

000 001 DYNAMICAL SYSTEM RECONSTRUCTION FROM PAR- 002 TIAL OBSERVATIONS USING STOCHASTIC DYNAMICS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

011 Learning stochastic models of dynamical systems underlying observed data is of
012 interest in many scientific fields. Here we propose a novel method for this task,
013 based on the framework of variational autoencoders for dynamical systems. The
014 method estimates from the data both the system state trajectories and noise time
015 series. This approach allows to perform multi-step system evolution and supports
016 a teacher forcing strategy, alleviating limitations of autoencoder-based approaches
017 for stochastic systems. We demonstrate the performance of the proposed approach
018 on six test problems, covering simulated and experimental data. We further show
019 the effects of the teacher forcing interval on the nature of the internal dynamics,
020 and compare it to the deterministic models with equivalent architecture.

021 022 1 INTRODUCTION 023

024 Many scientific fields are concerned with building mathematical models of dynamical systems un-
025 derlying the observed data. Over the last decade, works using artificial neural networks to achieve
026 this goal in data-driven fashion have emerged, showing considerable promise (Durstewitz et al.,
027 2023; Legaard et al., 2023). Compared with the related task of time series prediction, the problem of
028 dynamical system reconstruction (DSR) is distinguished by the following aspects: focus on the long-
029 term dynamics of the trained system, an interest in the interpretable structure of the variables and the
030 latent space, or reasoning and analysis of the learned system by the conceptual and computational
031 tools from the dynamical system theory.

032 Many of the influential DSR approaches assume that the underlying dynamics is deterministic (Brun-
033 ton et al., 2016; Pandarinath et al., 2018; Chen et al., 2018; Hess et al., 2023). Indeed, finding a
034 deterministic model which can accurately predict future behavior, matches the long-term properties
035 observed in the data, and correctly generalizes to new conditions can be viewed as the ultimate goal
036 of DSR. However, reaching such a goal is often infeasible due to the complexity of the underlying
037 process or due to limited experimental data. Furthermore, reconstructing a deterministic model in
038 full might not even be desirable if the resulting model is too computationally demanding for the in-
039 tended purposes. In such cases, relying on a stochastic framework might be a preferred alternative.
040 Stochastic models include a source of noise in the dynamical equations, in addition to the noisy ob-
041 servations that are commonly used with deterministic models. This system-level noise can represent
042 the elements of the system not explicitly modeled by the deterministic part, thus potentially easing
043 the demands on the deterministic dynamics that needs to be learned. Training of such models, how-
044 ever, requires different approaches than for deterministic ones. And while various approaches were
045 explored (Linderman et al., 2017; Duncker et al., 2019; Kramer et al., 2022; Course & Nair, 2023;
046 Pals et al., 2024, and others discussed in the following section), research into robust and efficient
047 algorithms with demonstrated ability to learn noise-driven dynamics with various characteristics is
ongoing.

048 The difficulty of training the system dynamics is influenced by the available observations. In this
049 work we focus on systems that are only partially observed, that is, where the number of observed
050 variables is lower than the assumed dimensionality of the system. For such under-observed systems
051 the time delay embedding (Takens, 1981) in some of its variants (Kraemer et al., 2021) can be used to
052 reconstruct the trajectory in the state space. Despite some theoretical results on forced and stochastic
053 systems (Stark et al., 1997), such attempts are in practice mostly limited to autonomous determin-
istic systems. Alternatively, the state trajectory has to be estimated jointly with the training of the

054 dynamical system. This can be done either directly by optimizing the state variables (Koppe et al.,
 055 2019; Kramer et al., 2022), or, more commonly when the training dataset is large, through training
 056 an encoder network that maps the observation onto the latent space. Partially observed systems re-
 057 quire gathering information from across the temporal dimension, using recurrent neural networks
 058 (Pandarinath et al., 2018), temporal convolutional networks (Brenner et al., 2024; Pals et al., 2024),
 059 or their combination (García et al., 2022). But these approaches were not yet sufficiently explored,
 060 in particular their performance with algorithms used for training stochastic dynamical systems.

061 In this work, we propose a novel method for DSR of partially observed systems based on stochastic
 062 dynamics. To ensure robust training and the ability to capture long-timescale patterns, we introduce
 063 a double projection approach, where we map the observations to both the system states and noise
 064 time series, on which we train the dynamical model.

065 Our main contributions are the following:

066

- 067 • First, we propose a novel method for stochastic DSR from minimally observed systems,
 068 based on variational autoencoders for dynamical systems and using dual encoding of ob-
 069 servation into the state space and noise space.
- 070 • Second, we test the method on six test problems, including models of both deterministic
 071 chaos and noise-driven dynamics, and experimental data.
- 072 • Third, we analyze the nature of the learned dynamics, and investigate the role of the pre-
 073 diction window on it.

074

075 2 RELATED WORK

076

077 A range of methods to train a stochastic dynamical system from the data were developed and ex-
 078 plored in recent years. An approach commonly employed for this task - on which we also build in
 079 this work - merges probabilistic state space models with the framework of variational autoencoders
 080 (VAE, Kingma & Welling, 2014). The key components of the approach are an encoder (or recogni-
 081 tion) network mapping the observations into the time series of latent variables, a discrete-time state
 082 space model parameterized by a flexible neural network, a decoder (or observation) mapping from
 083 the latent space back to the observations, and a training method based on minimizing the evidence
 084 lower bound (ELBO). The pioneering works often focused on other applications than DSR, among
 085 them are Deep Kalman Filters (Krishnan et al., 2015; 2017) with latent variables being the states of
 086 the system, Stochastic Recurrent Networks (Bayer & Osendorfer, 2015) with latent variables being
 087 the noise time series, or Variational Recurrent Neural Networks (Chung et al., 2016). These and
 088 other related works are reviewed by Girin et al. (2021).

089 In the explicit context of dynamical system reconstruction, Kramer et al. (2022) used the VAE
 090 framework to integrate multimodal data. Brenner et al. (2024) applied it together with a teacher
 091 forcing strategy proposed by Hess et al. (2023) for more robust training. Hernandez et al. (2020)
 092 trained state-dependent linear networks, reusing the generative model inside the recognition model.
 093 Sip et al. (2023) used coupled stochastic models to learn a model of brain network dynamics.

094 Different methods than dynamical VAEs for training neural network based models of stochastic dy-
 095 namics exist. Koppe et al. (2019) used the expectation-maximization (EM) method, where the latent
 096 states of the dynamical model are optimized directly, as opposed to through an encoder. Kramer
 097 et al. (2022) compared the EM method with a VAE-based approach, concluding better performance
 098 of EM compared to VAE for smaller problems, but at a cost of limited flexibility. Pals et al. (2024)
 099 trained stochastic low-rank recurrent neural networks using the variational sequential Monte Carlo
 100 method (Naesseth et al., 2018).

101 Other works use a different parameterization than neural networks for the stochastic generative
 102 model. Linderman et al. (2017) rely on a collection of linear dynamical systems with state-dependent
 103 probabilistic switching between them, leading to an interpretable representation. Using Gaussian
 104 processes to represent the dynamics of the generative model (Doerr et al., 2018; García et al., 2022)
 105 allows to naturally introduce the notion of uncertainty of the dynamics, but it requires careful choice
 106 of the kernel and is limited to state spaces of low dimensions.

107 The works reviewed so far use discrete-time models that can be viewed as an Euler-Maruyama
 108 discretization of a continuous stochastic differential equation. The gradients of the loss function

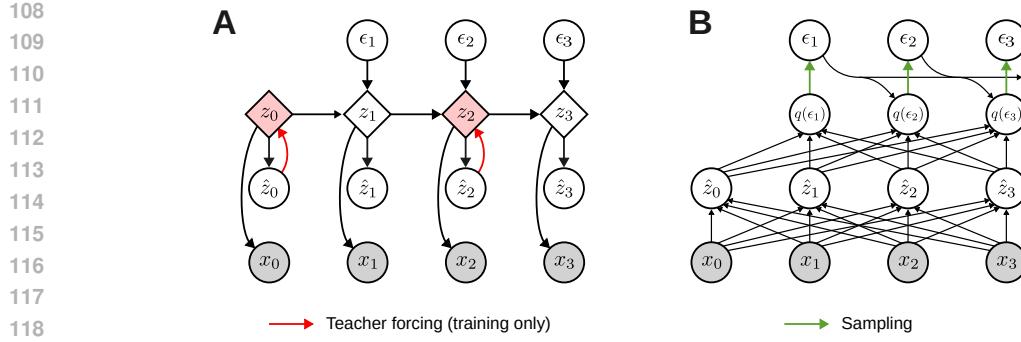


Figure 1: Graphical summary of the DPDSR method. (A) Generation, visualized for teacher forcing interval $\tau = 2$. (B) Encoding. For brevity, we use a shorthand for the posterior distributions $q(\epsilon_t)$ instead of $q(\epsilon_t | x, \hat{z}, \epsilon_{1:t-1})$.

necessary for the training are calculated from this discrete formulation, typically using automatic differentiation from modern ML toolboxes: an approach described as discretize-then-differentiate. Continuous-time models proceed in the other direction, differentiate-then-discretize, by formulating differential equations for the gradients, which can then be solved using any appropriate numerical scheme. Among their strengths is the ability to naturally deal with irregularly-spaced observations and better memory scaling properties. Li et al. (2020) generalized the neural ODE framework (Chen et al., 2018) for training continuous-time neural network based stochastic equations. Course & Nair (2023) proposed an efficient way to train the continuous-time stochastic model that avoids solving the differential equation through an ELBO reparameterization.

Dynamical models based on Gaussian processes were also applied for continuous-time approaches. Duncker et al. (2019) used Gaussian processes conditioned on the position of the fixed points and the associated Jacobians, leading to easily interpretable models. Hu et al. (2024) later extended the framework and integrated the ideas of switching linear models (Linderman et al., 2017) by introducing a novel smoothly-switching linear kernel.

3 METHODS

3.1 DOUBLE PROJECTION DYNAMICAL SYSTEM RECONSTRUCTION

In this study, we consider a dataset of N system observations with dimension d_x over T time steps $\{x_{1:T}^{(i)} \in \mathbb{R}^{T \times d_x}\}_{i=1}^N$. Our aim is to learn an underlying dynamical system

$$\begin{aligned} z_t &= F(z_{t-1}, \epsilon_t), \\ x_t &= g(z_t) + \Sigma_\eta \eta_t, \end{aligned}$$

with system states $z_t \in \mathbb{R}^{d_z}$, system noise $\epsilon_t \in \mathbb{R}^{d_\epsilon}$ of possibly lower dimension $d_\epsilon \leq d_z$, and observation noise $\eta_t \in \mathbb{R}^{d_x}$.

For this goal, we introduce a novel DSR method. We motivate it by the following reasoning: VAE-based approaches are powerful and flexible, but suffer from different drawbacks. For methods where system states $z_{1:T}$ are considered the latent variables (e.g. Deep Kalman Filter (Krishnan et al., 2015; 2017)), the system dynamics are in the loss function expressed by one step transition probability $p(z_t | z_{t-1})$. The system is never let to evolve more than one step from the estimated states, which limits its ability to learn long-term dependencies. On the other hand, for methods where the noise $\epsilon_{1:T}$ are the latent variables (e.g. STORN (Bayer & Osendorfer, 2015)), the absence of known system states prohibits the use of teacher forcing strategies, crucial for training deterministic systems Mikhaeil et al. (2022); Hess et al. (2023).

To circumvent this, we propose a method that uses trained encoders to estimate both the system states and the system noise: Double Projection Dynamical System Reconstruction (DPDSR, Fig. 1). Given one sample of the observed time series $x_{1:T}$, it first estimates the (possibly partial) state space

162 trajectory $\hat{z}_{1:T}$, and subsequently also the noise time series $\epsilon_{1:T}$. Then, starting from the estimated
 163 initial conditions, the system is evolved according to the trained dynamical system and using the
 164 estimated noise time series. Every τ steps, the state of the dynamical system is set to the estimated
 165 state \hat{z}_t . To calculate the loss function, the match of the generated trajectories to the observations
 166 $\mathbf{x}_{1:T}$ and the estimated state space trajectory $\hat{z}_{1:T}$ is combined with the Kullback-Leibler (KL)
 167 divergence of the latent variables $\epsilon_{1:T}$ from the white noise prior.

