Towards Learning Self-Organized Criticality of Rydberg Atoms using Graph Neural Networks

Anonymous Authors[1]

Abstract

Self-Organized Criticality (SOC) is a ubiquitous dynamical phenomenon believed to be responsible for the emergence of universal scale-invariant behavior in many, seemingly unrelated systems, such as forest fires, virus spreading or atomic excitation dynamics. SOC describes the buildup of large-scale and long-range spatio-temporal correlations as a result of only local interactions and dissipation. The simulation of SOC dynamics is typically based on Monte-Carlo (MC) methods, which are however numerically expensive and do not scale beyond certain system sizes. We investigate the use of Graph Neural Networks (GNNs) as an effective surrogate model to learn the dynamics operator for a paradigmatic SOC system, inspired by an experimentally accessible physics example: driven Rydberg atoms. To this end, we generalize existing GNN simulation approaches to predict dynamics for the internal state of the node. We show that we can accurately reproduce the MC dynamics as well as generalize along the two important axes of particle number and particle density. This paves the way to model much larger systems beyond the limits of traditional MC methods. While the exact system is inspired by the dynamics of Rydberg atoms, the approach is quite general and can readily be applied to other systems.

1. Introduction

Deep Learning methods deliver impact for many applications ranging from Computer Vision, Natural Language Processing, Audio and Speech Processing to Optimal Control. In recent years we also see these methods drive innovation in areas of scientific interest, such as drug discovery (Jim'enez-Luna et al., 2020), protein folding (Jumper et al., 2021), molecule dynamics (Noé et al., 2020; Gilmer et al., 2017; Henderson et al., 2021) as well as the dynamics of classic multi-particle systems (Battaglia et al., 2016; Sanchez-Gonzalez et al., 2020) and many more. In the context of particle dynamics the class of Graph Neural Networks (GNNs) (Bronstein et al., 2021; Sanchez-Gonzalez et al., 2020; Shlomi et al., 2021; Lemos et al., 2022) are particularly successful. We build up on this success and investigate the application of GNNs to learn the time-evolution of a complex many-body system in the regime of so-called self-organized criticality (SOC).

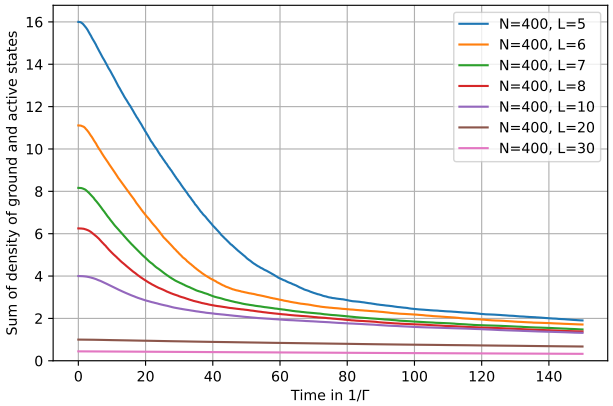

Figure 1. Results from the 60000-timestep rollout of the GNN, after being trained only on single timestep data. The available density (ground and active state atoms) shows the characteristic SOC behavior: all initial densities above the critical value of $n_c \approx 1$ converge to the same point, whereas all curves that start below n_c show only a slow decay.

SOC represents a widespread phenomenon observed in many different fields, ranging from noise in electric circuits (Bak et al., 1987), earthquakes (Sornette & Sornette, 1989), forest fires (Malamud et al., 1998; Drossel et al., 1993) over the spread of information or diseases (Gleeson et al., 2014; Rhodes & Anderson, 1996) to atomic excitation dynamics (Helmrich et al., 2020). The term refers to the property of a complex system to drive itself towards the critical point be-

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

tween two dynamical phases in the relevant parameter space during its time evolution without any external control or fine-tuning. Remarkably, even though the microscopic processes in these systems may be very different, certain properties such as the scaling laws of the critical behaviour remain the same, a phenomenon termed *universality*. This allows to study the properties of SOC dynamics in one system, e.g. forest fires, through controlled simulation and experimentation of other systems, e.g. atoms. An important challenge in studying SOC-behavior is the efficient simulation of the dynamics of large systems. State-of-the-art methods employ Markov-Chain Monte-Carlo (MCMC) methods that are limited in terms of number of particles and system size (Wintermantel et al., 2021). This renders the study of large-scale systems, required for reliable predictions of key SOC properties such as critical exponents, hard. Additionally, the simulation requirements in time and memory increase with the system size and make the simulations prohibitively expensive in the absence of high-performance computation infrastructure.

In this study, we investigate the use of GNNs to accurately approximate the time-evolution operator of a complex system inspired by the dynamics of a classical ensemble of Rydberg atoms in the SOC regime. The GNN acts as an effective surrogate model and replaces the costly MCMC step operator. We model the atoms as nodes in a GNN with internal degrees of freedom encoded in the node attribute and the Euclidean distances between the endpoints attached as edge attributes. We learn the one-timestep update function of the node attributes from data simulated using classical MCMC methods. We demonstrate that the GNN is capable to efficiently learn this operator and reliably reproduce the characteristic properties of the system such as the facilitation dynamics, Rydberg blockade effects and SOC excitation trajectories. Moreover, we show early indications that the model generalizes over several orders of magnitudes along two important axes: particle number and system size (particle density). This property enables the application of this model to system sizes beyond what can be simulated using classical MCMC methods, while also enabling computations on moderate hardware.

In addition, we are, to the best of our knowledge, the first to extend the simulation capabilities of GNNs to systems with dynamics in internal as well as external degrees of freedom. While previous work was mostly focused on simulating the position dynamics of the nodes (Battaglia et al., 2016; Kipf et al., 2018; Sanchez-Gonzalez et al., 2020), an external degree of freedom, we also discuss the dynamics of the node attribute, i.e. the internal degree of freedom, based on the interaction of a node with its neighbors.

