Deep Learning Based Integrators for Solving Newton's Equations with Large Timesteps

JCS Kadupitiya,[1, *] Geoffrey C. Fox,[1, †] and Vikram Jadhao[1, ‡]

[1]*Intelligent Systems Engineering, Indiana University, Bloomington IN 47408, USA*

Classical molecular dynamics simulations are based on Newton's equations of motion and rely on numerical integrators to solve them. Using a small timestep to avoid discretization errors, Verlet integrators generate a trajectory of particle positions as solutions to Newton's equations. We introduce an integrator based on deep neural networks that is trained on trajectories generated using the Verlet integrator and learns to propagate the dynamics of particles with timestep up to $4000\times$ larger compared to the Verlet timestep. We demonstrate significant net speedup of up to 32000 for 1 - 16 particle 3D systems and over a variety of force fields.

The time evolution of a diverse array of systems is governed by Newton's equations of motion [1]. These equations are the basis of powerful computational methods such as classical molecular dynamics (MD) that are used to understand the microscopic origins of a wide range of material and biological phenomena [2–5]. In the MD method, Newton's equations are integrated for a system of many particles using numerical integrators such as Verlet to produce trajectories of particles. The time evolution is generally performed one small timestep at a time for sufficiently long times in order to accurately sample enough representative configurations and extract useful structural and dynamical information. Consider the ordinary Verlet integrator

$$\vec{x}(t + \Delta) = 2\vec{x}(t) - \vec{x}(t - \Delta) + \Delta^2 \frac{\vec{f}}{m} + O(\Delta^4) \quad (1)$$

that updates the current position $\vec{x}(t)$ of a particle of mass m at time t to position $\vec{x}(t + \Delta)$ after timestep Δ using the previous position $\vec{x}(t - \Delta)$ and the force $\vec{f}$ at time t. With $\vec{x}(t + \Delta)$ as the current position, and $x(t)$ as the previous, the propagation is continued further. While the term $O(\Delta^4)$ denotes the incurred local error, the global (accumulated) error is $O(\Delta^2)$ and Equation 1 expresses a 2nd order integrator [2, 6]. The discretization errors are minimized by choosing a sufficiently small Δ which often makes the simulations computationally expensive.

The ordinary Verlet integrator requires a sequence of 2 position coordinates $(\vec{x}_t, \vec{x}_{t-\Delta})$ in order to update the position using other quantities such as $\vec{f}$ and m. However, it is possible to infer these other quantities using the information encoded in a longer stream of position data such that the time evolution can be done with only the history of positions as input. We illustrate this with a 1D example of a particle experiencing simple harmonic motion governed by the force $f = -kx$. It can be shown that the particle position can be accurately evolved to $t + \Delta$ via the function $\mathcal{V}(x_t, x_{t-\Delta}, x_{t-2\Delta}) = x_{t-\Delta}^{-1} \left(x_t^2 - x_{t-\Delta}^2 + x_t x_{t-2\Delta} \right)$ that uses a sequence of 3 positions. $\mathcal{V}$ also incurs a global error of $O(\Delta^2)$.

This idea generalizes for more accurate integrators and systems of many particles such that the time evolution can be performed via the operator $\mathcal{V}$: $\vec{x}(t + \Delta) = \mathcal{V}(\vec{x}_t, \vec{x}_{t-\Delta}, \ldots \vec{x}_{t-n_V \Delta})$ that takes a sequence of n_V positions and evolves the system forward by Δ. For example, one can derive a 4th order Adams-Bashforth-Moulton operator that uses a sequence of 4 positions to evolve the dynamics with a global error of $O(\Delta^4)$. The longer history of input positions enables such higher-order integrators to accurately perform time evolution with larger Δ. However, this advantage generally comes at the expense of higher computing costs per timestep that offset the speedup gained via the use of a larger Δ.

Here, we adopt a data-driven approach and show that accurate integrators based on a sequence of past configurations can be derived using deep learning. The deep neural network based operators are trained using the ground truth results obtained with the 2nd order Verlet integrator. They possess a complex mathematical structure described with $\sim 100,000$ parameters compared to the relatively simple functional forms of the traditional higher-order integrators. We demonstrate that the network complexity enables the integrators to perform the time evolution of systems of many particles for a wide range of force fields using a timestep up to $4000\times$ the baseline Δ associated with the 2nd order Verlet integrator. The relatively low time for inferring the positions as predictions of the deep learning model ensures that the overhead costs are low and we demonstrate that speedups up to 32000 can be obtained using larger timesteps.

RELATED WORK

Recent years have seen a dramatic rise in the use of machine learning (ML) including deep learning (DL) to enhance the performance of simulations of materials [7–22]. In MD simulations, ML has been used to accelerate the sampling of systems with rugged free-energy landscapes [14], generate new configuration updates in reduced-dimensional space [15], auto-tune simulation timestep [16], classify particle assembly landscapes [17], and derive "surrogates" of MD simulations [18–22]. However, the use of DL for advancements in the

FIG. 1. RNN-based $\mathscr{R}$ integrator propagates the single timestep (Δ_R) evolution of the N particle system characterized by input features of size d. Input system and LSTM parameters are shown for $N = 16$ in 3D. Symbols are defined in the main text.

design of integrators driving the MD simulation has been largely unexplored. Related efforts have focused on using DL to replicate the outputs of numerical integrators using baseline timestep values [23–29]. These integrators have been evaluated on relatively simple 1D systems governed by ordinary differential equations, with some studies probing the extension to learn partial differential equations [23, 25, 30, 31]. More recently, data-driven approaches aimed at time evolution with higher than baseline timesteps have been pursued [24, 32]. For example, Shen et. al [32] used an artificial neural network as a corrector to reduce the local truncation error in the Euler integrator, and demonstrated the success on simple 1D / 2-body problems by performing dynamics with timestep up to $200\times$ the baseline Euler approach. For the 1D problems, a $100\times$ increase in the timestep led to over 3 orders of magnitude increase in the error. Other related work focuses on using DL to learn the physics of simple systems such as 1D spring-mass problems and respect conservation laws in an unsupervised manner [29, 33–37]. Here, instead of performing explicit time evolution, neural networks observe the positions and momenta and produce the Hamiltonian as output [33].

We build on this research with the goal to produce integrators based on DL that enable simulations of 3D systems of particles interacting with commonly employed potentials describing soft-matter systems (e.g., Lennard-Jones potential). We show that recurrent neural networks (RNNs), using the high-dimensional particle trajectory data, can learn the interaction potentials and the New-

ton's equations of motion driving the particle dynamics. We demonstrate that the RNN-based integrator performs accurate time evolution of 1D and 3D systems over a variety of force fields using large timesteps, up to $4000\times$ the baseline, limiting the rise in error with increasing timestep to within an order of magnitude.

DATA-DRIVEN DESIGN OF INTEGRATORS FOR SOLVING NEWTON'S EQUATIONS OF MOTION

The use of deep learning (DL) in sequence processing and time series prediction problems has been well studied by the industry for different applications including voice recognition and translation [38], pattern recognition in stock market data [39], and ride-hailing [40]. Recurrent neural networks (RNN) are established DL tools in these applications. RNNs process input sequence data and maintain a vector $\vec{h}_t$ known as the "hidden state" for each recurrent cell to model the temporal behavior of sequences through directed cyclic connections between its cells. $\vec{h}_t$ is updated by applying a function f to the previous hidden state ($\vec{h}_{t-1}$) and the current input ($\vec{x}_t$). The cells are arranged in a fashion where they fire when the right sequence is fed. A common choice for f is the Long Short Term Memory networks (LSTM). An LSTM unit contains a cell which is the memory of the unit, and three gates (input gate, output gate, and forget gate) which regulate the flow of information into and out of the cells. This gated architecture allows the LSTM to remember longer dependencies of the sequences fed into the network and deal with the exploding and vanishing gradient problems encountered while training RNNs (SI).