168 In the following text, we write $\mathbf{x} = \mathbf{x}_{1:T}$ (and analogously for other variables) for readability.
 169

170 **Generative model** We consider a generative model of the following form:
 171

$$\begin{aligned} z_t &= \tanh(f(z_{t-1}) + B\epsilon_t), \\ x_t &= g(z_t) + \Sigma_\eta \eta_t. \end{aligned} \quad (1)$$

172 The evolution function has the form of a two-layer perceptron with residual connection,
 173

$$f(z_t) = z_t + W_2 \sigma(W_1 z_t + b_1) + b_2, \quad (2)$$

174 with $\sigma(z)$ being the ReLU function. The tanh nonlinearity is added to stabilize the dynamics and
 175 training by constraining the states to a finite interval. The observation function is a two-layer per-
 176 ceptron, $g(z_t) = W_2^g \sigma(W_1^g z_t + b_1^g) + b_2^g$. In the examples presented here, we use a one-dimensional
 177 noise ($d_\epsilon = 1$) injected into the last dimension only, so that $B = [0, \dots, 0, \sigma_\epsilon^2]^T \in \mathbb{R}^{d_z \times 1}$. The
 178 observation covariance matrix is diagonal and isotropic, $\Sigma_\eta = \sigma_\eta^2 I$.
 179

180 **Encoder** The encoding process has two steps. First, from the observed timeseries $\mathbf{x} \in \mathbb{R}^{T \times d_x}$ we
 181 compute a deterministic estimation of the system state timeseries $\hat{z} \in \mathbb{R}^{T \times d_z}$. These can be possibly
 182 only partial estimation of some dimensions of the states, $d_{\hat{z}} \leq d_z$. We use a WaveNet architecture
 183 (van den Oord et al., 2016), which is based on 1D dilated convolutional networks. Unless specified
 184 otherwise, we use a single stack of seven dilated convolutional layers (Tab. S1). The output of the
 185 last layer is linearly projected to an estimation \hat{z} . We are using non-causal layers, although we also
 186 train an auxilliary causal stack, which we use for prediction tasks (Sec. A.3.1).
 187

188 In the second step, we estimate the posterior distribution $q(\epsilon | \mathbf{x}, \hat{z})$ from the observed time series
 189 \mathbf{x} and estimated states \hat{z} . The noise ϵ serve as the latent variable in our VAE framework. We use an
 190 autoregressive Gaussian posterior, $q(\epsilon_t | \mathbf{x}, \hat{z}, \epsilon_{1:t-1}) = N(\epsilon_t | \mu_t(\mathbf{x}, \hat{z}, \epsilon_{1:t-1}), \sigma_t(\mathbf{x}, \hat{z}, \epsilon_{1:t-1}))$.
 191 The means μ_t and variances σ_t are computed by passing the input timeseries through a WaveNet
 192 block with the same architecture as in the first step, followed by an autoregressive LSTM with
 193 probabilistic output. This autoregressive form of the posterior distribution serves to increase the
 194 expressivity of the encoder of the posterior distribution.
 195

196 **Training and loss function** For each time series \mathbf{x} , the data is first projected to estimate of state
 197 \hat{z} and noise ϵ via the described encoders. If the projection to state space is only partial, the initial
 198 conditions are completed by a trainable linear projection to the remained states, $\tilde{z}_0 = [\hat{z}_0; f_{init}(\hat{z}_0)]$.
 199 Otherwise, the initial projected state is used, $\tilde{z}_0 = \hat{z}_0$. The system is then evolved according to
 200 the generative model (1), using a random sample ϵ from the posterior distribution. Using teacher
 201 forcing, every τ -th step the simulated state is replaced by the estimated state,
 202

$$\tilde{z}_{t+1} = \begin{cases} \tanh(f(\tilde{z}_t) + B\epsilon_t) & \text{if } t \bmod \tau \neq 0, \\ \tanh(f([\hat{z}_t; T_{d_{\hat{z}}:d_z} \tilde{z}_t]) + B\epsilon_t) & \text{if } t \bmod \tau = 0. \end{cases} \quad (3)$$

203 Here, $T_{d_{\hat{z}}:d_z}$ is a truncation matrix selecting the elements from the $d_{\hat{z}}$ -th to the last d_z -th dimension;
 204 that is, the remaining states for which the state was not estimated are left to evolve freely.
 205

206 For each sample \mathbf{x} , the loss function is composed from the reconstruction loss of \mathbf{x} and \hat{z} , and the
 207 KL term of the latent noise variables ϵ
 208

$$L = L_x^{\text{rec}} + L_{\hat{z}}^{\text{rec}} + L^{\text{KL}}. \quad (4)$$

209 The reconstruction loss of the observation is
 210

$$L_x^{\text{rec}} = \mathbb{E}_{\epsilon \sim q(\epsilon | \mathbf{x}, \hat{z})}[-\log p(\mathbf{x} | \hat{z})], \quad (5)$$

216 where $p(\mathbf{x} \mid \tilde{\mathbf{z}}) = \prod_t p(x_t \mid g(\tilde{z}_t), \Sigma_\eta)$ using the observation function g and covariance Σ_η from
 217 (1). We treat the estimated states as additional observations with identity mapping, we find that such
 218 an approach helps to stabilize the training. Therefore, the reconstruction loss of the estimated partial
 219 states is similar to the observation reconstruction loss, apart from replacing the observation operator
 220 by the identity function on the estimated states:

$$221 \quad L_{\tilde{\mathbf{z}}}^{\text{rec}} = \mathbb{E}_{\epsilon \sim q(\epsilon \mid \mathbf{x}, \tilde{\mathbf{z}})}[-\log p(\tilde{\mathbf{z}} \mid \tilde{\mathbf{z}})], \quad (6)$$

223 where $p(\tilde{\mathbf{z}} \mid \tilde{\mathbf{z}}) = \prod_t p(\tilde{z}_t \mid T_{1:d_{\tilde{\mathbf{z}}}} \tilde{z}_t, \Sigma_{\tilde{\mathbf{z}}})$ with $T_{1:d_{\tilde{\mathbf{z}}}}$ being the truncation matrix selecting the first
 224 $d_{\tilde{\mathbf{z}}}$ states, and the covariance matrix being diagonal and isotropic, $\Sigma_{\tilde{\mathbf{z}}} = \sigma_{\tilde{\mathbf{z}}}^2 I$. Finally, the KL term
 225 in the loss follows the standard formulation of variational autoencoders,

$$226 \quad L^{\text{KL}} = D_{\text{KL}}(q(\epsilon \mid \mathbf{x}, \tilde{\mathbf{z}}) \parallel p(\epsilon)). \quad (7)$$

227 We use a prior of standard normal distribution $p(\epsilon) = N(\epsilon \mid 0, I)$. In addition, two regularization
 228 terms are added to the loss. First, to favour sparse observation models, it is L1 regularization on
 229 the weights of projection g . Second, to encourage desired scale of the state trajectories, it is a
 230 regularization term for the scale and position of the estimated states. The loss function is minimized
 231 using Adam optimizer. Further details of the architecture and training are given in Sec. A.3.1.

233 3.2 COMPARED METHODS

235 **Single projection DSR (SPDSR)** is a deterministic variant of DPDSR. It uses the same architecture
 236 and teacher forcing training method, but it assumes that there is no noise in the dynamical system.
 237 The noise encoder is therefore absent, and the KL divergence term does not appear in the loss
 238 function (4). **Generalized teacher forcing** (Hess et al., 2023) is a method for reconstruction of
 239 deterministic dynamical system, and has been shown to outperform other deterministic methods.
 240 We test two variants: first, using partial forcing (GTF-PF), where only the one observed variable
 241 is used as an incomplete teacher signal. And second, using the time-delay embedding (GTF-TD)
 242 so that the full state can be forced. We used PECUZAL method for the time delay embedding
 243 (Kraemer et al., 2021), and where this failed, we defaulted to a delay embedding with constant
 244 offset and predetermined dimension $d = 8$. **Deep Kalman Filter** (DKF) (Krishnan et al., 2015;
 245 2017) is a method for reconstruction of stochastic systems. Using an encoder, it estimates the states
 246 of the system, while the one-step stochastic prediction of the generative model forms the basis of
 247 the loss function. Our implementation here mirrors the architecture of our proposed method where
 248 possible (in the architecture of the encoder and of the generative model), with the main difference
 249 being the form of the loss function. **Autoregressive LSTM** (AR-LSTM) (Graves, 2014) uses a
 250 standard LSTM network whose probabilistic output is fed back to the network at the next step, thus
 251 forming a stochastic dynamical system. During training, the original time series are used to feed
 252 the network, or can be replaced by the model generated output following the ideas of scheduled
 253 sampling (Bengio et al., 2015). The generative models of all methods were approximately matched
 254 in number of parameters (Tab. S2).

255 4 RESULTS

257 **Example: Double well model** First, we demonstrate the potential of our method and the limits of
 258 the alternatives on an example of stochastic dynamics: a noise driven double well model (Fig. 2).
 259 The double well dataset (Sec. A.1.3) is generated by numerical integration of a stochastic differential
 260 equation of bistable dynamics followed by four layers of exponential smoothing,

$$261 \quad \dot{z}_1 = -z_1^3 + z_1 + \sigma \eta(t), \\ 262 \quad \dot{z}_i = \alpha(z_{i-1} - z_i) \quad \text{for } i \in \{2, 3, 4, 5\},$$

263 with parameters σ scaling the noise amplitude and α being the temporal constant of exponential
 264 smoothing. We assume only the last variable z_5 is observable. The model has two stable fixed
 265 points, $z_i = \pm 1$, $i \in \{1, \dots, 5\}$.

266 As shown on Fig. 2C, reconstruction methods based on deterministic dynamics do not perform
 267 well on this problem. A good time delay embedding cannot be found for time series generated by
 268 stochastic dynamics, and GTF-TD fails to reproduce the bistable nature of the time series. Also par-
 269 tial forcing (GTF-PF) is not sufficient. The embedding-based method using deterministic dynamics

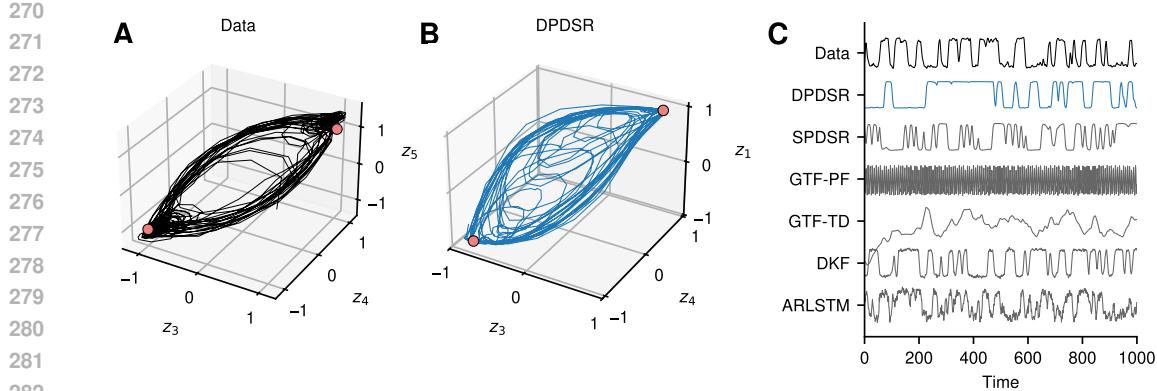


Figure 2: Results for the double well dataset. (A) Trajectory in the original state space of the model. The dots represent the stable fixed points of the model. (B) Simulated trajectory in the state space of the DPDSR model (hand-picked dimensions). (C) Original time series, and time series generated by the trained models. Here, and in all other figures, the time is represented in sample indices and not the original model time.

(SPDSR) performs better and is able to reproduce the bimodal distribution of the data via chaotic dynamics, but a quantitative evaluation in the following section will reveal the inferior match compared to the stochastic method. Among the stochastic methods, both DPDSR and DKF reproduce the bimodal nature of the data well, but DKF suffers from other problems, as we show next. Autoregressive LSTM performs better than the deterministic methods, but worse than embedding-based stochastic methods.

Quantitative evaluation For a more detailed evaluation of the DPDSR method, we test it on six datasets (Sec. A.1): 1) Lorenz, the well-known classic example of deterministic chaos (Lorenz, 1963). Even though training models of deterministic chaos is not the main target of the work, it is an important special case for stochastic methods too; 2) Cell cycle, a six-dimensional model of cell division cycle, generated by a deterministic chaotic model (Romond et al., 1999); 3) Double well, the dataset described above; 4) RNN, time series generated by a large recurrent neural network in chaotic regime, of which only one variable is observed. Due to the complexity of the model and minimal observation, we expect that the deterministic dynamics cannot be recovered, and stochastic model would prove to be a better alternative; 5) Neuron, an experimental dataset of somatic voltage of a rat pyramidal neuron driven by random and unknown stimulus current; 6) ECG, an electrocardiogram signal of a healthy adult.

We evaluate the quality of the reconstruction based on three criteria (Sec. A.4). First, the distribution distance D_d : the similarity of the spatial distributions of the original and model-generated data. Specifically, we generate long time series with the trained model, collapse the data across time, and compare the distributions using Wasserstein distance. Second, the spectral distance D_s : the similarity of the original and model-generated data in the frequency domain. We compute the power spectra of the original and model-generated data, and compare those using the Hellinger distance. And third, the short-term prediction error PE_{20} of the model-generated trajectories when starting from a state estimated from past data. For the ECG dataset, we also include the distance of interspike intervals D_{ISI} as an important measure of the reconstruction quality. We summarize these measures in a cumulative score, which we define as a weighted sum of the described measures,

$$S = w_1 D_d + w_2 D_s + w_3 PE_{20} + w_4 D_{ISI}.$$

Considering the different nature of the datasets, we set the weights differently for the different datasets in order to reflect the importance of the measures for each specific dataset (Sec. A.4).

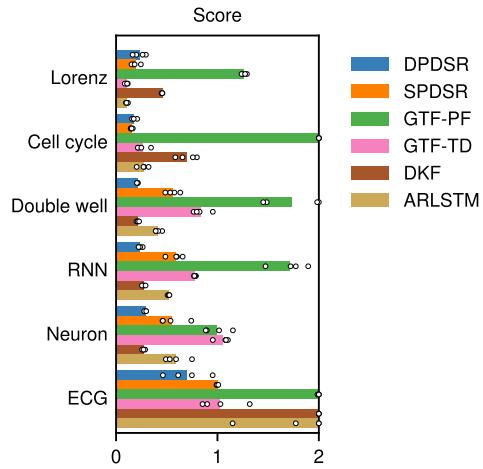
We train the models as described in Sec. 3 and Sec. A.3. For each method, we perform a parameter sweep over selected hyperparameters, and for each parameter combination we train the model with four random initializations. We then select the best model as the one with the lowest average score

324 across initializations. Examples of generated time series are on Fig. S4. The resulting scores for all
 325 problems are presented on Fig. 3, and the separate measures on Fig. S5.
 326

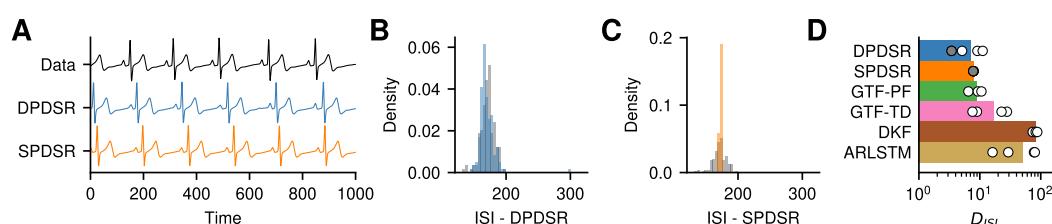
327 We summarize the results in three main points.
 328 First, for the datasets generated by low-
 329 dimensional deterministic models (Lorenz and
 330 Cell cycle), we see a good performance of de-
 331 terministic models, either based on time de-
 332 lay embedding (GTF-TD) or trained projection
 333 (SPDSR). Our stochastic model is, however,
 334 competitive in all three measures. DKF per-
 335 forms badly, mainly due to its limited predic-
 336 tion capacity.