We make both the code for our experiments as well as the Monte Carlo code to create our datasets publicly available upon acceptance.

2. Related Work

GNNs have been used in a multitude of applications. Most relevant to our study is the work by Sanchez-Gonzalez et al. (2020) to learn the dynamics of challenging physics motions such as viscous fluids and elastic bodies simulated using third party tools. They do so by breaking the objects into a discrete set of volumes that are subsequently represented as nodes in a graph. The position update of the volume is learned by predicting the next position using an encoder-decoder system with several message passing (MP) layers.

Pfaff et al. (2021) expand on this work to simulate the dynamics of wings and flags under external forces. They do so by introducing two different meshes, one for the surface of the object under study and a second one encoding the real-world, Euclidean distance between the nodes of the surface mesh. Using MP they learn to predict the surface mesh positions in the next time-step from simulated data.

Battaglia et al. (2016) study the use of graph representations in the data to enable a multi-layer perceptron to reason about complex interactions of physical systems. By making the relationships between physical objects explicit in the graph data structure, they are able to predict the orbital motion of planets as well as the dynamics of solids under non-trivial, external constraints.

Shlomi et al. (2021) give a review of the use of GNNs in high-energy particle physics. They draw attention to the benefits of using GNNs to model the different particle interactions through clever design of the node, edge and vertex attributes when representing the data as relationship graphs.

Lemos et al. (2022) use GNN simulators in conjunction with symbolic regression to rediscover Newton's law of gravity from simulated data of multi-planetary motions. This serves as a demonstration to infer physical laws from complex dynamics as multi-planetary motion can behave chaotically under certain circumstances.

Martinkus et al. (2021) discuss the construction of a hierarchical GNN in order to scale the simulation of particle dynamics from a few hundreds to tens of thousands. By recursively partitioning the space and compute interactions between the cells at each hierarchy they are able to reduce the $\mathcal{O}(N^2)$ interaction computation to $\mathcal{O}(N \log N)$ complexity.

Prakash & Tucker (2021) investigate the use of multi-task learning in GNNs to simultaneously learn the dynamics and the unknown physical parameters of an ensemble of objects with Newtonian dynamics. They show that using this

approach one can circumvent the design of inductive biases in the GNN construction, making the GNN architecture more versatile across datasets.

In contrast to these works, we are the first to discuss the dynamics of a complex system whose entities have an additional internal state, i.e. where the node attributes can change depending on the neighbors. This extends the use of GNNs to efficiently simulate systems that have internal as well as external degrees of freedom. On top of this we also demonstrate the ability of the GNNs to learn non-trivial spatio-temporal correlations, despite being trained on a spatially local, one-step prediciton task.

3. Self-Organized Criticality in Rydberg Atoms

A typical SOC system consists of many identical entities, each taking one of three internal states. The entity's state can vary over time as a function of the system's dynamics. Since in our case the entities are Rydberg atoms, we label these three states as the *ground*, *active* and *inactive* states, respectively (see Figure 2(b)). For illustration, in disease spread models these states are commonly referred to as susceptible, infected and recovered. The key ingredient to produce the characteristic SOC dynamics lies in the transitions between the internal states of each atom, as shown in Figure 2(a): the ground state $|g\rangle$ can only become active via the neighborhood-dependent rate $p_{g\to a}(N)$, called the *facilitation* process. The active state $|a\rangle$ can undergo the reverse process and return to the ground state again with a neighborhood-dependent rate $p_{a\to g}(N)$, or transition into the inactive state $|0\rangle$, which is a terminal state in the process.

The physical realization of the SOC system we study here is based on a gas of Rydberg atoms subject to external laser fields. The microscopic transition rates for each atom i per time interval δt of the model are given by

$$p_{i,g\to a} = 1 - \mathrm{e}^{-R_i \delta t} \tag{1}$$

$$p_{i,a\to g} = (1-\epsilon_i)\left(1 - \mathrm{e}^{-(R_i+\gamma)\delta t}\right) \tag{2}$$

$$p_{i,a\to 0} = \epsilon_i \left(1 - \mathrm{e}^{-(R_i+\gamma)\delta t}\right) \tag{3}$$

$$R_i = \frac{2\Omega^2\gamma}{\gamma^2 + \Delta_i^2} \tag{4}$$

$$\epsilon_i = \frac{b}{1 + R_i/\Gamma} \tag{5}$$

$$\Delta_i = \Delta_0 \left(-1 + \sum_{\ell\in N} \frac{1}{r_{i,\ell}^6}\kappa_\ell\right) \tag{6}$$

$$\kappa_\ell = \begin{cases} 1, & \text{if } s_\ell = |a\rangle \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

The parameters $\Delta_0, \Gamma, b, \gamma,$ and Ω are details of the system and can be treated as hyper-parameters to the problem. N denotes the total number of atoms in the system. The symbol κ_ℓ in Equation (6) ensures that the sum includes only the active atoms (state $s_\ell = |a\rangle$). Additionally, $r_{i,\ell}$ denotes the Euclidean distance between atom i and ℓ in natural units, where we set the length scale $R_F = 1$ (see Figure 2(b)). A more detailed introduction to this system is presented in Appendix A. It is worth pointing out that Equation (6) displays a resonance behavior dependent on the number and distance of the neighboring active atoms, giving rise to the neighborhood-dependent state transition rates: If $|\Delta|$ is large, the transition rate is suppressed and no change in the dynamics is happening. If an active atom is present at the right distance to an atom i, then $\Delta_i = 0$ and the transition is facilitated. The facilitated atom can in turn facilitate the excitation of another atom. Above a critical density this leads to an explosive spread of excitations. The system hence shows two distinct phases:

- The *active* phase is characterized by a large number of active atoms that spread through the entire system. The density of inactive atoms is small, and each active atom facilitates on average more or equal to one ground state atom before it ceases to be active.
- The *absorbing* phase is a state where the available density of ground state atoms is so low that transitions into the active state typically become inactive quickly before they can spread, i.e. facilitate the transition of another ground state atom into the active state.