Here, we show that RNNs can be used to design integrators to solve Newton's equations of motion and drive MD simulations. For the ease of exposition, we describe the DL-based integrator evolving positions of particles (akin to the ordinary Verlet integrator). The process can be adapted to describe an integrator that evolves both positions and velocities (SI). Each component of the position vector of a particle is identified as a feature, with the size of the input and output features being the same. The total feature size for N particles in $\mathcal{D}$ physical dimensions is $d = N \times \mathcal{D}$. We introduce an operator $\mathscr{R}$ based on RNN derived using LSTMs (Figure 1) that employs a sequence $\{x\} = \vec{x}_t, \vec{x}_{t-\Delta_R}, \ldots, \vec{x}_{t-S_R\Delta_R}$ of length S_R of current and past positions up to time t to predict the future position at time $t+\Delta_R$: $\vec{x}(t + \Delta_R) = \mathscr{R}[\{x\}]$. $\mathscr{R}$ can be written as $\mathscr{R}[\{x\}] = \mathscr{D}[\mathscr{L}_2[\mathscr{L}_1[\{x\}]]]$, where $\mathscr{D}, \mathscr{L}_1, \mathscr{L}_2$ are operators associated with the dense layer, first LSTM layer, and second LSTM layer of the RNN respectively. The layers are stacked up on each other such that the output of one (e.g., $\mathscr{L}_1$) becomes the input for another ($\mathscr{L}_2$) as shown in Figure 1. Each LSTM layer

FIG. 2. Overview of the DL approach based on recurrent neural networks (RNN) for integrating Newton's equations of motion in order to drive molecular dynamics (MD) simulations. The LSTM model in the RNN integrator block is described in Figure 1.

consists of n number of LSTM units and contains a set of parameters in the form of weights, biases, and activation functions. For example, $\mathscr{L}_1$ has n_1 LSTM units and is characterized with weights W, U, and bias b. It takes input feature vector $\{x\}$ and outputs hidden state vectors $\{h\}$. These hidden states are fed as input to the $\mathscr{L}_2$ layer characterized with its own set of weights and biases. A similar connection is made between $\mathscr{L}_2$ and the dense layer $\mathscr{D}$ which consists of n_D units in the form of weights and biases. Post training, these layers acquire well-determined values for all the parameters and the integrator $\mathscr{R}$ emerges as:

$$\mathscr{R}[\{x\}] = \mathscr{D}[\mathscr{L}_2[\mathscr{L}_1[\{x\}, \{P_1\}], \{P_2\}], \{P_D\}], \quad (2)$$

where $\{P_1\}, \{P_2\}, \{P_D\}$ are trained parameters associated with LSTM layer 1, LSTM layer 2, and the dense layer respectively. While training, $\{x\}$ is fed as a vector of size $B \times S_R \times d$, where B is a training parameter denoting the batch size. When used as an integrator (testing phase), $B = 1$. Characterized with up to $100,000$ parameters, $\mathscr{R}$ can be considered as a reformulation of the ordinary Verlet integrator $\mathscr{V}$ (Equation 1). This computational complexity accounts for its ability to handle larger timesteps in performing MD simulations.

Figure 2 describes the DL approach using $\mathscr{R}$ to evolve the dynamics of an N particle system with timestep Δ_{R}. $\mathscr{R}$ is trained using particle trajectories generated via $\mathscr{V}$ with small $\Delta = 0.001$; these Verlet results provide the ground truth. Input system attributes and Δ are fed to $\mathscr{V}$ to simulate the dynamics and obtain $S_{\mathrm{V}} = \Delta_{\mathrm{R}}(S_R - 1)/\Delta$ number of configurations. Out of S_{V}, S_R number of configurations (frames) separated by Δ_{R} are distilled to feed the $\mathscr{R}$ operator (S_{R} is the sequence length). $\mathscr{R}$ predicts the time evolution of the system after Δ_{R}. Until the end of the simulation, the input sequence to $\mathscr{R}$ is left shifted to discard the oldest time frame, and the latest frame predicted by $\mathscr{R}$ is appended to the right of the sequence. The adjusted input sequence is fed to $\mathscr{R}$ to evolve the system Δ_{R} further in time. The types of training and testing datasets generated and the LSTM architecture details are provided in the next section.

METHODS

The training and testing datasets are generated and processed for $N \in (1, 16)$ by changing the parameters characterizing the potential energies and the physical attributes describing the particles. Four different potentials are considered: 1) simple harmonic oscillator (SHO), 2) double well (DW), 3) Lennard-Jones (LJ), and 4) rugged (SI). In these initial studies, we select datasets generated by sweeping parameters that mainly shift the initial configuration (e.g., x_0 in LJ) or scale the particle attributes (e.g., mass m in LJ), and in some cases change the shape of the potential energy (e.g., k in SHO). Details on the data preparation and preprocessing are in the SI.

$\mathscr{R}$ integrator with LSTM layer 1, LSTM layer 2, and final dense layer (Figure 1) is implemented in TensorFlow for regression of the particle trajectories using n_1, n_2, and n_D number of hidden units respectively. $\mathscr{R}$ takes a $B \times S_{\mathrm{R}} \times d$ dimensional vector as input where B and d refer to batch and feature size respectively. All the parameters (weights, biases) describing the layers are trained with an error backpropagation algorithm, implemented via a stochastic gradient descent. Adam optimizer is used to optimize the error backpropagation. Outputs of the LSTM layers are wrapped with the $\tanh$ activation function and no activation functions are used in the dense layer. The L2 error (mean square loss) between target and predicted trajectories is used for error calculation. LSTM implementation, training, and testing are programmed using scikit-learn, Keras, and TensorFlow [41–43]. Scikit-learn is used for grid search and scaling, Keras is used to save and load models, and TensorFlow is used to create and train $\mathscr{R}$. Prototype implementation written in Python/C++ is available on GitHub [44].

n_1, n_2, and n_D are chosen depending on the problem complexity and data dimensions. We discuss these choices and other details of the feature extraction and regression process for the most complex case of 16 particles interacting with LJ potential in 3D with PBC. For this system the feature size $d = 96$. By performing a grid search, hyper-parameters such as the number of units for each of the two LSTM layers (n_1, n_2), number of units in the final dense layer (n_D), batch size (B), and
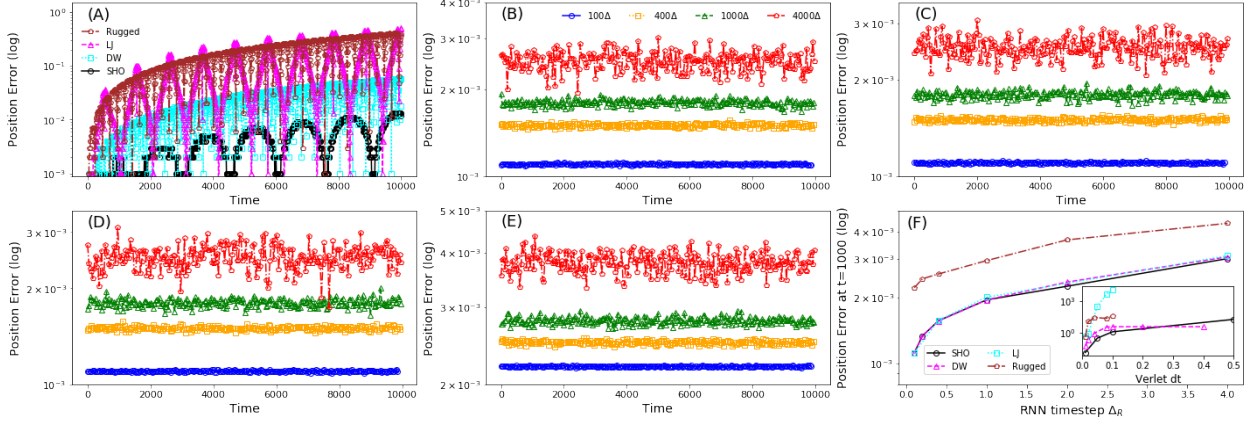
FIG. 3. Errors (log scale) in position updates for 1D simulations. Results are shown for 4 potentials: SHO with $m = 10$, $k = 1$, DW with $m = 1$, $x_0 = -2$, LJ with $m = 1$, $x_0 = 2$, and rugged with $m = 1$, $x_0 = -6$. Trajectories obtained using $\mathscr{V}$ with $\Delta = 0.001$ are taken as the ground truth. (A) Errors using $\mathscr{V}$ with timestep $dt = 10\Delta$. (B), (C), (D), and (E) show errors using $\mathscr{R}$ for SHO, DW, LJ, and rugged potentials respectively for $\Delta_R = 100\Delta$ (circles), 400Δ (squares), 1000Δ (triangles), and 4000Δ (pentagons). (F) Error incurred in position updates using $\mathscr{R}$ at $t = 1000$ vs. timestep Δ_R for the 4 potentials; inset shows corresponding error using $\mathscr{V}$ vs. Verlet timestep dt.

the number of epochs are optimized to 32, 32, 96, 256, and 2500 respectively. Here we have not done a thorough hyper-parameter optimization but have identified reasonable values for key choices – the architecture and the size of the network, and the length of time series. The learning rate of Adam optimizer is set to 0.0005 and the dropout layer is set to 0.15 to prevent overfitting. Both learning and dropout rates were selected using a trial-and-error process. The weights in the hidden layers and in the output layer are initialized for better convergence using a Xavier normal distribution at the beginning [45].

