

---

# Learning Dissipative Dynamics in Chaotic Systems

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Chaotic systems are notoriously challenging to predict because of their sensitivity to  
2 perturbations and errors due to time stepping. Despite this unpredictable behavior,  
3 for many dissipative systems, the long term trajectories converge to an invariant  
4 measure supported on a low dimensional set, known as the global attractor. For  
5 Markovian systems, the statistical properties of long-term trajectories are uniquely  
6 determined by a Markov operator that maps the evolution of the system over an  
7 infinitesimal time step. In this work, we propose a machine learning framework to  
8 learn the underlying Markov operator of dissipative chaotic systems which captures  
9 their statistical behavior without the need to predict the exact trajectories. Using  
10 this framework, for the first time, we are able to predict various statistics of the  
11 invariant measure for the turbulent Kolmogorov Flow dynamics with Reynolds  
12 numbers up to 500.

## 13 1 Introduction

14 **Machine learning methods for chaotic systems.** Chaotic systems are characterized by strong  
15 instabilities. Small changes in the initialization or errors during time-stepping accumulate and  
16 lead to vastly diverging trajectories. Such instability makes chaotic systems challenging, both for  
17 mathematical analysis and numerical simulation. Because of the intrinsic instability, it is infeasible  
18 for any method to capture the exact trajectory of a chaotic system for long periods. Therefore, prior  
19 works either fit recurrent neural networks (RNN) on extremely short trajectories or only learn a  
20 step-wise projection from a randomly generated evolution using reservoir computing (RC) [1–4].  
21 These previous attempts are able to push the limits of faithful prediction to moderate periods on  
22 low dimensional ordinary differential equations (ODEs), e.g. the Lorenz-63 system, or on one-  
23 dimensional partial differential equations (PDEs), e.g. the Kuramoto-Sivashinsky (KS) equation.  
24 However, they are less effective at modeling more complicated turbulent systems such as the Navier-  
25 Stokes equation (NS), especially over long time periods. Indeed, predicting long trajectories of such  
26 chaotic systems is an ill-posed problem and we cannot expect such attempts to be successful. Instead,  
27 we take a new perspective: we aim to capture statistical properties of long trajectories, even if we  
28 cannot precisely predict them.

29 **Invariants in chaos.** Despite their instability, many chaotic systems exhibit certain reproducible  
30 statistical properties, such as the auto-correlation and, for PDEs, the energy spectrum. Such properties  
31 remain the same for different realizations of the initial condition [5]. This is provably the case for the  
32 Lorenz-63 model [6, 7] and empirically holds for many dissipative PDEs, such as the KS equation  
33 and the two-dimensional Navier-Stokes equation (Kolmogorov flows) [8]. Dissipativity is a physical  
34 property of many natural systems. Intuitively higher-energy flows in such systems dissipate more  
35 strongly. Mathematically, there exists a global **attractor** which is a compact set towards which  
36 the system tends to evolve. The dissipativity property implies that for a given system there is a set  
37 which any given trajectory enters in finite time, depending on the initial condition, and thereafter

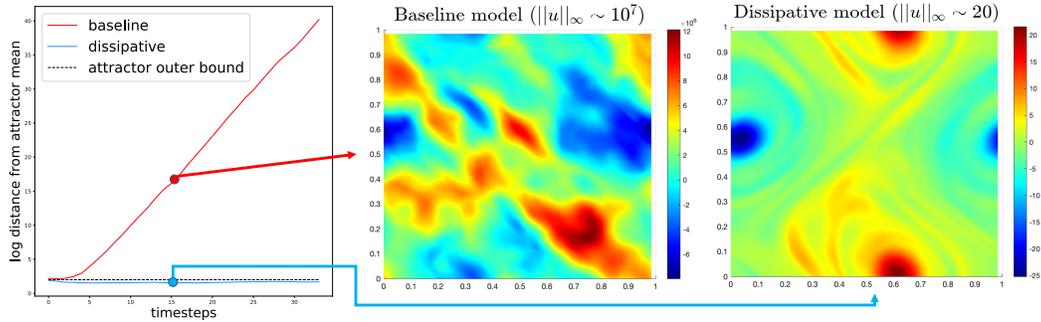


Figure 1: Dynamic evolution of the Markov neural operator for Kolmogorov Flow systems, from initial conditions near the attractor, with and without the enforced dissipativity constraints. The baseline model in the middle has no dissipativity constraint while the dissipative model on the right has the constraint enforced during its time evolution. Baseline model blows up, whereas the dissipative model returns to the attractor. The dissipative models is trained using the Fourier neural operator architecture in the manner shown in Figure 2.

38 remains inside. The global attractor is defined by mapping all initial conditions from bounded sets  
 39 forward in time. Furthermore there is strong empirical evidence that many dissipative systems are  
 40 ergodic i.e. there exists an invariant measure which charges the global attractor. While learning  
 41 infinite time-horizon trajectories is intractable, it is possible to approximate the attractor and invariant  
 42 measure using a **Markov operator**, a memoryless deterministic operator that captures the evolution  
 43 of the dynamics along an infinitesimal time step. The dissipativity property helps to make this  
 44 problem tractable [9, 10]. For Markovian systems, e.g. autonomous differential equations, samples  
 45 from the invariant measure can be obtained through repeatedly composing the Markov operator.  
 46 This property is implied since the family of Markov operators defined for any fixed time forms a  
 47 semigroup [9]. By learning a Markov operator, we are able to quickly and accurately generate an  
 48 approximate attractor and estimate its invariant measure for a variety of chaotic systems that are of  
 49 interest to the physics and applied mathematics communities [11–17].

50 **Neural operators.** To learn the Markov operators for PDEs, we need to model the time-evolution  
 51 of functions in infinite-dimensional function spaces. This is especially challenging when we need  
 52 to generate long trajectories since even a small error accumulates over multiple compositions of  
 53 the learned operator, potentially causing an exponential build-up or a collapse due to the high  
 54 dimension of the space. Because we study the evolution of functions in time, we propose to use  
 55 the recently developed operator learning method known as the neural operator [18, 19]. The neural  
 56 operator remedies the mesh-dependent nature of finite-dimensional operator methods such as RNNs,  
 57 CNNs, and RC. Neural operators are guaranteed to universally approximate any operator in a  
 58 mesh independent manner, and hence, can capture the Markov operator of chaotic systems. This  
 59 approximation guarantee and the absorption of trajectories by the global attractor makes it possible to  
 60 accurately follow it over long time horizons, allowing us access to the invariant measure of chaotic  
 61 systems.

62 **Our contributions.** In this work, we formulate a machine learning framework for chaotic systems  
 63 exploiting their dissipativity and Markovian properties. We propose the Markov neural operator  
 64 (MNO) and train it given only one-step evolution data from a chaotic system. By composing the  
 65 learned operator over a long horizon, we accurately approximate the global attractor of the system [20].  
 66 Our architecture is outlined in Figure 2. In order to assess its performance, we study the statistics of  
 67 the associated invariant measure such as the Fourier spectrum, the spectrum of the proper orthogonal  
 68 decomposition (POD), the point-wise distribution, the auto-correlation, and other domain-specific  
 69 statistics such as the turbulence kinetic energy and the dissipation rate. Furthermore we study the  
 70 behavior of our leaned operator over long horizons and ensure that it does not blow up or collapse but  
 71 rather accurately follows the global attractor. In this work:

- 72 • We theoretically prove that, under suitable conditions, the MNO can approximate the underlying  
 73 Markov operator of chaotic PDEs, while conventional neural networks lack such strong guarantees.

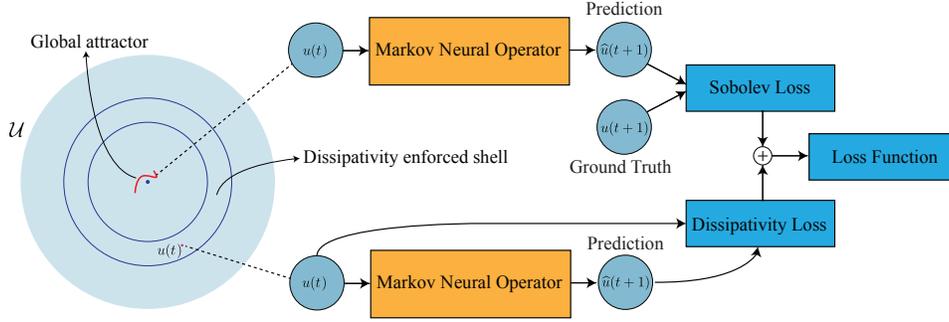


Figure 2: Markov neural operator (MNO): learn global dynamics from local data  
 Learn the MNO from the local time-evolution data with the Sobolev and dissipativity losses .  $u(t)$  is  $t$ 'th time step of the chaotic system. Sobolev losses of various order are used to compute the loss of prediction next time step  $\hat{u}(t+1)$  as of  $u(t+1)$ . Dissipativity loss is computed by drawing a random sample  $u(t)$  from the dissipativity shell to make sure that in expectation next time step prediction  $\hat{u}(t+1)$  dissipates and has a smaller norm as of  $v$ .

- 74 • We impose dissipativity by augmenting the data on an outer shell to enforce that the dynamic  
 75 evolution stays close to the attractor. We show this is a crucial for learning in a chaotic regime as  
 76 demonstrated by Figure 1. The resulting system remains stable against large perturbations.
- 77 • We study the choice of time steps for training the MNO, demonstrating that the error follows  
 78 a valley-shaped phenomenon. This gives rise to a recipe for choosing the optimal time step for  
 79 accurate learning.
- 80 • We show that standard mean square error (MSE) type losses for training are not adequate, and the  
 81 models often fail to capture the higher frequency information induced from the derivatives.
- 82 • We investigate various Sobolev losses in operator learning. We show that using the Sobolev  
 83 norms for training captures higher-order derivatives and moments, as well as high frequency  
 84 details [21, 22]. This is similar in spirit to pre-multiplying the spectrum by the wavenumber, an  
 85 approach commonly used in fluid analysis [23].
- 86 • We investigate multiple exiting deep learning architectures, including U-Net [24], long short-term  
 87 memory convolution neural networks (LSTM-CNN) [25], and gated recurrent unit (GRU) [26], in  
 88 place of the neural operator to learn the predictive operator. We show MNO provides an order of  
 89 magnitude lower error on all loss functions studied in this paper. Furthermore, we show that the  
 90 MNO desirably outperforms the above mentioned neural network models previous models, on all  
 91 statistics mentioned.

92 Our experiments show that Sobolev norms are crucial in training the MNO, allowing for significant  
 93 performance improvement over standard loss functions (e.g. MSE) and for preserving invariant  
 94 statistics such as the spectrum. Thus we propose a principled approach for learning chaotic systems  
 95 that incorporates operator learning and dissipativity.

## 96 2 Problem setting

97 We consider potentially infinite dimensional dynamical systems where the phase space  $\mathcal{U}$  is a Banach  
 98 space and, in particular, a function space on a Lipschitz domain  $D \subset \mathbb{R}^d$  (for finite dimensional  
 99 systems,  $\mathcal{U}$  will be a Euclidian space). We are interested in the initial-value problem

$$\frac{du}{dt}(t) = F(u(t)), \quad u(0) = u_0, \quad t \in (0, \infty) \quad (1)$$

100 for initial conditions  $u_0 \in \mathcal{U}$  where  $F$  is usually a non-linear operator. We will assume, given some  
 101 appropriate boundary conditions on  $\partial D$  when applicable, the solution  $u(t) \in \mathcal{U}$  exists and is unique  
 102 for all times  $t \in (0, \infty)$ . When making the spatial dependence explicit, if it is present, we will  
 103 write  $u(x, t)$  to indicate the evaluation  $u(t)|_x$  for any  $x \in D$ . We define the family of operators

104  $S_t : \mathcal{U} \rightarrow \mathcal{U}$  as mapping  $u_0 \mapsto u(t)$  for any  $t \geq 0$ , and note that, since (1) is autonomous,  $S_t$  satisfies  
 105 the Markov property i.e.  $S_t(S_s(u_0)) = u(s+t)$  for any  $s, t \geq 0$ . We adopt the viewpoint of casting  
 106 time-dependent PDEs into function space ODEs (1), as this leads to the semigroup approach to  
 107 evolutionary PDEs which underlies our learning methodology.

