
000 PANDA: A PRETRAINED FORECAST MODEL FOR 001 002 CHAOTIC DYNAMICS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Chaotic systems are intrinsically sensitive to small errors, challenging efforts to con-
011 struct predictive data-driven models of real-world dynamical systems such as fluid
012 flows or neuronal activity. Prior efforts comprise either specialized models trained
013 separately on individual time series, or foundation models trained on vast time
014 series databases with little underlying dynamical structure. Motivated by dynamical
015 systems theory, we present *Panda*, Patched Attention for Nonlinear DynAmics.
016 We train *Panda* on a novel synthetic, extensible dataset of 2×10^4 chaotic dy-
017 namical systems that we discover using an evolutionary algorithm. Trained purely
018 on simulated data, *Panda* exhibits emergent properties: zero-shot forecasting of
019 unseen chaotic systems preserving both short-term pointwise accuracy and distri-
020 butional measures. Despite having been trained only on low-dimensional ordinary
021 differential equations, *Panda* spontaneously develops the ability to predict partial
022 differential equations without retraining. We also demonstrate a neural scaling
023 law for differential equations, underscoring the potential of pretrained models for
024 probing abstract mathematical domains like nonlinear dynamics.

025 026 1 INTRODUCTION 027

028 Nonlinear dynamical systems test the limits of scientific machine learning (SciML). When an
029 approximate model is constructed of a chaotic nonlinear system, any small error grows exponentially
030 over time, precluding long-term forecasting. This intrinsic property underscores the practical difficulty
031 of accurately forecasting systems like weather fronts, neural activity, or economic markets (Li et al.,
032 2022; Mikhaeil et al., 2022; Price et al., 2025).

033 Recent empirical studies show surprising progress on the classical problem of forecasting chaos,
034 including the ability to predict these systems well-beyond the classical predictability timescale for
035 nonlinear systems (Gilpin, 2021; 2023; Pathak et al., 2018). These approaches construct *local* forecast
036 models trained on past observations of a single dynamical system, and then forecast future, unseen
037 states of the same system. For dynamical systems, this represents an in-domain generalization task,
038 because future timepoints are drawn from the same underlying differential equations. This problem
039 thus reduces to learning the numerical propagator for the true underlying governing equations.

040 However, a frontier in SciML is out-of-domain generalization (Göring et al., 2024; Wang et al., 2022):

041 *Can a dynamics model effectively forecast unseen dynamical systems?*

042 This task requires a *global* forecast model, which combines training on a large body of background
043 knowledge with local adaptation to generate meaningful forecasts of unseen systems (Sen et al.,
044 2019). Moreover, what kind of data is required to train a forecasting model for dynamical systems in
045 order to achieve generalization? A global nonlinear forecast model has intrinsic theoretical interest
046 in SciML, which has long questioned the degree to which complexity can be "transformed out" i.e.
047 whether the predictability of a system is determined by its intrinsic properties or by the choice of
048 measurement coordinates (Brunton et al., 2022; Mezić, 2013).

051
052
053 *Equal contribution. †Corresponding author.

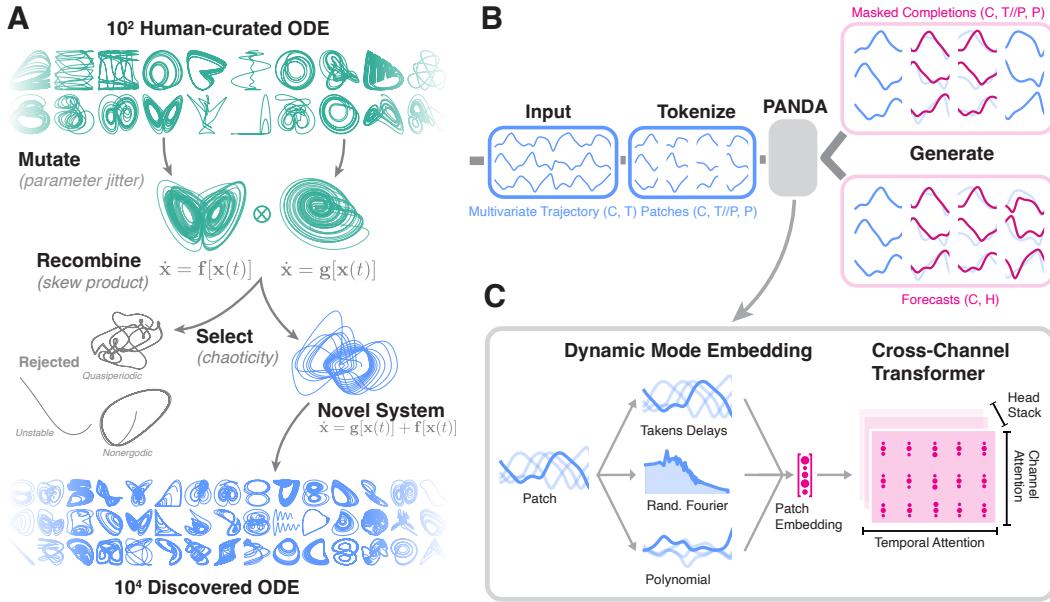


Figure 1: **A large-scale chaotic dynamics dataset and dynamics-informed forecast model.** (A) Evolutionary creation of a large dataset of chaotic ODEs through mutation and recombination of known systems. (B) Patch model architecture with forecasting and masked completion output modes. (C) The dynamics-informed time series embedding and attention modules.