EXPERIMENTS AND RESULTS

We now compare the results obtained from simulations using $\mathscr{R}$ against the baseline velocity Verlet integrator $\mathscr{V}$ (unless otherwise noted, $\mathscr{V}$ will denote the velocity Verlet integrator). For 1D systems, results are shown for four different potentials: 1) simple harmonic oscillator (SHO), 2) Lennard-Jones (LJ), 3) double well (DW), and 4) rugged (see SI Text for details). For the 3D many particle system, experiments are performed on systems of particles interacting with LJ potential. In each case we adopt units such that the values of the system parameters and predicted quantities are around 1. For example, in the case of 3D many particle system, we use reduced LJ units common in MD simulations.

1D Simulations Our first experiments are on testing $\mathscr{R}$ to predict the time evolution of single particle systems in 1D. Training datasets of trajectories simulated using $\mathscr{V}$ with $\Delta = 0.001$ up to $t = 200$ are used. We

find that unlike the traditional Verlet integrator with large timesteps, the errors in positions (trajectory errors) predicted by $\mathscr{R}$ do not increase over time for t up to 10000, and are $O(10^{-3})$ for all $\Delta_R \in (100\Delta, 4000\Delta)$ (Figure 3 B, C, D, E). These bounded errors can be attributed to the unique functional form of the $\mathscr{R}$ integrator produced via the connected LSTM layers, $\mathscr{L}_1$ and $\mathscr{L}_2$, and the dense layer $\mathscr{D}$ with each layer parameterized using fixed (time-independent) weights. The outputs of $\mathscr{L}_1$ and $\mathscr{L}_2$ are wrapped with a tanh function and are bounded regardless of the range of values associated with the input positions. The matrix multiplication of the bounded outputs of $\mathscr{L}_2$ and the fixed weights of the linear layer $\mathscr{D}$ with no wrapping function produces the updated positions.

The errors rise with increase in Δ_R and the complexity of the potential (e.g., higher for rugged than LJ) but remain within an order of magnitude (Figure 3F). The average error across different potentials for $\Delta_R = [100\Delta, 400\Delta, 1000\Delta, 4000\Delta]$ are $\approx$ $0.001, 0.002, 0.003, 0.004$ respectively. On the other hand, the trajectory errors incurred using $\mathscr{V}$ with timestep 10Δ show exponential increase with t and with rising timestep for all potentials considered (Figure 3A and Figure 3F). Similarly, velocity v and position x predictions produced by $\mathscr{R}$ exhibit small trajectory errors, and the phase diagram (v vs. x) and the total energy track the ground truth results for t up to 10000 and Δ_R up to 4000Δ (SI Figure 8). $\mathscr{R}$ learns the energy conservation feature associated with the dynamics governed by Newton's equations of motion.

Lyapunov Instability Next set of experiments involve assessing the stability of the solutions (trajectories) pre-
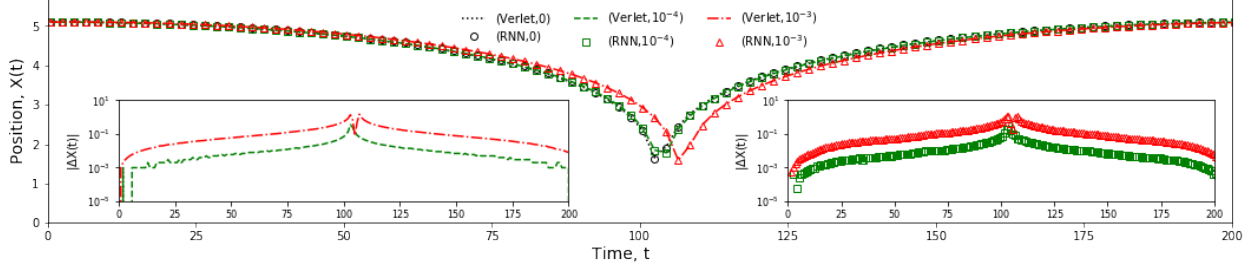
FIG. 4. Lyapunov instability tests in 1D simulations of a particle in an LJ potential with $m = 1, x_0 = 5.1$ using $\mathscr{V}$ with timestep $\Delta = 0.001$ and $\mathscr{R}$ with timestep $\Delta_{\rm R} = 0.1 = 100\Delta$. Legend indicates (integrator, δp) where δp represents the small momentum shift introduced at the start of trajectory. Lines and markers are the results of simulations driven by $\mathscr{V}$ and $\mathscr{R}$ integrators respectively. Outset shows the trajectory for $\delta p = 0, 10^{-4}, 10^{-3}$. Insets show the absolute difference between positions generated with $\delta p \neq 0$ and $\delta p = 0$ ($\Delta x(t) = |x(t, 0) - x(t, \delta p)|$) as produced by $\mathscr{V}$ (left inset) and $\mathscr{R}$ (right inset).

dicted by the integrator $\mathscr{R}$. This is typically discussed in terms of Lyapunov instability which states that close-by trajectories diverge exponentially [46, 47]. Lyapunov instability is present in standard Verlet integrators for many potentials. Trajectories of a particle of $m = 1$ in 1D LJ potential generated using $\mathscr{V}$ are used for these experiments. A small perturbation in the momentum δp is introduced at the start to investigate its effects on long-time evolution of the trajectory. We find that the RNN surrogate inherits the characteristic Lyapunov instability of $\mathscr{V}$ (Figure 4). As δp increased from 10^{-4} to 10^{-3}, the trajectories predicted by $\mathscr{R}$ ($\mathscr{V}$) exhibit similar average divergences from 0.0153 (0.0152) to 0.149 (0.148).


FIG. 5. Energy conservation using $\mathscr{R}$ in a $N = 16$ particle 3D simulation with LJ interactions ($\epsilon = 2$ and $m = 1$) in PBC with initial positions randomly selected and initial velocities set to 0. Black lines are the ground truth results obtained using $\mathscr{V}$ with $\Delta = 0.001$. Outset shows the energy deviation δE (log scale) vs. time t (log scale) predicted by $\mathscr{R}$ for $\Delta_{\rm R} = 100\Delta, 400\Delta, 1000\Delta, 4000\Delta$. Inset shows the same for time evolution using $\mathscr{V}$ with timestep $10\Delta, 40\Delta$, and 100Δ. Note the order of magnitude difference between the outset and inset y-axis values.

Many Particle Simulations in 3D Our final set of experiments probe the extension of the proposed DL approach to systems of $N = 3, 8, 16$ particles interacting with LJ potential. $\mathscr{R}$ is trained using trajectories generated by $\mathscr{V}$ ($\Delta = 0.001$) in a cubic box with periodic boundary conditions (PBC) or in a spherical hard-wall confinement with reflective boundaries. In the former case, $\mathscr{R}$ observed the positions of the particles governed by Newton's equations of motion and the re-mapping dictated by the use of PBC. We find that $\mathscr{R}$ successfully evolved the positions and velocities of particles up to time $t = 10^6$ for all the N-particle systems studied with $\Delta_{\rm R}$ up to 4000Δ. In the interest of brevity, we discuss the results of experiments on the $N = 16$ system in PBC. Results with the spherical hard-wall confinement exhibited similar position errors ($\delta r \sim O(10^{-3})$) (SI Figure 10). The training dataset for this study involved time evolution up to $t = 2000$ and the characteristic LJ potential energy $\epsilon = 1$. Results are shown for time evolution using $\mathscr{R}$ for $\epsilon = 2$ (not present in the entire dataset); slightly better results were obtained for $\epsilon = 1$. The random initial configuration for testing $\mathscr{R}$ was selected outside the training dataset. We find that the trajectory error (computed as an average over 16 particles) associated with the positions predicted by $\mathscr{R}$ for $\Delta_{\rm R} = 100\Delta$ to 4000Δ is $O(10^{-3})$ for all times (SI Figure 9), exhibiting a behavior similar to the 1D results (Figure 3). The total energy $E_t \equiv E(t)$ and the energy deviation $\delta E(t) = (E_t - E_0)/E_0$ of the system show that $\mathscr{R}$ conserves the total energy of the system ($\delta E(t) \lesssim 10^{-3}$) for up to $t = 10^6$ for all $\Delta_{\rm R}$ (from 100Δ to 4000Δ). In strike contrast, the baseline $\mathscr{V}$ integrator suffers from growing accumulated error of $> 10^{-1}$ when simulated with timestep of 40Δ and exhibits a rapid divergence in the energy for $t > 10^5$ (with associated average trajectory error of 10^9 at $t \approx 10^5$; SI Figure 9).