108 **Dissipativity.** Systems for which there exists some bounded, positively-invariant set  $E$  such that for  
 109 any bounded  $B \subset \mathcal{U}$ , there is some time  $t^* = t^*(E, B)$  beyond which the dynamics of any trajectory  
 110 starting in  $B$  enters and remains in  $E$  are known as **dissipative systems** ([9]). The set  $E$  is known as  
 111 the **absorbing set** of the system. For such systems, the global attractor  $A$ , defined subsequently, is  
 112 characterized as the  $\omega$ -limit set of  $E$ . In particular, for any initial condition  $u_0 \in \mathcal{U}$ , the trajectory  
 113  $u(t)$  approaches  $A$  as  $t \rightarrow \infty$ . In this work, we consider dissipative dynamical systems where there  
 114 exist some  $\alpha \geq 0$  and  $\beta > 0$  such that

$$\frac{1}{2} \frac{d}{dt} \|u\|^2 \leq \alpha - \beta \|u\|^2 \quad (2)$$

115 for all  $u \in \mathcal{U}$ . It can be shown that systems which satisfy this inequality are dissipative ([9]) with  
 116 the absorbing set  $E$ , an open ball of radius  $\sqrt{\alpha/\beta + \varepsilon}$  for any  $\varepsilon > 0$ . There are several well-known  
 117 examples of dynamical systems that satisfy the above inequality. In this paper we consider the  
 118 finite-dimensional Lorenz-63 system and the infinite-dimensional cases of the Kuramoto-Sivashinsky  
 119 and 2D incompressible Navier-Stokes equations, in the form of Kolmogorov flows ([8]).

120 **Global Attractors.** The long time behavior of the solution to (1) is characterized by the set  
 121  $U = U(u_0) \subset \mathcal{U}$  which is **invariant** under the dynamic i.e.  $S_t(U) = U$  for all  $t \geq 0$ , and the orbit  
 122  $u(t)$  converges

$$\inf_{v \in U} \|u(t) - v\|_{\mathcal{U}} \rightarrow 0 \quad \text{as} \quad t \rightarrow \infty.$$

123 When it exists,  $U$  is often identified as the  $\omega$ -limit set of  $u_0$ . The chaotic nature of certain dynamical  
 124 systems arises due to the complex structure of this set because  $u(t)$  follows  $U$  and  $U$  can be, for  
 125 example, a fractal set. A compact, invariant set  $A$  is called a **global attractor** if, for any bounded  
 126 set  $B \subset \mathcal{U}$  and any  $\epsilon > 0$  there exists a time  $t^* = t^*(\epsilon, A, B)$  such that  $S_t(B)$  is contained within  
 127 an  $\epsilon$ -neighborhood of  $A$  for all  $t \geq t^*$ . Many PDEs arising in physics such as reaction-diffusion  
 128 equations describing chemical dynamics or the Navier-Stokes equation describing the flow of fluids  
 129 are dissipative and possess a global attractor which is often times finite-dimensional [8]. Therefore,  
 130 numerically characterizing the attractor is an important problem in scientific computing with many  
 131 potential applications.

132 **Data distribution.** For many applications, an exact form for the possible initial conditions to (1)  
 133 is not available; it is therefore convenient to use a stochastic model to describe the initial states. To  
 134 that end, let  $\mu_0$  be a probability measure on  $\mathcal{U}$  and assume that all possible initial conditions to (1)  
 135 come as samples from  $\mu_0$  i.e.  $u_0 \sim \mu_0$ . Then any possible state of the dynamic (1) after some time  
 136  $t > 0$  can be thought of as being distributed according to the pushforward measure  $\mu_t := S_t^\# \mu_0$  i.e.  
 137  $u(t) \sim \mu_t$ . Therefore as the dynamic evolves, so does the type of likely functions that result. This  
 138 further complicates the problem of long time predictions since training data may only be obtained  
 139 up to finite time horizons hence the model will need the ability to predict not only on data that is  
 140 out-of-sample but also out-of-distribution.

141 **Ergodic systems.** To alleviate some of the previously presented challenges, we consider **ergodic**  
 142 systems. Roughly speaking, a system is ergodic if there exists an **invariant measure**  $\mu$  such that after  
 143 some time  $t^* > 0$ , we have  $\mu_t \approx \mu$  for any  $t \geq t^*$  (in fact,  $\mu$  can be defined without any reference to  
 144  $\mu_0$  or its pushforwards, see [27] for details). That is, after some large enough time, the distribution of  
 145 possible states that the system can be in is fixed for any time further into the future. Indeed,  $\mu$  charges  
 146 the global attractor  $A$ . Notice that ergodicity is a much more general property than having **stationary**  
 147 **states** which means that the system has a fixed period in time, or having **steady states** which means  
 148 the system is unchanged in time.

149 Ergodicity mitigates learning a model that is able to predict out-of-distribution since both the input  
 150 and the output of  $\hat{S}_h$ , an approximation to  $S_h$ , will approximately be distributed according to  $\mu$ .  
 151 Furthermore, we may use  $\hat{S}_h$  to learn about  $\mu$  since sampling it simply corresponds to running the  
 152 dynamic forward. Indeed, we need only generate data on a finite time horizon in order to learn  $\hat{S}_h$ ,

153 and, once learned, we may use it to sample  $\mu$  indefinitely by simply repeatedly composing  $\hat{S}_h$  with  
 154 itself. Having samples of  $\mu$  then allows us to compute statistics which characterize the long term  
 155 behavior of the system and therefore the global attractor  $A$ . Note further that this strategy avoids the  
 156 issue of accumulating errors in long term trajectory predictions since we are only interested in the  
 157 property that  $\hat{S}_h(u(t)) \sim \mu$ .

158 Notably, the existence of a global attractor does not imply the existence of an invariant measure.  
 159 Indeed, the only deterministic and chaotic systems that are proven to possess an invariant measure  
 160 are certain ODEs such as the Lorenz-63 system [7]. On the other hand, proving the existence of an  
 161 invariant measure for deterministic and chaotic PDEs such as the KS or KF equations are still open  
 162 problems, despite ergodic behavior being observed empirically.

### 163 3 Learning the Markov neural operator in chaotic dynamics

164 We propose the Markov neural operator, a method for learning the underlying Markov operators of  
 165 chaotic dynamical systems. In particular, we approximate the operator mapping the solution from  
 166 the current step to the next step  $\hat{S}_h : u(t) \mapsto u(t+h)$ . We approximate the Markov operator  $S_h$ , an  
 167 element of the underlying continuous time semigroup  $\{S_t : t \in [0, \infty)\}$ , using a neural operator as  
 168 detailed in Figure 2. See Appendix B.1 for background on the Markov operator and semigroup.

169 **Long-term predictions.** Having access to the map  $\hat{S}_h$ , its semigroup properties allow for approxi-  
 170 mating long time trajectories of (1) by repeatedly composing  $\hat{S}_h$  with its own output. Therefore, for  
 171 any  $n \in \mathbb{N}$ , we compute  $u(nh)$  as follows,

$$u(nh) \approx \hat{S}_h^n(u_0) := \underbrace{(\hat{S}_h \circ \dots \circ \hat{S}_h)}_{n \text{ times}}(u_0). \quad (3)$$

172 The above semigroup formulation can be applied with various choices of the backbone model for  $\hat{S}_h$ .  
 173 In general, we prefer models that can be evaluated quickly and have approximation guarantees so the  
 174 per-step error can be controlled. Therefore, we choose the standard feed-forward neural network [28]  
 175 for ODE systems, and the Fourier neural operator [29] for infinite dimensional PDE systems.

176 For the the neural operator parametric class, we prove the following theorem regarding the Markov  
 177 neural operator. The result states that our construction can approximate trajectories of infinite-  
 178 dimensional dynamical systems arbitrary well. The proof is given in the Appendix.

179 **Theorem 1.** *Let  $K \subset \mathcal{U}$  be a compact set and assume that, for some  $h > 0$ , the Markov operator*  
 180  *$S_h : \mathcal{U} \rightarrow \mathcal{U}$  associated to the dynamic (1) is locally Lipschitz. Then, for any  $n \in \mathbb{N}$  and  $\epsilon > 0$  there*  
 181 *exists a neural operator  $\hat{S}_h : \mathcal{U} \rightarrow \mathcal{U}$  such that*

$$\sup_{u_0 \in K} \sup_{k \in \{1, \dots, n\}} \|u(kh) - \hat{S}_h^k(u_0)\|_{\mathcal{U}} < \epsilon.$$

182 Theorem 1 indicates that of choice of backbone model is rich enough to approximate many chaotic  
 183 dynamical systems for a arbitrarily long period. For finite-dimensional systems, the same theorem  
 184 holds with feed-forward neural networks instead of neural operators. We note that standard neural  
 185 networks such as RNNs and CNNs *do not possess* such approximation theorems in the infinite-  
 186 dimensional setting.

187 **Invariant statistics.** A useful application of the Markov operators is to estimate statistics of the  
 188 invariant measure of a chaotic system. Assume the target system is ergodic and there exists an  
 189 invariant measure  $\mu$  such that  $u(t) \sim \mu$  for any  $t$  as discussed in Section 2. An invariant statistic is  
 190 defined as

$$T_G := \int_{\mathcal{U}} G(u) d\mu(u) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T G(u(t)) dt \quad (4)$$

191 for any functional  $G : \mathcal{U} \rightarrow \mathbb{R}^d$ . Examples include the  $L^2$  norm, any spectral coefficients, and the  
 192 spatial correlation, as well as problem-specific statistics such as the turbulence kinetic energy and  
 193 dissipation rates in fluid flow problems. Given the property (3) and using the ergodicity from (4), the  
 194 approximate model  $\hat{S}_h$  can be used to estimate any invariant statistic simply by computing  $T_G \approx$   
 195  $\frac{h}{T} \sum_{k=1}^n G(\hat{S}_h^k(u_0))$  for some  $n = T/h$  and  $T > 0$  large enough. Examples of fast approximation  
 196  $\hat{S}_h$  which accurately predict invariant statics are given in Section 4.

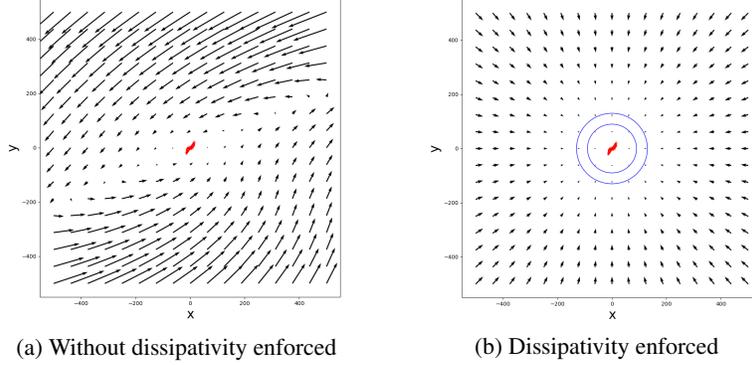


Figure 3: Enforcing dissipativity on the Lorenz 63 system – extrapolated flow maps. The red points are training data on the attractor. The dissipativity is imposed by augmenting the data on the blue shell. The dissipativity enforcing training results in a learned dissipative Markov operator.

197 **Enforcing dissipativity.** In practice, training data for dynamical systems is typically drawn from  
 198 trajectories that lie close to the global attractor of the system, so a priori there is no guarantee of a  
 199 learned model’s behavior far from the attractor. Thus, if we seek to learn the global attractor and  
 200 invariant statistics of a dynamical system, it is crucial that we place constraints on the model to enforce  
 201 that it learns a dissipative dynamical system.