The critical point of the system occurs at the point where the two phases meet, i.e. where an active atom facilitates exactly one ground state atom on average before it decays. Since the inactive state is a dead-end, the available density of ground and active state atoms can only decrease over time. Consequently, the active phase automatically moves towards the critical point, which constitutes a self-organization of the critical phase termed SOC. Since SOC systems at the critical point share certain universal properties, being able to simulate a single SOC system with high accuracy and scalability allows investigations of many different physical systems.

4. Learning the Rydberg SOC Dynamics with a GNN

Training data. We model the atoms as the nodes of a GNN with the internal states of the atoms mapped to node attributes and the distance between them as the edge attributes (see Figure 2(b)). An advantage of the GNN structure over conventional discretization schemes is the independence of system density from the discretization length and thus an easier computational description of the problem at hand.

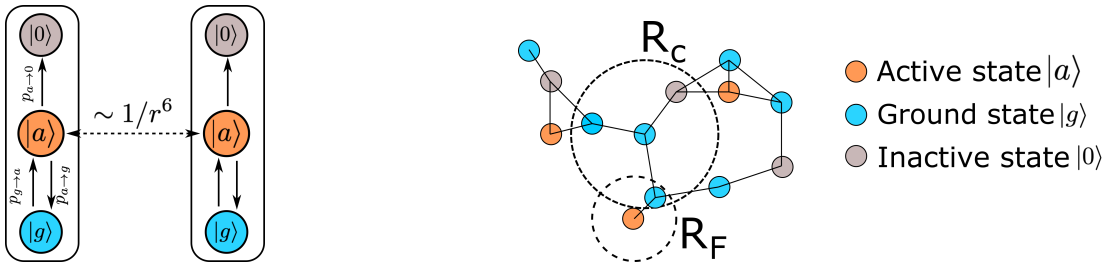

Figure 2. Schematics for the SOC problem construction. Figure 2(a): The model entities (atoms) have each three internal states. The ground state $|g\rangle$, the active state $|a\rangle$ and the inactive state $|0\rangle$. The atoms have a distance-dependent interaction scaling as $1/r^6$. Figure 2(b): When constructing the input graph there are two length scales that are relevant: the facilitation radius R_F denoting the relevant interaction distances and the cutoff radius R_C used to limit the number of edges in the graph to avoid all-to-all connectivity.

Note that throughout this study we always use training data for two-dimensions, where all atoms are placed into a square box in the xy-plane.

The training data is generated using classical Monte-Carlo methods (see Appendix B) and we create a large dataset covering a broad range of particle numbers and densities to expose the model to a diverse set of scenarios during training. Each training sample consists of X, a matrix of shape (N, N_f), and the ground truth Y, a matrix of shape $(N, 3)$. Here, N_f denotes the number of node features, which are the position and internal state of each node. Note that X is not directly the input to the GNN, since we replace the absolute positions with relative distances. The ground truth matrix $Y_i = (P_i(g|s_i), P_i(a|s_i), P_i(0|s_i))$ contains for each node i with the internal state s_i the probabilities to change to or remain in each of the three states in the next timestep. Before feeding a graph into the GNN, we apply the following physically motivated processing steps to facilitate easier learning for the model:

- We mask out all nodes in the inactive state $|0\rangle$, since they do not contribute to the dynamics and can be neglected without loss of accuracy.
- The edges between nodes are created by a radius-graph scheme: For two nodes to be connected their relative (Euclidean) distance needs to be smaller than a critical radius R_c, which is a hyperparameter for our model. Increasing R_c yields higher accuracy, but also increases the computational burden due to a quadratic rise in the number of neighbors. However, as the interaction drops off as $1/r^6$ the benefit of a large R_c is minimal. For all our calculations we use $R_c = 2$ (in natural units).
- Two nodes in the ground state $|g\rangle$ do not interact with each other, allowing us to remove edges between them. This causes no loss of accuracy and reduces the total number of edges in a graph by up to 80%, since the fraction of active atoms is typically lower than 15%.
- Each edge is given the Euclidean distance between its nodes as an edge attribute. The absolute position data is not encoded to encourage translation and rotational invariance of the model, since only relative distances are relevant for the SOC behavior.

Network architecture. To construct the network architecture we first note some characteristics of the probability distribution for the state transition probabilities. As the data is generated using Monte-Carlo simulations of a continuous process, the state transition probabilities are small in order to minimize the numerical error of the simulation. As a result, the probability of each atom to remain in its current state is large in order to preserve probability. This results in a bimodal distribution of the state transition probabilities with the actual transitions orders of magnitude smaller, i.e. $\mathcal{O}(10^{-7} - 10^{-1})$ than the probabilities to remain in the current state $\mathcal{O}(1)$ (For more details see Appendix B). However, the dynamics of the system are encoded in the changes in the state transition probabilities and without accounting for the sizable differences in the probabilities the model will simply learn the large background probabilities and not pay much attention to the subtle fluctuations in the transition probabilities that are caused by neighborhood interaction effects.