To address these questions, we introduce *Panda*¹ — Patched Attention for Nonlinear DynAmics. Our key contributions are as follows:

1. We introduce a framework for generating novel chaotic dynamical systems, allowing us to create a dataset of $\sim 2 \times 10^4$ ODEs, algorithmically-discovered based on evolutionary recombination of 129 chaotic systems such as the Lorenz attractor, double pendulum, etc.
2. We pretrain a global forecast model for nonlinear dynamics purely on chaotic trajectories integrated from our dataset. Our model exhibits competitive zero-shot forecasts for real systems including mechanical motion of *C. Elegans*, electronic circuits, and turbulent flows.
3. We demonstrate the effectiveness of features motivated by dynamical systems theory: (a) masked pretraining for dynamical continuity (b) channel attention for dynamical coupling, (c) kernelized patch embeddings based on dynamic mode decomposition.
4. Despite being trained only on low-dimensional ODEs, *Panda* develops emergent ability to zero-shot forecast high-dimensional PDEs.

2 RELATED WORK

Machine learning for dynamical systems. Machine learning models for dynamical systems (MLDS) leverage as inductive biases the unique properties of dynamical systems, relative to traditional time series. These include: (1) *Strong channel coupling*: The evolution of system variables is governed by deterministic differential or difference equations, implying coupled functional dependencies among variables rather than statistically correlations. Several MLDS approaches perform large-scale multivariate dynamical modeling, or infer interactions networks among measurement channels (Bhaskar et al., 2024; Brunton et al., 2022; Chen et al., 2018; Li et al., 2022). (2) *Invariant statistical measures*: Ergodic dynamical systems possess invariant probability measures supported on non-wandering sets, such as limit cycles or strange attractors, resulting in well-defined long-term statistical distributions for all observables. Recent works incorporate these properties as inductive biases in modern methods in MLDS settings (Cheng et al., 2025; Koppe et al., 2019; Pedersen et al., 2025).

¹Code available: https://anonymous.4open.science/r/anonymous_panda-3AE0

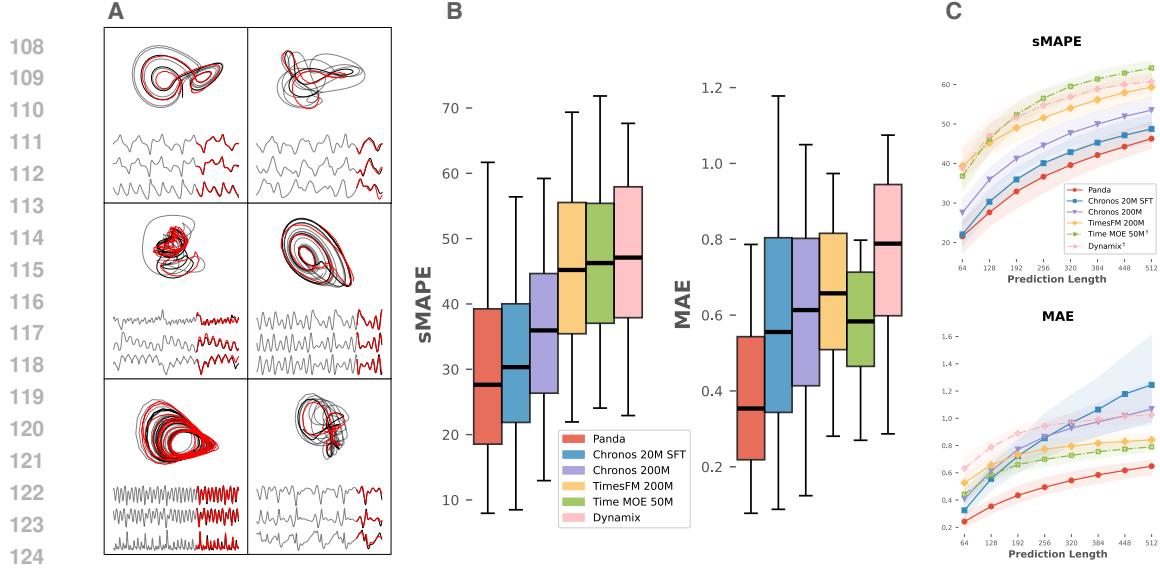


Figure 2: **Panda** zero-shot forecasts unseen nonlinear dynamics. (A) Example zero-shot forecasts on novel chaotic skew-systems. (B) sMAPE and MAE of *Panda* compared to zero-shot time series models over a 128 timepoint prediction horizon. (C) Error versus forecast horizon. Error ranges correspond to median and semi-interquartile range across 9.3×10^3 held-out dynamical systems, 6 forecasts per system. Note: † indicates some NaNs present in forecasts (more examples in Appendix C; dataset description in Section 3). See Table 7 in Appendix D for statistical significance tests.

Discovering new dynamical systems. Small datasets of dynamical systems have previously been curated from the published literature (Gilpin, 2021; 2023; La Cava et al., 2021). Several pretrained models, particularly for partial differential equations (PDE), generate new equations for training by randomly-perturbing parameters or initial conditions from known systems (Chen et al., 2024; Herde et al., 2024; Sun et al., 2025; Tripura & Chakraborty, 2023). Others construct *de novo* systems by combining terms from a fixed function library (Ziegler et al., 2024), or leveraging language models to create candidate symbolic expressions (d’Ascoli et al., 2023; Du et al., 2024). However, these approaches do not address the harder task of sampling based on whether a system exhibits a unique dynamical attractor. Richer sampling requires *post-hoc* analysis of candidate dynamics, akin to intrinsically-motivated discovery previously used in domains such as cellular automata and coupled oscillators (Crutchfield & Mitchell, 1995; Falk et al., 2024; Kumar et al., 2024; Reinke et al., 2020). Some foundation models generate synthetic time series using stochastic dynamics like Gaussian processes (Ansari et al., 2024; Das et al., 2024), or simulated physics environments (Lin et al., 2025; Wang et al., 2024).