202 In particular, given some Markov operator mapping between time-steps  $\hat{S}_h : u(t) \mapsto u(t + h)$  and  
 203 cost functional  $C_D : \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}$ , we supplement the loss function with the additional term

$$\mathbb{E}_{u \sim \nu} \left[ C_D \left( \hat{S}_h(u), u \right) \right] = \int_{\mathcal{U}} \|\hat{S}_h(u) - \lambda u\|_{\mathcal{U}}^2 d\nu(u), \quad (5)$$

204 up to some multiplicative constant with respect to the other terms in the loss function, where  
 205  $0 < \lambda < 1$  is some constant factor for scaling down (i.e., enforcing dissipativity) inputs  $u$  drawn  
 206 from a probability measure  $\nu$ . We choose  $\nu$  to be a uniform probability distribution supported on  
 207 some shell with a fixed inner and outer radii from the origin in  $\mathcal{U}$ . Our choice of cost functional  $C_D$   
 208 as given in eq. 5 scales down  $u$  by some constant factor  $\lambda$ , but in principle alternative dissipative cost  
 209 functionals can be used.

210 We find that enforcing this dissipativity constraint on a shell at a sufficiently large radius encourages  
 211 the learned Markov operator to produce dissipative predictions arbitrarily far away from the shell. In  
 212 Section 4 we demonstrate that the constraint in eq. 5 prevents blow-up of a Markov operator trained  
 213 on the turbulent Kolmogorov flow system.

## 214 4 Experiments

215 We evaluate the efficacy of our approach on the finite-dimensional, chaotic Lorenz-63 system as well  
 216 as chaotic regimes of the 1D Kuramoto-Sivashinsky and 2D Navier-Stokes equations. In all cases we  
 217 show that enforcing dissipativity is crucial for capturing the global attractor and evaluating statistics  
 218 of the invariant measure. To the best of our knowledge, we showcase the first machine learning  
 219 method able to predict the statistical behavior of a highly turbulent regime of the Navier-Stokes.

### 220 4.1 Lorenz-63 system

221 To motivate and justify our framework for learning chaotic systems in the infinite-dimensional setting  
 222 (e.g., Navier-Stokes equations), we first apply our framework on the simple yet still highly chaotic  
 223 Lorenz-63 ODE system.

224 The Lorenz-63 system constitutes a simplified climate model and is described by the following ODEs,

$$\dot{u}_x = \alpha(u_y - u_x), \quad \dot{u}_y = -\alpha u_x - u_y - u_x u_z, \quad \dot{u}_z = u_x u_y - b u_z - b(r + \alpha). \quad (6)$$

225 In this paper, we use the canonical parameters  $(\alpha, b, r) = (10, 8/3, 28)$  [30]. Since the Markov  
 226 operator of the Lorenz-63 system is finite-dimensional, we learn the Markov operator by training a

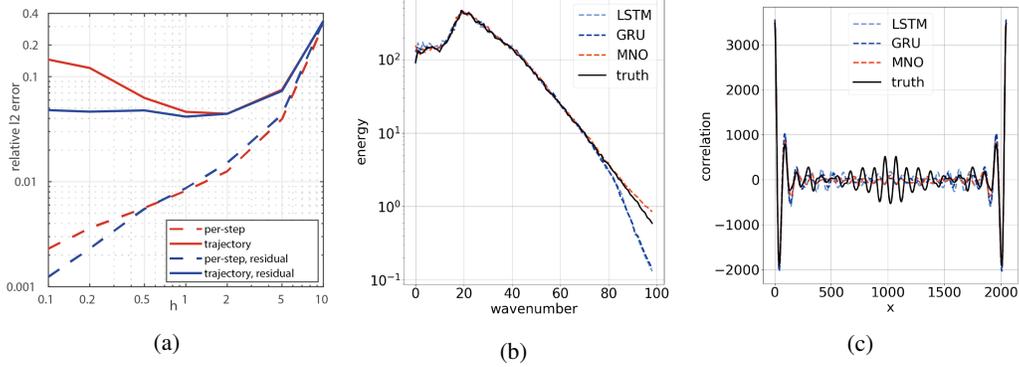


Figure 4: (a) Choice of time step for directly learning the Markov operator and learning its identity residual. We observe that both models induce smaller per step error for smaller  $h$ . When learn models are composed to generate longer trajectories, we observe that learning residual is advantageous. (b) Fourier spectrum of the predicted attractor. All models are able to capture the Fourier modes with magnitude larger than  $O(1)$ , while MNO is more accurate on the tail. (c) Spatial correlation of the attractor, averaged in the time dimension. MNO is more accurate on the near-range correlation, but all models miss the long-range correlation.

227 feedforward neural network on a single trajectory with  $h = 0.05s$  on the Lorenz attractor. Figure 3  
 228 shows that enforcing dissipativity produces predictions that isotropically point towards the attractor,  
 229 implying that the attractive properties of the Lorenz attractor are learned in the process. Observe  
 230 that the our network is also dissipative outside the shell in which dissipativity was enforced during  
 231 training. We conjecture that this is a property of ReLU networks that prevents model blow-up even in  
 232 more difficult learning problems (see Figure 1).

233 We empirically find that enforcing dissipativity does not reduce the relative  $L^2$  error compared to  
 234 the baseline neural network. Further, the invariant statistics of the dissipative model align well with  
 235 the ground-truth data distribution. Our results suggest that dissipativity can be enforced without  
 236 significantly affecting the model’s step-wise error and the learned statistical properties of the attractor.  
 237 See Appendix A.1 for more details.

## 238 4.2 Kuramoto-Sivashinsky equation

239 We consider the following one-dimensional KS equation,

$$\frac{\partial u}{\partial t} = -u \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, \text{ on } [0, L] \times (0, \infty), \text{ with initial condition: } u(\cdot, 0) = u_0, \text{ on } [0, L] \quad (7)$$

240 where the spatial domain  $[0, L]$  is equipped with periodic boundary conditions. We study the impact  
 241 the time step  $h$  has on learning. Our study shows that when the time steps are too large, the correlation  
 242 is chaotic and hard to capture. But counter-intuitively, when the time steps are too small, the evolution  
 243 is also hard to capture. In this case, the input and output of the learned operator will be very close,  
 244 and the identity map will be a local minimum. We thus propose to use the MNO to also learn the  
 245 time-derivative or residual directly. Figure 4a shows the results for varying  $h$  and when MNO is used  
 246 to learn either the identity residual or the Markov operator itself. We observe that the residual model  
 247 has a better per-step error and accumulated error at smaller  $h$ . When the time step is large, there is no  
 248 difference in modeling the residual. This idea can generalize to other integrators as an extension of  
 249 Neural ODEs to PDEs [31].

250 As shown in Figure 4b and 4c, we compare the performance of the MNO model against LSTM and  
 251 GRU that we use to model the evolution operator of the KS equation with  $h = 1s$ . We observe  
 252 that MNO model accurately recovers the Fourier spectrum of KS equation. We also present various  
 253 other invariant statistics of the KS equation in Appendix A.2. For other statistics, all models perform  
 254 similarly, except on the velocity distribution where MNO outperforms the LSTM and GRU. We  
 255 emphasize that some these statistics are very challenging to capture and most machine learning  
 256 approaches in the literature thus far fail to do so. (see Appendix A.2 for the details).

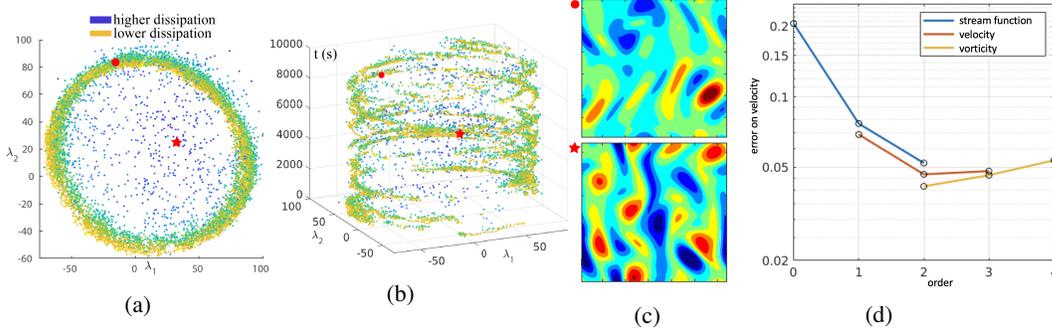


Figure 5: The learned attractor of the Kolmogorov flow and choice of Sobolev loss for KF. (a-c) The 10000 time steps trajectory generated by MNO projected onto the first two components of PCA. Each point corresponds to an snapshot on the attractor. Two points are selected for further visualization of vorticity field. (d) The velocity error in models trained on stream function, velocity and vorticity using Sobolev loss of different orders. We observe that Sobolev norm of second order provides the best performing model across different learning regimes.

### 257 4.3 Kolmogorov Flow

258 We consider two-dimensional Kolmogorov flow (a form of the Navier-Stokes equations) for a viscous,  
 259 incompressible fluid,

$$\frac{\partial u}{\partial t} = -u \cdot \nabla u - \nabla p + \frac{1}{Re} \Delta u + \sin(ny)\hat{x}, \quad \nabla \cdot u = 0, \quad \text{on } [0, 2\pi]^2 \times (0, \infty) \quad (8)$$

260 with initial condition  $u(\cdot, 0) = u_0$  where  $u$  denotes the velocity,  $p$  the pressure, and  $Re > 0$  is the  
 261 Reynolds number. We enforce dissipativity during training with the criterion described in eq. 5, with  
 262  $\lambda = 0.5$  and  $\nu$  being a uniform probability distribution supported on a shell around the origin. We  
 263 test the effect of enforcing dissipativity in the highly turbulent (and blow-up prone)  $Re = 500$  setting,  
 264 where we observe that a non-dissipative model blows up when composed with itself multiple times if  
 265 the initial condition is perturbed slightly from the attractor (Figure 1), even though the model achieves  
 266 relatively low  $L^2$  error. In contrast, we empirically observe that the dissipative MNO does not blow  
 267 up and its composed predictions returns to the attractor even when the initial condition is perturbed.

Model	training	loss	$L^2$ error	$H^1$ error	$H^2$ error	TKE error	$\epsilon$ error
MNO	$L^2$ loss	0.0166	0.0187	0.0474	0.1729	<b>0.0136</b>	0.0303
	$H^1$ loss	0.0184	0.0151	0.0264	0.0656	0.0256	<b>0.0017</b>
	$H^2$ loss	0.0202	<b>0.0143</b>	<b>0.0206</b>	<b>0.0361</b>	0.0226	0.0193
U-Net	$L^2$ loss	0.0269	0.0549	0.1023	0.3055	0.0958	0.0934
	$H^1$ loss	0.0377	0.0570	0.0901	0.2164	0.1688	0.1127
	$H^2$ loss	0.0565	0.0676	0.0936	0.1749	0.0482	0.0841
ConvLSTM	$L^2$ loss	0.2436	0.2537	0.3854	1.0130	0.0140	24.1119
	$H^1$ loss	0.2855	0.2577	0.3308	0.5336	0.6977	6.9167
	$H^2$ loss	0.3367	0.2727	0.3311	0.4352	0.8594	4.0976

Table 1: Benchmark on vorticity for the Kolmogorov flow with  $Re = 40$

268 **Accuracy with respect to various norms.** We study performance of MNO along with other  
 269 predictive architectures, U-Net [24] and LSTM-CNN [25], on modeling the vorticity  $w$  in the  
 270 relatively non-turbulent  $Re = 40$  case 1. The MNO achieves one order of magnitude better accuracy  
 271 compared to this architectures. As shown in Table 1, we train each model using the balanced  
 272  $L^2 (= H^0)$ ,  $H^1$ , and  $H^2$  losses, defined as the sum of the relative  $L^2$  loss grouped by each order of  
 273 derivative. We measure the error with respect to the standard (unbalanced) norms.

