the lowest RMS error, PNN C predicts the melting temperature as:

$$T_m^{PNN C} = 11.9034 \theta_0 + 0.000499 \theta_3 + 0.00796 \frac{\theta_0^2}{\theta_1} \quad (13)$$

Quite interestingly, this model combines the Lindemann expression with Debye temperature and bulk (not shear) modulus. This combination is not surprising given the expressions above, but the selection of bulk over shear modulus is not clear at this point and should be explored further.

In summary, we proposed parsimonious neural networks that are capable of learning interpretable physics models from data; importantly, they can extract underlying symmetries in the problem at hand and provide physical insight. This is achieved by balancing accuracy with parsimony, an adjustable parameter is used to control the relative importance of these two terms and generate a family of models that are Pareto optimal. We quantify parsimony by ranking individual activation functions and favoring fixed weights over adjustable ones.

Future work should explore other measures of complexity such as the complexity of a polynomial expansion of the resulting PNN expression<sup>52</sup>, or the curvature of the PNN expression evaluated over the training data<sup>53</sup>. The combination of genetic optimization with neural networks enables PNNs to explore a large function space and obviate the need for estimating numerical derivatives or matching a library of candidate functions, as was done in prior efforts<sup>17–19</sup>. Additionally, PNNs perform complex composition of functions in contrast to sparse regression, which combine functions linearly. The libraries of activation functions in our first examples of PNNs are relatively small and based on physics intuition, the application of PNNs in areas where less is known about the problem at hand requires more extensive sets of activation functions, at increased computational cost. As opposed to most efforts attempting to discover differential equations, such as DLGA-PDE, PNN can discover laws in situations where an underlying differential equation may not exist. The state-of-art solutions PNNs provide two quite different problems attest to the power and versatility of the approach.

From data describing the classical evolution of a particle in an external potential, the PNN produces integration schemes that are accurate, conserve energy and satisfy time reversibility. Furthermore, they can be easily interpretable as discrete versions of Newton's equations of motion. Quite interestingly, the PNNs learn the non-trivial need to evaluate the force at the half step for time reversibility. The optimization could have learned the first order Runge–Kutta algorithm, which is not reversible, but it favored central-difference based integrators. Furthermore, parsimony favors Verlet-type integrators over more complex expressions that describe the training data more accurately but do not exhibit good stability. We note that other high-order integrators are not compatible with our initial network, but these can easily be incorporated by starting with a more complex network. As discussed above, the resulting algorithms would not come as a surprise to experts in molecular dynamics simulations as this community has developed, over decades, accurate algorithms to integrate Newton's equations of motion. The fact that such knowledge and algorithms can be extracted automatically from observational data has, however, deep implications in other problems and fields. This is confirmed with a second example that shows the ability of PNNs to extract interpretable melting laws from experimental data. We discover a family of expressions with varying tradeoffs between accuracy vs. parsimony and our results show that the widely used Lindeman law, proposed over a century ago, is remarkably close to the pareto front but we find PNNs that outperform it. The PNN models highlight the relationships between melting and a materials Debye's temperature as well as shear moduli, providing insight into the processes that determine melting.

## Methods

**Training data.** To discover integration schemes, training data was generated using molecular dynamics simulations under an NVE ensemble, using the velocity Verlet scheme with a short timestep (about a tenth of what is required for accurate integration), see section S1 of the SM for additional details. Training and validation data are obtained from trajectories with four different total energies, 20% of which is used as a validation set. Our test set is a separate trajectory with a different total energy. For the damped dynamics cases, a frictional force proportional to negative the velocity is added, with frictional coefficient  $\gamma = 0.004 \text{ eVps}/\text{\AA}^2$ . To discover novel melting laws, we queried the Pymatgen and Wolfram alpha databases for experimental melting temperatures. We obtained fundamental material properties such as volume and shear modulus by querying the Materials Project. Additional details are provided in section S8 of the SM.

**Evolutionary optimization.** We used populations of 200 and 500 individuals and a two-point crossover scheme and random mutations to evolve the population (weights and activation functions)<sup>54</sup>. For each generation, individual networks in the population are trained using backpropagation using the same protocols as for the feed-forward networks; only adjustable weights are optimized in this operation. The populations were evolved over 50 generations, additional details of the genetic algorithm are included in section S5 of the SM.