To address this challenge, we make use of a residual network architecture as depicted in Figure 3, where we give the network the opportunity to model the large background distribution, termed the *mean-field* contribution, and the neighborhood effects of the changes in the state transition probabilities, termed *fluctuations*. This terminology is in accordance with corresponding techniques in the traditional modelling of dynamic systems, where one tries to describe the system in terms of small variations, the fluctuations, around a stationary point, the mean-field. More specifically, we embed the state of each node into a higher-dimensional embedding space. Note that we removed the inactive state $|0\rangle$ from the data. We then feed each node state individ-

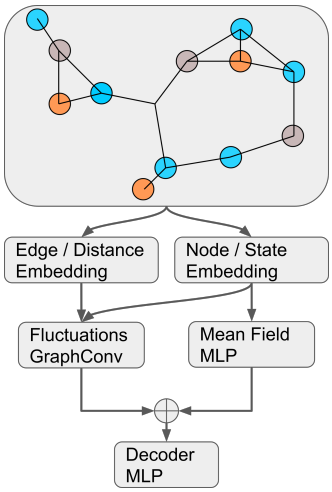

Figure 3. We use a residual architecture for the network structure to encourage the model to learn the background mean-field and the fluctuations on top of the mean field separately. To this end we feed the node embeddings into a simple MLP that learns the factorized node distribution. The fluctuations are learned using a Graph Convolution network that receives the node as well as edge embeddings as input. Finally, we add the mean-field and fluctuations before feeding them through the final MLP to predict the update probabilities.

ually through a simple feed forward network. Since this network does not take into account the edge attribute, it effectively learns only the factorized, non-correlated features of the nodes, which corresponds to the mean-field part of the dynamics. To capture the correlations, we combine the embedded node states and edges and feed them through a Message Passing Graph Convolution layer (Gilmer et al., 2017). We then add the outputs of the mean-field and fluctuation modules before passing them through a final MLP network to predict the state transition probabilities of each node. We use ReLU nonlinearities as activation funtions and normalize the activations using LayerNorm for the MLPs.

Loss function. As noted before the transition probability distribution changes over many orders of magnitude and hence we need a loss function that works over such a large interval. We experimented with mean-squared-error (MSE) and Kullback-Leibler divergence (KLD) and found the latter to be more stable. This can be motivated by the fact that the logarithmic dependency of the KLD penalizes differences in the exponents of the probabilities much more severely than the MSE loss that ignores such differences if the probabilities are small in favor of minimizing the differences for large probabilities.

5. Experiments & Results

Having established the modelling approach, we now turn to the experimental validation. We implement the GNN based on PyTorch (Paszke et al., 2019) using the PyTorch Lightning framework (Falcon & the PyTorch Lightning team, 2019) as well as PyTorch Geometric (Fey & Lenssen, 2019). The network consists of a $2D$ embedding module for the internal node state with an output dimension of 10, a state encoder network (mean-field component) consisting of 2 fully connected layers with hidden dimension of 50 and a final output dimension of 32, a Graph Convolution layer with a single message passing step and a final decoder with 2 fully connected layers and hidden dimension 50. The Graph Convolution uses the NNConv module from (Gilmer et al., 2017) that is implemented in PyTorch Geometric and has an output dimension of 64. The MLP inside of the NNConv has a single hidden layer with 50 units. Inside the fluctuations component we apply another MLP of 2 hidden layers with 100 units each and an output dimension of 32, before adding both components. In total the network has approximately 81000 parameters.

For all experiments we use the Adam optimizer (Kingma & Ba, 2015) with a learning rate of 10^{-3}, beta values of $(0.9, 0.999)$ and batch size 15. We set the cutoff radius $R_c = 2$ (in natural units) and train the model for 25 epochs with 6000 training instances each on an Intel® Xeon® Gold Processor 6154 with a peak memory usage of about 240MB. The training instances are drawn from three datasets that we created that all contain $N \approx 150$ atoms per instance but had varying densities. In two datasets we used $dx = 0.25$ and $dx = 0.5$, respectively. Here, we define dx as $dx = L/\sqrt{N}$, where L is the length of one edge of the square box in which we place the atoms. Intuitively, dx represents the distance between the atoms if they were placed in an ideal square lattice. In the third we artificially increased the fraction of active state atoms to 50% at a density of $dx = 1.25$ to address the fact that active atoms are otherwise rare in the training data. Finally, for reasons of numerical stability we transform the probabilities using a square-root transformation. This ensures the probabilities stay in the range $[0, 1]$ while mitigating the extreme bimodal nature of the ground-truth probability. See the Appendix B for more details.

Transition rates of internal node dynamics. To the best of our knowledge there is no comparable publication that predicts transition probabilities in SOC systems or considers the dynamics of internal degrees of freedom of a node. Hence we developed metrics to assess the quality and validate the learning of the GNN. We base these metrics on key aspects of the Rydberg-SOC dynamics.

The fastest process in the SOC system is the *facilitation* pro-

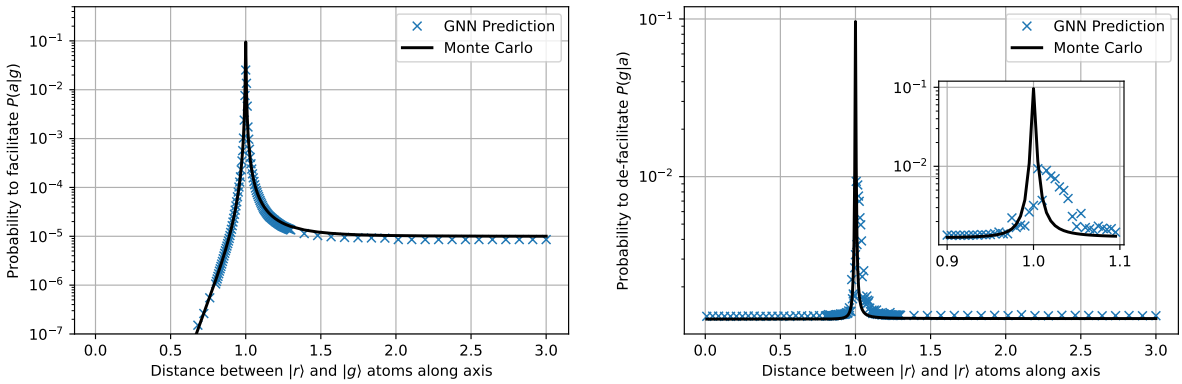

(a) Model prediction of the facilitation process of a ground state atom as well as the ground-truth curve from Monte Carlo.