274 The MNO with  $H^2$  loss consistently achieves the smallest error on vorticity for all of the  $L^2$ ,  $H^1$ ,  
 275 and  $H^2$  norms. However, the  $L^2$  loss achieves the smallest error on the turbulence kinetic energy  
 276 (TKE), while the  $H^1$  loss achieves the smallest error on the dissipation  $\epsilon$ .

277 The KF equation with  $Re = 40$  is shown in Figure 10. (a) shows the ground truth vorticity field  
278 where each column represents a snapshot at  $t = 100s$  with different initial conditions. (b), (c), and  
279 (d) show the predicted trajectory of MNO vorticity, using the  $L^2$ ,  $H^1$ , and  $H^2$  losses respectively.  
280 The five columns represent  $t = 1000s, 2000s, 3000s, 4000s, 5000s$  respectively. The figure indicates  
281 that the predicted trajectories (b) (c) (d) share the same behaviors as in the ground truth (a), indicating  
282 that the MNO model is stable.

283 **Estimating the attractor and invariant statistics.** We compose MNO 10000 times to obtain the  
284 global attractor, and we compute the PCA (POD) basis of these 10000 snapshots and project them  
285 onto the first two components. As shown in Figure (a), we obtain a cycle-shaped attractor. The true  
286 attractor has dimension on the order of  $\mathcal{O}(100)$  [8]. If the attractor is a high-dimensional sphere, then  
287 most of its mass concentrates around the equator. Therefore, when projected to low-dimension, the  
288 attractor will have a ring shape. We see that most of the points are located on the ring, while a few  
289 other points are located in the center. The points in the center have high dissipation, implying they are  
290 intermittent states. In Figure (b) we add the time axis. While the trajectory jumps around the cycle,  
291 we observe there is a rough period of 2000s. We perform the same PCA analysis on the training  
292 data, showing the same behavior. Furthermore, in Figure 11, we present invariant statistics for the  
293 KF equation ( $Re = 40$ ), computed from a large number of samples from our approximation of the  
294 invariant measure. We find that the MNO consistently outperforms all other models in accurately  
295 capturing these statistics.

296 **Derivative orders.** Roughly speaking, vorticity is the derivative of velocity; velocity is the derivative  
297 of the stream function. Therefore we can denote the order of derivative of vorticity, velocity, and  
298 stream function as 2, 1, and 0 respectively. Combining vorticity, velocity, and stream function, with  
299  $L^2$ ,  $H^1$ , and  $H^2$  loss, we have in total the order of derivatives ranging from 0 to 4. We observe, in  
300 general, it is best practice to keep the order of derivatives in the model at a number slightly higher  
301 than that of the target quantity. For example, as shown in Figure 5d, when querying the velocity  
302 (first-order quantity), it is best to use second-order (modeling velocity plus  $H^1$  loss or modeling  
303 vorticity plus  $L^2$  loss). This is further illustrated in Table 3 for the  $Re = 500$  case. In general, using  
304 a higher order of derivatives as the loss will increase the power of the model and capture the invariant  
305 statistics more accurately. However, a higher-order of derivative means higher irregularity. It in turn  
306 requires a higher resolution for the model to resolve and for computing the discrete Fourier transform.  
307 This trade-off again suggests it is best to pick a Sobolev norm not too low or too high.

## 308 5 Discussion and future works

309 In this work, we propose a machine learning framework that trains from local data and predicts the  
310 global attractor and invariant statistics of chaotic systems. By enforcing dissipativity, we learn a  
311 Markov operator that empirically does not collapse or blow up over a long or infinite time horizon.  
312 Experiments also show the MNO predicts the attractor which shares the same distribution and  
313 statistics as the true function space trajectories. The simulations achieved by our MNO model  
314 have the potential to further our understanding of many physical phenomena and the mathematical  
315 models that underline them. Furthermore, the MNO shows great potential for modeling partially  
316 observed systems or studying systems that exhibit bifurcations both of which are of great interest for  
317 engineering applications. We also expect to generalize our theoretical work to show convergence to  
318 the global attractor and invariant measure by adapting ideas from the standard theory of numerical  
319 integrators [32].

320 This work provides a method for fast computation with applications to many scientific computing  
321 problems. Our methods have two main long-term impacts beyond the immediate interests of the  
322 scientific computing community. Since they are orders of magnitude faster than the traditional  
323 solvers currently employed in supercomputers, edge devices, and servers, the deployment of our  
324 methods significantly reduces the carbon footprint caused by scientific studies. Furthermore, the  
325 proposed methods are extremely flexible. The off-the-shelf usage of our methods allows scientists  
326 from a variety of disciplines, e.g. chemistry, biology, ecology, epidemiology, physics, and applied  
327 mathematics, to deploy them on their complex systems without the need to build elaborate numerical  
328 methods.

## References

- 329
- 330 [1] Pantelis R Vlachas, Wonmin Byeon, Zhong Y Wan, Themistoklis P Sapsis, and Petros Koumout-  
331 sakos. Data-driven forecasting of high-dimensional chaotic systems with long short-term  
332 memory networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineer-*  
333 *ing Sciences*, 474(2213):20170844, 2018.
- 334 [2] Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction  
335 of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical*  
336 *review letters*, 120(2):024102, 2018.
- 337 [3] Pantelis R Vlachas, Jaideep Pathak, Brian R Hunt, Themistoklis P Sapsis, Michelle Girvan,  
338 Edward Ott, and Petros Koumoutsakos. Backpropagation algorithms and reservoir computing  
339 in recurrent neural networks for the forecasting of complex spatiotemporal dynamics. *Neural*  
340 *Networks*, 2020.
- 341 [4] Huawei Fan, Junjie Jiang, Chun Zhang, Xingang Wang, and Ying-Cheng Lai. Long-term  
342 prediction of chaotic systems with machine learning. *Physical Review Research*, 2(1):012080,  
343 2020.
- 344 [5] James Gleick. *Chaos: Making a new science*. Open Road Media, 2011.
- 345 [6] Warwick Tucker. The lorenz attractor exists. *Comptes Rendus de l'Académie des Sciences -*  
346 *Series I - Mathematics*, 328(12):1197–1202, 1999.
- 347 [7] Mark Holland and Ian Melbourne. Central limit theorems and invariance principles for Lorenz  
348 attractors. *Journal of the London Mathematical Society*, 76(2):345–364, 10 2007.
- 349 [8] Roger Temam. *Infinite-dimensional dynamical systems in mechanics and physics*, volume 68.  
350 Springer Science & Business Media, 2012.
- 351 [9] A. Stuart and A.R. Humphries. *Dynamical Systems and Numerical Analysis*. Cambridge  
352 Monographs on Applie. Cambridge University Press, 1998.
- 353 [10] AR Humphries and AM Stuart. Runge–kutta methods for dissipative and gradient dynamical  
354 systems. *SIAM journal on numerical analysis*, 31(5):1452–1485, 1994.
- 355 [11] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Machine-learning-based spatio-temporal super  
356 resolution reconstruction of turbulent flows. *Journal of Fluid Mechanics*, 909, 2021.
- 357 [12] José I Cardesa, Alberto Vela-Martín, and Javier Jiménez. The turbulent cascade in five dimen-  
358 sions. *Science*, 357(6353):782–784, 2017.
- 359 [13] Gary J Chandler and Rich R Kerswell. Invariant recurrent solutions embedded in a turbulent  
360 two-dimensional kolmogorov flow. *Journal of Fluid Mechanics*, 722:554–595, 2013.
- 361 [14] Péter Koltai and Stephan Weiss. Diffusion maps embedding and transition matrix analysis of  
362 the large-scale flow structure in turbulent rayleigh–bénard convection. *Nonlinearity*, 33(4):1723,  
363 2020.
- 364 [15] Jason J Bramburger, J Nathan Kutz, and Steven L Brunton. Data-driven stabilization of periodic  
365 orbits. *IEEE Access*, 9:43504–43521, 2021.
- 366 [16] Jacob Page, Michael P Brenner, and Rich R Kerswell. Revealing the state space of turbulence  
367 using machine learning. *Physical Review Fluids*, 6(3):034402, 2021.
- 368 [17] Victor Churchill and Dongbin Xiu. Deep learning of chaotic systems from partially-observed  
369 data. *arXiv preprint arXiv:2205.08384*, 2022.
- 370 [18] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya,  
371 Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial  
372 differential equations. *arXiv preprint arXiv:2003.03485*, 2020.
- 373 [19] Nikola B. Kovachki, Zongyi Li, Liu Burigede, Kamyar Azizzadenesheli, Kaushik Bhattacharya,  
374 Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function  
375 spaces. *In Preparation*.

- 376 [20] Georgy Anatolevich Sviridyuk. On the general theory of operator semigroups. *Russian*  
377 *Mathematical Surveys*, 49(4):45–74, 1994.
- 378 [21] Wojciech Marian Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Świrszcz, and Razvan  
379 Pascanu. Sobolev training for neural networks. *arXiv preprint arXiv:1706.04859*, 2017.
- 380 [22] Alex Beatson, Jordan Ash, Geoffrey Roeder, Tianju Xue, and Ryan P Adams. Learning com-  
381 posable energy surrogates for pde order reduction. *Advances in Neural Information Processing*  
382 *Systems*, 33, 2020.
- 383 [23] Alexander J Smits, Beverley J McKeon, and Ivan Marusic. High–reynolds number wall  
384 turbulence. *Annual Review of Fluid Mechanics*, 43:353–375, 2011.
- 385 [24] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for  
386 biomedical image segmentation. In *International Conference on Medical image computing and*  
387 *computer-assisted intervention*, pages 234–241. Springer, 2015.
- 388 [25] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun  
389 Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting.  
390 *arXiv preprint arXiv:1506.04214*, 2015.
- 391 [26] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation  
392 of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*,  
393 2014.
- 394 [27] Grigorios Pavliotis and Andrew Stuart. *Multiscale Methods: Averaging and Homogenization*,  
395 volume 53. 01 2008.
- 396 [28] Jooyoung Park and Irwin W Sandberg. Universal approximation using radial-basis-function  
397 networks. *Neural computation*, 3(2):246–257, 1991.
- 398 [29] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya,  
399 Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differen-  
400 tial equations. *arXiv preprint arXiv:2010.08895*, 2020.
- 401 [30] Edward N Lorenz. Deterministic nonperiodic flow. *Journal of atmospheric sciences*, 20(2):130–  
402 141, 1963.
- 403 [31] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary  
404 differential equations. In *Advances in neural information processing systems*, pages 6571–6583,  
405 2018.
- 406 [32] A. R. Humphries and A. M. Stuart. Runge–kutta methods for dissipative and gradient dynamical  
407 systems. *SIAM Journal on Numerical Analysis*, 31(5):1452–1485, 1994.
- 408 [33] Aly-Khan Kassam and Lloyd N. Trefethen. Fourth-order time-stepping for stiff pdes. *SIAM*  
409 *Journal on Scientific Computing*, 26(4):1214–1233, 2005.
- 410 [34] Gabriel J Lord, Catherine E Powell, and Tony Shardlow. *An introduction to computational*  
411 *stochastic PDEs*, volume 50. Cambridge University Press, 2014.
- 412 [35] Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction  
413 with lstm. 1999.
- 414 [36] Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and  
415 error bounds for fourier neural operators. *Journal of Machine Learning Research*, 22:1–76,  
416 2021.
- 417 [37] Andrew Stuart. Perturbation theory for infinite dimensional dynamical systems. 1995.

## 418 Checklist

419 The checklist follows the references. Please read the checklist guidelines carefully for information on  
420 how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or  
421 **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing  
422 the appropriate section of your paper or providing a brief inline description. For example:

- 423 • Did you include the license to the code and datasets? **[Yes]**
- 424 • Did you include the license to the code and datasets? **[No]** The code and the data are  
425 proprietary.
- 426 • Did you include the license to the code and datasets? **[N/A]**

427 Please do not modify the questions and only use the provided macros for your answers. Note that the  
428 Checklist section does not count towards the page limit. In your paper, please delete this instructions  
429 block and only keep the Checklist section heading above along with the questions/answers below.