Pretrained models for SciML. Pretrained foundation models for dynamics enable transfer learning and zero-shot inference. One study trains transformers across diverse PDEs to create a shared multiphysics embedding space (McCabe et al., 2024). Another study proposes supervised pretraining to enable out-of-domain generalization for scientific foundation models, and derives scaling laws for transfer learning on PDEs (Subramanian et al., 2023). Several recent studies evaluate the zero-shot performance of time series and language models in MLDS, and observe performance only comparable to standard time series tasks (Liu et al., 2024; Zhang & Gilpin, 2025b). Several studies apply pretrained transformers to control or symbolic equation discovery tasks (Becker et al., 2023; d’Ascoli et al., 2023; Lee et al., 2023; Zhang et al., 2024). One work generates pretraining data by randomizing the parameters of four named ODE (Song et al., 2024), similar to the first step of our evolutionary algorithm described below, with a small founder pool. Another work samples systems from a fixed function space, selecting based on total variation over time (Ziegler et al., 2024), while another study uses latent ODE as a prior for zero-shot imputation (Seifner et al., 2025). A contemporaneous work to our study, DynaMix, is a multivariate mixture-of-experts model for zero-shot dynamical systems reconstruction (Hemmer & Durstewitz, 2025), built from Almost-Linear RNN experts and trained on the founder pool for our dataset (Brenner et al., 2024). Our work is distinguished by (1) a rich data generation process, which discovers novel chaotic flows with diverse properties, and (2) a multivariate patched-based architecture which demonstrates emergent forecasting capabilities like zero-shot PDE inference.

162 **3 DATASET**
163

164 **Evolutionary search.** We discover 2×10^4 novel chaotic ODEs (schematic in Fig. 1A, example
165 systems in Appendix A). **1. Founding population:** We start from a human-curated dataset of 129
166 previously-published low-dimensional chaotic systems (Gilpin, 2021), consisting of curated ODEs
167 from the literature (e.g. the Lorenz equations or blinking vortex flow) of the form $\dot{\mathbf{x}} = \mathbf{f}_\theta(\mathbf{x}, t)$. The
168 default parameters of each system θ and initial conditions $\mathbf{x}(0)$ were hand-tuned to the chaotic regime,
169 and the integration timescales were standardized based on calculations of invariant mathematical
170 properties of the underlying equations, such as the largest Lyapunov exponent. **2. Mutation:** We
171 randomly sample pairs of systems $\mathbf{f}_a, \mathbf{f}_b$. For each ODE's default system parameters, we add random
172 Gaussian noise, $\theta'_a \sim \mathcal{N}(\theta_a, \sigma), \theta'_b \sim \mathcal{N}(\theta_b, \sigma)$. **3. Recombination:** We combine the mutated
173 parents using an additive skew-product coupling:

174
$$\dot{\mathbf{x}} = \mathbf{f}_a(\mathbf{x}, t) \quad (1)$$

175
$$\dot{\mathbf{y}} = \kappa_b \mathbf{f}_b(\mathbf{y}, t) + \kappa_a \mathbf{f}_a(\mathbf{x}, t) \quad (2)$$

176 This coupling between flows is asymmetric, and thus we refer to \mathbf{f}_a as the driver and \mathbf{f}_b as the
177 response. In general, skew-product coupling maps can be symmetric and nonlinear, but may be harder
178 to integrate as a result. This particular recombination scheme, for appropriate scale factors, preserves
179 chaoticity because the response system either synchronizes to the chaotic driver or continues to
180 exhibit chaotic dynamics (Gilpin, 2025; Pecora & Carroll, 1990). For the scale factors, we compute
181 the inverse RMS norm $\kappa = 1/\sqrt{\mathbb{E}||\mathbf{f}(\mathbf{x}, t)||^2}$ for each individual flow over a representative trajectory.

182 **4. Selection:** We integrate trajectories from multiple initial conditions using a 5th order implicit
183 Runge-Kutta integrator (see Appendix A), and use a suite of *attractor tests* to cull systems that
184 fail to exhibit chaotic behavior. First, transient systems that converge to a fixed point or diverge to
185 infinity are filtered. Then, we apply the chaos 0-1 test, which distinguishes quasiperiodic dynamics
186 from true chaos (Falconer et al., 2007). We also apply a near-recurrence test to reject limit cycles, a
187 power spectrum test to reject trajectories with only a few distinct sharp peaks, and the data-driven
188 Rosenstein estimator (Rosenstein et al., 1993) to ensure a positive maximum Lyapunov exponent.
189 Finally, we filter for stationarity using the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) (Kwiatkowski
190 et al., 1992) and augmented Dickey-Fuller (ADF) (Dickey & Fuller, 1979) statistical tests.

191 **Augmentations.** On top of the integrated trajectories, we expand the training data by applying
192 dynamics-inspired augmentations that preserve the property that the transformed timeseries arise
193 from a closed nonlinear dynamical system. Our augmentations are: *Random time-delay embedding*
194 $x_i(t) \rightarrow x_i(t - \tau_i)$, $\tau_i \sim \mathcal{U}(1, d_{\text{embed}})$. This augmentation produces dynamics diffeomorphic to the
195 original trajectory due to Takens' embedding theorem (Packard et al., 1980; Takens, 1981). Given
196 $X \in \mathbb{R}^{C \times T}$ and $d \sim \mathcal{U}(d_{\min}, d_{\max})$, *Convex combinations* take random linear combinations of
197 coordinates with coefficients sampled from a Dirichlet distribution: $X \leftarrow MX \in \mathbb{R}^{d \times T}$; $M \in \mathbb{R}^{d \times C}$,
198 $M_{i,:} \sim \text{Dir}(\alpha \mathbf{1}_C)$. *Affine transforms* implement $X \leftarrow AX + b$, $[A \ b] \in \mathbb{R}^{d \times (c+1)}$, $[A \ b]_{ij} \sim$
199 $\mathcal{N}(0, \sigma^2)/\sqrt{d}$. We set $d_{\min} = 3$, $d_{\max} = 10$, and $d_{\text{embed}} = 10$ for our experiments.