(b) Model prediction of the facilitated de-excitation of an active atom as well as the ground-truth curve from Monte Carlo. The inset shows a magnification of the peak around $r = 1$.

Figure 4. Benchmark and validation curves to assess correct learning of key interaction features for the SOC dynamics in Rydberg atoms.

cess by which an active atom promotes surrounding ground state atoms to the active state, according to the resonance feature of Equation (6). The facilitation probability of a ground state atom at a distance r to an active atom is characterized by an extremely sharp peak at $r = 1$ (in natural units), several orders of magnitudes larger than the non-interacting ($r \to \infty$) case. Additionally, as $r \to 0$ the facilitation probability decreases rapidly to zero, due to the *Rydberg blockade* (Lukin et al., 2001) (also see (Saffman et al., 2010) for a review). For $r \gg R_F$, the facilitation probability tends to a small, but finite value. To check this behaviour we construct a custom graph where we place an active state atom at the origin and a ground state atom at the position $(x, y) = (r, 0)$. Additionally, we randomly place other ground state atoms in the graph to keep the total particle number and density close to those seen during training. The value we record is then $P_i(a|g)$ of the ground state atom to become active at position $(x, y) = (r, 0)$. We then sweep this setting across different values of r, creating a new random graph each time. The result of this experiment is shown in Figure 4(a) and shows excellent agreement with the exact solution of the facilitation rate around the facilitation peak $r = 1$ (in natural units) and in the $r \to \infty$ limit. In the blockade regime for $r \ll 1$ the GNNs prediction converges to $P_i(a|g) \sim 10^{-11}$, six orders of magnitude smaller than the limit $r \to \infty$. The true convergence to zero however is not captured. This is an expected result given the fact that even making large relative errors on very small ground-truth probabilities incurs a small loss penalty for the algorithm. Nonetheless, since the predicted blockaded facilitation probability is sufficiently small compared to all SOC timescales, we hypothesize that this deviation from the ground truth does not lead to qualitative changes in the GNNs predictions.

Task	Configuration	KLD
num. particles	N=50, dx=0.5, d=2	1.17e-4
	N=150, dx=0.5, d=2	1.50e-4
	N=500, dx=0.5, d=2	1.83e-4
	N=1000, dx=0.5, d=2	1.97e-4
	N=1500, dx=0.5, d=2	2.04e-4
density	N=150, dx=0.125, d=2	3.00e-5
	N=150, dx=0.25, d=2	8.17e-5
	N=150, dx=0.5, d=2	1.50e-4
	N=150, dx=1, d=2	7.54e-5
	N=150, dx=2.5, d=2	1.99e-5

Table 1. Test loss averaged over 2000 graphs each from different datasets, where either the particle number or density was varied. For the case of particle number generalization, we test on up to 10 times as many particles as seen during training and observe that the loss rises slowly with the particle number. For the case of varying densities we go a factor of 2 smaller than the smallest density seen during training (including the engineered dataset with more active atoms), as well as a factor of two above. In both cases, the loss is reduced compared to $dx = 0.5$, which we attribute to less facilitation events and more blockade (for higher density, low dx) or more non-interacting cases (for lower density, high dx).

The inverse process is called *de-facilitation* and is, to the best of our knowledge, a process unique to Rydberg atoms. In this case, a de-excitation (return to the ground state) of an active atom is facilitated by the presence of a second active atom. If these two active atoms are close to each other, the probability of one atom to return to the ground state increases sharply at the relative distance $r = 1$. Even though both the facilitation and the de-facilitation share the same physical origin, their ground-truth curves in Figures 4(a) and 4(b) look dissimilar on short distances. The is due to the second process that allows active atoms to return to

the ground state, which is independent of the neighborhood (see γ in eqn. (2)). We check this behaviour the same way as for the facilitation metric, but in this case the atom at position $(x, y) = (r, 0)$ is also active and we record $P_i(g|a)$. The results are shown in Figure 4(b). The defacilitation is much harder to learn as the case of two active atoms at close distance is very rare in the dataset, especially in the already very narrow region around $r = 1$ where the behavior changes drastically. We find that the GNN accurately predicts the non-resonant case as well as the position of the peak ar $r = 1$. However, the interval between $r \in [0.995, 1.005]$, where the ground truth values increase by a factor of 10, is not reproduced, as the GNNs prediction does not exceed 0.009, an order of magnitude short of the ground truth.

Particle number and density generalization. In addition to the SOC specific metrics we investigate the prediction accuracy on graphs consisting of more nodes, i.e. higher particle number, or denser, i.e. shorter average distance between particles, than those seen during training. The results are summarized in Table 1. We see that when trained on $N \approx 150$ atoms, the model suffers no significant loss of accuracy when predicting on a varying number of atoms up to $N = 1500$. The lowest test loss occurs at smaller N than during training, which might be the result of boundary effects that the model learns when trained on small system sizes. This scalability of the GNN has been observed in other particle-based regression tasks (Sanchez-Gonzalez et al., 2020) and is most likely due to the local scope of the message passing, which encodes spatial translation invariance of the dynamics operator.

In addition, we see that the test loss of the GNN predicting on datasets with different densities is also small across a range of average dx distances. During training, the model is confronted with densities $dx = 0.5$ and $dx = 0.25$ as well as artificial graphs with $dx = 1.25$. The test case of $dx = 0.5$ on an unseen dataset shows the highest test loss, which is most likely due to the fact that more facilitation events happen at this density than at lower or higher densities. In the case of lower densities, the chance of a ground state atom to be located in the small interval around $r = 1$ to an active atom is reduced, and in the case of larger densities the case of $r \ll 1$ occurs more frequently, where we see extremely small state-change probabilities (blockade).