337 Second, for the Double well, RNN, and Neu-
 338 ron datasets, where the time series are dom-
 339 inated by random (or seemingly random) ef-
 340 fects on short time scales, the stochastic mod-
 341 els DPDSR and DKF show the best per-
 342 formance. The deterministic models (SPDSR
 343 and GTF) have to approximate the random transi-
 344 tions through deterministic chaos; the SPDSR
 345 projection method does it better than the GTF-
 346 TD method based on the time embedding. We
 347 note that for these three dataset, the state of the
 348 art PECUZAL embedding failed and we used
 349 embedding with fixed delay and number of di-
 350 mension instead. Thus, a worse performance of
 351 the time embedding method can be expected.
 352 The stochastic AR-LSTM method stands in be-
 353 tween the two groups.

354 Third, for the ECG dataset, several method are capable of learning a dynamical system that can
 355 generate the periodic ECG signal (Fig. S4). However, closer inspection reveals the variations of the
 356 interspike intervals (ISI) in the data (Fig. 4). Unlike our stochastic DPDSR method (Fig. 4B), the
 357 deterministic methods do not reproduce the ISI variations (Fig. 4C,D), leading to higher spectral
 358 distance and higher overall score. Our stochastic method, however, can exhibit different undesirable
 359 behavior. The trained models can “skip a beat” (Fig. 4B); the prevalence of this behavior depends
 360 on the random initialization of the training, as demonstrated by the variable results in Fig. 4D.
 361



362 Figure 3: Main results showing the score (lower
 363 = better) for all datasets and methods. Each circle
 364 represents one of four initialization of the training
 365 method, and the bar shows the mean. The mea-
 366 sures from which the score is computed are shown
 367 on Fig. S5. For visualization purposes, the values
 368 are clipped to the upper limit of the shown range.



368 Figure 4: Results for the ECG dataset. (A) Example time series: original, and simulated by the
 369 proposed stochastic model (DPDSR) and its deterministic variant (SPDSR). (B) Distribution of the
 370 interspike intervals (ISI) in the data (gray) and generated by the DPDSR model. The two dis-
 371 tributions mostly overlap, but note the non zero bin near interval 300 indicating skipped beat in the
 372 simulated data. (C) Distribution of the ISI in the data (gray) and generated by the deterministic
 373 SPDSR model. The model generated data show periodic behavior with all ISIs concentrated on
 374 interval 168. (D) Distance between the data and model-generated ISI distributions for all methods.
 375 Each circle represent a different training initialization, bar shows their mean. The gray circle repre-
 376 sent the model used in panels A-C. The match of ISI distributions in the DPDSR model is strongly
 377 dependent on the initialization, which affects the proportion of skipped beats.

378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
Internal dynamics of the trained models To better understand the behavior of the proposed DPDSR method, we now analyze the trained models using tools from dynamical systems theory. Specifically, we look at the nature of the attractors of the deterministic part of the trained generative models. For contrast, we compare the results with those from the deterministic SPDSR models. Apart from the absence of the noise, to SPDSR method is identical to the DPDSR in the architecture and training procedure. As such, it provides a useful comparison point for the differences of the stochastic and deterministic methods.

In Fig. 5A we show the maximal Lyapunov exponent of the attractors in the trained models for all six datasets. For each dataset and method, we first detect the attractors by evolving the system from randomly initialized points, and then estimated the maximal Lyapunov exponent (Sec. A.5). The results show that the DPDSR method learns dynamical system exhibiting deterministic chaos for the Lorenz and Cell cycle dataset, both of which are indeed generated by a chaotic system. As shown also on Fig. 2, the model trained on the Double well problem has two stable fixed points (with noise driven transitions between them). The similar RNN problem results in a noise-driven fluctuations around a single fixed point. The variations in the interspike intervals in the ECG model are due to noisy fluctuations around a stable limit cycle.

In contrast, the best performing deterministic models trained with the SPDSR method exhibit chaotic dynamics in all cases. Indeed, in the absence of noise, deterministic chaos is the only option to model the random transitions of the Double well and RNN datasets, or spikes of the Neuron datasets.

Next, we explore the role of the teacher forcing interval τ on the performance and nature of the dynamics of the trained models. The double well model (Fig. 5B,C) illustrates the phenomena that are to a large extent consistent across all datasets (Fig. S6). In the stochastic DPDSR model, we observe that for smaller intervals ($\tau \leq 40$) the trained models contain chaotic attractors, and they do not rely on the noise. We quantify this by the KL divergence of the estimated posterior distribution of noise from the prior distribution, $KL(q(\epsilon | \mathbf{x}) \parallel p(\epsilon))$, averaged across all samples. If the two distributions overlap, the KL divergence is zero, and no information is encoded in the posterior distribution. As the interval τ increases, the models transition into the regime of noise-driven dynamics with two stable fixed points. This is complemented by the increasing importance of the noise, so that the observed dynamics can be generated by noise-driven fluctuations.

Such behavior is mostly consistent for DPDSR models across datasets, with variations in the position of the transition into the noise-driven regime or the number and nature (fixed point or limit cycle) of the stable attractors. However, the position of the optimal model can differ: it is located in the deterministic chaos regime for Lorenz and Cell cycle, or stochastic regime for Double well,

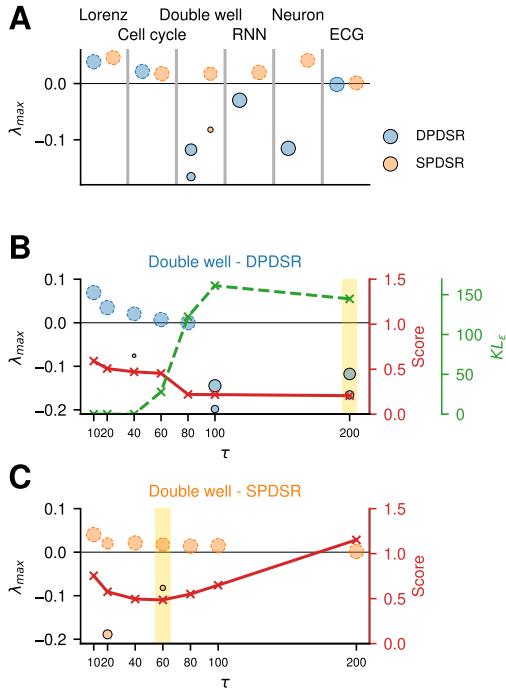


Figure 5: Analysis of the attractors of the trained models. (A) Maximal Lyapunov exponent λ_{\max} for the models trained on the six datasets. Each point correspond to one attractor. Dashed circle outline represents a chaotic attractor ($\lambda_{\max} > 0$), colored solid outline represents a limit cycle, and black solid outline a fixed point. Size of the circle corresponds to the size of the basin of attraction. (B) Influence of the teacher forcing interval τ on the attractors in the DPDSR models. Circles show the maximal Lyapunov exponent as in (A). Solid red line shows the score (lower = better). Dashed green line shows the KL divergence of the estimated posterior distribution of noise to the prior distribution $KL_{\epsilon} = D_{KL}(q(\epsilon | \mathbf{x}) \parallel p(\epsilon))$. The yellow band indicates the optimal τ value for the dataset. (C) As in (B), but for the deterministic SPDSR models.

432 RNN, Neuron, and ECG problems. In the deterministic SPDSR models, contrastingly, the stochastic
 433 regime cannot exist by construction. The models therefore mainly stay in the chaotic regime. The
 434 best performing models are found for lower teacher forcing intervals τ , and the score worsens for
 435 increasing values. This behavior matches the results reported by Mikhaeil et al. (2022) when training
 436 deterministic models using a teacher forcing scheme on the data from Lorenz system, forced Duffing
 437 oscillator, and empirical temperature time series.

438

439

440

441 5 DISCUSSION

442

443

444 In this work we have introduced a novel method for reconstructing stochastic dynamical models
 445 from the data. The method is based on a double projection approach, where the time series of the
 446 observations are projected onto the estimates of the system states and the estimates of the driving
 447 noise using trained encoders. The estimated system states are used for teacher forcing, while the
 448 driving noise is used as the latent variable in a variational autoencoder framework. We benchmark
 449 the method on six test problems, demonstrating its performance for data generated by noise-driven
 450 models, deterministic chaotic models, and empirical data (Fig. 3). While other methods, both based
 451 on deterministic and stochastic models, can provide equal or even better performance for some of
 452 the problems, only our method performs competitively for all tested problems. We then analyze
 453 the nature of the attractors in the learned dynamics in the examined problems, and evaluate the
 454 role of the teacher forcing interval τ on it (Fig. 5 and Fig. S6). We identify the existence of two
 455 regimes occurring across the test problems: a deterministic regime for lower values of τ with chaotic
 456 attractors, and a noise-driven regime for higher values of τ .

457 The proposed method has some notable shortcomings and limitations. First, it is the dependency
 458 of the behavior of the trained model on the teacher forcing interval. Performing parameter sweeps
 459 across a range of values to find the best performing model, while robust, is computationally de-
 460 manding. As analyzed in Sec. B.1, our results indicate that an adaptive scheme for τ , akin to the
 461 scheme of Hess et al. (2023), might not be sufficient to converge to the optimally performing model.
 462 The question of optimal strategy for setting the teacher forcing interval thus remains open for future
 463 studies.

464 We have chosen a parameterization of the dynamical system using fully connected neural networks.
 465 Such approach offers flexibility in terms of the dynamics it can produce, but limited interpretability,
 466 desirable for dynamical system reconstruction. Other works chose different strategies and trade-
 467 offs. Symbolic approaches recovering sparse models using predefined function basis (Brunton et al.,
 468 2016; Champion et al., 2019) lead to models that are open to pen-and-paper manipulation, but it
 469 can be at a cost of reduced performance on empirical datasets (Hess et al., 2023). For specific
 470 forms of neural networks with piecewise linear ReLU activation, the fixed points and limit cycles
 471 can be found analytically in low dimensions (Schmidt et al., 2021; Brenner et al., 2022; Hess et al.,
 472 2023; Pals et al., 2024), greatly simplifying their analysis. Alternatively, Duncker et al. (2019) used
 473 Gaussian processes conditioned on the position of fixed points, which are then directly available
 474 after training. Switching linear dynamical systems (Linderman et al., 2017; Hu et al., 2024) offer
 475 some interpretability by decomposing the state space into regions of linear dynamics.

476 The ability of a dynamical system reconstruction method to robustly learn dynamics with diverse
 477 timescales is of great importance for many problems, and improvements could be made to our pro-
 478 posed methods in this direction. Gated variants of recurrent neural networks, most notably LSTM
 479 (Hochreiter & Schmidhuber, 1997) and GRU (Cho et al., 2014), were designed to deal with long-
 480 range temporal dependencies. Other approaches with explicit time scale separation were suggested
 481 for DSR with deterministic models, often outperforming the traditional architectures. Among the
 482 proposed approaches were: using multiple coupled RNNs operating with different temporal resolu-
 483 tion (Liu et al., 2022; Farooq et al., 2024), regularization of the parameters of the neural network
 484 to introduce multiple time scales (Schmidt et al., 2021), separating time scales using dynamic mode
 485 decomposition before applying the DSR algorithm (Bramburger et al., 2020), or using echo state
 486 networks with different leak rates of leaky integrator neurons (Tanaka et al., 2022). Applying these
 487 ideas to stochastic models could open a way to more robust methods for problems with disparate
 488 time scales.