Speedup The DL approach uses Verlet integrator to kickstart the MD simulation and $\mathscr{R}$ to evolve the dynam-

ics forward in time. This detail should be incorporated in calculating the net speedup S_p:

$$S_p = t_V S_T / \left[S_V t_V + (S_T - S_V) t_R \Delta / \Delta_R \right] \quad (3)$$

where S_T is the total number of steps needed if the time evolution is performed only using $\mathcal{V}$, t_V is the time for one forward step propagation using $\mathcal{V}$, and t_R is the time for one forward step propagation using $\mathcal{R}$. We have not accounted for the time spent on creating training datasets ($<$ 24 hours for the experiments shown). S_p is 1 if $S_R = 0$ (no time evolution using $\mathcal{R}$). As most MD simulations are expected to run for long times, in the limit $S_T \gg S_V$, we obtain $S_p \approx t_V \Delta_R / (t_R \Delta)$. For example, for the system of 16 LJ particles with $t_V \approx 0.04392$ seconds, $\Delta = 0.001$, $t_R \approx 0.0026$ seconds, and $\Delta_R = 4.0$, we find $S_p > 10^4$. Table I shows S_p for different experiments at different Δ_R. For time evolution up to $t = 10^6$

TABLE I. Net speedup S_p using the $\mathcal{R}$ integrator

Expt. / Δ_R	100	200	400	1000	2000	4000
SHO	0.5	1.3	3.2	8.6	20.0	45.0
Double-well	0.6	1.2	2.8	8.7	17.3	38.0
LJ	0.9	1.5	3.9	12.8	22.5	42.3
Rugged	0.4	0.8	2.1	4.7	9.7	20.6
LJ, 8	600	1000	1500	5500	8300	12000
LJ, 16	3000	4900	7100	20000	28000	32000

requiring $S_T = 10^9$ steps, we find S_p up to $45\times$ for 1D systems, and from $230\times$ to $32000\times$ for the many particle simulations with LJ forces. $S_p < 1$ data is generally for 1D systems with $\Delta_R = 100\Delta$, 200Δ and when $t_V < t_R$ (e.g., for the 1D cases where the computational cost for time evolution using $\mathcal{V}$ is smaller than that for $\mathcal{R}$).

DISCUSSION

Role of Sequence Length The DL-based integrator $\mathcal{R}$ developed here learns both the interaction potentials and the dynamics of the particles based on its experience with the solutions (trajectories) of Newton's equations of motion via the training data generated using the 2nd order Verlet integrator. The $\mathcal{R}$ integrator employs a sequence of 5 past configurations (e.g., positions of particles) to predict the future configuration using a timestep up to 4000Δ with high accuracy (where Δ is the baseline timestep associated with the 2nd order Verlet integrator). We hypothesize that the long input sequence length endows $\mathcal{R}$ with properties that enable it to perform the time evolution with large timesteps. To test this hypothesis, we trained $\mathcal{R}$ with an input sequence length of 3 and 4

positions, and used the resulting integrator to evolve the dynamics of the 1D system of a particle in an LJ potential. We extracted the position error δr incurred in the time evolution up to $t = 1000$ using the ground truth results obtained with the 2nd order baseline ordinary Verlet integrator $\mathcal{V}$ with $\Delta = 0.001$.

We find that for smaller sequence length $S_R = 3$ or 4, $\mathcal{R}$ is only able to accurately propagate the dynamics with $\delta r \sim O(\Delta^{-3})$ for $\Delta_R \lesssim 10\Delta$ (SI Figure 11). For $S_R = 3$, the error associated with $\mathcal{R}$ rises steeply for $\Delta_R > 10\Delta$ and spans over 4 orders of magnitude, similar to the results for the 2nd order Verlet integrator. For $S_R = 4$, the accuracy improves and error scaling is similar to that produced by a traditional 4th order integrator such as Runge-Kutta or Adams-Bashforth-Moulton method (with global error $\sim O(\Delta^4)$). $\mathcal{R}$ integrator trained with sequence length $S_R = 5$ shows a much weaker rise in error limited to within an order of magnitude as Δ_R rises up to 4000Δ. Thus, the $\mathcal{R}$ operator with increasing sequence length S_R learns a progressively more accurate integrator employing the trajectory data generated by the 2nd order Verlet ($\mathcal{V}$) integrator.

Comparison with Traditional Large-Timestep Integrators The main source of speedup shown in Table I is the use of large timestep $\Delta_R \gg \Delta$ in performing particle dynamics with $\mathcal{R}$. Extending the approximate expression for speedup to a general higher-order / multi-step traditional integrator $\mathcal{I}$, we find $S_p^I \approx (t_I/t_R) \times (\Delta_R/\Delta_I) = S_p \times (t_I/t_V) \times (\Delta/\Delta_I)$, where t_I and Δ_I are the time spent and timestep used in the single-step forward propagation using $\mathcal{I}$. The use of $\mathcal{I}$ as the baseline (reference) integrator can reduce the reported net speedup numbers tabulated above as $\Delta/\Delta_I < 1$. However, the increase of timestep with the use of $\mathcal{I}$ generally accompanies a rise in the computational overhead costs as the time for single step propagation typically rises, i.e. $t_I > t_V$. For example, in the case of a 3D system of 16 particles interacting with LJ potential, a 4th order integrator (e.g., Runge-Kutta, Adams-Bashforth-Moulton) is expected to yield $\Delta/\Delta_I \approx 0.1$ but also take $2t_V$ to $5t_V$ longer along with a much higher memory utilization ($4\times$ to $8\times$). Moreover, t_I/t_V is expected to rise with the number of particles N following the $O(N^2)$ scaling in several applications due to the higher number of force computations and memory utilization. Thus, with increasing N, the reduction in Δ/Δ_I is expected to be offset by the increase in t_I/t_V to yield S_p^I values similar to S_p. The $\mathcal{R}$ integrator uniquely bypasses many of these computational challenges at the design stage and enjoys a relatively small forward step propagation (inference) time t_R that is largely independent of the timestep Δ_R and exhibits $O(N)$ scaling. In the case of the above example, $t_R = 1.48 \times 10^{-5}$ seconds (compared to $t_V \approx 0.04$ seconds) is roughly the same for over 3 orders of magnitude change in Δ_R and is expected

to scale linearly with N.

Challenges and Future Work While the use of the DL approach to design integrators has many unique benefits as evidenced by the initial results, to make the approach more generalizable and suitable for practical applications, further research is needed. One key challenge is predicting the dynamics of systems that exhibit rare-event characteristics such that times to yield the low-energy configurations are longer due to the kinetic bottlenecks that can trap the particles in metastable pathways. In these cases, training the RNN model to "see" a variety of distinct phase space explorations may prove too time-consuming. Another challenge is the scaling of the approach to a larger number of particles N. Future work will seek to address these challenges. We will investigate the design of DL-based models for predicting the trajectories of particles in systems where self-organization into aggregates occurs at longer times depending on the combination of particle and potential attributes. We plan to address the issue of scaling with N in part by the use of hierarchical RNNs [48] and alternative DL approaches [49] including transformers [50]. Also, we plan to explore the "recurrent" versions of physics-informed neural network architectures [33, 34] where in addition to the use of sequence-to-sequence mapping, the infusion of conservation laws in the design of loss functions of the LSTM units can help scale the approach to larger N. Future work will also explore the training of $\mathscr{R}$ using trajectories produced via simulations in different ensembles (e.g., NVT) in order to test its accuracy in predicting configuration updates where thermal effects (often stochastic in nature) are incorporated. The implications of the DL-based integrators for other MD methods such as discontinuous MD, which enable large gains in speedup with small loss of accuracy, will also be explored [51].

CONCLUSION

We introduced an integrator based on recurrent neural networks which can learn the dynamics of many particle systems described by Newton's equations of motion using trajectory data generated via the traditional Verlet integrators. The data-driven approach provides a new way to design integrators that are accurate for large timesteps by controlling only a few parameters (e.g., the sequence length). The deep learning model enables an accurate and rapid determination of the large set of 100,000 trainable parameters characterizing the mathematically complex functional form of the neural network based integrator. This complexity makes the integrator robust enough to be applicable over a wide range of potentials and generalizable to systems of many particles. To evaluate our integrator, we showed that it exhibits excellent energy conservation and produces accurate predictions of the trajectories of systems with up to 16 particles using up to $4000\times$ larger timestep than the Verlet integrators. We also showed that our integrator inherits the Lyapunov instabilities of the Verlet integrator. Finally, we demonstrated significant net speedup over a variety of interaction potentials and total number of particles. The deep learning based integrator introduced here and the associated results illustrate an important approach to learn the time evolution operators which can be applied across a variety of fields including fluid dynamics [31], robotics [52], and fusion-energy science [53].