200 **Standardization.** For all trajectories, we apply instance-normalization to standardize the scales
201 per channel. For integration, we standardize the integration horizon and granularity based on the
202 number of timepoints (4096) and dominant timescale; note, however, that the numerical integrator
203 ultimately dictates the stepsizes (Gilpin, 2021; Rosenstein et al., 1993). We observe no decrease in the
204 range of invariant properties (maximum Lyapunov exponents, fractal dimension) across generations,
205 suggesting that the starting population is sufficiently large and diverse (see Appendix A more details).

206 **Held-out systems.** For our zero-shot test metrics, we evaluate on an unseen set of 9.3×10^3 systems.
207 We form the test set by holding out a random subset of 20 systems from the 129 founding system
208 population and ensure that none of these systems or their descendants (systems where the parent is
209 the driver or response) appear in the training set. We then evolve these systems into the test set and
210 include all skew product systems descended from these held-out systems.

211 **4 MODEL ARCHITECTURE**
212

213 Dynamical systems differ from traditional time series, and so we introduce a novel architecture
214 motivated by dynamical systems theory (Fig. 1B). Time series foundation models with causal
215 decoders that tokenize time series on a per-observation basis tend to *parrot* motifs from their context,

216 leading to over-confident predictions on out-of-domain tasks (Olsson et al., 2022; Zhang & Gilpin,
217 2025a;b). Parroting is a useful emergent inductive bias when modeling invariant properties in long
218 forecasts is prioritized over accuracy — otherwise known as forecasting the *climate*. However, we
219 opt for an encoder-only, fixed prediction horizon forecast model that maximizes short-term pointwise
220 accuracy, known as predicting the *weather* in SciML.

221 *Panda* generalizes PatchTST, a transformer for univariate forecasting that tokenizes timeseries on a
222 per-patch basis (Nie et al., 2022). Section 5.2 shows that univariate-only architectures underperform
223 on dynamical systems, motivating channel attention. Moreover, patching admits an inductive bias for
224 dynamical systems due to Takens’ theorem which states that time-delayed copies of a low-dimensional
225 measurement of a dynamical system result in a multivariate time series that preserves key topological
226 features of the true attractor (Packard et al., 1980; Takens, 1981).

227 **Patching.** We tokenize a length T trajectory $\mathcal{T} \in \mathbb{R}^{C \times T}$ by patching it into a token sequence of size
228 P patches with stride S so that in general, $\mathcal{T}_{P,S} \in \mathbb{R}^{C \times (\lfloor \frac{T-P}{S} \rfloor + 1) \times P}$. We choose stride $S = P$ so
229 that the token sequences are $\mathcal{T}_P \in \mathbb{R}^{C \times (T/P) \times P}$. We choose $P = 16$, unless stated otherwise.
230

231 **Dynamics Embedding.** We lift the patched multivariate timeseries to a higher-dimensional em-
232 bedding space (d_{model}) by concatenating each patch token $\mathcal{P} \in \mathbb{R}^{C \times P}$ with random poly-
233 nomial and random Fourier features. For random polynomial features with degree d , we sam-
234 ple a fixed index set $\mathcal{I} \subset \{1, \dots, P\}^d$ of $|\mathcal{I}| = N_{\text{poly}}$ (number of features) d -tuples such that
235 for $I \in \mathcal{I}$: $\Phi_{c,i}(\mathcal{P}) := \prod_{j=1}^d \mathcal{P}_{c,I_j} = \mathcal{P}_{c,I_1} \cdot \dots \cdot \mathcal{P}_{c,I_d}$. The random Fourier features sam-
236 ple parameters $W \in \mathbb{R}^{P \times (N_{\text{rff}}/2)}$, $b \in \mathbb{R}^{N_{\text{rff}}/2}$ such that $W_{ij}, b_i \sim \mathcal{N}(0, \sigma^2)$ and $\mathcal{F}(\mathcal{P}) :=$
237 $[\sin(PW + b) \ \cos(PW + b)] \in \mathbb{R}^{C \times N_{\text{rff}}}$, where b added across channels (Rahimi & Recht, 2007).
238 The overall patch embedding is $\mathcal{E}(\mathcal{P}) := [P \ \Phi(\mathcal{P}) \ \mathcal{F}(\mathcal{P})] \in \mathbb{R}^{C \times (P + N_{\text{poly}} + N_{\text{rff}})}$. We use degrees
239 $d \in \{2, 3\}$ and choose N_{poly} and N_{rff} such that $d_{\text{model}} = P + N_{\text{poly}} + N_{\text{rff}} = 512$. The use of
240 polynomial and Fourier features as a lifted dynamics embedding is motivated by prior approximations
241 of the Koopman operator via extended dynamic mode decomposition (eDMD) (Kutz et al., 2016;
242 Williams et al., 2015) and next-generation reservoir computers, which use polynomial features to
243 forecast chaotic systems (Gauthier et al., 2021). See Appendix B for our choices of hyperparameters.
244