486 REPRODUCIBILITY STATEMENT
487488 The data and code associated with the study are available in the supplementary material.
489490 REFERENCES
491492 Justin Bayer and Christian Osendorfer. Learning Stochastic Recurrent Networks, March 2015. URL
493 <http://arxiv.org/abs/1411.7610>.494 Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL
495 https://proceedings.neurips.cc/paper_files/paper/2015/hash/e995f98d56967d946471af29d7bf99f1-Abstract.html.
496497 Jason J. Bramburger, Daniel Dylewsky, and J. Nathan Kutz. Sparse identification of slow timescale
498 dynamics. *Physical Review E*, 102(2):022204, August 2020. doi: 10.1103/PhysRevE.102.022204.
499 URL <https://doi.org/10.1103/PhysRevE.102.022204>.500 Manuel Brenner, Florian Hess, Jonas M. Mikhaeil, Leonard Bereska, Zahra Monfared, Po-Chen
501 Kuo, and Daniel Durstewitz. Tractable Dendritic RNNs for Reconstructing Nonlinear Dynamical
502 Systems. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162,
503 pp. 2292–2320. PMLR, July 2022. doi: 10.48550/arXiv.2207.02542. URL <http://arxiv.org/abs/2207.02542>.
504505 Manuel Brenner, Florian Hess, Georgia Koppe, and Daniel Durstewitz. Integrating Multimodal
506 Data for Joint Generative Modeling of Complex Dynamics. In *Proceedings of the 41st International
507 Conference on Machine Learning*, pp. 4482–4516. PMLR, July 2024. URL <https://proceedings.mlr.press/v235/brenner24a.html>.
508509 Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations
510 from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National
511 Academy of Sciences*, 113(15):3932–3937, April 2016. doi: 10.1073/pnas.1517384113. URL
512 <https://doi.org/10.1073/pnas.1517384113>.513 Kathleen Champion, Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Data-driven discovery
514 of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116
515 (45):22445–22451, November 2019. doi: 10.1073/pnas.1906995116. URL <https://doi.org/10.1073/pnas.1906995116>.
516517 Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural Ordinary
518 Differential Equations. In *Advances in Neural Information Processing Systems*, volume 31. Curran
519 Associates, Inc., 2018. URL https://papers.nips.cc/paper_files/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html.
520521 Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the Properties
522 of Neural Machine Translation: Encoder–Decoder Approaches. In Dekai Wu, Marine
523 Carpuat, Xavier Carreras, and Eva Maria Vecchi (eds.), *Proceedings of SSST-8, Eighth Workshop
524 on Syntax, Semantics and Structure in Statistical Translation*, pp. 103–111, Doha, Qatar,
525 October 2014. Association for Computational Linguistics. doi: 10.3115/v1/W14-4012. URL
526 <https://doi.org/10.3115/v1/W14-4012>.527 Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron Courville, and Yoshua Bengio.
528 A Recurrent Latent Variable Model for Sequential Data, April 2016. URL <http://arxiv.org/abs/1506.02216>.
529530 Kevin Course and Prasanth Nair. Amortized Reparametrization: Efficient and
531 Scalable Variational Inference for Latent SDEs. *Advances in Neural Information Processing
532 Systems*, 36:78296–78318, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/f72d4fdfd5eb425cd81df9fe6272a533-Abstract-Conference.html.
533

540 Andreas Doerr, Christian Daniel, Martin Schiegg, Nguyen-Tuong Duy, Stefan Schaal, Marc Tous-
 541 saint, and Trimpe Sebastian. Probabilistic Recurrent State-Space Models. In *Proceedings of the*
 542 *35th International Conference on Machine Learning*, pp. 1280–1289. PMLR, July 2018. URL
 543 <https://proceedings.mlr.press/v80/doerr18a.html>.

544 Lea Duncker, Gergo Bohner, Julien Boussard, and Maneesh Sahani. Learning interpretable
 545 continuous-time models of latent stochastic dynamical systems. In *Proceedings of the 36th Inter-*
 546 *national Conference on Machine Learning*, volume 97, pp. 1726–1734. PMLR, May 2019. URL
 547 <http://proceedings.mlr.press/v97/duncker19a.html>.

548 Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. Reconstructing computational system
 549 dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24
 550 (11):693–710, November 2023. ISSN 1471-0048. doi: 10.1038/s41583-023-00740-7. URL
 551 <https://doi.org/10.1038/s41583-023-00740-7>.

552 Junaid Farooq, Danish Rafiq, Pantelis R. Vlachas, and Mohammad Abid Bazaz. RefreshNet:
 553 Learning multiscale dynamics through hierarchical refreshing. *Nonlinear Dynamics*, 112(16):
 554 14479–14496, August 2024. ISSN 1573-269X. doi: 10.1007/s11071-024-09813-3. URL
 555 <https://doi.org/10.1007/s11071-024-09813-3>.

556 Constantino A. García, Paulo Félix, Jesús M. Presedo, and Abraham Otero. Stochastic embeddings
 557 of dynamical phenomena through variational autoencoders. *Journal of Computational Physics*,
 558 454:110970, April 2022. ISSN 0021-9991. doi: 10.1016/j.jcp.2022.110970. URL <https://doi.org/10.1016/j.jcp.2022.110970>.

559 William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling.
 560 In *Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Bench-*
 561 *marks Track (Round 2)*, August 2021. URL <https://openreview.net/forum?id=enYjtbjYJrf>.

562 Laurent Girin, Simon Leglaive, Xiaoyu Bie, Julien Diard, Thomas Hueber, and Xavier Alameda-
 563 Pineda. Dynamical Variational Autoencoders: A Comprehensive Review. *Foundations and*
 564 *Trends® in Machine Learning*, 15(1-2):1–175, 2021. ISSN 1935-8237, 1935-8245. doi:
 565 10.1561/2200000089. URL <http://arxiv.org/abs/2008.12595>.

566 Alex Graves. Generating Sequences With Recurrent Neural Networks, June 2014. URL <http://arxiv.org/abs/1308.0850>.

567 Daniel Hernandez, Antonio K. Moretti, Shreya Saxena, Ziqiang Wei, John Cunningham, and Liam
 568 Paninski. Nonlinear Evolution via Spatially-Dependent Linear Dynamics for Electrophysiology
 569 and Calcium Data. *Neurons, Behavior, Data analysis, and Theory*, 3(3), June 2020. URL <https://arxiv.org/abs/1811.02459v3>.

570 Florian Hess, Zahra Monfared, Manuel Brenner, and Daniel Durstewitz. Generalized Teacher Forc-
 571 ing for Learning Chaotic Dynamics. In *Proceedings of the 40th International Conference on*
 572 *Machine Learning*, pp. 13017–13049. PMLR, July 2023. URL <https://proceedings.mlr.press/v202/hess23a.html>.

573 Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Comput.*, 9(8):1735–
 574 1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.

575 Amber Hu, David Zoltowski, Aditya Nair, David Anderson, Lea Duncker, and Scott Linder-
 576 man. Modeling Latent Neural Dynamics with Gaussian Process Switching Linear Dynami-
 577 cal Systems. *Advances in Neural Information Processing Systems*, 37:33805–33835, Decem-
 578 ber 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/3b64416915026a6744bf10a819571041-Abstract-Conference.html.

579 Renaud Jolivet, Alexander Rauch, Hans-Rudolf Lüscher, and Wulfram Gerstner. Predicting spike
 580 timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational*
 581 *Neuroscience*, 21(1):35–49, August 2006. ISSN 1573-6873. doi: 10.1007/s10827-006-7074-5.
 582 URL <https://doi.org/10.1007/s10827-006-7074-5>.

594 Renaud Jolivet, Felix Schürmann, Thomas K. Berger, Richard Naud, Wulfram Gerstner, and Arnd
 595 Roth. The quantitative single-neuron modeling competition. *Biological Cybernetics*, 99(4):417–
 596 426, November 2008. ISSN 1432-0770. doi: 10.1007/s00422-008-0261-x. URL <https://doi.org/10.1007/s00422-008-0261-x>.

597

598 Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes, May 2014. URL <http://arxiv.org/abs/1312.6114>.

599

600 Georgia Koppe, Hazem Toutounji, Peter Kirsch, Stefanie Lis, and Daniel Durstewitz. Identifying
 601 nonlinear dynamical systems via generative recurrent neural networks with applications
 602 to fMRI. *PLOS Computational Biology*, 15(8):e1007263, August 2019. ISSN 1553-7358.
 603 doi: 10.1371/journal.pcbi.1007263. URL <https://dx.plos.org/10.1371/journal.pcbi.1007263>.

604

605 K H Kraemer, G Datseris, J Kurths, I Z Kiss, J L Ocampo-Espindola, and N Marwan. A unified and
 606 automated approach to attractor reconstruction. *New Journal of Physics*, 23(3):033017, March
 607 2021. ISSN 1367-2630. doi: 10.1088/1367-2630/abe336. URL <https://dx.doi.org/10.1088/1367-2630/abe336>.

608

609

610 Daniel Kramer, Philine L. Bommer, Carlo Tombolini, Georgia Koppe, and Daniel Durstewitz. Re-
 611 constructing Nonlinear Dynamical Systems from Multi-Modal Time Series. In *Proceedings of
 612 the 39th International Conference on Machine Learning*, pp. 11613–11633. PMLR, June 2022.
 613 URL <https://proceedings.mlr.press/v162/kramer22a.html>.

614

615 Rahul Krishnan, Uri Shalit, and David Sontag. Structured Inference Networks for Nonlinear State
 616 Space Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), February
 617 2017. ISSN 2374-3468. doi: 10.1609/aaai.v31i1.10779. URL <https://ojs.aaai.org/index.php/AAAI/article/view/10779>.

618

619 Rahul G. Krishnan, Uri Shalit, and David Sontag. Deep Kalman Filters, November 2015. URL
 620 <http://arxiv.org/abs/1511.05121>.

621

622 Christian Legaard, Thomas Schranz, Gerald Schweiger, Ján Drgoňa, Basak Falay, Cláudio Gomes,
 623 Alexandros Iosifidis, Mahdi Abkar, and Peter Larsen. Constructing Neural Network Based Mod-
 624 els for Simulating Dynamical Systems. *ACM Comput. Surv.*, 55(11):236:1–236:34, February
 625 2023. ISSN 0360-0300. doi: 10.1145/3567591. URL <https://doi.org/10.1145/3567591>.

626

627 Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable Gradients
 628 for Stochastic Differential Equations. *arXiv:2001.01328 [cs, math, stat]*, October 2020. URL
 629 <http://arxiv.org/abs/2001.01328>.

630

631 Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski.
 632 Bayesian Learning and Inference in Recurrent Switching Linear Dynamical Systems. In
 633 *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, vol-
 634 ume 54, pp. 914–922. PMLR, April 2017. URL <https://proceedings.mlr.press/v54/linderman17a.html>.

635

636 Yuying Liu, J. Nathan Kutz, and Steven L. Brunton. Hierarchical deep learning of multiscale differ-
 637 ential equation time-steppers. *Philosophical Transactions of the Royal Society A: Mathematical,
 638 Physical and Engineering Sciences*, 380(2229):20210200, June 2022. doi: 10.1098/rsta.2021.
 639 0200. URL <https://doi.org/10.1098/rsta.2021.0200>.

640

641 Edward N. Lorenz. Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*,
 642 20(2):130–141, March 1963. ISSN 0022-4928, 1520-0469. doi: 10.1175/1520-0469(1963)
 643 020<0130:DNF>2.0.CO;2. URL [https://doi.org/10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).

644

645 Jonas Mihkaeil, Zahra Monfared, and Daniel Durstewitz. On the difficulty of learning chaotic
 646 dynamics with RNNs. *Advances in Neural Information Processing Systems*, 35:11297–11312,
 647 December 2022. URL https://papers.neurips.cc/paper_files/paper/2022/hash/495e55f361708bedbab5d81f92048dcd-Abstract-Conference.html.

648 Christian Naesseth, Scott Linderman, Rajesh Ranganath, and David Blei. Variational Sequential
 649 Monte Carlo. In *Proceedings of the Twenty-First International Conference on Artificial Intel-*
 650 *ligence and Statistics*, pp. 968–977. PMLR, March 2018. URL <https://proceedings.mlr.press/v84/naesseth18a.html>.

651

652 Matthijs Pals, A. Erdem Sağtekin, Felix Pei, Manuel Gloeckler, and Jakob H. Macke.
 653 Inferring stochastic low-rank recurrent neural networks from neural data. *Advances*
 654 *in Neural Information Processing Systems*, 37:18225–18264, December 2024. URL
 655 https://proceedings.neurips.cc/paper_files/paper/2024/hash/209423f076b6479ab3a4f45886e30306-Abstract-Conference.html.

656

657

658 Chethan Pandarinath, Daniel J. O’Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D. Stavisky,
 659 Jonathan C. Kao, Eric M. Trautmann, Matthew T. Kaufman, Stephen I. Ryu, Leigh R. Hochberg,
 660 Jaimie M. Henderson, Krishna V. Shenoy, L. F. Abbott, and David Sussillo. Inferring single-trial
 661 neural population dynamics using sequential auto-encoders. *Nature Methods*, 15(10):805–815,
 662 October 2018. ISSN 1548-7091, 1548-7105. doi: 10.1038/s41592-018-0109-9. URL <https://doi.org/10.1038/s41592-018-0109-9>.

663

664 Alexander Rauch, Giancarlo La Camera, Hans-Rudolf Lüscher, Walter Senn, and Stefano Fusi.
 665 Neocortical Pyramidal Cells Respond as Integrate-and-Fire Neurons to In Vivo-Like Input Cur-
 666 rents. *Journal of Neurophysiology*, 90(3):1598–1612, September 2003. ISSN 0022-3077. doi:
 667 10.1152/jn.00293.2003. URL <https://doi.org/10.1152/jn.00293.2003>.

668

669 Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. Deep PPG: Large-Scale
 670 Heart Rate Estimation with Convolutional Neural Networks. *Sensors*, 19(14):3079, January
 671 2019. ISSN 1424-8220. doi: 10.3390/s19143079. URL <https://doi.org/10.3390/s19143079>.

672

673 Pierre-Charles Romond, Mauro Rustici, Didier Gonze, and Albert Goldbeter. Alternating Oscilla-
 674 tions and Chaos in a Model of Two Coupled Biochemical Oscillators Driving Successive Phases
 675 of the Cell Cycle. *Annals of the New York Academy of Sciences*, 879(1):180–193, 1999. ISSN
 676 1749-6632. doi: 10.1111/j.1749-6632.1999.tb10419.x. URL <https://doi.org/10.1111/j.1749-6632.1999.tb10419.x>.

677

678 Dominik Schmidt, Georgia Koppe, Zahra Monfared, Max Beutelspacher, and Daniel Durstewitz.
 679 Identifying nonlinear dynamical systems with multiple time scales and long-range dependencies.
 680 In *International Conference on Learning Representations*, March 2021. URL <http://arxiv.org/abs/1910.03471>.

681

682

683 Viktor Sip, Meysam Hasbemi, Timo Dickscheid, Katrin Amunts, Spase Petkoski, and Viktor Jirsa.
 684 Characterization of regional differences in resting-state fMRI with a data-driven network model of
 685 brain dynamics. *Science Advances*, 9(11):eabq7547, March 2023. doi: 10.1126/sciadv.abq7547.
 686 URL <https://doi.org/10.1126/sciadv.abq7547>.

687

688 H. Sompolinsky, A. Crisanti, and H. J. Sommers. Chaos in Random Neural Networks. *Physical*
 689 *Review Letters*, 61(3):259–262, July 1988. doi: 10.1103/PhysRevLett.61.259. URL <https://doi.org/10.1103/PhysRevLett.61.259>.