### 430 1. For all authors...

- 431 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s  
432 contributions and scope? **[Yes]**
- 433 (b) Did you describe the limitations of your work? **[Yes]** See Section C
- 434 (c) Did you discuss any potential negative societal impacts of your work? **[No]**
- 435 (d) Have you read the ethics review guidelines and ensured that your paper conforms to  
436 them? **[Yes]**

### 437 2. If you are including theoretical results...

- 438 (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** Statement of  
439 Theorem 1
- 440 (b) Did you include complete proofs of all theoretical results? **[Yes]** See Appendix

### 441 3. If you ran experiments...

- 442 (a) Did you include the code, data, and instructions needed to reproduce the main experi-  
443 mental results (either in the supplemental material or as a URL)? **[Yes]** we plan to do  
444 so.
- 445 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
446 were chosen)? **[Yes]** Specified in Section 4
- 447 (c) Did you report error bars (e.g., with respect to the random seed after running ex-  
448 periments multiple times)? **[No]** The error is consistent in the operator learning for  
449 PDEs.
- 450 (d) Did you include the total amount of compute and the type of resources used (e.g., type  
451 of GPUs, internal cluster, or cloud provider)? **[Yes]** Specified in Section 4

### 452 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- 453 (a) If your work uses existing assets, did you cite the creators? **[Yes]**
- 454 (b) Did you mention the license of the assets? **[Yes]** Supplement
- 455 (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]**  
456 Supplement
- 457 (d) Did you discuss whether and how consent was obtained from people whose data you’re  
458 using/curating? **[N/A]**
- 459 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
460 information or offensive content? **[N/A]**

### 461 5. If you used crowdsourcing or conducted research with human subjects...

- 462 (a) Did you include the full text of instructions given to participants and screenshots, if  
463 applicable? **[N/A]**
- 464 (b) Did you describe any potential participant risks, with links to Institutional Review  
465 Board (IRB) approvals, if applicable? **[N/A]**
- 466 (c) Did you include the estimated hourly wage paid to participants and the total amount  
467 spent on participant compensation? **[N/A]**

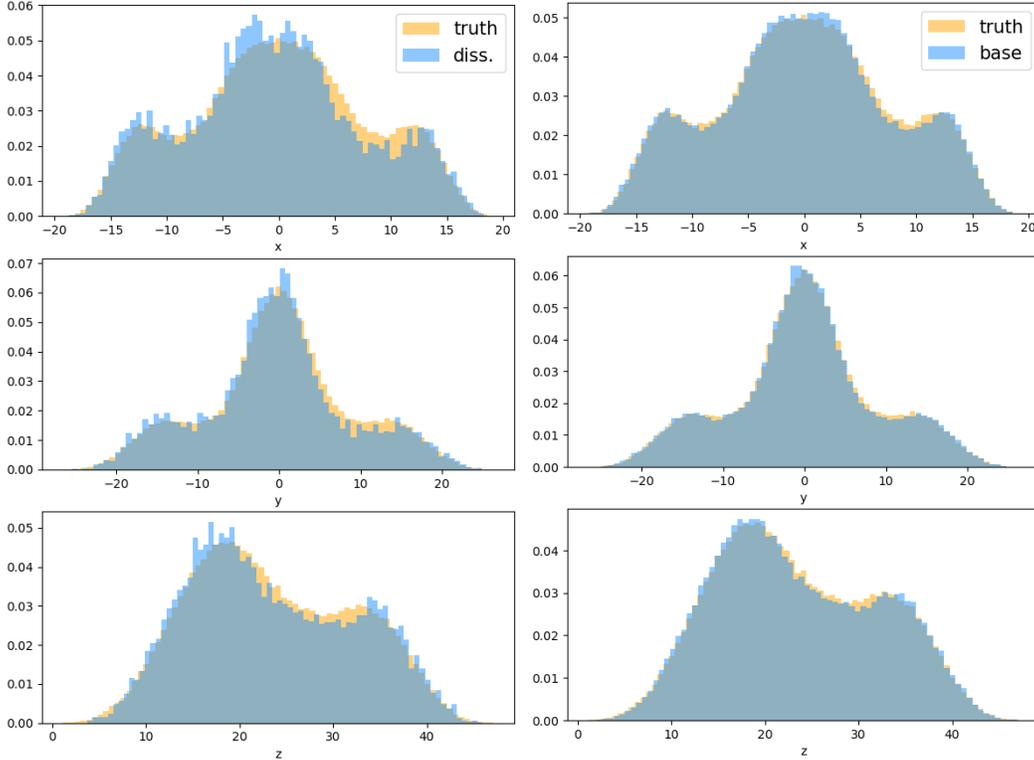


Figure 6: Lorenz-63 histograms of  $x, y, z$ -coordinates for the baseline network (no dissipativity enforced) and the MNO (dissipativity enforced).

## 468 A Full Experiments

469 In this section, we present the full numerical experiments on the Lorenz-63 system, the Kuramoto-  
470 Sivashinsky equation, and the Kolmogorov flow.

### 471 A.1 Lorenz-63

472 The Lorenz-63 system is three dimensional ODE system:

$$\begin{aligned}\dot{u}_x &= \alpha(u_y - u_x), \\ \dot{u}_y &= -\alpha u_x - u_y - u_x u_z, \\ \dot{u}_z &= u_x u_y - b u_z - b(r + \alpha).\end{aligned}$$

473 In this paper, we use the canonical parameters  $(\alpha, b, r) = (10, 8/3, 28)$  [30].

474 **Learning the Markov operator.** We use a simple feedforward neural network with 6 hidden layers  
475 and 150 neurons per layer to learn the Markov operator for the Loren-63 system. We discretize the  
476 training trajectory into time-steps of 0.05 seconds.

477 **Enforcing dissipativity.** We enforce dissipativity during training with the criterion described in eq.  
478 5, with  $\lambda = 0.5$  and  $\nu$  being a uniform probability distribution supported on a shell around the origin  
479 with inner radius 90 and outer radius 130.

480 **Evaluation metrics.** We evaluate the effectiveness of our technique for enforcing dissipativity in  
481 three ways:

- 482 • **Relative L2 error:** We evaluate the learned Markov operators on relative (i.e., normalized)  
483 L2 error over 1 second by composing the model 20 times to gauge their ability to predict

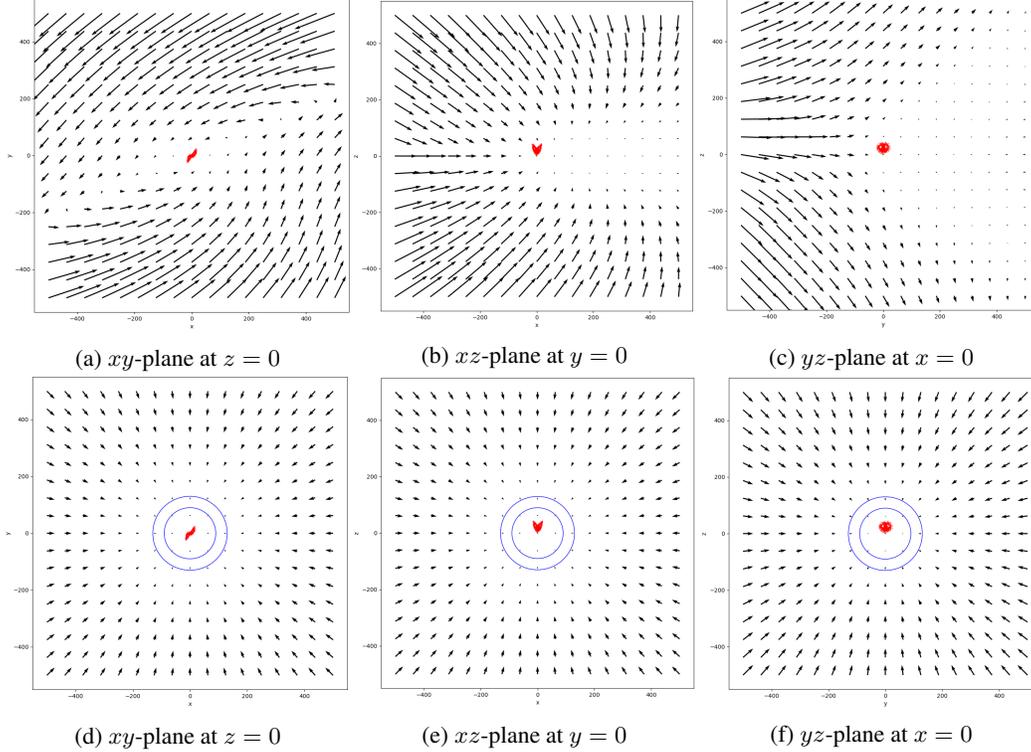


Figure 7: Predicted flow field of dissipative network on Lorenz-63. The red points are training points on the attractor. Dissipativity is enforced uniformly within the blue shell.

Model	Per-step error	Per-second error	Dissipativity error
No diss. enforced	0.000570	0.0300	-
Diss. enforced	0.000564	0.0264	0.000667

Table 2: Relative  $L^2$  error rates on the Lorenz-63 system. Per-step error is the error on the time-scale used in training. Per-second error is the error of the model composed with itself 21 times. Dissipativity error is the error of following the enforced dissipativity constraint on each step.

484 longer trajectories. As seen in Table 2, enforcing dissipativity does not cause a decrease in  
 485 relative L2 error when compared to the baseline network, even though both models have the  
 486 same number of parameters.

487 • **Vector field far from the attractor:** We also qualitatively evaluate the learned vector  
 488 fields far from the attractor both with and without enforced dissipativity. We observe that  
 489 enforcing dissipativity produces predictions that isotropically point towards the attractor,  
 490 implying that the attractive properties of the Lorenz attractor are learned in the process. See  
 491 Figure 7. Observe that the dissipative network is also dissipative outside the shell in which  
 492 dissipativity was enforced during training.

493 • **Invariant statistics of the attractor:** To justify learning the Markovian map between  
 494 time-steps, we also compare the coordinate histograms of the ground-truth and the learned  
 495 attractors (with and without dissipativity enforced) after composing for 200000 time steps.  
 496 Both models match the ground-truth coordinate histograms. Since the coordinate-wise  
 497 histograms of the models matches the ground-truth, this indicates that the models are  
 498 learning the distribution of the attractor and thus the invariant measure of the system.

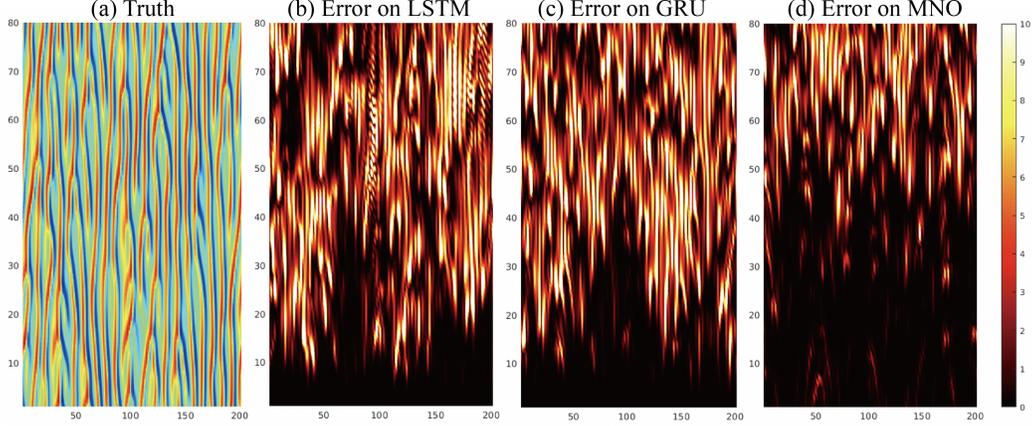


Figure 8: Trajectory and error on the KS equation

The x-axis is the spatial domain; the y-axis is the temporal domain. The figure shows that LSTM and GRU start to diverge at  $t = 20s$  while MNO is able to keep up with the exact trajectory until  $t = 50s$ .