Received: 15 January 2021; Accepted: 31 May 2021

Published online: 17 June 2021

## References

1. K. Simonyan and A. Zisserman, *Very Deep Convolutional Networks for Large-Scale Image Recognition*, ArXiv Preprint <https://arxiv.org/abs/1409.1556> (2014).
2. A. Krizhevsky, I. Sutskever, and G. E. Hinton, *ImageNet Classification with deep convolutional neural networks*, in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., 2012), pp. 1097–1105.
3. Bengio, Y., Ducharme, R., Vincent, P. & Jauvin, C. A neural probabilistic language model. *J. Mach. Learn. Res.* **3**, 1137 (2003).
4. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431 (2017).
5. Senior, A. W. *et al.* Improved protein structure prediction using potentials from deep learning. *Nature* **577**, 7792 (2020).
6. Meredig, B. *et al.* Combinatorial screening for new materials in unconstrained composition space with machine learning. *Phys. Rev. B* **89**, 094104 (2014).
7. Carrete, J., Li, W., Mingo, N., Wang, S. & Curtarolo, S. Finding unprecedentedly low-thermal-conductivity half-heusler semiconductors via high-throughput materials modeling. *Phys. Rev. X* **4**, 011019 (2014).
8. Bassman, L. *et al.* Active learning for accelerated design of layered materials. *NPJ Comput. Mater.* **4**, 1 (2018).
9. Kaufmann, K. *et al.* Discovery of high-entropy ceramics via machine learning. *NPJ Comput. Mater.* **6**, 1 (2020).
10. Snyder, J. C., Rupp, M., Hansen, K., Müller, K.-R. & Burke, K. Finding density functionals with machine learning. *Phys. Rev. Lett.* **108**, 253002 (2012).
11. Li, Z., Kermodé, J. R. & De Vita, A. Molecular dynamics with on-the-fly machine learning of quantum-mechanical forces. *Phys. Rev. Lett.* **114**, 096405 (2015).

12. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
13. Jacobsen, T. L., Jørgensen, M. S. & Hammer, B. On-the-fly machine learning of atomic potential in density functional theory structure optimization. *Phys. Rev. Lett.* **120**, 026102 (2018).
14. M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics informed deep learning (part i): data-driven solutions of nonlinear partial differential equations, <https://arxiv.org/abs/1711.10561> [Cs, Math, Stat] (2017).
15. M. Raissi, P. Perdikaris, and G. E. Karniadakis, Physics informed deep learning (part ii): data-driven discovery of nonlinear partial differential equations, <https://arxiv.org/abs/1711.10566> [Cs, Math, Stat] (2017).
16. Ling, J., Kurzawski, A. & Templeton, J. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *J. Fluid Mech.* **807**, 155 (2016).
17. Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **324**, 81 (2009).
18. Rudy, S. H., Brunton, S. L., Proctor, J. L. & Kutz, J. N. Data-driven discovery of partial differential equations. *Sci. Adv.* **3**, 1602614 (2017).
19. Champion, K., Lusch, B., Kutz, J. N. & Brunton, S. L. Data-driven discovery of coordinates and governing equations. *Proc. Natl. Acad. Sci.* **116**, 22445 (2019).
20. S. A. Harp, T. Samad, and A. Guha, Designing application-specific neural networks using the genetic algorithm, in *Advances in Neural Information Processing Systems* (1990), pp. 447–454.
21. Miller, G. F., Todd, P. M. & Hegde, S. U. Designing neural networks using genetic algorithms. *ICGA* **89**, 379–384 (1989).
22. Stepniewski, S. W. & Keane, A. J. Pruning backpropagation neural networks using modern stochastic optimisation techniques. *Neural Comput. Appl.* **5**, 76 (1997).
23. Yao, X. & Liu, Y. A new evolutionary system for evolving artificial neural networks. *IEEE Trans. Neural Netw.* **8**, 694 (1997).
24. Montana, D. J. & Davis, L. Training feedforward neural networks using genetic algorithms. *IJCAI* **89**, 762–767 (1989).
25. F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning, ArXiv Preprint <https://arxiv.org/abs/1712.06567> (2017).
26. A. Costa, R. Dangovski, S. Kim, P. Goyal, M. Soljačić, and J. Jacobson, Interpretable neuroevolutionary models for learning non-differentiable functions and programs, ArXiv Preprint <https://arxiv.org/abs/2007.10784> (2020).
27. S. Greydanus, M. Dzamba, and J. Yosinski, *Hamiltonian neural networks*, <https://arxiv.org/abs/1906.01563> [Cs] (2019).
28. Z. Chen, J. Zhang, M. Arjovsky, and L. Bottou, Symplectic recurrent neural networks, <https://arxiv.org/abs/1909.13334> [Cs, Stat] (2019).
29. M. J. Eslamibidgoli, M. Mokhtari, and M. H. Eikerling, Recurrent Neural network-based model for accelerated trajectory analysis in AIMD simulations, <https://arxiv.org/abs/1909.10124> [Physics] (2019).
30. Udrescu, S.-M., Tegmark, M. & Feynman, A. I. A physics-inspired method for symbolic regression. *Sci. Adv.* **6**, 2631 (2020).
31. Iten, R., Metger, T., Wilming, H., Del Rio, L. & Renner, R. Discovering physical concepts with neural networks. *Phys. Rev. Lett.* **124**, 010508 (2020).
32. G. Martius and C. H. Lampert, *Extrapolation and learning equations*, ArXiv Preprint <https://arxiv.org/abs/1610.02995> (2016).
33. S. Kim, P. Lu, S. Mukherjee, M. Gilbert, L. Jing, V. Ceperic, and M. Soljacic, Integration of neural network-based symbolic regression in deep learning for scientific discovery, ArXiv Preprint <https://arxiv.org/abs/1912.04825> (2019).
34. H. Xu, H. Chang, and D. Zhang, DLGA-PDE: discovery of PDEs with incomplete candidate library via combination of deep learning and genetic algorithm, *J. Comput. Phys.* 109584 (2020).
35. Lindemann, F. A. Über Die Berechnung Molekularer Eigenfrequenzen. *Z. Physik* **11**, 609 (1910).
36. S. Desai and A. Strachan, *Discovering Discretized Classical Equations of Motion Using Parsimonious Neural Networks*, (2020).
37. Z. Long, Y. Lu, and B. Dong, *PDE-Net 2.0: Learning PDEs from Data with a Numeric-Symbolic Hybrid Deep Network*, *Journal of Computational Physics* **399**, 108925 (2019).
38. L. Verlet, *Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules*, *Phys. Rev.* **159**, 98 (1967).
39. Tuckerman, M., Berne, B. J. & Martyna, G. J. Reversible multiple time scale molecular dynamics. *J. Chem. Phys.* **97**, 1990 (1992).
40. Yoshida, H. Construction of higher order symplectic integrators. *Phys. Lett. A* **150**, 262 (1990).
41. Rowlands, G. A numerical algorithm for hamiltonian systems. *J. Comput. Phys.* **97**, 235 (1991).
42. Izaguirre, J. A., Reich, S. & Skeel, R. D. Longer time steps for molecular dynamics. *J. Chem. Phys.* **110**, 9853 (1999).
43. R. De Vogelaere, *Methods of Integration Which Preserve the Contact Transformation Property of the Hamilton Equations*, Technical Report (University of Notre Dame. Dept. of Mathematics) (1956).
44. Trotter, H. F. On the product of semi-groups of operators. *Proc. Am. Math. Soc.* **10**, 545 (1959).
45. Eshet, H., Khaliullin, R. Z., Kühne, T. D., Behler, J. & Parrinello, M. Ab initio quality neural-network potential for sodium. *Phys. Rev. B* **81**, 184107 (2010).
46. Behler, J., Martoňák, R., Donadio, D. & Parrinello, M. Metadynamics simulations of the high-pressure phases of silicon employing a high-dimensional neural network potential. *Phys. Rev. Lett.* **100**, 185501 (2008).
47. Chmiela, S., Sauceda, H. E., Müller, K.-R. & Tkatchenko, A. Towards exact molecular dynamics simulations with machine-learned force fields. *Nat. Commun.* **9**, 1 (2018).
48. Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M. & Gagné, C. DEAP: evolutionary algorithms made easy. *J. Mach. Learn. Res.* **13**, 2171 (2012).
49. F. Chollet, *Keras* (2015).
50. Hairer, E., Lubich, C. & Wanner, G. Geometric numerical integration illustrated by the Stormer-Verlet method. *Acta Numer.* **12**, 399 (2003).
51. Born, M. Thermodynamics of crystals and melting. *J. Chem. Phys.* **7**, 591 (1939).
52. Stinstra, E., Rennen, G. & Teeuwen, G. Metamodeling by symbolic regression and pareto simulated annealing. *Struct. Multidiscip. Optim.* **35**, 315 (2008).
53. L. Vanneschi, M. Castelli, and S. Silva, *Measuring bloat, overfitting and functional complexity in genetic programming*, in *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation* (2010), pp. 877–884.
54. T. Bäck, D. B. Fogel, and Z. Michalewicz, *Evolutionary Computation 1: Basic Algorithms and Operators* (CRC press, 2018).

## Acknowledgements

Partial support from the Network for Computational Nanotechnology, Grant EEC-1227110, is acknowledged.

## Author contributions

S.D. Investigation, visualization, writing – original draft A.S. writing – reviewing and editing, supervision, funding acquisition.

## Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-021-92278-w>.

**Correspondence** and requests for materials should be addressed to A.S.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021