690

691 Julien Clinton Sprott. *Chaos and Time-Series Analysis*. Oxford University Press, January 2003.
 692 ISBN 978-0-19-850839-7. doi: 10.1093/oso/9780198508397.001.0001. URL <https://doi.org/10.1093/oso/9780198508397.001.0001>.

693

694 J. Stark, D. S. Broomhead, M. E. Davies, and J. Huke. Takens embedding theorems for forced
 695 and stochastic systems. *Nonlinear Analysis: Theory, Methods & Applications*, 30(8):5303–5314,
 696 December 1997. ISSN 0362-546X. doi: 10.1016/S0362-546X(96)00149-6. URL [https://doi.org/10.1016/S0362-546X\(96\)00149-6](https://doi.org/10.1016/S0362-546X(96)00149-6).

697

698

699 Floris Takens. Detecting strange attractors in turbulence. In David Rand and Lai-Sang Young
 700 (eds.), *Dynamical Systems and Turbulence, Warwick 1980*, pp. 366–381, Berlin, Heidelberg,
 701 1981. Springer. ISBN 978-3-540-38945-3. doi: 10.1007/BFb0091924. URL <https://doi.org/10.1007/BFb0091924>.

702 Gouhei Tanaka, Tadayoshi Matsumori, Hiroaki Yoshida, and Kazuyuki Aihara. Reservoir computing
703 with diverse timescales for prediction of multiscale dynamics. *Physical Review Research*, 4(3):
704 L032014, July 2022. doi: 10.1103/PhysRevResearch.4.L032014. URL <https://doi.org/10.1103/PhysRevResearch.4.L032014>.

705

706 Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves,
707 Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. WaveNet: A Generative Model for
708 Raw Audio, September 2016. URL <http://arxiv.org/abs/1609.03499>.

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A SUPPLEMENTARY METHODS
757758 A.1 DATASETS
759760 All datasets were generated or obtained and processed as described below, and all were normalized
761 across time by subtracting the variable mean and dividing by the variable standard deviation.
762763 A.1.1 LORENZ SYSTEM
764765 The Lorenz system (Lorenz, 1963) is a three-dimensional model exhibiting chaotic behavior, and is
766 among the most used benchmarks in the dynamical system reconstruction field. It is described by
767 three equations,
768

769
$$\begin{aligned}\dot{x} &= s(y - x), \\ \dot{y} &= rx - y - xz, \\ \dot{z} &= xy - bz,\end{aligned}$$

770
771

772 with parameters $s = 10$, $r = 28$, $b = 2.667$. The system was simulated for $T = 10000$ using RK45
773 method implemented in Scipy library with the default relative and absolute tolerances of 10^{-3} and
774 10^{-6} respectively. The simulated time series are exported with sampling period $\Delta t = 0.05$, leaving
775 200 000 time points, divided equally into the train and test time series. Only the first variable x is
776 considered to be observed.
777778 A.1.2 CELL CYCLE
779780 The model of cell division cycle (Romond et al., 1999) represents an example of deterministic chaos
781 in a six-dimensional state space. It models the evolution of two coupled biochemical oscillators
782 using six differential equations,
783

784
$$\begin{aligned}\dot{C}_1 &= v_{i1} \frac{K_{im1}}{K_{im1} + M_2} - v_{d1} X_1 \frac{C_1}{K_{d1} + C_1} - k_{d1} C_1, \\ \dot{M}_1 &= V_1 \frac{1 - M_1}{K_1 + (1 - M_1)} - V_2 \frac{M_1}{K_2 + M_1}, \\ \dot{X}_1 &= V_3 \frac{1 - X_1}{K_3 + (1 - X_1)} - V_4 \frac{X_1}{K_4 + X_1}, \\ \dot{C}_2 &= v_{i2} \frac{K_{im2}}{K_{im2} + M_1} - v_{d2} X_2 \frac{C_2}{K_{d2} + C_2} - k_{d2} C_2, \\ \dot{M}_2 &= U_1 \frac{1 - M_2}{H_1 + (1 - M_2)} - U_2 \frac{M_2}{H_2 + M_2}, \\ \dot{X}_2 &= U_3 \frac{1 - X_2}{H_3 + (1 - X_2)} - U_4 \frac{X_2}{H_4 + X_2},\end{aligned}$$

785
786
787
788
789
790
791
792
793
794
795
796
797
798
799

with

800
$$\begin{aligned}V_1 &= \frac{C_1}{K_{c1} + C_1} V_{M1}, & V_3 &= M_1 \cdot V_{M3}, \\ U_1 &= \frac{C_2}{K_{c2} + C_2} U_{M1}, & U_3 &= M_2 \cdot U_{M3},\end{aligned}$$

801
802
803
804
805

806 and $V_{M1} = U_{M1} = 0.3$, $v_{i1} = v_{i2} = 0.05$, $K_{1,2,3,4} = H_{1,2,3,4} = 0.01$, $V_2 = U_2 = 0.15$,
807 $V_{M3} = U_{M3} = 0.1$, $V_4 = U_4 = 0.05$, $K_{c1} = K_{c2} = 0.5$, $v_{d1} = v_{d2} = 0.025$, $K_{d1} = K_{d2} =$
808 0.02 , $k_{d1} = k_{d2} = 0.001$. Following Gilpin (2021), who included the model in a chaotic system
809 collection, we set the value of the bifurcation parameter $K_{im1} = K_{im2} = 0.65$, for which a single
chaotic attractor exist in the state space.

810 The system was simulated for $T = 800000$ using RK45 method implemented in Scipy library with
 811 the default relative and absolute tolerances of 10^{-3} and 10^{-6} respectively, and maximum time step
 812 0.04. The simulated time series are exported with sampling period $\Delta t = 5$, leaving 160 000 time
 813 points, divided equally into the train and test time series. Only the first variable C_1 is considered to
 814 be observed.

815 **A.1.3 DOUBLE WELL**

816 The double well model represents a system with two fixed points and noise driven switches between
 817 the two basins of attractions. It is described by a cubic stochastic differential equation with fixed
 818 points at -1 and 1, followed by four exponential smoothing equations,

$$821 \quad \dot{z}_1 = -z_1^3 + z_1 + \sigma\eta(t), \\ 822 \quad \dot{z}_i = \alpha(z_{i-1} - z_i) \quad \text{for } i \in \{2, 3, 4, 5\},$$

823 with $\alpha = 0.4$ and additive Gaussian noise $\eta(t)$ with variance $\sigma^2 = 0.2$. The last variable z_5 is
 824 considered the observation of the system. The system is simulated with Euler-Maruyama method
 825 with $\Delta t = 0.2$ for time $T = 400000$, and then downsampled by a factor of 10, leaving 200 000 time
 826 points with sampling period 2, equally divided into a train and test set.

827 **A.1.4 CHAOTIC RNN**

828 The chaotic recurrent neural network model (Sompolinsky et al., 1988) describes the activity of
 829 a population of n randomly connected neurons. The dynamics of the synaptic current of the i -th
 830 neuron is given by

$$833 \quad \dot{h}_i = -h_i + \sum_{j=1}^n J_{ij}\phi(h_j),$$

834 where $\phi(h) = \tanh(h)$, and the connectivity matrix \mathbf{J} contains independent random elements J_{ij}
 835 following Gaussian distribution with mean 0 and variance g/n^2 . We choose the factor $g = 2$, for
 836 which the model exhibits chaotic behavior.

837 We set the number of neurons $n = 1000$, and we solve the system using RK45 method implemented
 838 in Scipy library with the default relative and absolute tolerances of 10^{-3} and 10^{-6} respectively. We
 839 solve the system for $T = 100000$ and export with sampling period $\Delta t = 0.5$, leaving 200 000 time
 840 points, divided equally into the train and test time series. Only the firing rate of the first neuron
 841 $\phi(h_1)$ is taken as observed.

842 **A.1.5 NEURON**

843 The dataset represents the voltage time series of an *in vitro* cortical pyramidal neuron from rat
 844 barrel cortex stimulated by a randomly generated fluctuating current (Rauch et al., 2003; Jo-
 845 livet et al., 2006); the dataset was also used in a spike-timing prediction competition (Jolivet
 846 et al., 2008) from which we obtained the data (<https://lcnwww.epfl.ch/gerstner/QuantNeuronMod2007/challenge.html>). In the experiment, the neuron was stimulated
 847 by a current generated by an Ornstein-Uhlenbeck process. Although available in the data, here we
 848 have assumed that the input current is unknown, and aimed to estimate a stochastic dynamical model
 849 of both the neuron and the noise source.

850 From the available data, we have used only one recording (00-0). The signal was sampled at 5 kHz
 851 and contained 34 000 data points, we have discarded the initial 600 and final 1200 time points where
 852 the stimulus current was not applied. The time series were smoothed with a Gaussian filter with
 853 $\sigma = 0.2$ ms, normalized over time, and divided equally into the train and test time series of 16 100
 854 time points each.

855 **A.1.6 ELECTROCARDIOGRAM (ECG)**

856 The ECG signal captures the heart’s electrical activity over time. We have used human ECG record-
 857 ing from the PPG-DaLiA dataset (Reiss et al., 2019), as preprocessed and used by Hess et al. (2023).
 858 They performed smoothing and normalization of the time series, and used the signal of length

864 100 000 time points at sampling rate 700 Hz (duration 143 s) for both training and test dataset. We
 865 have further downsampled the time series by a factor of 4 (25 000 time points, sampling rate 175 Hz,
 866 duration 143 s for each train and test time series) for reasons of reduced computational costs.
 867

868 For the sake of consistency with other datasets in this study, we have not used the time embedding
 869 provided by Hess et al. (2023), but performed the time embedding ourselves (described below),
 870 although using the same embedding method as in the original study. Our approach resulted in 5-
 871 dimensional embedding, consistent with the results of Hess et al. (2023).
 872

872 A.2 TIME DELAY EMBEDDING

873 For all time series we first attempted to perform the time delay embedding using the PECUZAL
 874 algorithm (Kraemer et al., 2021) using the implementation in DelayEmbedding Julia module. Our
 875 settings allowed for possible delay values between 0 and 100 time points, the Theiler window w was
 876 set to the first minimum of mutual information of the signal with itself, and we used the economy-
 877 mode for L-statistic computation, while all other arguments were set to default. The PECUZAL
 878 embedding was successfully achieved for the Lorenz system (dimension $d = 3$), cell cycle model
 879 ($d = 5$), and ECG signal ($d = 5$). The PECUZAL embedding failed for the double well sys-
 880 tem, chaotic RNN, and the neuron recording; in these cases we used a time delay embedding with
 881 repeated delays equal to the minimum of mutual information of the signal with itself, and hand-
 882 selected embedding dimension $d = 8$.
 883

884 A.3 ARCHITECTURE AND TRAINING DETAILS

885 For all methods, the same procedure was followed to train the models. For each dataset, a param-
 886 eter sweep was performed over selected hyperparameters (differing across methods), and for each
 887 parameter combination, four models were trained with four different random weight initializations.
 888 The models were checkpointed during training. The models were evaluated on a separate test set
 889 in terms of the score, detailed below. The best parameter combination was chosen by the lowest
 890 score averaged across initializations. Where only one model was used for a visualization, the best
 891 performing model from the four initializations of the best performing parameter combination was
 892 used unless stated otherwise.
 893

894 A.3.1 DOUBLE PROJECTION DYNAMICAL SYSTEM RECONSTRUCTION (DPDSR)

895 The DPDSR method is described in the main text (Sec. 3). Here we describe further details. The
 896 regularization term for the observation function is an L1 regularization on the weights of the pro-
 897 jection g with coefficient $\alpha_g = 0.3$. The regularization term for the scale and position of the es-
 898 timated states aims to weakly enforce the desired scale of the state trajectories. It has the form
 899 $\alpha_{\hat{z}} D_{\text{KL}}(N(\mu_{\hat{z}}, \text{diag}(\sigma_{\hat{z}}^2) \parallel N(0, I))$, where the mean and variances of the states are calculated
 900 across time and samples in a batch, but not features. Denoting the sample in a batch by a super-
 901 script, $\mu_{\hat{z}} = \mathbb{E}_{b,t}[\hat{z}_t^b]$ and $\sigma_{\hat{z}}^2 = \text{Var}_{b,t}(\hat{z}_t^b)$. We are using value $\alpha_{\hat{z}} = 0.001$.
 902

903 Tab. S1 shows the parameters of the encoder model, and Tab. S2 shows the parameters of the gener-
 904 ative model. We divide the time series into chunks of $T = 300$, and use batch size 16. To avoid
 905 the boundary effects of the convolutional encoders, we evaluate the losses (5-7) on shortened tra-
 906 jectories $\mathbf{x}_{1+a:T-a}$ (and equally shortened ϵ , \tilde{z} , and \hat{z}) instead of the full trajectory $\mathbf{x} = \mathbf{x}_{1:T}$ with
 907 $a = 50$.
 908

909 For all test problems, the parameter grid exploration was performed over two parameters: teacher
 910 forcing interval $\tau \in [1, 10, 20, 40, 60, 80, 100, 200]$, and observation noise variance $\log \sigma_{\eta}^2 \in$
 911 $[-4, -2, 0]$. The noise variance of the estimated states was set to $\log \sigma_{\hat{z}}^2 = \log \sigma_{\eta}^2 + 2$ to avoid
 912 introducing more free parameters and to reflect the secondary importance of reconstructing the es-
 913 timated states compared to reconstructing the original observations.

914 The loss function is minimized using Adam optimizer. The optimizer is ran for 30k batch evalua-
 915 tions, with learning rate starting at 0.001 and reduced by a factor of 0.3 at 10k and 20k steps. The
 916 reparameterization trick of variational autoencoders is used to sample from the posterior, with 4 sam-
 917 ples used to evaluate the expectations in (5) and (6). We perform gradient clipping with threshold
 100. The model is saved every 5000 iterations.