SUPPORTING INFORMATION

Reformulation of Velocity Verlet Integrator using Deep Learning

Many numerical integration methods can perform the dynamics of particles determined by Newton's equation $\vec{F} = m\vec{a}$. Among these, the velocity Verlet integrator $\mathscr{V}$ is a popular choice that is also used in our work to generate the ground truth for many experiments. $\mathscr{V}$ is expressed as:

$$\vec{x}(t + \Delta) = \vec{x}(t) + \Delta\vec{v}(t) + \frac{\Delta^2}{2m}\vec{f}(t), \tag{4}$$

$$\vec{v}(t + \Delta) = \vec{v}(t) + \frac{\Delta}{2m}\left(\vec{f}(t) + \vec{f}(t + \Delta)\right). \tag{5}$$

Here we have adopted the notation for a single particle of mass m; generalization to many particles is straightforward. $\mathscr{V}$ updates the current position $\vec{x}(t)$ at time t to the future position $\vec{x}(t + \Delta)$ after timestep Δ using the current velocity $\vec{v}$ and the force $\vec{f}$ at time t. The current velocity $\vec{v}(t)$ at time t is updated to velocity $\vec{v}(t + \Delta)$ after timestep Δ using the force $\vec{f}(t)$ at time t and the force $\vec{f}(t + \Delta)$ at time $t + \Delta$. Note that after the position update in Equation 4, the force needs to be re-evaluated to update the velocity in Equation 5. With $\vec{x}(t + \Delta)$ as the current position and $\vec{v}(t + \Delta)$ as the current velocity, the propagation is continued further.

We describe a data-driven approach to design an integrator that takes a sequence of past positions and velocities as input and generates the future position and velocity of the particle. Each component of the position and the velocity vectors of a particle is identified as a feature. Inputs and outputs have the same feature size. The total feature size for N particles in $\mathcal{D}$ dimensions is $d = N \times \mathcal{D} \times 2$. For example, for $N = 16$ particles interacting in 3D, the size of both the input and output features is $d = 96$. We introduce an operator $\mathscr{R}$ based on RNN derived using LSTMs (Figure 1) that employs a sequence of current and past positions and velocities (treated as input features) up to time t to predict the future position and velocity at time $t+\Delta_{\mathrm{R}}$: $(\vec{x}_{t+\Delta_{\mathrm{R}}}, \vec{v}_{t+\Delta_{\mathrm{R}}}) = \mathscr{R}\left[\vec{x}_t, \vec{x}_{t-\Delta_{\mathrm{R}}}, \ldots, \vec{x}_{t-S_R\Delta_{\mathrm{R}}}, \vec{v}_t, \vec{v}_{t-\Delta_{\mathrm{R}}}, \ldots, \vec{v}_{t-S_R\Delta_{\mathrm{R}}}\right]$ where Δ_{R} is the timestep associated with the operator $\mathscr{R}$ and S_R is the length of the sequence. $\mathscr{R}$ can be written as $\mathscr{R}\{x,v\} = \mathscr{D}[\mathscr{L}_2[\mathscr{L}_1[\{x,v\}]]]$, where $\mathscr{D}$, $\mathscr{L}_1$, $\mathscr{L}_2$ are operators associated with the dense layer, first LSTM layer, and second LSTM layer of the RNN respectively, and $\{x,v\}$ is a shorthand for the sequence $\vec{x}_t, \vec{x}_{t-\Delta_{\mathrm{R}}}, \ldots, \vec{x}_{t-S_R\Delta_{\mathrm{R}}}, \vec{v}_t, \vec{v}_{t-\Delta_{\mathrm{R}}}, \ldots, \vec{v}_{t-S_R\Delta_{\mathrm{R}}}$. Similar to the derivation shown in the main text, these layers are stacked up on each other such that the output of one (e.g., $\mathscr{L}_1$) becomes the input for another ($\mathscr{L}_2$). Post training, these layers acquire well-determined values for all the parameters and the integrator $\mathscr{R}$ emerges with the more explicit expression:

$$\mathscr{R}[\{x,v\}] = \mathscr{D}[\mathscr{L}_2[\mathscr{L}_1[\{x,v\}, \{P_1\}], \{P_2\}], \{P_D\}], \tag{6}$$

where $\{P_1\}$, $\{P_2\}$, $\{P_D\}$ are trained parameters associated with LSTM layer 1, LSTM layer 2, and the dense layer respectively. While training, $\{x,v\}$ is fed as a vector of size $B \times S_R \times d$, where B is a training parameter denoting the batch size. When used as an integrator (testing phase), $B = 1$. $\mathscr{R}$ is trained on trajectories produced by the using traditional velocity Verlet integrator $\mathscr{V}$ to evolve the dynamics dictated by Newton's equations of motion. Characterized with up to $100,000$ parameters, $\mathscr{R}$ can be considered as a reformulation of the velocity Verlet integrator $\mathscr{V}$ (Equation 4 and 5).

LSTM Networks

There are several architectures of LSTM units. A common architecture is composed of a cell (the memory part of the LSTM unit) and three "regulators", usually called gates, that regulate the flow of information inside the LSTM unit. An input gate (i_t) controls how much new information is added from the present input (x_t) and past hidden state (h_{t-1}) to our present cell state (c_t). A forget gate (f_t) decides what is removed or retained and carried forward to the current cell state (c_t) from the previous cell state (c_{t-1}). An output gate (o_t) decides what to output as the current hidden state (h_t) from the current cell state (c_t). The LSTM formulation can be expressed as:

$$
\begin{aligned}
f_t &= \sigma_g(W_f x_t + U_f h_{t-1} + b_f) \\
i_t &= \sigma_g(W_i x_t + U_i h_{t-1} + b_i) \\
o_t &= \sigma_g(W_o x_t + U_o h_{t-1} + b_o) \\
\tilde{c}_t &= \sigma_h(W_c x_t + U_c h_{t-1} + b_c) \\
c_t &= f_t \circ c_{t-1} + i_t \circ \tilde{c}_t \\
h_t &= o_t \circ \sigma_h(c_t).
\end{aligned}
\tag{7}
$$

Here, $x_t \in \mathbf{R}^d$ is the input vector to the LSTM unit, $f_t \in \mathbf{R}^h$ is the forget gate's activation vector, $i_t \in \mathbf{R}^h$ is the input gate's activation vector, $o_t \in \mathbf{R}^h$ is the output gate's activation vector, $h_t \in \mathbf{R}^h$ is the hidden state vector also known as the output vector of the LSTM unit, $c_t \in \mathbf{R}^h$ is the cell state vector, and $\circ$ is the Hadamard product operator. $W \in \mathbf{R}^{h \times d}$ and $U \in \mathbf{R}^{h \times h}$ are the weight matrices and $b \in \mathbf{R}^h$ are the bias vector parameters which need to be learned during training. σ_g and σ_h represent sigmoid function and hyperbolic tangent functions respectively. The superscripts d and h refer to the number of input features and the number of hidden units respectively.

Data Generation, Preparation, and Preprocessing

In most cases, prior domain experience and backward elimination using the adjusted R squared method is employed to determine the important input parameters that significantly change the desired output. In some cases, such as particles in Lennard-Jones (LJ) potential, only a subset of the important input parameters are varied to create the dataset to train the RNN-based $\mathscr{R}$ integrator. Unlike traditional deep neural networks where the physical inputs are mapped to outputs generally distinct from inputs, the inputs and outputs for the RNN approach are both trajectory data as shown in Figure 6. MD simulation driven using the velocity Verlet algorithm generally produces time series data separated by $\Delta = 0.001$ in time. The data is characterized by multiple features such as position and velocity vectors. To train the integrator $\mathscr{R}$, we prepare the inputs and target data as explained in Figure 6. Since integrator $\mathscr{R}$ will operate with different timestep Δ_R (e.g., 100Δ, 1000Δ), we filter the time series data with a filtering factor defined as $\Delta_{\mathrm{R}}/\Delta$ to make the data separated with Δ_{R}. Next, we do windowing through the time series data with the window length of S_R (the RNN model sequence length), and frame overlap length of $S_R - 1$ to generate a sequence of length S_R

as input and the future configuration as the output (target). We do this repeatedly until the time series is ended and then move to the next choice for Δ_R and a new associated filtering factor. We do the same process for all simulation data and append all the input sequences and target vectors as shown in Figure 6.
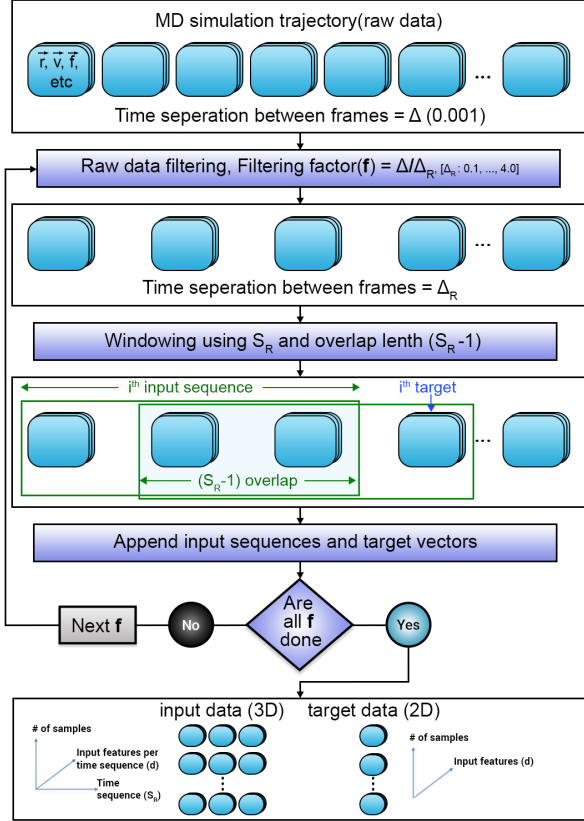

FIG. 6. Data processing for training and testing of integrator $\mathscr{R}$ with sequence length S_{R} as a key input parameter.