245 **Temporal Attention.** We mix information over the temporal dimension by taking the channel
246 dimension as a batch dimension and performing self-attention with p -RoPE (Barbero et al., 2025) (a
247 modification of rotary positional encoding, RoPE (Su et al., 2023)) over the T/P univariate patches
248 of dimension d_{model} . For all experiments, we use a RoPE wavelength of 500 and $p = 75\%$.
249

250 **Multivariate Attention.** Several time series foundation models are univariate, and thus, channel-
251 independent; they solely employ temporal attention for information mixing (Nie et al., 2022).
252 However, chaotic dynamical systems exhibit strong channel coupling. We demonstrate this em-
253 pirically for the electronic circuits dataset in Fig. 4D, where we show the benefit of channel
254 attention as the coupling strength increases. We interleave channel attention layers without po-
255 sitional encoding after each temporal attention layer. Each layer transposes the token sequence,
256 treating the token dimension as a batch dimension and the channels as a set before self-attention
257 $\text{ChannelAttention}(\mathcal{T}_P) := \text{SelfAttention}(\mathcal{T}_P^\top), \mathcal{T}_P^\top \in \mathbb{R}^{T/P \times C \times d_{\text{model}}}$. Temporal attention is fol-
258 lowed by a feedforward residual network, GeLU activations (Hendrycks & Gimpel, 2016), and
259 RMSNorm (Zhang & Sennrich, 2019). In the prediction head, processed tokens are aggregated along
260 the sequence dimension T/P and mapped with a linear layer into a length H channel-wise forecast.
261 The architecture is further described in Appendix B.
262

263 5 RESULTS

264 5.1 *Panda* ZERO-SHOT FORECASTS UNSEEN NONLINEAR DYNAMICS

265 To evaluate the quality of the generated forecasts, we measure (1) short-term forecast accuracy via
266 mean squared error (MSE), mean absolute error (MAE), symmetric mean absolute percentage error
267 (sMAPE), and Spearman correlation, as well as (2) attractor reconstruction accuracy via correlation
268 dimension, spectral Hellinger distance, and Kullback-Leibler (KL) divergence from the ground-truth
269 attractor. For brevity, we report only the sMAPE and MAE (short-term), and KL divergence and
270 spectral Hellinger distance (global) in the main text; the other metrics show similar results and are

270 included in the Appendix D and Appendix C. We compute all metrics for forecasts generated from
 271 zeroshot (held-out) systems never seen during training. Specifically, these are $N_{\text{test}} = 9.3 \times 10^3$
 272 unique skew-product dynamical systems found using the methodology described in section 3. We
 273 additionally include results for scaling up model size and training dataset size in Appendix K.
 274

275 **Comparison to baseline models.** We train *Panda* with $21M$ parameters and evaluate against several
 276 time series foundation models of comparable or larger scale: *Chronos 20M*, a causal univariate
 277 model which was recently shown to produce competitive forecasts of chaos systems (Ansari et al.,
 278 2024; Zhang & Gilpin, 2025b). *Chronos 20M SFT*: Chronos supervised-finetuned on our entire
 279 chaotic systems dataset (Section 3). *Time MOE 50M*: A $50M$ parameter univariate model based
 280 on sparse mixture of experts (Shi et al., 2024). *TimesFM 200M*: A patch-based $200M$ parameter
 281 decoder-only univariate model (Das et al., 2024). *DynaMix*: A multivariate pretrained dynamical
 282 systems model based on RNNs trained with teacher-forcing, enabling efficiency ($10k$ parameters)
 283 (Hemmer & Durstewitz, 2025). For univariate baselines, each dimension is forecast independently.
 284

285 Across 9.3×10^3 held-out systems, we find *Panda* outperforms the baselines across a variety of
 286 prediction horizons and error metrics (Fig. 2). While we train our model exclusively on $d = 3$ -
 287 dimensional dynamical systems, the evaluation set includes arbitrary dimension systems, indicating
 288 that channel attention enables multivariate generalization. Moreover, we use window autoregression
 289 to extend our evaluation forecast horizon well beyond the forecast horizon used during training.
 290 Our model maintains its performance advantage, indicating that it learns an effective dynamical
 291 propagator independent of a single timescale. In Appendix D, we show that our results are robust to
 292 the choice of metric (see Fig. 15).

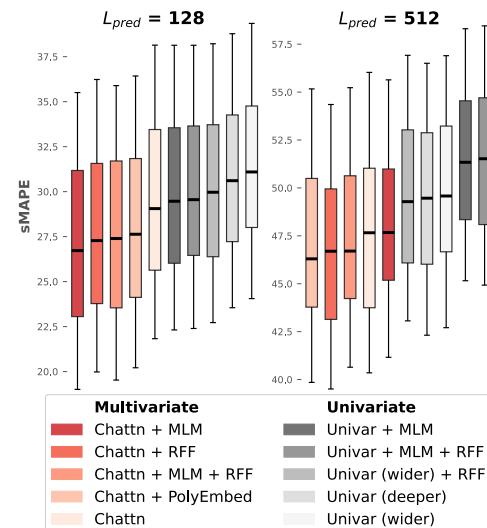
293 **Ablations.** We also ablate several features of *Panda*,
 294 in order to verify the contributions of our dynamics-
 295 based architectural choices. These include (1) Chan-
 296 nel Attention, (2) Dynamics Embedding, and (3)
 297 Masked Language Modeling (MLM) Pretraining.
 298

299 We observe a significant improvement due to chan-
 300 nel attention and MLM pretraining (See Section E
 301 for example zero-shot completions). However, the
 302 combined effect of the MLM with the dynamics em-
 303 bedding appears to be more complex: with no MLM,
 304 the dynamics embedding helps, but with MLM, it
 305 reduces performance. Moreover, the dynamics em-
 306 bedding improves the error on autoregressive rollout,
 307 whereas MLM reduces performance on rollout. We
 308 conclude that using the dynamics embedding with
 309 polynomial features (*PolyEmbed*) gives us the best
 310 model for long prediction horizons.