918

919

Table S1: Parameters of the encoder models of the DPDSR method.

920

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

STATE ENCODER	NUMBER OF LAYERS	7
	KERNEL SIZE	7
	DILATION	[1,2,4,8,16,32,64]
	CHANNELS	24
NOISE ENCODER	NUMBER OF LAYERS	7
	KERNEL SIZE	7
	DILATION	[1,2,4,8,16,32,64]
	CHANNELS	24
		LSTM STATE SIZE 32

Table S2: Parameters of the generative models.

		LORENZ	CELL CYCLE	DOUBLE WELL
		ECG	RNN	NEURON
DPDSR	STATE SIZE	5	8	8
	F HIDDEN SIZE	256	256	256
	G HIDDEN SIZE	32	32	32
	# PARAMETERS	3015	4650	4650
SPDSR	STATE SIZE	5	8	8
	F HIDDEN SIZE	256	256	256
	G HIDDEN SIZE	32	32	32
	# PARAMETERS	3046	4681	4681
GTF-PF	STATE SIZE	5	8	8
	F HIDDEN SIZE	256	256	256
	# PARAMETERS	2826	4368	4368
GTF-TD	STATE SIZE	3	5	8
	F HIDDEN SIZE	424	384	256
	# PARAMETERS	2974	4234	4368
DKF	STATE SIZE	5	8	8
	F HIDDEN SIZE	256	256	256
	# PARAMETERS	2832	4377	4377
AR-LSTM	STATE SIZE	24	32	32
	INITIAL CONDITIONS SIZE	5	8	8
	# PARAMETERS	2642	4546	4546

Causal encoder The DPDSR method uses encoders based on dilated convolutional neural networks. These encoders are non-causal, that is, the estimated states \hat{z}_t at time t are computed from both past and future observations, and therefore \hat{z}_t depends on all elements of $\mathbf{x}_{1:T}$. Such approach allows to effectively gather information from the whole provided sample. However, evaluating the predictive performance of such model is difficult, as the initial state would need to be estimated also from future data, invalidating the results. For this reason, we have trained also an auxiliary causal encoder, which uses the same architecture of dilated convolutional networks, but with all connections to future observations set to zero, so that \hat{z}_t depends only on $\mathbf{x}_{1:t}$.

The causal encoder is trained in parallel with the main model. In each iteration, the loss and the gradients of the main model are first computed and the update is performed. After that, an update of the causal encoder is performed. Denoting the non-causal encoder as F and causal encoder as F_c , the loss function is given by

$$L^{ce} = \|F(\mathbf{x}) - F_c(\mathbf{x})\|,$$

that is, the causal encoder is trained so its output matches the output of the non-causal encoder. We use Adam optimizer with the same learning rate and same batch sizes as for the main model. The causal encoder is used only for the prediction tasks, while for all other purposes we are using the non-causal variant. The noise posterior distribution is always estimated using the non-causal

972 encoder, since for the prediction tasks the noise is sampled from the prior and not the posterior
 973 distribution, and future information is thus not leaked through it.
 974

975 A.3.2 SINGLE PROJECTION DYNAMICAL SYSTEM RECONSTRUCTION (SPDSR) 976

977 The architecture of DPDSR allows to consider a special case of a noise-free model by setting $B = 0$
 978 in the generative model (1). In this case, no noise is used in the simulations, and the noise encoder
 979 is not used; for this reason we call this special case a Single Projection Dynamical System Re-
 980 construction (SPDSR). In effect, SPDSR is a variant of sparse teacher forcing methods for training
 981 deterministic systems, with the teacher signal estimated using the encoder based on convolutional
 982 neural networks. We use this variant as a useful comparison point, as it allows us to directly evaluate
 983 the effects of the stochastic formulation of DPDSR against a method equivalent in the architecture
 984 and training methods.

985 Indeed, the architecture and training of SPDSR is equal to DPDSR, with the difference that the
 986 noise encoder is not employed, and the KL divergence term of the loss function (7) is zero. Same
 987 as for DPDSR, the parameter grid exploration was performed over teacher forcing interval $\tau \in$
 988 $[1, 10, 20, 40, 60, 80, 100, 200]$, and observation noise variance $\log \sigma_\eta^2 \in [-4, -2, 0]$. As for the
 989 DPDSR method, a secondary causal encoder (Sec. A.3.1) is trained and used for the prediction
 990 tasks.

991 A.3.3 GENERALIZED TEACHER FORCING (GTF) 992

993 Generalized teacher forcing is a training method developed for training deterministic models of
 994 chaotic dynamics, and has shown a superior performance on range of problems (Hess et al., 2023).
 995 We are using the method to compare our method with state of the art method for deterministic
 996 dynamical system reconstruction. Strictly speaking, “generalized teacher forcing” refers only to the
 997 training method, but in this paper we use it as a shorthand for both the training method and the
 998 specific architecture of the generative model used in the original paper of Hess et al. (2023).

999 The generative model has the form

$$1000 z_t = Az_{t-1} + W_1\sigma(W_2z_{t-1} + h_2) + h_1$$

1001 similar to our parameterization (2) apart from the diagonal matrix A , the tanh transformation in our
 1002 method, and absent noise. The model is trained using a generalized teacher forcing scheme. In every
 1003 step, the state is replaced by a linear interpolation between the simulated state \tilde{z}_t and the data z_t ,

$$1004 \hat{z}_t = (1 - \alpha)\tilde{z}_t + \alpha z_t, \quad (8)$$

1005 with coefficient $\alpha \in [0; 1]$. Such approach was shown theoretically to rectify to problem of exploding
 1006 gradients in learning chaotic dynamics.

1007 We tested two variants of this approach. In the first variant, the full states of the system are estimated
 1008 using time delay embedding (GTF-TD), and these states are used in the forcing scheme (8). The
 1009 observations are then equal to the states, $x_t = z_t$. In the second variant we avoid performing time
 1010 delay embedding and rely on partial forcing instead (GTF-PF). The original time series is used, and
 1011 only one dimension of the state is forced to the teacher signal in (8), while other variables are left to
 1012 evolve freely. The observations are then equal to the first state of the system, $x_t = z_{t,1}$.

1013 We used the original implementation in Julia language provided by authors (<https://github.com/DurstewitzLab/GTF-shPLRNN>).

1016 We followed the example in (Hess et al., 2023) to set the method parameters. We divided the time
 1017 series into chunks of size $T = 200$ and used batch size 16. The parameters were optimized using
 1018 RADAM optimizer with 5000 epochs, 50 batches per epoch, and exponential decay schedule with
 1019 initial and final learning rate 10^{-3} and 10^{-6} respectively. Every 1000 epochs the model was saved
 1020 and evaluated, and best performing model was kept. For all test problems, the parameter sweep
 1021 was performed across the value of teacher forcing parameter $\alpha \in [0.0, 0.1, 0.2, 0.3, 0.4, 0.5]$. The
 1022 parameters of the generative model are given in Tab. S2.

1023 A.3.4 DEEP KALMAN FILTER (DKF) 1024

1025 Deep Kalman Filter (Krishnan et al., 2015; 2017) is based upon the framework of variational au-
 1026 toencoders applied to dynamical systems, where the states of the system z are considered to be the

1026 latent variables. The generative model of DKF is
 1027

$$1028 \quad z_t = f(z_{t-1}) + \Sigma_\epsilon \epsilon_t, \\ 1029 \quad x_t = g(z_t) + \Sigma_\eta \eta_t.$$

1030 We use the parameterization of the evolution function
 1031

$$1032 \quad f(z_t) = z_t + W_2 \sigma(W_1 z_t + b_1) + b_2$$

1033 and linear observation function g . This formulation differs from the one used in the DPDSR method
 1034 (2) by the absence of the \tanh nonlinearity; we have found that this stabilization is not needed for
 1035 training the DKF models.
 1036

1037 During training, for an observation \mathbf{x} , DKF first projects the observation to a posterior distribution
 1038 state trajectories $q(\mathbf{z} \mid \mathbf{x})$. Then a state trajectory is sampled from the posterior, $\mathbf{z} \sim q(\mathbf{z} \mid \mathbf{x})$, and
 1039 the loss function for the data point is calculated as

$$1040 \quad L(\mathbf{x}) = D_{\text{KL}}(q(\mathbf{z} \mid \mathbf{x}) \parallel p(\mathbf{z})) - \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z} \mid \mathbf{x})} [\log p(\mathbf{x} \mid \mathbf{z})].$$

1041 The prior for the system states is
 1042

$$1043 \quad p(\mathbf{z}) = p(z_0) \prod_{t=1}^T p(z_{t+1} \mid z_t),$$

1044 with $p(z_0) = N(z_0 \mid 0, I)$ and $p(z_{t+1} \mid z_t) = N(z_{t+1} \mid f(z_t), \Sigma_\epsilon)$. The probability in the
 1045 reconstruction loss is
 1046

$$1047 \quad p(\mathbf{x} \mid \mathbf{z}) = \prod_{t=1}^T N(x_t \mid g(z_t), \Sigma_\eta).$$

1048 The system and observation noise covariance matrices were assumed to be diagonal and isotropic,
 1049 $\Sigma_\epsilon = \sigma_\epsilon^2 I$ and $\Sigma_\eta = \sigma_\eta^2 I$. Compared to our method, the evolution of the dynamical system is
 1050 represented in the prior term in the KL divergence and the observation function in the reconstruction
 1051 loss, while in our method both are represented in the reconstruction loss. The disadvantage of the
 1052 method is that it does not allow to evolve the system more than one step from the estimated states,
 1053 leading to reduced capacity to learn long-term dependencies.
 1054

1055 In the experiments, we are using the same architecture for the state encoder as for the noise encoder
 1056 in the DPDSR method, that is, the encoder is composed of a stack of dilated convolutional networks
 1057 followed by an autoregressive LSTM network. The same parameters as for the DPDSR noise en-
 1058 coder are used (Tab. S1). The loss function is minimized using Adam optimizer. The optimizer is
 1059 ran for 30k batch evaluations, with learning rate starting at 0.001 and reduced by a factor of 0.3 at
 1060 10k and 20k steps. The models were saved every 5k steps. The parameter sweep was performed over
 1061 the observation noise variance, $\log \sigma_\eta^2 \in [-4, -2, 0]$ and initial values of the system noise variance,
 1062 $\log \sigma_\epsilon^2 \in [-8, -6, -4, -2]$. As for the DPDSR method, a secondary causal encoder (Sec. A.3.1) is
 1063 trained and used for the prediction tasks.
 1064

1065 A.3.5 AUTOREGRESSIVE LSTM (AR-LSTM)

1066 Long short-term memory (LSTM) network (Hochreiter & Schmidhuber, 1997) is an established
 1067 architecture of recurrent neural networks designed to handle long-term dependencies in the input
 1068 data. Autoregressive LSTM model (Graves, 2014) represents one approach to introduce stochasticity
 1069 in the standard LSTM network. The generative model is given by
 1070

$$1071 \quad h_t, c_t = \text{LSTM}(h_{t-1}, c_{t-1}, x_{t-1}), \quad (9) \\ 1072 \quad x_t \sim N(\mu_t, \sigma_t^2),$$

1073 where h_t and c_t are the hidden state and cell state vectors at time t , and $\text{LSTM}(h, c, x)$ represents
 1074 the standard LSTM cell. The observation x_t at each step are drawn from a normal distribution,
 1075 and fed back to LSTM as an input in the next step. The parameters of the normal distribution are
 1076 computed by a linear projection from the hidden state
 1077

$$1078 \quad [\mu_t; \log \sigma_t^2] = A h_t + b.$$

1080 For training the model, the timeseries are split into chunks \mathbf{x} with length $T = T_{\text{past}} + T_{\text{pred}}$. The
 1081 initial conditions $[h_0; c_0; x_0]$ are first estimated from the past observation of length $x_{1:T_{\text{past}}}$. This
 1082 is done using a stack of dilated convolutional neural networks, mirroring the architecture of the
 1083 state encoder in our method. From the last layer and the last element in time the low-dimensional
 1084 representation of the initial conditions $z_0 \in \mathbb{R}^d$ is computed via linear projection. Then the full
 1085 dimension initial conditions are computed through two layer MLP with ReLU nonlinearity.

1086 From the initial conditions the system is evolved according to (9) to obtain the simulated observa-
 1087 tions $\tilde{\mathbf{x}}_{T_{\text{past}}+1:T}$ and means and variances $\mu_{T_{\text{past}}+1:T}, \sigma_{T_{\text{past}}+1:T}^2$. Using the principle of scheduled
 1088 sampling, during the evolution the observations entering the LSTM cell are either randomly sampled
 1089 from the last step prediction ($\tilde{x}_t \sim N(\mu_t, \sigma_t^2)$) with probability γ , or replaced by the data x_t with
 1090 probability $1 - \gamma$.

1091 The loss function for one sample \mathbf{x} is given by

$$1093 L(\mathbf{x}) = -\log p(\mathbf{x}_{T_{\text{past}}+1:T} \mid \mu_{T_{\text{past}}+1:T}, \sigma_{T_{\text{past}}+1:T}^2).$$

1094 The loss function is minimized using Adam optimizer. The optimizer is ran for 30k batch eval-
 1095 uations, with learning rate starting at 0.001 and reduced by a factor of 0.3 at 10k and 20k steps.
 1096 The models were saved every 5k steps. The parameter sweep was performed for parameter
 1097 $\gamma \in [0., 0.2, 0.4, 0.6, 0.8, 1.0]$ and for prediction length $T_{\text{pred}} \in [20, 50, 100, 200]$.