In the end, all the input sequences are expected to form a 3D matrix (number of samples $\times$ feature size (d) $\times$ S_R) and the associated output vectors form a 2D matrix (number of samples $\times$ d). The entire data set is randomly shuffled (along the axis of the number of samples) and separated into training and testing sets using a ratio of 0.8:0.2. We now describe the datasets associated with systems characterized by different potential energies and particle attributes used in the experiments for training and testing $\mathscr{R}$. The potential energy functions are shown in Figure 7. For the 1D systems, the timestep and the total time associated with generating the ground truth using the velocity Verlet integrator were $\Delta = 0.001$ and $t_f = 100$. For the 3D many particle systems, $\Delta = 0.001$ and $t_f = 2000$.


FIG. 7. Potential energies associated with the 1D experiments. Dash-dotted, dotted, dashed, and solid lines represent SHO, DW, LJ, and rugged potential respectively.

Simple harmonic oscillator (SHO) For this system, the potential energy is given by

$$U = \frac{1}{2}kx^2, \tag{8}$$

and the associated force is

$$f = -kx. \tag{9}$$

The dataset is generated by varying three input parameters: mass of the particle $m \in [1, 10]$, spring constant $k \in [1, 10]$, and initial position of the particle $x_0 \in [-10, -1]$. The parameter sweep generated a dataset of 500 simulations, each having 50,000 position and velocity values.

Particle in a double well (DW) For this system, the potential energy is given by

$$U = \frac{1}{4}x^4 - \frac{1}{2}x^2, \tag{10}$$

and the associated force is

$$f = x - x^3. \tag{11}$$

The dataset is generated by varying two input parameters: mass of the particle $m \in [1, 10]$ and the initial position of the particle $x_0 \in [-3.1, 3.1]$. The parameter sweep generated a dataset of 500 simulations, each having 50,000 position and velocity values.

1D Lennard-Jones (LJ) system For this system, the potential energy is given by

$$U(x) = 4\epsilon \left(\left(\frac{1}{x} \right)^{12} - \left(\frac{1}{x} \right)^{6} \right), \tag{12}$$

and the associated force is

$$f = \frac{48\epsilon}{x} \left(\left(\frac{1}{x} \right)^{12} - 0.5 \left(\frac{1}{x} \right)^{6} \right). \tag{13}$$

The dataset is generated by varying two input parameters: mass of the particle $m \in [1, 10]$ and the initial position of the particle $x_0 \in [1.0, 3.0]$. The parameter sweep generated a dataset of 500 simulations, each having 50,000 position and velocity values.

Particle in a rugged potential For this system, the potential energy is given by

$$U(x) = \frac{x^4 - x^3 - 16x^2 + 4x + 48}{50} + \frac{\sin\left(30(x+5)\right)}{5}, \tag{14}$$

and the associated force is

$$f = \frac{-4x^3 + 3x^2 + 32x - 4}{50} - 6\cos\left(30(x+5)\right). \tag{15}$$

The dataset is generated by varying two input parameters: mass of the particle $m \in [1, 10]$ and the initial position of the particle $x_0 \in [-6.1, 6.1]$. The parameter sweep generated a dataset of 640 simulations, each having 64,000 position and velocity values.

Many particles interacting with LJ potential For this 3D system, the interaction potential energy between any two particles is given by the LJ potential:

$$U(r) = 4\epsilon \left(\left(\frac{1}{r}\right)^{12} - \left(\frac{1}{r}\right)^6 \right) + 0.0163\epsilon \quad \text{for } r \leq 2.5, \tag{16}$$

and the associated force is

$$\vec{f} = \frac{48\epsilon}{r^2} \left(\left(\frac{1}{r}\right)^{12} - 0.5\left(\frac{1}{r}\right)^6 \right) \vec{r} \quad \text{for } r \leq 2.5. \tag{17}$$

For $r > 2.5$, both U and $\vec{F}$ are 0. Note that the unit of distance is chosen to be the particle diameter σ and the unit of energy is chosen such that the LJ potential well depth ϵ has values around 1. We prepared two different types of simulation boxes to generate the datasets: cubic box with periodic boundary conditions (PBC) and spherical box with reflective boundary. In each box, we performed simulations with $N = 3, 8$ and 16 particles. Each of these simulations was created as a separate dataset that yielded 6 different datasets to train and test $\mathscr{R}$. The dataset is generated by varying two input parameters: mass of the particle $m \in [1, 10]$ and the initial position of the particles (x_0, y_0, z_0) with each Cartesian coordinate chosen between -3.0 and 3.0. The well depth was held constant to $\epsilon = 1$ in creating the training dataset. Initial velocities were chosen to be zero. The parameter sweep generated a dataset of 5000 simulations for each of the aforementioned cases (or 30,000 simulations in total).

Additional Experiments and Empirical Results

Phase diagram for LJ This experiment focused on testing $\mathscr{R}$ to predict the time evolution of the 1D LJ sys-

tem. Training dataset comprised of trajectories (positions and velocities) generated using $\mathscr{V}$ with $\Delta = 0.001$ for up to $t = 200$. Figure 8 shows the 1st period (top row) from $t = 1$ to $t = 13$ and the 75th period (bottom row) from $t = 987$ to $t = 1000$ of the simulation. The 75th period is towards the end of the simulation. The result is shown for a particle of mass $m = 1$, initial position $x_0 = 2.5$, and $\epsilon = 1$. For either time periods characterizing the oscillations, the positions predicted by $\mathscr{R}$ using a wide range of timestep $\Delta_\mathrm{R} \in [100\Delta, 4000\Delta]$ remain close to the numerically exact result. On the other hand, the trajectory simulated using $\mathscr{V}$ shows errors rising with t for timesteps $\gtrsim 10\Delta$. Similarly, velocities v produced by $\mathscr{R}$ exhibit very small deviations from the ground truth result for all Δ_R investigated as evidenced by the phase space plots (v vs. x) in Figure 8. On the other hand, the phase space plot generated using trajectories evolved with $\mathscr{V}$ deviated from the ground truth for timesteps $\gtrsim 10\Delta$. Finally, the total energy produced by $\mathscr{R}$ tracks the ground truth result for t up to 1000 and $\Delta_\mathrm{R} \in [100\Delta, 4000\Delta]$, in stark contrast with results using $\mathscr{V}$ as the integrator.

Many particle experiments Figure 9 shows the total error associated with the positions of many particles as a function of time for dynamics driven by $\mathscr{R}$ and $\mathscr{V}$. The errors are computed using:

$$\delta r = \frac{1}{N} \sum_{i=1}^{N} |\vec{r}_{i,\mathrm{GT}} - \vec{r}_{i,\mathscr{R}}| \tag{18}$$

where the formula is shown for $\mathscr{R}$. Here N is the number of particles, $\vec{r}_{i,\mathrm{GT}}$ is the ground truth result for the position of the i^th particle at time t, $\vec{r}_{i,\mathscr{R}}$ is the position of the i^th particle predicted by $\mathscr{R}$ at the same time t. The ground truth (GT) results for the trajectories are generated using $\mathscr{V}$ with $\Delta = 0.001$. Note that the training set consisted of dynamics up to $t = 2000$ and we intentionally kept the random initial configuration tested here outside the training and testing datasets. The simulated system is $N = 16$ particles interacting in 3D with LJ potentials under PBC. We find that the trajectory error δr associated with the positions predicted by $\mathscr{R}$ for $\Delta_\mathrm{R} = 100\Delta$ to 4000Δ is $O(10^{-3})$ up to $t = 10^6$. On the other hand, δr for the system evolved using $\mathscr{V}$ with timestep 10Δ increases to $O(10^{-1})$ at around $t \approx 1000$. For timestep values of $40\Delta, 100\Delta$, $\mathscr{V}$ does worse as δr rises to very large values $> 10^9$, rendering the time evolution of the system meaningless.