SOC excitation trajectories. Having established that the model captures key aspects of the SOC dynamics and is able to generalize over particle numbers as well as densities, we can use the trained model to iteratively predict the time evolution of an initial ensemble of atoms. To this end we randomly create an initial 2D distribution of atoms, each in the ground state. Additionally, we assign a velocity to

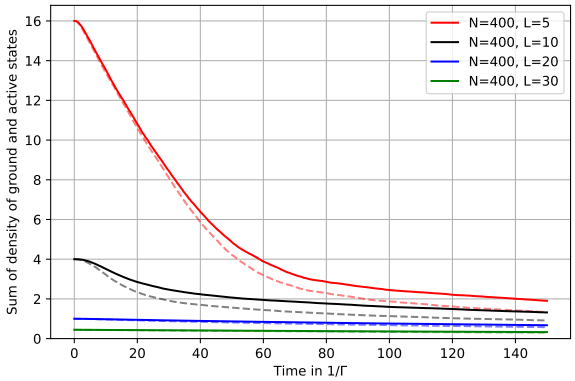

Figure 5. Comparison between the 60000-timestep predictions of the GNN (solid) and the Monte-Carlo simulation (dashed) from which the training data was obtained. We do not show all initial conditions in the interest of readability.

each atom which we sample from a Maxwell-Boltzmann distribution with the most probable velocity $\hat{v} = 1$. We input this initial state into the graph network and use the output probabilities to roll new internal states for each node and construct a new graph, where we record the fraction of atoms not in the inactive state. Then, we input the new graph into the GNN and repeat this procedure for $60,000$ steps. In each step, we update the positions of each node using the velocity: $\vec{x} \to \vec{x} + \vec{v}\delta t$. We use repelling boundary conditions, where the corresponding velocity component is reflected when a particle collides with the boundary. Lastly we average the resulting excitation trajectories of 50 runs with equivalent initial conditions (N, L). The number of steps is chosen such that the total predicted time T is equivalent to $T = 60,000\,\delta t = 150\Gamma^{-1}$ since it takes approximately $T = 150\Gamma^{-1}$ for the system to reach the critical state. Here $\Gamma = 1$ (the decay rate of the active state) serves as the natural unit of (inverse) time in the system. The results are shown in Figure 1. All curves that have an initial density above the critical density $n_c \approx 1$ (in the active phase) feature three distinct regimes: (i) a fast initial decay, (ii) a convergence to the same critical density n_c and (iii) a slow decay beyond that. All curves that start below the critical density (absorbing phase) exhibit only feature (iii). This is a hallmark characteristic of SOC behavior. When comparing the GNN predictions to the MCMC results in Figure 5 we see that the GNN shows a slower decay in phase (i). This effect is less pronounced the lower the initial density, which suggests the deviation stems from imperfect reproduction of the facilitation peak in Figure 4(a). Indeed, since the GNN underestimates the transition probability very close to $r = 1$, less facilitation events occur here than in the Monte Carlo baseline, causing a slower initial spread of excitations.

Discussion of the results. There are a few key implications we want to discuss based on the experimental results. We showed that the GNN learns the one-step time-evolution operator successfully and can display the hallmark characteristics of SOC dynamics. However, by construction, the GNN convolution operators are spatially translation invariant and local in nature. This implies that the long-range correlation of the SOC phase requires a time-dimension and is truly spatio-temporal as the spatial locality of the Graph Convolution does not allow to learn long-range spatial correlations.

A second aspect is that the spatial translation invariance of the Graph Convolution operator allows us to scale the time-evolution of the system to very large particle numbers, as all that matters is the local connectivity of the graph, not the overall particle number. Hence, the traditional $\mathcal{O}(N^2)$ scaling for MCMC does not apply.

Finally, the use of a graph data structure enables us to also scale over many densities, if the model has learned the facilitation metrics correctly. This is in contrast to standard simulations based on the discretization of the real-space coordinate. Due to the necessity of having a fine-grained grid, the simulated densities are typically small in that case in order to approximate continous systems. However, this limitation does not exist for GNNs as the relevant scale is the distance encoded in the edge attribute r as a fraction of the natural distance R_F. Taken together with the particle number scaling, it might allow the study of dense systems with large particle numbers.

6. Summary

We present a new method to model the time dynamics of SOC systems using machine learning on graphs. We use a physically motivated, lightweight GNN architecture to learn the one-step time-evolution operator of each atom of the SOC system. In contrast to previous studies, we extend the GNN simulation toolkit to learn the dynamics of internal degrees of freedom, i.e. the node-state, in addition to the external degrees of freedom, i.e. the distance between atoms. We demonstrate high accuracy of the predictions via excellent agreement of the learned facilitation rate compared to exact ground-truth rates. Additionally, we show that the model generalizes to at least one order of magnitude more particles than seen during training without significant loss of accuracy, as well as higher and lower densities than it was trained on. Moreover, we showed that the model can successfully reproduce the decay to a common critical density and the separation of time-scales in the excitation trajectories, a hallmark characteristic of SOC, despite having been trained on single-timestep data only.

We leave the extraction of critical exponents and other SOC elements as well as scaling to even larger particle numbers to future work. Additionally, our model could most likely benefit from hierarchical graph methods as presented in (Martinkus et al., 2021), which could reduce the computational complexity of larger systems. Lastly, the Monte Carlo training data that we used could be replaced by real observations of SOC systems. In that case, the GNN model could learn directly the interaction laws and circumvent the time discretization of differential equations, thereby providing an independent check of the underlying theory.

References

Bak, P., Tang, C., and Wiesenfeld, K. Self-organized criticality: An explanation of the 1/f noise. *Phys. Rev. Lett.*, 59:381–384, Jul 1987. doi: 10.1103/PhysRevLett.59.381. URL https://link.aps.org/doi/10.1103/PhysRevLett.59.381.

Battaglia, P. W., Pascanu, R., Lai, M., Rezende, D. J., and Kavukcuoglu, K. Interaction networks for learning about objects, relations and physics. *ArXiv*, abs/1612.00222, 2016.