## 499 A.2 Kuramoto-Sivashinsky equation

500 We consider the following one-dimensional Kuramoto-Sivashinsky equation,

$$\begin{aligned} \frac{\partial u}{\partial t} &= -u \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, & \text{on } [0, L] \times (0, \infty) \\ u(\cdot, 0) &= u_0, & \text{on } [0, L] \end{aligned}$$

501 where  $L = 32\pi$  or  $64\pi$  and the spatial domain  $[0, L]$  is equipped with periodic boundary conditions.  
 502 We assume the initial condition  $u_0 \in \dot{L}_{\text{per}}^2([0, L]; \mathbb{R})$ , where  $\dot{L}_{\text{per}}^2([0, L]; \mathbb{R})$  is the space of all mean  
 503 zero  $L^2$ -functions that are periodic on  $[0, L]$ . Existence of the semigroup  $S_t : \dot{L}_{\text{per}}^2([0, L]; \mathbb{R}) \rightarrow$   
 504  $\dot{L}_{\text{per}}^2([0, L]; \mathbb{R})$  is established in [8, Theorem 3.1]. Data is obtained by solving the equation using the  
 505 exponential time-differencing fourth-order Runge-Kutta method from [33]. Random initial conditions  
 506 are generated according to a mean zero Gaussian measure with covariance  $L^{-2/\alpha} \tau^{\frac{1}{2}(2\alpha-1)} (-\Delta +$   
 507  $(\tau^2/L^2)I)^{-\alpha}$  where  $\alpha = 2$ ,  $\tau = 7$ , and periodic boundary conditions on  $[0, L]$ ; for details see [34].

508 **Benchmarks for Kuramoto-Sivashinsky.** We compare MNO with common choices of recurrent  
 509 neural networks including the long short-term memory network (LSTM)[35] and gated recurrent unit  
 510 (GRU)[26]. All models use the time-discretization  $h = 1s$ . The training dataset consists of 1000  
 511 different realizations of trajectories on the time interval  $t \in [50, 200]$  (the first 50s is truncated so  
 512 the dynamics reach the ergodic state), which adds up to  $1000 \times 150 = 150,000$  snapshots in total.  
 513 Another 200 realizations are generated for testing. Every single snapshot has the resolution 2048. We  
 514 use Adam optimizer to minimize the relative  $L^2$  loss with learning rate = 0.001, and step learning  
 515 rate scheduler that decays by half every 10 epochs for 50 epochs in total. **LSTM and GRU:** having  
 516 tested many different configurations, we choose the best hyper-parameters: the number of layers = 1,  
 517 width = 1000. During the evaluation, we additionally provide 1000 snapshots as a warm-up of the  
 518 memory. **MNO:** we choose 1-d Fourier neural operator as our base model with four Fourier layers  
 519 with 20 frequencies per channel and width = 64. Experiments run on Nvidia V100 GPUs.

520 **Accuracy with respect to the true trajectory.** In general, MNO has a smaller per-step error  
 521 compared to RNN. As shown in Figure 8, the MNO model captures a longer period of the exact  
 522 trajectory compared to LSTM and GRU. LSTM and GRU start to diverge at  $t = 20s$  while FNO is  
 523 able to keep up with the exact trajectory until  $t = 50s$ .

524 **Invariant statistics for the KS equations** As shown in Figure 9, we present enormous invariant  
 525 statistics for the KS equation. We use 150000 snapshots to train the MNO, LSTM, and GRU to model  
 526 the evolution operator of the KS equation with  $h = 1s$ . We compose each model for  $T = 1000$  time  
 527 steps to obtain a long trajectory (attractor), and estimate various invariant statistics from them.

- 528 • **(a) Fourier spectrum:** the Fourier spectrum of the predicted attractor. All models are able  
529 to capture the Fourier modes with magnitude larger than  $O(1)$ , while MNO is more accurate  
530 on the tail.
- 531 • **(b) Spatial correlation:** the spatial correlation of the attractor, averaged in the time di-  
532 mension. MNO is more accurate on the near-range correlation, but all models miss the  
533 long-range correlation.
- 534 • **(c) Auto-correlation of the Fourier mode:** the auto-correlation of the  $10^{th}$  Fourier mode.  
535 Since the Fourier modes are nearly constant, the auto-correlation is constant too.
- 536 • **(d) Auto-correlation of the PCA mode:** the auto-correlation of the first PCA mode (with  
537 respect to the PCA basis of the ground truth data). The PCA mode oscillates around  
538  $[-40, 40]$ , showing an ergodic state.
- 539 • **(e) Distribution of kinetic energy:** the distribution of kinetic energy with respect to the  
540 time dimension. MNO captures the distribution most accurately.
- 541 • **(f) Pixelwise distribution of velocity:** since the KS equation is homogeneous, we can  
542 compute the distribution of velocity with respect to pixels. MNO captures the pixelwise  
543 distribution most accurately too.

544 **Choice of time discretization.** We further study the choice of time steps  $h$ . As shown in Figure 4a,  
545 when the time steps are too large, the correlation is chaotic and hard to capture. But counter-intuitively,  
546 when the time steps are too small, the evolution is also hard to capture. In this case, the input and  
547 output of the network will be very close, and the identity map will be a local minimum. An easy fix  
548 is to use MNO to learn the time-derivative or residual. This is shown in the figure, where the residual  
549 model (blue line) has a better per-step error and accumulated error at smaller  $h$ . When the time step  
550 is large, there is no difference in modeling the residual. This idea can generalize to other integrators  
551 as an extension of Neural ODEs to PDEs [31].

### 552 A.3 Kolmogorov Flow

553 We consider the two-dimensional Navier-Stokes equation for a viscous, incompressible fluid,

$$\begin{aligned}
 \frac{\partial u}{\partial t} &= -u \cdot \nabla u - \nabla p + \frac{1}{Re} \Delta u + \sin(ny)\hat{x}, & \text{on } [0, 2\pi]^2 \times (0, \infty) \\
 \nabla \cdot u &= 0 & \text{on } [0, 2\pi]^2 \times [0, \infty) \\
 u(\cdot, 0) &= u_0 & \text{on } [0, 2\pi]^2
 \end{aligned}$$

554 where  $u$  denotes the velocity,  $p$  the pressure, and  $Re > 0$  is the Reynolds number. The domain  
555  $[0, 2\pi]^2$  is equipped with periodic boundary conditions. The specific choice of forcing  $\sin(ny)\hat{x}$   
556 constitutes a Kolmogorov flow; we choose  $n = 4$  in all experiments. We define  $\mathcal{U}$  to be the closed  
557 subspace of  $L^2([0, 2\pi]^2; \mathbb{R}^2)$ ,  $\mathcal{U} = \{u \in \dot{L}_{\text{per}}^2([0, 2\pi]^2; \mathbb{R}^2) : \nabla \cdot u = 0\}$  and assume  $u_0 \in \mathcal{U}$ .  
558 We define the vorticity  $w = (\nabla \times u)\hat{z}$  and the stream function  $f$  as the solution to the Poisson  
559 equation  $-\Delta f = w$ . Existence of the semigroup  $S_t : \mathcal{U} \rightarrow \mathcal{U}$  is established in [8, Theorem 2.1]. We  
560 denote turbulence kinetic energy (TKE)  $\langle (u - \bar{u})^2 \rangle$ , and dissipation  $\epsilon = \langle w^2 \rangle / Re$ . Data is obtained  
561 by solving the equation in vorticity form using the pseudo-spectral split step method from [13].  
562 Random initial conditions are generated according to a mean zero Gaussian measure with covariance  
563  $7^{3/2}(-\Delta + 49I)^{-2.5}$  with periodic boundary conditions on  $[0, 2\pi]^2$ .

564 **Benchmarks for the 2d Kolmogorov flow.** We compare MNO with common standard two-  
565 dimensional dynamic models including U-Net[24] and LSTM-CNN[25] on modeling the vorticity  
566  $w$  in the relatively non-turbulent  $Re = 40$  case. We choose the discretization  $h = 1s$ . The training  
567 dataset consists of 180 realizations of trajectories on time interval  $t \in [100, 500]$  (the first 100 seconds  
568 are discarded) which adds up to  $180 \times 400 = 72,000$  snapshots in total. Another 20 realizations  
569 are generated for testing. Each single snapshot has resolution  $64 \times 64$ . We use the Adam optimizer  
570 to minimize the relative  $L^2$  loss with learning rate = 0.0005, and step learning rate scheduler that  
571 decays by half every 10 epochs for 50 epochs in total. **U-Net:** we use five layers of convolution and  
572 deconvolution with width from 64 to 1024. **LSTM-CNN:** we use one layer of LSTM with width  
573 = 64. **MNO:** we parameterize the 2-d Fourier neural operator consists of four Fourier layers with 20  
574 frequencies per channel and width = 64.