Testing $\mathscr{R}$ with different sequence lengths In this experiment, we tested the dependence of the accuracy of time evolution performed by $\mathscr{R}$ integrator on the sequence length of input configurations used to train the operator. We trained $\mathscr{R}$ with different input sequence length S_R experimenting with $S_R = 3, 4$ and 5. The training was performed using the same dataset employed
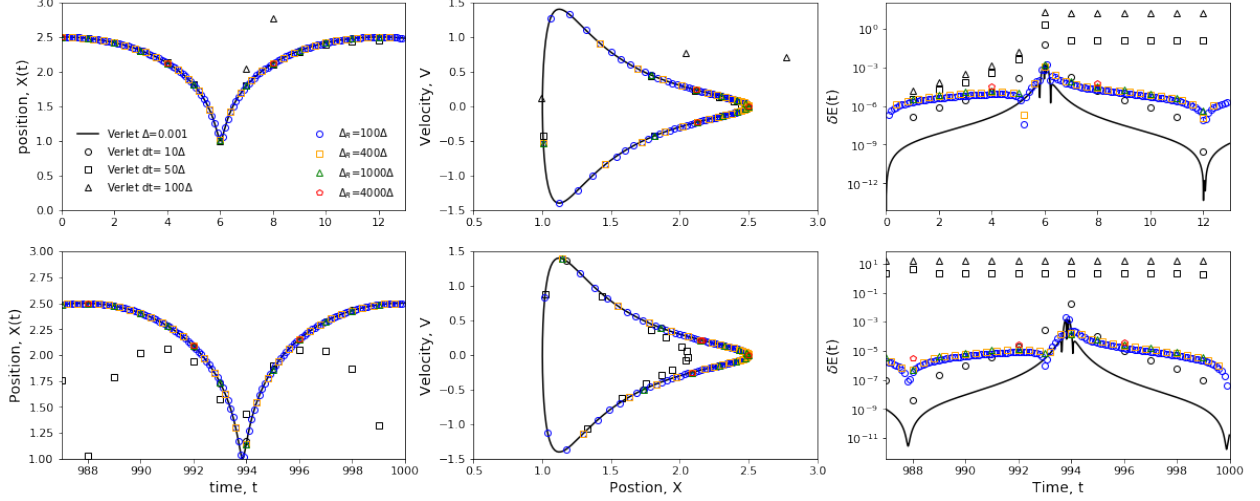
FIG. 8. Dynamics of a particle in a 1D LJ potential ($m = 1, x_0 = 2.5$) captured by tracking the position x vs. time t (left column), phase space plot of velocity v vs. x (center column), and the deviation in the total energy (right column). Results are shown for the dynamics produced by $\mathscr{V}$ and $\mathscr{R}$. Top row is the dynamics from $t = 0$ to $t = 13$, and bottom row is the dynamics from $t = 987$ to $t = 1000$. Black lines are the ground truth results obtained using $\mathscr{V}$ with $\Delta = 0.001$. Black color circles, squares and triangles represent the results obtained using $\mathscr{V}$ with timestep $10\Delta, 50\Delta$, and 100Δ. Blue circles, orange squares, green triangles and red pentagons represent the results obtained using $\mathscr{R}$ with $\Delta_R = 100\Delta, 400\Delta, 1000\Delta$, and 4000Δ respectively.


FIG. 9. Average error δr (log scale) in evolved positions vs. time t incurred in a 3D simulation of $N = 16$ particles interacting with LJ potentials ($\epsilon = 2$ and $m = 1$) in periodic boundary conditions (PBC). Initial positions are randomly selected and velocities are initialized to zero. All results are compared with the ground truth result obtained using $\mathscr{V}$ with $\Delta = 0.001$. Left plot shows δr for the time evolution up to $t = 1$ million (log scale) for $\mathscr{V}$ using different $dt = 10\Delta, 40\Delta, 100\Delta$ values. Right plot shows δr for the same time evolution using $\mathscr{R}$ with $\Delta_R = 100\Delta, 400\Delta, 1000\Delta$, and 4000Δ.

in the original 1D LJ system experiment. The resulting integrator was used to evolve the dynamics of the 1D system of a particle in an LJ potential. Figure 11 shows the position error δr incurred in the time evolution up to $t = 1000$ vs. Δ using different integrators. δr is evaluated using the ground truth results obtained with the 2nd order baseline Verlet integrator $\mathscr{V}$ with $\Delta = 0.001$.

We find that for smaller sequence length $S_R = 3$ or 4, $\mathscr{R}$ fails to accurately perform the time evolution for $\Delta_R \gtrsim 10\Delta$ incurring errors $\delta r \gtrsim O(\Delta^{-2})$ (SI Figure 11). For $S_R = 3$, the error associated with $\mathscr{R}$ rises steeply for $\Delta_R > 10\Delta$ and spans over 4 orders of magnitude, similar to the results for the 2nd order Verlet inte-

grator (or its equivalent $\mathscr{V}$ operator that takes a sequence of 3 positions as input). For $S_R = 4$, the accuracy improves and error scaling is similar to that produced by a traditional 4th order integrator such as Runge-Kutta or Adams-Bashforth-Moulton method (with global error $\sim O(\Delta^4)$). $\mathscr{R}$ integrator trained with sequence length $S_R = 5$ shows a much weaker rise in error limited to within an order of magnitude as Δ_R rises up to 4000Δ.

───────────

* kadu@iu.edu
† gcf@iu.edu

FIG. 10. Average error δr (log scale) in evolved positions vs. time t incurred in a 3D simulation of $N = 16$ particles interacting with LJ potentials ($\epsilon = 2$ and $m = 1$) in spherical hard-wall confinement. Initial positions are randomly selected and velocities are initialized to zero. All results are compared with the ground truth result obtained using $\mathscr{V}$ with $\Delta = 0.001$. Left plot shows δr for the time evolution up to $t = 1$ million (log scale) for $\mathscr{V}$ using different $dt = 10\Delta, 40\Delta, 100\Delta$ values. Right plot shows δr for the same time evolution using $\mathscr{R}$ with Δ_R=100Δ, 400Δ, 1000Δ, and 4000Δ.
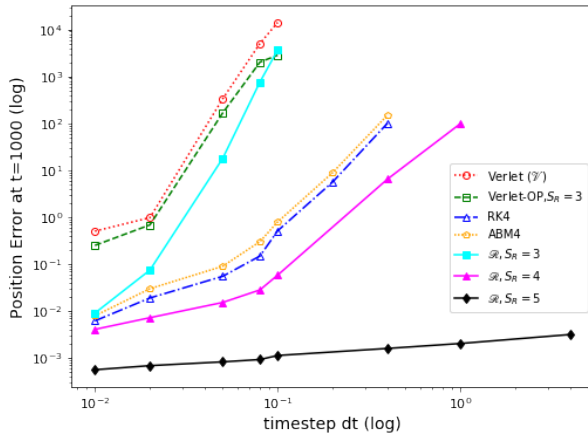

FIG. 11. Error in position at time $t = 1000$ for different timestep $dt \geq 10^{-2}$ in 1D simulations of a particle of mass $m = 1$ in an LJ potential with initial position $x_0 = 2.0$. The 2nd order baseline Verlet integrator $\mathscr{V}$ with $dt = 10^{-3}$ is used to generate the ground truth results. Open symbols correspond to traditional numerical integrators and closed symbols are results from the $\mathscr{R}$ integrator. Open circles and squares represent the results of using the traditional Verlet integrator ($\mathscr{V}$) and its equivalent operator (Verlet-OP or $\mathcal{V}$) with $dt \geq 10^{-2}$. Open triangles and pentagons represent the errors obtained with 4th order Runge-Kutta (RK4) and 4th order Adams-Bashforth-Moulton (ABM4) respectively. Closed squares, triangles, and diamonds represent errors incurred using $\mathscr{R}$ with sequence length $S_R = 3, 4$ and 5 respectively.

‡ vjadhao@iu.edu

[1] Issac Newton, Philosophiae Naturalis Principia Mathematica (1687).

[2] Daan Frenkel and Berend Smit, Understanding Molecular Simulation, 2nd ed. (Academic Press, 2001).

[3] Sharon C Glotzer, "Assembly engineering: Materials design for the 21st century (2013 pv danckwerts lecture)," Chemical Engineering Science 121, 3–9 (2015).

[4] Nicholas E Brunk and Vikram Jadhao, "Computational studies of shape control of charged deformable nanocontainers," Journal of Materials Chemistry B 7, 6370 (2019).

[5] Oliver Beckstein, Kaihsu Tai, and Mark SP Sansom, "Not ions alone: barriers to ion permeation in nanopores and channels," Journal of the American Chemical Society 126, 14694–14695 (2004).

[6] Hans C Andersen, "Rattle: A velocity version of the shake algorithm for molecular dynamics calculations," Journal of Computational Physics 52, 24–34 (1983).