1099 A.4 EVALUATION CRITERIA

1100 Dynamical system reconstruction aims at training models that can robustly reproduce the temporal
 1101 patterns observed in the training data on long-term scale. In this spirit we evaluate the models using
 1102 two measures evaluating long-term behavior, distribution distance D_d and spectral distance D_s , and
 1103 one measure of short-term prediction capacity, 20-step prediction error PE_{20} . The first two are
 1104 evaluated by a comparison of long model-generated time series with the original data. To generate
 1105 the data, we take a point on the embedded state trajectory (via trained projection embedding method,
 1106 or time delay embedding depending on the model). Using this point as initial conditions, we evolve
 1107 the system for 40000 steps, using random noise for stochastic models.

1108 The distribution distance D_d measures the similarity of the distribution of the original and generated
 1109 data in the observation space. To calculate it, we take the original and simulated data, collapse them
 1110 across time, and compare the distributions using the Wasserstein distance (also known as Earth
 1111 mover's distance). Loosely speaking, the Wasserstein distance correspond to the cost of reshaping
 1112 one distribution into the other by transporting the mass. In one dimension, the Wasserstein distance
 1113 between two probability distributions u and v with cumulative distribution functions F_u and F_v is
 1114 defined as:

$$1116 D_d(u, v) = \int_{-\infty}^{\infty} |F_u(x) - F_v(x)| \, dx. \quad (10)$$

1117 We use the SciPy implementation for computations.

1118 The spectral distance D_s measures the similarity of the long time series in frequency space. To
 1119 calculate it, we compute the power spectral density of the original and simulated data using Welch's
 1120 method (using the SciPy implementation) with segment length equal to 4096 points. We smooth
 1121 the frequency spectra using a Gaussian filter with $\sigma = 2$ time points, normalize them, and calculate
 1122 their Hellinger distance. The Hellinger distance for two discrete distributions U and V is given by

$$1124 D_s(u, v) = \frac{1}{\sqrt{2}} \sqrt{\sum_i (\sqrt{u_i} - \sqrt{v_i})^2}.$$

1127 We also consider a measure of short term prediction capability, the n -step prediction error. For
 1128 data chunk $\mathbf{x} = (x_1, x_2, \dots, x_T)$ we use the first k time points to estimate the latent state at time
 1129 k -th step. We then repeatedly simulate the next n steps to with random noise obtain predictions
 1130 $\tilde{\mathbf{x}} = (\tilde{x}_{k+1}, \dots, \tilde{x}_{k+n})$. The prediction error is then

$$1132 \text{PE}_n = \frac{1}{n} \sum_{i=1}^n \|x_{k+i} - \tilde{x}_{k+i}\|,$$

1134 which we average over 20 random noise samples (for probabilistic models only) and 2000 chunks
 1135 from the test dataset.

1136 For the ECG dataset, we also measure the distance between the original and simulated distributions
 1137 of the interspike intervals D_{ISI} . To do so, use the long time series generated as described above, and
 1138 we detect the peaks in the signal using the SciPy `find_peaks` tool with height 2 and prominence 1.
 1139 We then then set D_{ISI} to be the Wasserstein distance (10) between the ISI distributions from the data
 1140 and simulated signal.

1141 We calculate the overall score as weighted sum of the distribution distance D_d , spectral distance D_s ,
 1142 20-step prediction error PE_{20} , and (for ECG) the distance of ISI distributions D_{ISI} .

1143

$$S = w_1 D_d + w_2 D_s + w_3 \text{PE}_{20} + w_4 D_{\text{ISI}}.$$

1144

1145 Given the different nature of the datasets, we set the weights differently across datasets. For the
 1146 Lorenz and Cell datasets, we use $\mathbf{w} = (w_1, w_2, w_3, w_4) = (1, 1, 1, 0)$. For the Double well,
 1147 RNN, and Neuron datasets, where the predictability is lower, we reduce the weight on PE_{20} ,
 1148 $\mathbf{w} = (1, 1, 1, 0.2, 0)$. For the ECG dataset, we include the distance of interspike intervals, D_{ISI}
 1149 as an important measure of the reconstruction quality, $\mathbf{w} = (1, 1, 1, 0.05)$. The weight is chosen
 1150 lower due to large magnitude of unnormalized values (Fig. 4). All measures are evaluated on a test
 1151 dataset which was not used for training the model.

1152

1153

1154 A.5 ANALYSIS OF THE ATTRACTORS

1155

1156 To identify the attractors in the state space of the model, we randomly select 100 points on the
 1157 state space trajectory obtained by projecting the training data into the state space using the trained
 1158 encoder. We then simulate the system forward for $T_{\text{warmup}} + T$ steps with $T_{\text{warmup}} = 1000$ and
 1159 $T = 20000$. For stochastic models, we set the noise in the simulation to zero. After discarding
 1160 the first T_{warmup} steps to allow the transients to decay, the remaining trajectories are analyzed to
 1161 detect distinct attractors. For each trajectory we compare it against previously identified attractors
 1162 by computing pairwise distances between trajectory points. Specifically, for each candidate new
 1163 attractor trajectory $\mathbf{z}^a \in \mathbb{R}^{T \times d_z}$ and an already identified attractor $\mathbf{z}^b \in \mathbb{R}^{T \times d_z}$, we compute two
 1164 way distances between the attractors,

1165

$$d_{a \rightarrow b} = Q_{t_a} \left(\min_{t_b} \|z_{t_a}^a - z_{t_b}^b\|, 0.8 \right),$$

1166

$$d_{b \rightarrow a} = Q_{t_b} \left(\min_{t_a} \|z_{t_a}^a - z_{t_b}^b\|, 0.8 \right),$$

1167

1168 where $Q(\cdot, q)$ denotes the q -th quantile. We define the attractor distance as
 $d_{a,b} = \max(d_{a \rightarrow b}, d_{b \rightarrow a})$, and consider the attractors to be distinct if $d_{a,b} > \text{tol}$, with $\text{tol} = 10^{-5}$
 1169 if the already identified attractor b is a fixed point, and $\text{tol} = 10^{-1}$ for a limit cycle or a chaotic
 1170 attractor. We use the quantile instead of maximum for computing the distances, and relatively high
 1171 tolerances, both for more robust solution when dealing with trajectories with finite length; this is at
 1172 a cost of possibly conflating close attractors.

1173 For each new attractor we calculate the maximal Lyapunov exponent λ_{max} numerically (Sprott,
 1174 2003). In the algorithm, we take a point on the attractor and a point with a small perturbation. We
 1175 repeatedly advance the trajectories from the initial points using the deterministic dynamics, while
 1176 rescaling the deviation of the perturbed trajectory to its original norm at every step. We advance the
 1177 system for 1000 time steps for, and calculate λ_{max} as an average from the estimates from all steps.
 1178 We consider the attractor chaotic if $\lambda_{\text{max}} > 0$, limit cycle if $\lambda_{\text{max}} \leq 0$ and the trajectory does not
 1179 converge to a single point (with $\text{tol} = 10^{-5}$), and fixed point otherwise.

1180

1181

1182 A.6 USE OF LARGE LANGUAGE MODELS

1183

1184 Large Language Models (LLMs) were used in preparing the manuscript to polish the writing. LLMs
 1185 were also used, together with specialized literature search tools, to identify relevant research works.

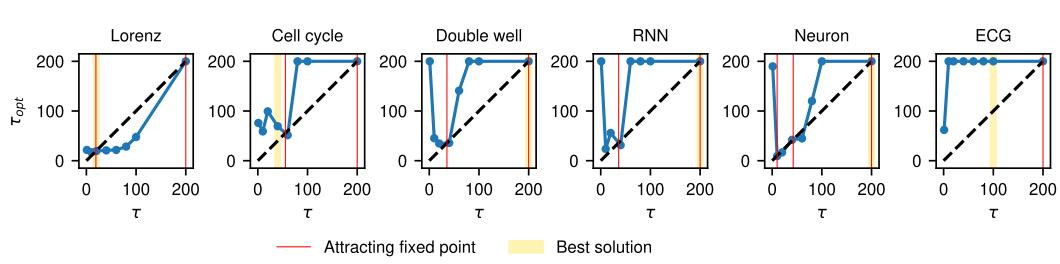


Figure S1: Optimal value of teacher forcing interval τ_{opt} across datasets. For each value of τ , we calculated the maximal Lyapunov exponent λ_{max} along the estimated trajectory using the final trained model. Following Mikhaeil et al. (2022), we calculate $\tau_{\text{opt}} = \log 2 / \lambda_{\text{max}}$ for $\lambda_{\text{max}} > 0$, and we set τ_{opt} equal to the maximal trajectory length (=200) otherwise. The red lines indicate the attracting fixed points of a relaxed fixed point iteration scheme $\tau_{i+1} = (1 - \alpha)\tau_i + \alpha\tau_{\text{opt},i}$ if ran on the visualized relation $\tau_{\text{opt}}(\tau)$ obtained from the final models, that is, they show the points of diagonal crossings with derivative smaller than 1. We assume linear interpolation between the evaluated points. The yellow line indicates the position of the best trained models, same as in Fig. S6. Note that here the Lyapunov exponents are evaluated along the estimated trajectories, and not for the deterministic attractors as in Fig. S6. For most datasets, two fixed points exist, one in the low τ (deterministic) regime, and one in the high τ (noise-driven) regime, indicating that the adaptive τ scheme might not robustly converge to the optimal solution.

B SUPPLEMENTARY RESULTS

B.1 CHOICE OF THE TEACHER FORCING INTERVAL

An important question for practical applications of the method is how to select the teacher forcing interval τ . In this work, we have performed a parameter sweep across a range of values. Such approach, while robust, is costly in terms of computational time. Existing works proposed several approaches to similar issues in deterministic models. Mikhaeil et al. (2022) used teacher forcing interval equal to the predictability time of a chaotic system

$$\tau_{\text{opt}} = \frac{\log 2}{\lambda_{\text{max}}} \quad (11)$$

with λ_{max} being the maximal Lyapunov exponent estimated ahead from the data; they showed that such estimates match closely the optimal values. Hess et al. (2023) introduced an adaptive scheme where the parameter of generalized teacher forcing is updated during the training based on the product of Jacobians of the trained system, evaluated along the forced trajectory.

In the first approach, estimating the Lyapunov exponent from the data is performed under the assumption of deterministic chaotic system, and therefore is unsuitable for our purposes. However, a possible option would be to combine the two methods, that is, an adaptive scheme where at each step i of the training we estimate the maximal Lyapunov exponent along the forced trajectory of our stochastic system, calculate the optimal interval $\tau_{\text{opt},i}$ using (11), and adapt the parameter following a relaxed fixed point iteration scheme

$$\tau_{i+1} = (1 - \alpha)\tau_i + \alpha\tau_{\text{opt},i} \quad (12)$$

with relaxation parameter $\alpha \in (0, 1]$.

For a preliminary investigation if such approach is feasible, we reanalyze the results from our parameter sweep, and calculate τ_{opt} from the final trained models for each constant value of τ (Fig. S1). Using this visualization we can see where the equation $\tau_{\text{opt}}(\tau) = \tau$ solved by scheme (12) has its attracting fixed points. These can provide an indication to which values of τ might the hypothetical adaptive scheme converge. While it is not guaranteed that the adaptive scheme would converge to the same solutions as the scheme with fixed τ , these results suggest that the scheme might not be sufficiently robust. On most datasets we see the existence of two attracting fixed points: one in the low range of τ in the deterministic regime with chaotic dynamics, and one in the noise-driven regime, where the negative Lyapunov exponent results in the maximal teacher forcing interval. Although

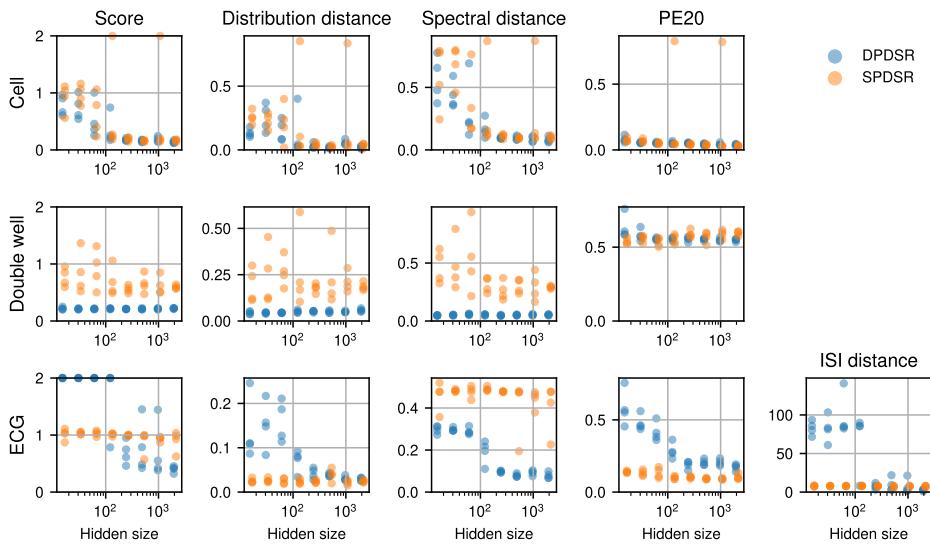


Figure S2: Effect of the number of units in the hidden layers. The effect on the model quality was investigated for three datasets (rows), the score and the separate measures of reconstruction quality are in the columns. For clearer visualization, the results were clipped at the upper limit of the shown range.

the best solution closely matches one of these fixed points for all but one problem, the adaptive scheme might not be able to correctly identify the optimal one. We therefore conclude that for practical purposes where robust results are required the parameter sweep remains the best option, and we highlight of task of finding the optimal teacher forcing strategy in stochastic approaches as an interesting research problem for future studies.

B.2 ARCHITECTURE VARIATIONS

B.2.1 NUMBER OF PARAMETERS IN THE GENERATIVE MODEL

We investigated how does the performance of the stochastic DPDSPR method change when the number of parameters in the generative model is increased or decreased. In particular, we aimed at a comparison with the behavior of its deterministic modification: the SPDSPR method. We considered three problems, the Cell, Double well, and ECG datasets. For each, we trained stochastic and deterministic models with 16 to 2048 units in the hidden layer of the generative model (2), and evaluated the models as before.