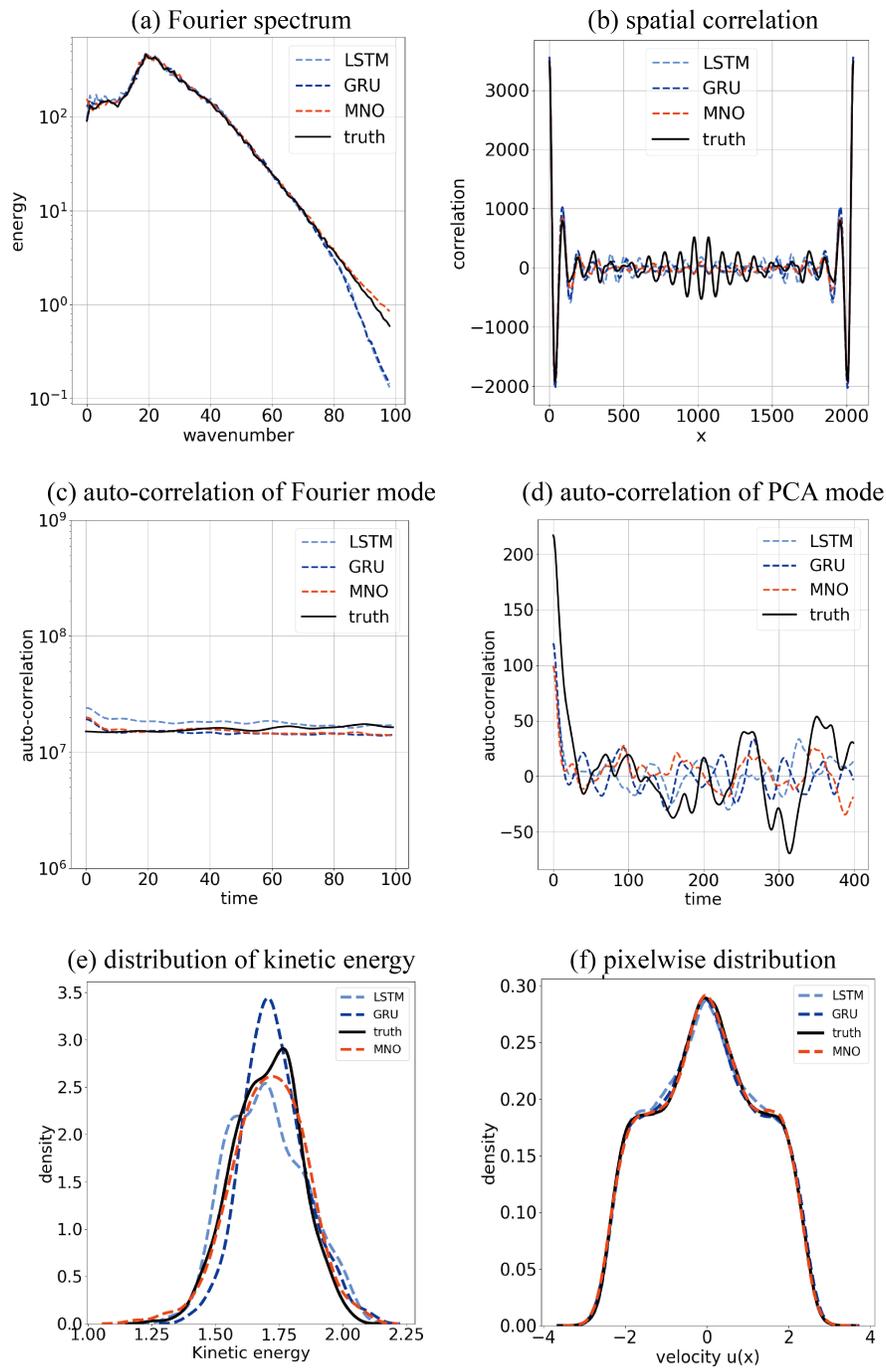


Figure 9: Invariant statistics for the KS equation

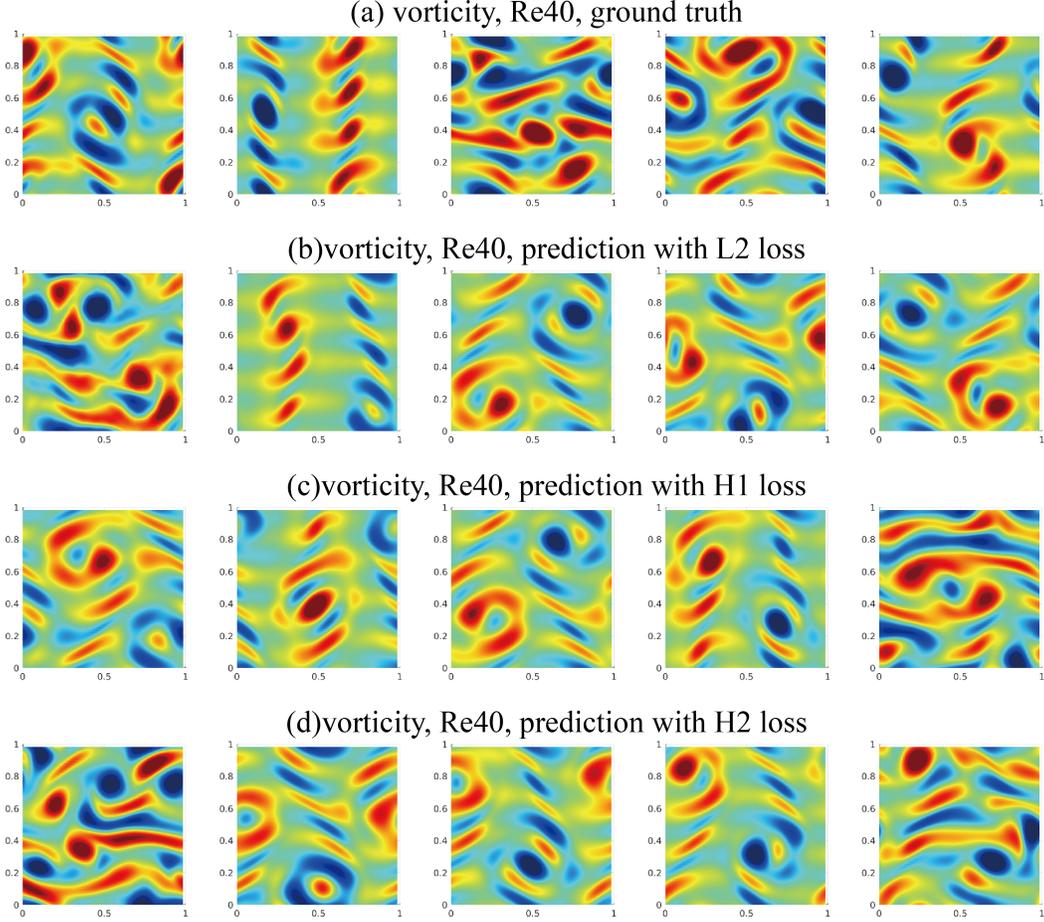


Figure 10: visualization of the KF equation, Re40

575 **Enforcing dissipativity.** We enforce dissipativity during training with the criterion described in eq.  
 576 5, with  $\lambda = 0.5$  and  $\nu$  being a uniform probability distribution supported on a shell around the origin.

577 **Accuracy with respect to various norms.** MNO shows near one order of magnitude better accuracy  
 578 compared to U-Net and LSTM-CNN. As shown in Table 1, we train each model using the balanced  
 579  $L^2 (= H^0)$ ,  $H^1$ , and  $H^2$  losses, defined as the sum of the relative  $L^2$  loss grouped by each order of  
 580 derivative. And we measure the error with respect to the standard (unbalanced) norms. The MNO  
 581 with  $H^2$  loss consistently achieves the smallest error on vorticity on all of the  $L^2$ ,  $H^1$ , and  $H^2$   
 582 norms. However,  $L^2$  loss achieves the smallest error on the turbulence kinetic energy (TKE);  $H^1$  loss  
 583 achieves the smallest error on the dissipation  $\epsilon$ .

584 The NS equation with Re40 is shown in Figure 10. (a) show the ground truth data of vorticity  
 585 field, each column represents a snapshot at  $t = 100s$  with a different initial condition. (b), (c),  
 586 (d) show the predicted trajectory of MNO on vorticity, using  $L^2$ ,  $H^1$ , and  $H^2$  losses respectively.  
 587 We are able to generate a long trajectory with the MNO model. The five columns represent  $t =$   
 588  $1000s, 2000s, 3000s, 4000s, 5000s$  respectively. As shown in the figure, the predicted trajectories  
 589 (b) (c) (d) share the same behaviors as in the ground truth (a). It indicates the MNO model is stable.

590 **Visualizing the attractor generated by MNO.** We compose MNO 10000 times to obtain the  
 591 global attractor, and we compute the PCA (POD) basis of these 10000 snapshots and project them  
 592 onto the first two components. As shown in Figure (a), we obtain a cycle-shaped attractor. The true  
 593 attractor has a degree of freedom around  $O(100)$  [8]. If the attractor is a high-dimensional sphere,  
 594 then most of the mass concentrates around its equator. Therefore, when projected to low-dimension,  
 595 the attractor will have the shape of a ring. Most of the points are located on the ring, while a few

Model	training loss	(order)	error on $f$	error on $u$	error on $w$
Stream function $f$	$L^2$ loss	0.0379 (0 <sup>th</sup> order)	0.0383	0.2057	2.0154
Stream function $f$	$H^1$ loss	0.0512 (1 <sup>st</sup> order)	0.0268	0.0769	0.3656
Stream function $f$	$H^2$ loss	0.0973 (2 <sup>nd</sup> order)	0.0198	0.0522	0.2227
Velocity $u$	$L^2$ loss	0.0688 (1 <sup>st</sup> order)	0.0217	0.0691	0.3217
Velocity $u$	$H^1$ loss	0.1246 (2 <sup>nd</sup> order)	<b>0.0170</b>	0.0467	0.1972
Velocity $u$	$H^2$ loss	0.2662 (3 <sup>rd</sup> order)	0.0178	0.0482	0.1852
Vorticity $w$	$L^2$ loss	0.1710 (2 <sup>nd</sup> order)	0.0219	<b>0.0415</b>	0.1736
Vorticity $w$	$H^1$ loss	0.3383 (3 <sup>rd</sup> order)	0.0268	0.0463	<b>0.1694</b>
Vorticity $w$	$H^2$ loss	0.4590 (4 <sup>th</sup> order)	0.0312	0.0536	0.1854

Table 3: Vorticity, velocity, and stream function for the Kolmogorov flow with  $Re = 500$

596 other points are located in the center. The points in the center have high dissipation, implying they are  
597 intermittent states. In Figure (b) we add the time axis. While the trajectory jumps around the cycle,  
598 we observe there is a rough period of 2000s. We perform the same PCA analysis on the training data,  
599 which shows the same behavior.

600 **Invariant statistics.** Similarly, we present enormous invariant statistics for the NS equation (Re40),  
601 as shown in Figure 11,. We use 72000 snapshots to train the MNO, UNet, and ConvLSTM to model  
602 the evolution operator of the KS equation with  $h = 1s$ . We compose each model for  $T = 10000$  time  
603 steps to obtain a long trajectory (attractor), and estimate various invariant statistics from them.

- 604 • **(a, d) Fourier spectrum of velocity and vorticity:** the Fourier spectrum of the predicted  
605 attractor. Again, all models are able to capture the Fourier modes with magnitude larger  
606 than  $O(1)$ , while MNO is more accurate on the tail. Using the Sobolev norm further helps  
607 to capture the tail.
- 608 • **(b, e) Pixelwise distribution of velocity and vorticity:** All models preserve the pixelwise  
609 distribution.
- 610 • **(c, f) Distribution of kinetic energy and dissipation rate:** the distribution of kinetic energy  
611 with respect to the time dimension. MNO captures the distribution most accurately.
- 612 • **(g) Auto-correlation of the Fourier mode:** the auto-correlation of the 10<sup>th</sup> Fourier mode.  
613 Since the Fourier modes are nearly constant, the auto-correlation is constant too. Notice it is  
614 very expensive to generate long-time ground truth data, so the figure does not include the  
615 ground truth. However, it is easy to obtain the auto-correlation by MNO.
- 616 • **(h) Auto-correlation of the PCA mode:** the auto-correlation of the first PCA mode (with  
617 respect to the PCA basis of the ground truth data). The PCA mode oscillates around  
618  $[-1000, 1000]$ , showing an ergodic state. The UNet oscillates around  $[0, 2000]$ .
- 619 • **(i) Spatial correlation:** the spatial correlation of the attractor, averaged in the time dimen-  
620 sion. The four columns represent the truth and MNO with different losses. As seen from the  
621 figure, there is a wave pattern matching the force term  $\sin(4y)$ .

622 **Order of derivatives.** Roughly speaking, vorticity is the derivative of velocity; velocity is the  
623 derivative of the stream function. Therefore we can denote the order of derivative of vorticity, velocity,  
624 and stream function as 2, 1, and 0 respectively. Combining vorticity, velocity, and stream function,  
625 with  $L^2$ ,  $H^1$ , and  $H^2$  loss, we have in total the order of derivatives ranging from 0 to 4. We observe,  
626 in general, it is best practice to keep the order of derivatives in the model at a number slightly higher  
627 than that of the target quantity. For example, as shown in Figure 5d, when querying the velocity  
628 (first-order quantity), it is best to use second-order (modeling velocity plus  $H^1$  loss or modeling  
629 vorticity plus  $L^2$  loss). This is further illustrated in Table 3. In general, using a higher order of  
630 derivatives as the loss will increase the power of the model and capture the invariant statistics more  
631 accurately. However, a higher-order of derivative means higher irregularity. It in turn requires a  
632 higher resolution for the model to resolve and for computing the discrete Fourier transform. This  
633 trade-off again suggests it is best to pick a Sobolev norm not too low or too high.