[7] Andrew L Ferguson, "Machine learning and data science in soft materials engineering," Journal of Physics: Condensed Matter 30, 043002 (2017).

[8] Pankaj Mehta, Marin Bukov, Ching-Hao Wang, Alexandre GR Day, Clint Richardson, Charles K Fisher, and David J Schwab, "A high-bias, low-variance introduction to machine learning for physicists," Physics reports (2019).

[9] Keith T Butler, Daniel W Davies, Hugh Cartwright, Olexandr Isayev, and Aron Walsh, "Machine learning for molecular and materials science," Nature 559, 547 (2018).

[10] Geoffrey Fox, James A. Glazier, JCS Kadupitiya, Vikram Jadhao, Minje Kim, Judy Qiu, James P. Sluka, Endre Somogyi, Madhav Marathe, Abhijin Adiga, Jiangzhuo Chen, Oliver Beckstein, and Shantenu Jha, "Learning everywhere: Pervasive machine learning for effective High-Performance computation," in HPDC Workshop at IPDPS 2019 (2019).

[11] Samuel S. Schoenholz, "Combining machine learning and physics to understand glassy systems," Journal of Physics: Conference Series 1036, 012021 (2018).

[12] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar, "Integrating physics-based modeling with machine learning: A survey," arXiv preprint arXiv:2003.04919 (2020).

[13] Matthew Spellings and Sharon C Glotzer, "Machine learning for crystal identification and discovery," AIChE Journal 64, 2198–2206 (2018).

[14] Ashley Z Guo, Emre Sevgen, Hythem Sidky, Jonathan K Whitmer, Jeffrey A Hubbell, and Juan J de Pablo, "Adaptive enhanced sampling by force-biasing using neural networks," The Journal of chemical physics 148, 134108

(2018).

[15] Venkatesh Botu and Rampi Ramprasad, "Adaptive machine learning framework to accelerate ab initio molecular dynamics," International Journal of Quantum Chemistry **115**, 1074–1083 (2015).

[16] JCS Kadupitiya, Geoffrey C Fox, and Vikram Jadhao, "Machine learning for parameter auto-tuning in molecular dynamics simulations: Efficient dynamics of ions near polarizable nanoparticles," The International Journal of High Performance Computing Applications (2020), 10.1177/1094342019899457.

[17] Andrew W Long, Jie Zhang, Steve Granick, and Andrew L Ferguson, "Machine learning assembly landscapes from particle tracking data," Soft Matter **11**, 8141–8153 (2015).

[18] Alireza Moradzadeh and Narayana R Aluru, "Molecular dynamics properties without the full trajectory: A denoising autoencoder network for properties of simple liquids," The journal of physical chemistry letters **10**, 7568–7576 (2019).

[19] Yangzesheng Sun, Robert F DeJaco, and J Ilja Siepmann, "Deep neural network learning of complex binary sorption equilibria from molecular simulation data," Chemical science **10**, 4377–4388 (2019).

[20] Florian Hse, Ignacio Fdez. Galvn, Aln Aspuru-Guzik, Roland Lindh, and Morgane Vacher, "How machine learning can assist the interpretation of ab initio molecular dynamics simulations and conceptual understanding of chemistry," Chem. Sci. **10**, 2298–2307 (2019).

[21] JCS Kadupitiya, Geoffrey C Fox, and Vikram Jadhao, "Machine learning for performance enhancement of molecular dynamics simulations," in International Conference on Computational Science (2019) pp. 116–130.

[22] JCS Kadupitiya, Fanbo Sun, Geoffrey Fox, and Vikram Jadhao, "Machine learning surrogates for molecular dynamics simulations of soft materials," Journal of Computational Science , 101107 (2020).

[23] Maziar Raissi and George Em Karniadakis, "Hidden physics models: Machine learning of nonlinear partial differential equations," Journal of Computational Physics **357**, 125–141 (2018).

[24] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis, "Multistep neural networks for data-driven discovery of nonlinear dynamical systems," arXiv preprint arXiv:1801.01236 (2018).

[25] Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong, "Pde-net: Learning pdes from data," arXiv preprint arXiv:1710.09668 (2017).

[26] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud, "Neural ordinary differential equations," in Advances in neural information processing systems (2018) pp. 6571–6583.

[27] Katsuhiro Endo, Katsufumi Tomobe, and Kenji Yasuoka, "Multi-step time series generator for molecular dynamics," in Thirty-Second AAAI Conference on Artificial Intelligence (2018).

[28] Philip G Breen, Christopher N Foley, Tjarda Boekholt, and Simon Portegies Zwart, "Newton vs the machine: solving the chaotic three-body problem using deep neural networks," arXiv preprint arXiv:1910.07291 (2019).

[29] Zhengdao Chen, Jianyu Zhang, Martin Arjovsky, and Léon Bottou, "Symplectic recurrent neural networks," arXiv preprint arXiv:1909.13334 (2019).

[30] Maziar Raissi, Paris Perdikaris, and George E Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics **378**, 686–707 (2019).

[31] Yohai Bar-Sinai, Stephan Hoyer, Jason Hickey, and Michael P Brenner, "Learning data-driven discretizations for partial differential equations," Proceedings of the National Academy of Sciences **116**, 15344–15349 (2019).

[32] Xing Shen, Xiaoliang Cheng, and Kewei Liang, "Deep euler method: solving odes by approximating the local truncation error of the euler method," arXiv preprint arXiv:2003.09573 (2020).

[33] Sam Greydanus, Misko Dzamba, and Jason Yosinski, "Hamiltonian neural networks," arXiv preprint arXiv:1906.01563 (2019).

[34] Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho, "Lagrangian neural networks," arXiv preprint arXiv:2003.04630 (2020).

[35] Stefan Chmiela, Alexandre Tkatchenko, Huziel E Sauceda, Igor Poltavsky, Kristof T Schütt, and Klaus-Robert Müller, "Machine learning of accurate energy-conserving molecular force fields," Science advances **3**, e1603015 (2017).

[36] Alvaro Sanchez-Gonzalez, Victor Bapst, Kyle Cranmer, and Peter Battaglia, "Hamiltonian graph networks with ode integrators," arXiv preprint arXiv:1909.12790 (2019).

[37] Yeonjong Shin, Jerome Darbon, and George Em Karniadakis, "On the convergence and generalization of physics informed neural networks," arXiv preprint arXiv:2004.01806 (2020).

[38] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al., "Google's neural machine translation system: Bridging the gap between human and machine translation," arXiv preprint arXiv:1609.08144 (2016).

[39] Eunsuk Chong, Chulwoo Han, and Frank C Park, "Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies," Expert Systems with Applications **83**, 187–205 (2017).

[40] X Huang, G C Fox, S Serebryakov, A Mohan, P Morkisz, and D Dutta, "Benchmarking deep learning for time series: Challenges and directions," in 2019 IEEE International Conference on Big Data (Big Data) (ieeexplore.ieee.org, 2019) pp. 5679–5682.

[41] François Chollet et al., "Keras," (2015).

[42] Lars Buitinck et al., "Api design for machine learning software: experiences from the scikit-learn project," arXiv:1309.0238 (2013).

[43] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al., "Ten-

sorflow: a system for large-scale machine learning." in OSDI, Vol. 16 (2016) pp. 265–283.

[44] Available at https://github.com/softmaterialslab/RNN-MD/.

[45] Xavier Glorot and Yoshua Bengio, "Understanding the difficulty of training deep feedforward neural networks," (2010) pp. 249–256.

[46] HA Posch and Wm G Hoover, "Lyapunov instability of dense lennard-jones fluids," Physical Review A **38**, 473 (1988).

[47] GD Venneri and William G Hoover, "Simple exact test for well-known molecular dynamics algorithms," Journal of computational physics (Print) **73**, 468–475 (1987).

[48] Yong Du, Wei Wang, and Liang Wang, "Hierarchical recurrent neural network for skeleton based action recognition," (2015) pp. 1110–1118.

[49] Deniz Eroglu, Norbert Marwan, Martina Stebich, and Jürgen Kurths, "Multiplex recurrence networks," Physical Review E **97**, 012312 (2018).

[50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention is all you need," in Advances in neural information processing systems (2017) pp. 5998–6008.

[51] Lisandro Hernández de la Peña, Ramses van Zon, Jeremy Schofield, and Sheldon B Opps, "Discontinuous molecular dynamics for rigid bodies: Applications," The Journal of chemical physics **126**, 074106 (2007).

[52] Peiyao Shen, Xuebo Zhang, and Yongchun Fang, "Essential properties of numerical integration for time-optimal path-constrained trajectory planning," IEEE Robotics and Automation Letters **2**, 888–895 (2017).

[53] Julian Kates-Harbeck, Alexey Svyatkovskiy, and William Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," Nature **568**, 526–531 (2019).