Fig. S2 shows that the behavior of the method depends on the dataset. For the Cell dataset, generated by a low-dimensional deterministic model, both DPDSPR and SPDSPR behave similarly, with increasing performance for increasing number of parameters. For the Double well dataset, generated by a noise-driven model, the stochastic model performs well even for minimal number of parameters, and consistently outperforms the deterministic model. For the ECG dataset we see that the deterministic models perform equally well across the range of parameters. The stochastic models, however, perform considerably worse for models with less than 100 units in the hidden layer, but outperform the deterministic models for larger hidden layers.

B.2.2 ENCODER ARCHITECTURE

We further investigated the role of the encoder architecture on the performance of the DPDSPR method (Fig. S3). Specifically, we compared the baseline architecture with two modifications: First, a variation with a noise encoder with the autoregressive RNN removed, and with the parameters otherwise kept equal as described in Tab. S1. Second, a variation with an autoregressive RNN in noise encoder also removed, but with increased number of channels (26 instead of 24) in the dilated convolutional neural network to keep the total number of parameters approximately equal. Degradation

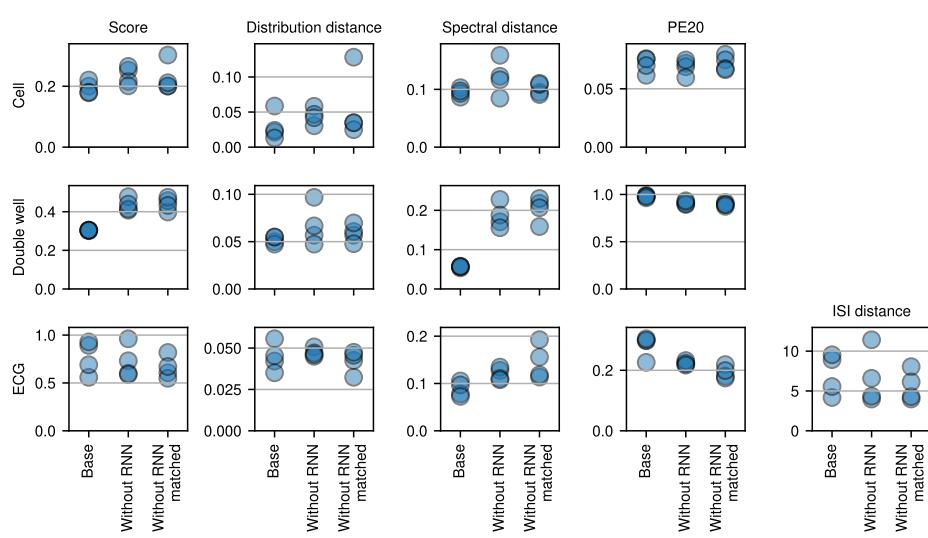


Figure S3: Role of the encoder architecture. For three datasets, three variations of the DPDSR model architecture were investigated: the baseline, a variation with the autoregressive RNN in the noise encoder removed, and a variation with the autoregressive RNN removed, but increased number of parameters in the convolutional neural network to match the total parameters. For each variant, four models with different random weight initializations were trained.

in performance with the removed RNN can be seen for the Double well problem, while for the Cell and ECG problems the modification have little impact.

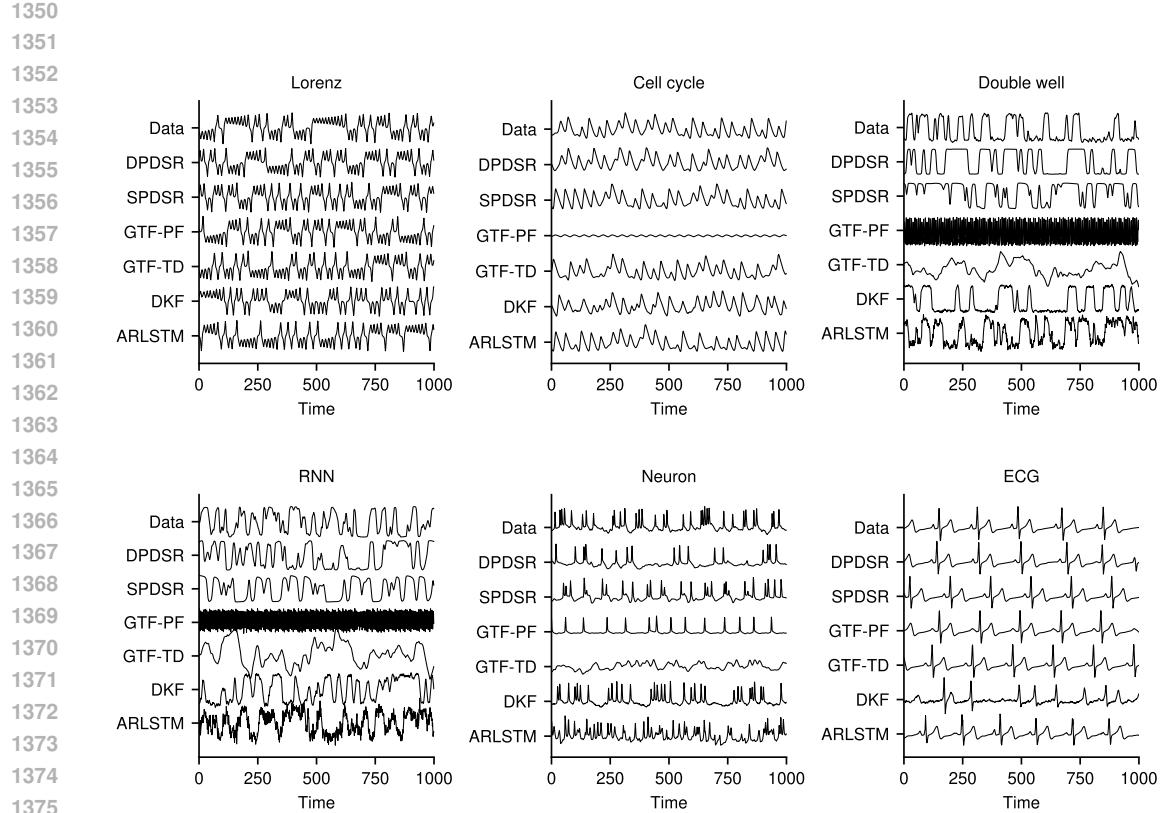


Figure S4: Examples of generated time series for all test problems and evaluated methods. The time is represented in sample indices and not the original model or real time.

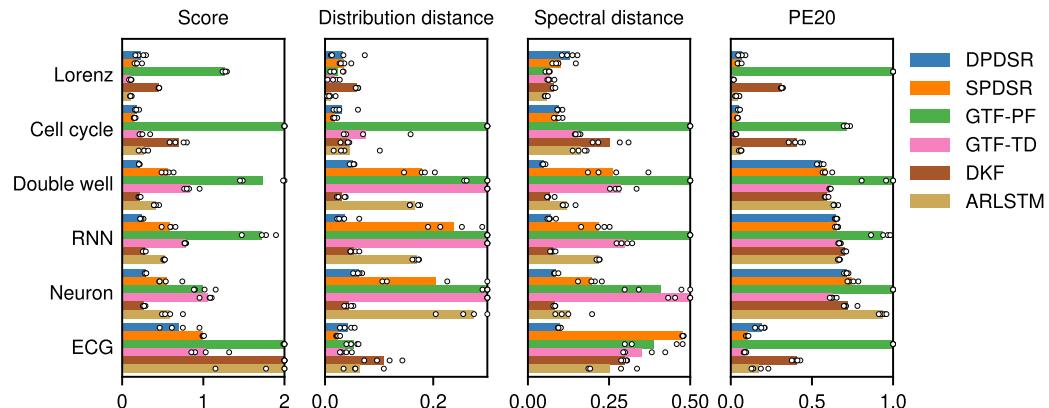


Figure S5: Reconstruction quality of the evaluated methods for the six datasets. The panels show the overall score, distribution distance D_d , spectral distance D_s , and the 20-step prediction error PE_{20} (lower = better for all measures). Each circle represent one of four initializations of the training method, and the bar shows the mean. For clearer visualization, the results were clipped to the shown interval.

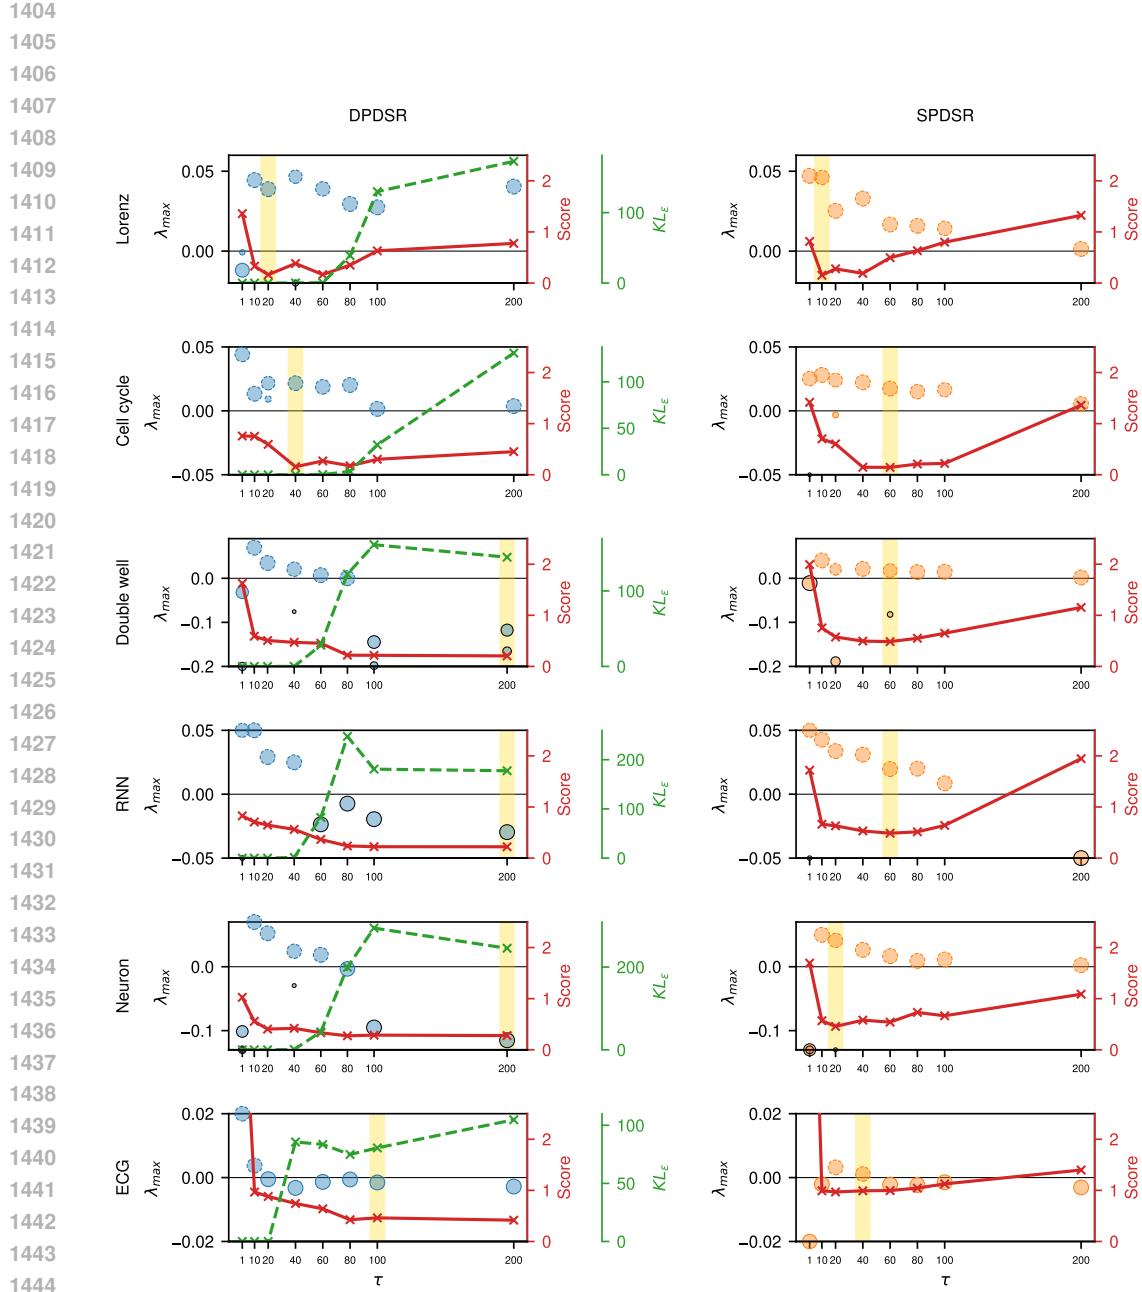


Figure S6: Analysis of the attractors of the trained models for all datasets, extending Fig. 5. Left column: DPDSR models (stochastic); right column: SPDSR models (deterministic). Circles show the maximal Lyapunov exponent. Dashed circle outline represents a chaotic attractor ($\lambda_{\max} > 0$), colored solid outline represents a limit cycle, and black solid outline a fixed point. Size of the circle corresponds to the size of the basin of attraction. Solid red line shows the score (lower = better). Dashed green line shows the KL divergence of the estimated posterior distribution of the noise to the prior distribution $KL_{\epsilon} = D_{KL}(q(\epsilon | x) \parallel p(\epsilon))$. The yellow band indicates the optimal τ value for the dataset; note that the optimal τ value is selected based on the mean score across all initializations, while for other visualizations we plot only the best model. For clearer visualization, the positions of the Lyapunov exponents are clipped to the shown interval.