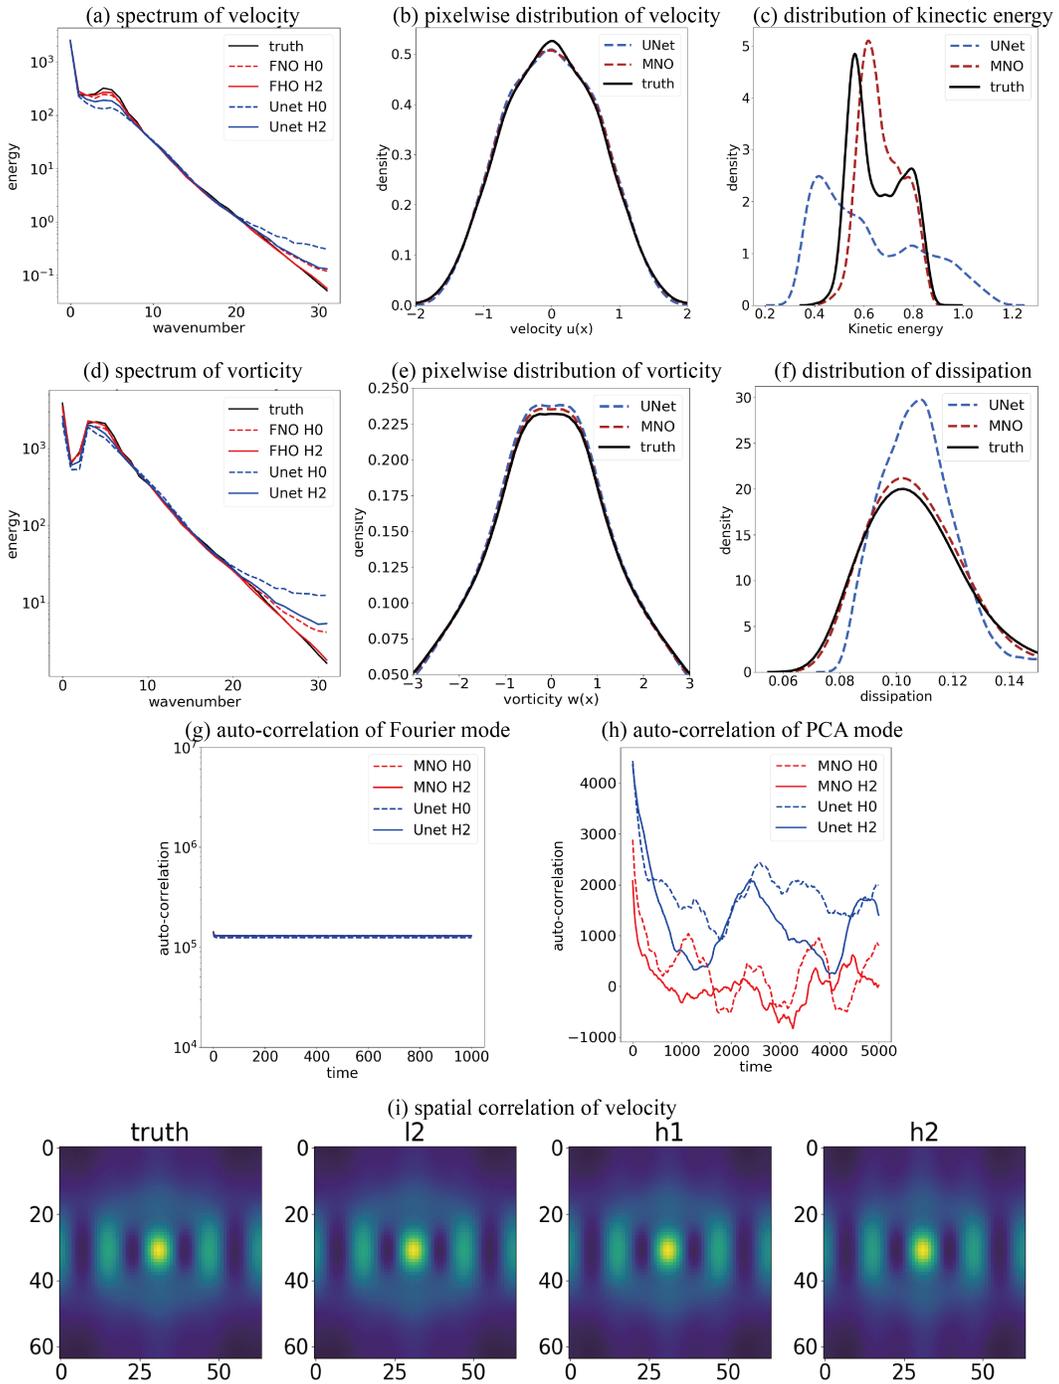


Figure 11: Invariant statistics for the NS equation

## 634 B Theoretical background

### 635 B.1 Markov operators and the semigroup

636 Since the system (1) is autonomous, that is,  $F$  does not explicitly depend on time, under the well-  
 637 posedness assumption, we may define, for any  $t \in [0, \infty)$ , a Markov operator  $S_t : \mathcal{U} \rightarrow \mathcal{U}$  such that  
 638  $u(t) = S_t u(0)$ . This map satisfies the properties

- 639 1.  $S_0 = I$ ,
- 640 2.  $S_t(u_0) = u(t)$ ,
- 641 3.  $S_t(S_s(u_0)) = u(t + s)$ ,

642 for any  $s, t \in [0, \infty)$  and any  $u_0 \in \mathcal{U}$  where  $I$  denotes the identity operator on  $\mathcal{U}$ . In particular, the  
 643 family  $\{S_t : t \in [0, \infty)\}$  defines a semigroup of operators acting on  $\mathcal{U}$ . Our goal is to approximate a  
 644 particular element of this semigroup associated to some fixed time step  $h > 0$  given observations of  
 645 the trajectory from (1). We build an approximation  $\hat{S}_h : \mathcal{U} \rightarrow \mathcal{U}$  such that

$$\hat{S}_h \approx S_h. \quad (9)$$

### 646 B.2 Proof of Theorem 1

647 *Proof of Theorem 1.* Since  $S_h$  is continuous and  $K$  is compact, the set

$$R = \bigcup_{t=0}^n S_h^t(K)$$

648 is compact. Therefore, there exist a set of representers  $\varphi_1, \varphi_2, \dots \in R$  such that

$$\lim_{m \rightarrow \infty} \sup_{v \in R} \inf_{u \in R_m} \|u - v\|_{\mathcal{U}} = 0$$

649 where  $R_m = \text{span}\{\varphi_1, \dots, \varphi_m\}$ . For any  $m \in \mathbb{N}$ , let  $P_m : \mathcal{U} \rightarrow R_m$  denote a projection of  $U$  to  
 650  $R_m$ . Since  $R$  is compact, the set

$$P = R \bigcup \left( \bigcup_{m=1}^{\infty} P_m(R) \right)$$

651 is compact [19, Lemma 14]. Since  $S_h$  is locally Lipschitz and  $P$  is compact, there exists a constant  
 652  $C = C(P) > 0$  such that

$$\|S_h(u_1) - S_h(u_2)\|_{\mathcal{U}} \leq C \|u_1 - u_2\|_{\mathcal{U}}, \quad \forall u_1, u_2 \in P.$$

653 Without loss of generality, assume  $C \neq 1$  and define

$$M = \left( C^n + \frac{1 - C^n}{1 - C} \right)^{-1}.$$

654 By the universal approximation theorem for neural operators [19, Theorem 4] or [36, Theorem 2.5],  
 655 there exists a neural operator  $\hat{S}_h : \mathcal{U} \rightarrow \mathcal{U}$  such that

$$\sup_{u_0 \in P} \|S_h(u_0) - \hat{S}_h(u_0)\|_{\mathcal{U}} < \epsilon M.$$

656 Perusal of the proof of the universal approximation theorem for neural operators shows that  $\hat{S}_h$  can  
 657 be chosen so that  $\hat{S}_h(P) \subseteq R_m$  for some  $m \in \mathbb{N}$  large enough, therefore  $\hat{S}_h(P) \subseteq P$ . Let  $u_0 \in K$ ,  
 658 then the triangle inequality implies

$$\begin{aligned} \|u(nh) - \hat{S}_h^n(u_0)\|_{\mathcal{U}} &= \|S_h^n(u_0) - \hat{S}_h^n(u_0)\|_{\mathcal{U}} \\ &= \|S_h(S_h^{n-1}(u_0)) - \hat{S}_h(\hat{S}_h^{n-1}(u_0))\|_{\mathcal{U}} \\ &\leq \|S_h(S_h^{n-1}(u_0)) - S_h(\hat{S}_h^{n-1}(u_0))\|_{\mathcal{U}} + \|S_h(\hat{S}_h^{n-1}(u_0)) - \hat{S}_h(\hat{S}_h^{n-1}(u_0))\|_{\mathcal{U}} \\ &\leq C \|S_h^{n-1}(u_0) - \hat{S}_h^{n-1}(u_0)\|_{\mathcal{U}} + \epsilon M. \end{aligned}$$

659 By the discrete time Grönwall lemma,

$$\begin{aligned} \|u(nh) - \hat{S}_h^n(u_0)\|_{\mathcal{U}} &\leq C^n \|S_h(u_0) - \hat{S}_h(u_0)\|_{\mathcal{U}} + \epsilon M \left( \frac{1 - C^n}{1 - C} \right) \\ &< \epsilon M \left( C^n + \frac{1 - C^n}{1 - C} \right) \\ &= \epsilon \end{aligned}$$

660 which completes the proof.

661

□

## 662 C Discussion and future work

663 In this work, we learn MNO from only local data and compose it to obtain the global attractor of  
664 chaotic systems, and by explicitly enforcing dissipativity, we empirically show that MNO predictions  
665 do not blow up or collapse even in the long-time horizon, while still achieving relatively low error  
666 running several orders of magnitude faster than traditional methods.

667 MNO has two major limitations. First, it assumes the target system is approximately Markovian. If  
668 the system is heavily path-dependent, then the MNO framework does not directly apply. Second,  
669 although we develop an approximation theorem for finite period, it does not hold for infinite time  
670 horizon.

671 As discussed previously, it is infeasible to track the exact trajectory of chaotic systems on an infinite  
672 time horizon. Even very small errors will accumulate in each step, and eventually cause the simulation  
673 to diverge from the true trajectory. However, it is possible to track the attractor of the system. An  
674 attractor is absorbing. If the simulated trajectory only makes a small error, the attractor will absorb  
675 it back, so that the simulated trajectory will never diverge from the true attractor. Therefore, it is  
676 possible to have the simulated trajectory capture the true attractor.

677 To obtain an infinite-time approximation error bound is non-trivial. Previously, [9, 10] (cf. Theorem  
678 3.12) show a result for (finite-dimensional) ODE systems. If the system is Lipschitz then there exists  
679 a numerical simulation that forms a dissipative dynamical system that does not blow up or collapse.  
680 And the the simulated attractor  $\mathcal{A}_h$  approximates the true attractor  $\mathcal{A}$  with the time step  $h$

$$\text{dist}(\mathcal{A}_h, \mathcal{A}) \rightarrow 0, \quad \text{as } h \rightarrow 0$$

681 To generalize such theorem to Markov neural operator (MNO), we need to overcome two difficulties  
682 (1) generalize the formulation from (finite-dimensional) ODE systems to (infinite-dimensional) PDE  
683 systems, and (2) show MNO can obtain a sufficient error rate with respect to the time step  $h$ .

684 The first aspect requires extending the theory from finite dimension to infinite dimension, which  
685 is non-trivial since the operator  $F$  in (1) is not compact or bounded. This makes it hard to bound  
686 the error with respect to the attractor [37]. The second aspect requires to formulate MNO slightly  
687 differently. In the current formulation, the evolution operator is chosen for a fixed time step  $h$ . To  
688 achieve  $O(h)$  error we need to formulate the evolution operator continuously for infinitesimal  $h$ .  
689 Especially, for a semi-linear PDE system

$$\frac{du}{dt} + Au = F(u)$$

690 where  $A$  is a linear, self-adjoint operator and  $F$  is a continuous but nonlinear operator (this formulation  
691 includes the KS and NS equations). The evolution can be written as

$$u(t+h) = e^{-Ah}u(t) + \int_0^h e^{-A(h-s)}F(u(t+s))ds$$

692 Where  $\Phi(u(t), A, t) := \int_0^h e^{-A(h-s)}F(u(s))ds$  is bounded despite  $F$  is not. If one can approximate  
693  $\Phi(u(t), A, h)$  by a neural operator, then MNO can potentially achieve the needed error rate. This  
694 shows hope to obtain an approximation error bound for infinite time zero. We leave this as a promising  
695 future direction.