311 We include additional forecast metrics in Fig. 16 in
 312 Appendix D. We continue the discussion and evalua-
 313 tion of MLM on the completions task in Appendix E.
 314 In particular, we compare the correlation dimension
 315 of the completions against that of the ground truth
 316 trajectories (Fig. 18) and show a strong match. Fur-
 317 thermore, we investigate the effect of patch size on
 318 *Panda*’s performance in Appendix J.

319 5.2 *Panda* ZERO-SHOT FORECASTS EXPERIMENTAL DATA

320 We next show that *Panda* generalizes to experimental time series from real-world dynamical systems.
 321 These experimental datasets have nonstationarity, missing values, noise, and other complexities not
 322 seen during training. Following prior works, we select systems in which the experimental data is
 323 known to have an underlying dynamical process generating it: the positions and momenta of the
 324 tips of the two rods in an experimental recording of a double pendulum (Asseman et al., 2018), the
 325 leading independent components of body posture from a light microscopy video of *C. elegans* worms
 326 crawling on agar gel (Ahamed et al., 2021), and voltage recordings from networks of 28 randomly



327 **Figure 3: Ablations of key architectural fea-
 328 tures of *Panda*: MLM pretraining, chan-
 329 nel attention (*Chattn*), and components of
 330 the dynamics embedding (*RFF* denotes ran-
 331 dom Fourier features and *PolyEmbed* includes
 332 polynomial features).**

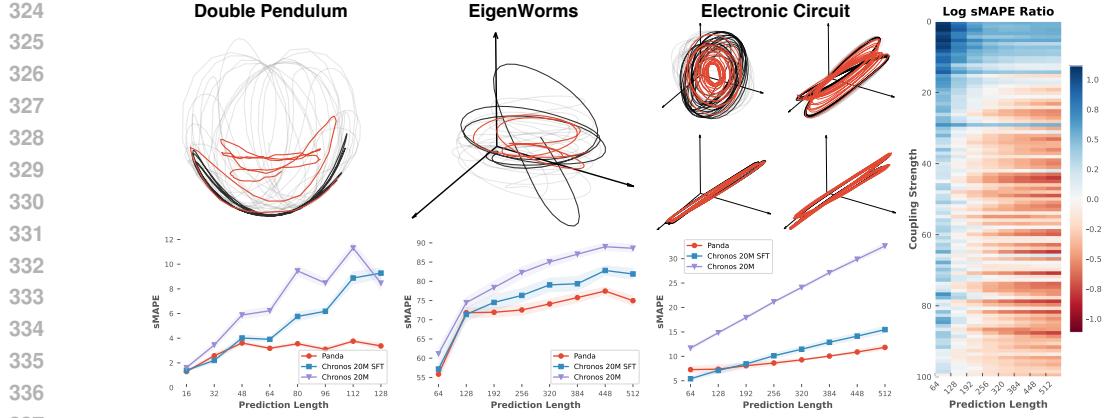


Figure 4: Zero-shot forecasts of experimental data from (a) Double Pendulum (Asseman et al., 2018), (b) Eigenworms (Ahamed et al., 2021), and (c) Electronic Circuit (Vera-Ávila et al., 2020). (d) Relative change in forecast error for *Panda* compared to Chronos-SFT (as measured in $\log(sMAPE_{Panda}/sMAPE_{Chronos-SFT})$), showing the advantage of our approach as the coupling strength between variables increases, for various prediction horizons.

connected electrical oscillators (Vera-Ávila et al., 2020). In all cases, the zero-shot performance of *Panda* outperforms Chronos-SFT (Fig. 4a).

For the circuit dataset in particular, we find that as the experimental coupling strength increases, the relative advantage of *Panda* over Chronos-SFT increases (red regions), particularly at long prediction horizons—leading to a visible Pareto front between the two models (Fig. 4b). This finding underscores the importance of channel attention for capturing nonlinear couplings typical in real world dynamical systems.

5.3 *Panda* EXHIBITS A DYNAMICAL SYSTEMS SCALING LAW

We create eight independent pretraining datasets that are subsets of the 2×10^4 unique systems generated using our methodology in Section 3. Across these eight datasets, we maintain a constant number of total timepoints while taking, at one extreme, a single trajectory (one initial condition) from each unique system, and at the other extreme, several trajectories (multiple initial conditions) from only a few unique systems. These datasets thus allow us to measure how dynamical diversity (unique systems versus initial conditions) affects generalization. We repeat our zero-shot evaluations on our set of 9.3×10^3 held-out systems for each model trained on the eight datasets.

In particular, let N_{ics} be the number of sampled initial conditions and N_{sys} the number of unique systems. Keeping $N_{ics} \times N_{sys}$ fixed, our eight dataset splits are constructed as $\{(N_{sys} \approx 2 \times 10^4, N_{ics} = 1), (N_{sys} \approx 10^4, N_{ics} = 2), \dots, (N_{sys} \approx 156, N_{ics} = 128)\}$, where each subsequent split uses a strict subset of the systems of the previous split, but with double the number of sampled initial conditions N_{ics} .