Bronstein, M. M., Bruna, J., Cohen, T., and Veličković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. 2021. doi: 10.48550/ARXIV.2104.13478. URL https://arxiv.org/abs/2104.13478.

Drossel, B., Clar, S., and Schwabl, F. Exact results for the one-dimensional self-organized critical forest-fire model. *Phys. Rev. Lett.*, 71:3739–3742, Dec 1993. doi: 10.1103/PhysRevLett.71.3739. URL https://link.aps.org/doi/10.1103/PhysRevLett.71.3739.

Falcon, W. and the PyTorch Lightning team. Pytorch lightning, 2019. URL https://www.pytorchlightning.ai.

Fey, M. and Lenssen, J. E. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. *ArXiv*, abs/1704.01212, 2017.

Gleeson, J. P., Ward, J. A., O'Sullivan, K. P., and Lee, W. T. Competition-induced criticality in a model of meme popularity. *Phys. Rev. Lett.*, 112:048701, Jan 2014. doi: 10.1103/PhysRevLett.112.048701. URL https://link.aps.org/doi/10.1103/PhysRevLett.112.048701.

Helmrich, S., Arias, A., Lochead, G., Wintermantel, T. M., Buchhold, M., Diehl, S., and Whitlock, S. Signatures of self-organized criticality in an ultracold atomic gas. *Nature*, 577:481–486, 2020.

Henderson, R., Clevert, D.-A., and Montanari, F. Improving molecular graph neural network explainability with orthonormalization and induced sparsity. In *ICML*, 2021.

Jim'enez-Luna, J., Grisoni, F., and Schneider, G. Drug discovery with explainable artificial intelligence. *Nat. Mach. Intell.*, 2:573–584, 2020.

Jumper, J. M., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Zídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., Back, T., Petersen, S., Reiman, D. A., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A. W., Kavukcuoglu, K., Kohli, P., and Hassabis, D. Highly accurate protein structure prediction with alphafold. *Nature*, 596:583 – 589, 2021.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.

Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pp. 2688–2697. PMLR, 2018.

Lemos, P., Jeffrey, N., Cranmer, M., Ho, S.-L., and Battaglia, P. W. Rediscovering orbital mechanics with machine learning. *ArXiv*, abs/2202.02306, 2022.

Lukin, M. D., Fleischhauer, M., Cote, R., Duan, L., Jaksch, D., Cirac, J. I., and Zoller, P. Dipole blockade and quantum information processing in mesoscopic atomic ensembles. *Physical review letters*, 87(3):037901, 2001.

Malamud, Morein, and Turcotte. Forest fires: An example of self-organized critical behavior. *Science*, 281 5384: 1840–2, 1998.

Martinkus, K., Lucchi, A., and Perraudin, N. Scalable graph networks for particle simulations. In *AAAI*, 2021.

Noé, F., Tkatchenko, A., Müller, K.-R., and Clementi, C. Machine learning for molecular simulation. *Annual review of physical chemistry*, 2020.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. W. Learning mesh-based simulation with graph networks. *ArXiv*, abs/2010.03409, 2021.

Prakash, S. K. A. and Tucker, C. S. Graph network for learning bi-directional physics. 2021.

Rhodes, C. J. and Anderson, R. M. Power laws governing epidemics in isolated populations. *Nature*, 381:600–602, 1996.

Saffman, M., Walker, T. G., and Mølmer, K. Quantum information with rydberg atoms. *Reviews of modern physics*, 82(3):2313, 2010.

Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. W. Learning to simulate complex physics with graph networks. *ArXiv*, abs/2002.09405, 2020.

Scully, M. O. and Zubairy, M. S. *Quantum Optics*. Cambridge University Press, 1 edition, September 1997. ISBN 9780521435956 9780521434584 9780511813993. doi: 10.1017/CBO9780511813993. URL https://www.cambridge.org/core/product/identifier/9780511813993/type/book.

Shlomi, J., Battaglia, P. W., and Vlimant, J. R. Graph neural networks in particle physics. *Mach. Learn. Sci. Technol.*, 2:21001, 2021.

Sornette, A. and Sornette, D. Self-organized criticality and earthquakes. *EPL*, 9:197–202, 1989.

Wintermantel, T., Buchhold, M., Shevate, S., Morgado, M., Wang, Y., Lochead, G., Diehl, S., and Whitlock, S. Epidemic growth and griffiths effects on an emergent network of excited atoms. *Nature Communications*, 12 (1):1–6, 2021.

A. Derivation of Rate Equations from Optical Bloch Equations / State-transition description of a Driven Rydberg Gas

In this section, we shortly summarize the derivation of the rate equations discussed in the main text. To help readers draw parallels with the physics literature, we will use the subscript r for the active (rather: Rydberg) state and the subscript g for the ground state. The derivation considers the density matrix ρ for a two-level system consisting of the ground and Rydberg states that is coupled to an external light field which we assume to be classical (not in the single-photon limit). Since we take the full quantum mechanics of the two-level atom into account, this is called the *semi-classical* approach. For a two level system with atom-light interaction in the semi-classical limit, the optical Bloch equations are then given as (Scully & Zubairy, 1997)

$$\dot{\rho}_{rr} = -i\Omega(\rho_{rg} - \rho_{gr}) - \Gamma\rho_{rr}, \tag{8}$$

$$\dot{\rho}_{rg} = -\gamma\rho_{rg} + i\Delta\rho_{rg} - i\Omega(\rho_{rr} - \rho_{gg}), \tag{9}$$

where ρ_{ij} represents the entries of the density matrix. For large dephasing γ, $\dot{\rho}_{rg}$ can be adiabatically eliminated. This gives

$$\rho_{rg} = \frac{i\Omega(\rho_{rr} - \rho_{gg})}{-\gamma + i\Delta} \tag{10}$$

$$\dot{\rho}_{rr} = -\frac{2\Omega^2\gamma}{\gamma^2 + \Delta^2}(\rho_{rr} - \rho_{gg}) - \Gamma\rho_{rr} \tag{11}$$

$$\equiv -\Gamma_{\uparrow}\rho_{rr} + \Gamma_{\downarrow}\rho_{gg}. \tag{12}$$

This gives the transition rates $\Gamma_{\uparrow}$ from the ground to excited state and $\Gamma_{\downarrow}$ from the excited to the ground state. To simulate these rates, we choose the microscopic parameters in such a way as to achieve a sufficient separation of timescales, see table 2. Using these values, we find the following separation of timescales:

parameter	value
Δ/Γ	2000
Ω/Γ	20
Γ	1 (Unit of time)
γ/Γ	20
$\delta t\Gamma$	0.0025
b	0.5