We observe clear scaling of zero-shot performance on unseen systems with the number of new dynamical systems encountered. We emphasize that this scaling law is distinct from traditional neural scaling laws for total training data, because we hold the number of timepoints constant while

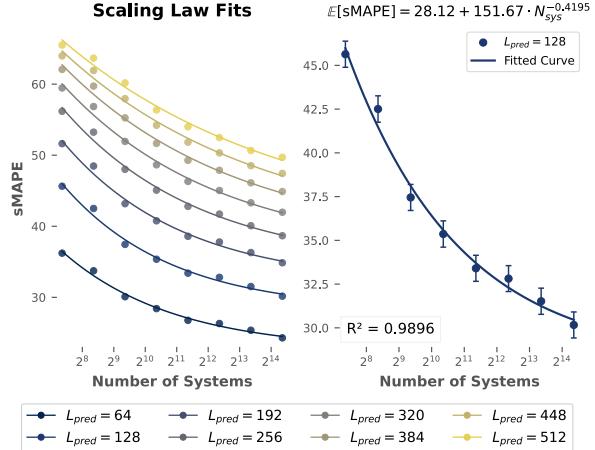


Figure 5: Scaling laws in zero-shot forecast error as the number of unique dynamical systems increases. The total amount of training timepoints is held constant.

controlling the diversity of the data (Kaplan et al., 2020). These results highlight the advantages of scaling with diverse synthetic data. This finding accords with classical nonlinear dynamics theory: additional on-attractor trajectories continuously produce new information about that particular attractor’s measure (a result of Pesin’s theorem), but beyond a certain point they fail to provide new topological information about winding, voids, etc (Gilmore, 1998; Pesin, 1977). The distinction between these information types partly explains the gap between in-domain and out-of-domain generalization in MLDS (Göring et al., 2024).

5.4 *Panda* EXHIBITS EMERGENT PDE FORECASTING CAPABILITY

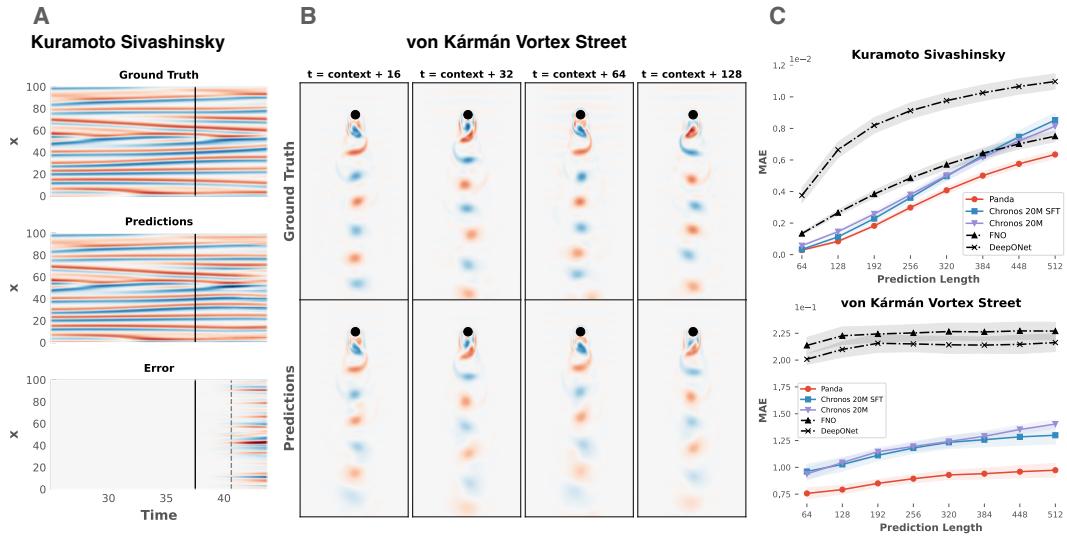


Figure 6: (A) Zero-shot forecasts of the Kuramoto-Sivashinsky equation. The time axis $t = 25$ to $t \approx 44$ contains 768 timepoints (512 context + 256 prediction): solid black line marks end of context window, and dashed gray line marks length 128 prediction horizon. (B) Zero-shot forecasts of the von Kármán vortex street. (C) The horizoned MAE (with standard errors bars) compared to baselines. We show point-wise MAE, due to sMAPE’s saturation at the upper bound. We include two baselines, Fourier Neural Operators and DeepONet (black traces) fully-trained on the context (see Appendix H).

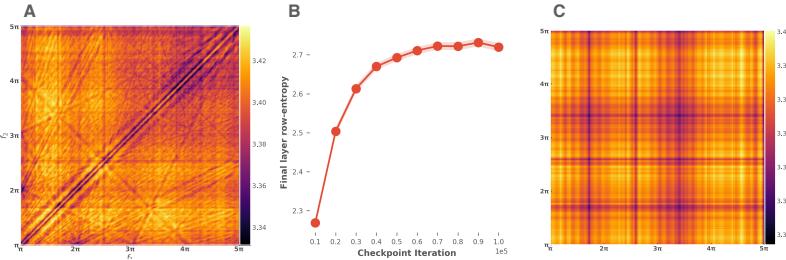
Partial differential equations (PDEs) are dynamical systems on continuous domains, with diverse applications in weather prediction or materials science (Kochkov et al., 2024). Conceptually, PDEs may be seen as coupled ordinary differential equations evolving in an infinite-dimensional space. We apply our trained model to the problem of forecasting two weakly-turbulent PDEs representing standard SciML benchmarks: the Von-Karman vortex street (VKVS) describes the unsteady motion of flow past a cylinder, and the Kuramoto-Sivashinsky (KS) models a laminar flame front (Cvitanović et al., 2010). More details on the PDE evaluation setup can be found in Appendix H.