Table 2. Microscopic parameters used in the creation of Monte-Carlo data. We use units where $\hbar = 1$.

$$\Gamma_{\text{facil}} = \frac{2\Omega^2}{\gamma} = 40\Gamma \tag{13}$$

$$b\Gamma = 0.5\Gamma, \tag{14}$$

$$\Gamma_{\text{non-int}} = \frac{2\Omega^2\gamma}{\gamma^2 + \Delta_0^2} \approx 0.004\Gamma, \tag{15}$$

$$\Gamma_{\text{facil}} = 80\, b\Gamma \tag{16}$$

$$b\Gamma \approx 125\Gamma_{\text{non-int}}. \tag{17}$$

To observe SOC behavior, a separation of approximately a factor of 100 between each rate is recommended.

B. Monte-Carlo simulation of the Training Data

B.1. Creation of Training Data

As outlined in the main text, the training data for the GNN is generated using classical Monte-Carlo methods where each training instance consists of X, a matrix of shape (N, N_f), and the ground truth Y, a matrix of shape $(N, 3)$. $Y_{i,s} = (P_i(g|s_i), P_i(a|s_i), P_i(0|s_i))$ contains for each node i with the internal state s_i the probabilities to change to or remain in each of the three states in the next timestep.

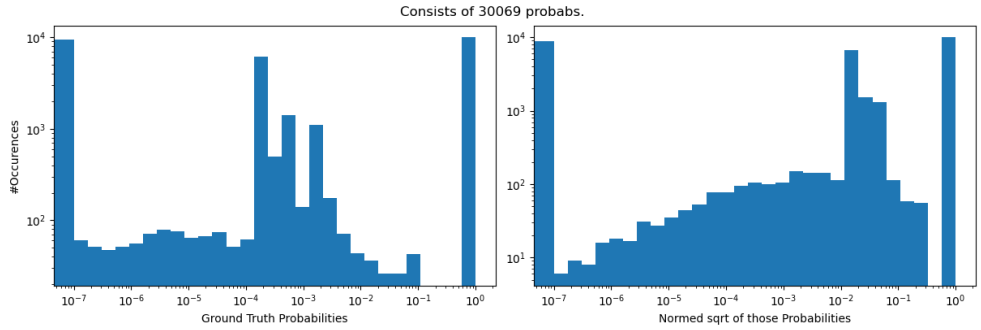

Figure 6. Left: Histogram over the target probabilities. The leftmost bin contains all values smaller than 10^{-7}. One notices the peaks at 0.9 and higher and at around 10^{-3}. Right: Histogram of the ground truth probabilities after the square-root transformation.

In each training instance, even though the total number of particles is fixed for every dataset, the distribution of the three internal node states may be different. Since over the course of a SOC simulation (see Figure 1) the fractions of ground, active and inactive states change drastically, we set bounds for the fraction of each state and draw instances at random. These bounds are

$$p(g) = [0.15, 1.0] \tag{18}$$
$$p(a) = [0, 0.15] \tag{19}$$
$$p(0) = [0, 0.9]\,. \tag{20}$$

Since the $|0\rangle$ atoms are masked out at every step, this leads to a varying particle number for each training instance.

B.2. Square Root Transformation

In defining a model architecture, we are guided by the expected distribution of transition probabilities $P(i|s)$, where $i, s \in \{\text{ground}, \text{active}, \text{inactive}\}$ states. The training data stems from fixed-timestep Monte Carlo, where the timestep has to be chosen to be small against the inverse rate at which the states of the atoms change, such that the error stemming from time discretization is small. As a result, the probability of a node to remain in its current state is always large, independent of the states of neighboring nodes. While neighborhood-effects can increase the probability of state-change by factors of ten-thousand, the absolute probability of any state-change is nonetheless $< 10\%$. Additionally, the SOC system imposes further constraints on the transition probabilities. To observe critical behaviour in the SOC regime, the three timescales in the system need to be separated by at least a factor of 100 each (see Appendix A). Since the time increment δt must be chosen sufficiently short such that the transition probability connected to the largest rate is about 10^{-1}, the transition probability connected to the smallest rate can maximally be 10^{-5} (see Figure 4(a) for $r \to \infty$). Lastly, to accurately capture the Rydberg blockade effect for small inter-atomic distances (see Figure 4(a) for $r \to 0$), we need to model transition probabilities to at least two orders of magnitudes below the slowest transition rate at maximally 10^{-7}. This results in a bi-modal target distribution, where most values are in the interval $[0.9, 1.0]$ as well as in the interval $\left[10^{-4}, 0.1\right]$, including a long tail for even smaller probabilities as shown in 6. Learning probabilities as small as 10^{-7} is very challenging, as noise during training becomes significant for such small numbers. To avoid this, we decided to perform a transformation on the ground truth data that increases the magnitude of very small numbers but keeps large ones effectively unchanged, while still producing numbers in the interval $[0, 1]$ as required by the loss function. For this purpose we chose the square root function, which we apply node-wise on all three probabilities before normalizing their sum to unity. To obtain real predictions from the model, it is then necessary to take the normalized square from the model's output.

We found that this transformation generally increased training stability as well as performance on the (de-)facilitation metrics.