Surprisingly, *Panda* outperforms baselines in zero-shot forecasting these systems (Fig. 6), *despite having never encountered PDE during training*. Unlike baselines, our model predicts nonlinear phenomena like merging of flame fronts in the KS equation or vortex pinchoff in the VKVS. While prior works require specially-trained models to forecast chaotic PDEs (Pathak et al., 2018), our zero-shot approach does not require extensive in-distribution training data, highlighting the advantages of cross-channel attention and multivariate training in generalization.

5.5 *Panda* DEVELOPS INTERPRETABLE INTERNAL REPRESENTATIONS OF COMPLEX DYNAMICS

To probe the role of channel attention in *Panda*, we feed two-tone sinusoids into the model and measure the response. The frequencies f_1, f_2 are each swept out over the range $[2\pi, 5\pi]$. Let \tilde{A} denote the attention rollout (Abnar & Zuidema, 2020) of the temporal attention matrices. Since \tilde{A} is the product of row-stochastic matrices, \tilde{A} remains row-stochastic. Thus, we can measure the response

432
433
434
435
436
437
438
439
440



441 Figure 7: (A) Nonlinear resonance structure measured by average row-wise entropy of temporal
442 attention rollout matrices. (B) Mean row-wise entropy of the final layer during training. (C) Lack of
443 nonlinear resonance structure in the temporal attention rollout entropy for our *univariate* ablation.

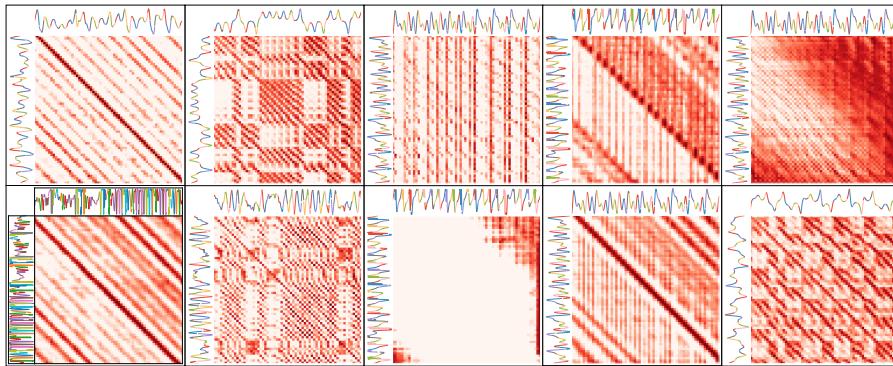
444
445
446
447
448
449
450

from "shaking" the model at frequency mixtures f_1, f_2 by measuring the average of the rowwise entropies of \hat{A} (c.f. Fig. 7A). The attention maps exhibit complex, multiscale structure indicating nonlinear resonance, a phenomenon in dynamical systems where a physical system (such as a kicked rotor or forced pendulum) exhibits gain with nonlinear dependence on the input frequencies. As a control, the frequency response of an equally trained univariate model does not exhibit the same nonlinear multiscale structure (Fig. 7C).

451
452
453
454
455
456
457
458

We next analyze *Panda*'s attention maps to probe its underlying forecast strategy. The attention maps largely concentrate mass away from the diagonals, which indicates that *Panda* effectively uses the context. In contrast, a model implementing a purely local rule (like a numerical integrator) would exhibit predominant diagonal structure, indicating that *Panda* performs more complex operations than few-step integration. For example, some attention maps form recurrence maps, which encode large-scale attractor geometry in classical nonlinear dynamics (Donner et al., 2010; Gilpin, 2025). Other layers show banding and circulant structure (Fig. 8), consistent with global integral transforms like Fourier series.

459
460
461
462
463
464
465
466
467
468
469
470



471
472
473
474

Figure 8: Temporal attention maps from *Panda* on context from different chaotic systems, showing Toeplitz, block, selector, and hybrid/combined structures (left to right). Appendix G further discusses spatiotemporal coupling and cross-channel maps (Fig. 22).

475
476

5.6 LIMITATION: REGRESSION TO THE MEAN

477
478
479
480
481
482
483
484
485

Divergence of error is inevitable when forecasting chaos with finite precision. Eventually, the prediction error grows to the point where a point forecast is useless, but invariant and geometric properties of the chaotic system can still be estimated in the long horizon regime. We quantify the utility of long horizon forecasts by measuring the geometry of forecasts much farther than $4 \times$ the training prediction horizon. Specifically, we compute: the maximum Lyapunov exponents (Appendix F.6); the forward KL divergence (Table 1) between the attractor and the predictions (Hess et al., 2023); and the spectral Hellinger distance (Table 2), an f -divergence between power spectra of the attractor and predictions (Mikhaeil et al., 2022).

For these experiments, we include *Dynamix*, a much smaller ($\sim 10K$ parameters) autoregressive dynamical systems foundation model that excels in capturing long term geometry (Hemmer &

486 Durstewitz, 2025). For Tables 1 and 2, we report mean \pm std. dev. across all test systems, averaged
 487 over 5 context windows for prediction horizons $L_{\text{pred}} \in \{512, 1024, 2048, 3072\}$. Only one context
 488 window is available for evaluating $L_{\text{pred}} = 3584$, as our dataset contains trajectories of length 4096.
 489

| $D_{KL}(\text{Ground Truth}(L_{\text{pred}}) \text{Model Prediction}(L_{\text{pred}}))$ | | | | | |
|--|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Model | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
| Panda | 3.93 ± 3.51 | 4.72 ± 3.64 | 5.63 ± 3.71 | 6.14 ± 3.68 | 6.39 ± 3.90 |
| Chronos 20M SFT | 4.72 ± 5.00 | 5.09 ± 4.90 | 5.62 ± 4.86 | 5.93 ± 4.84 | 6.05 ± 5.34 |
| Chronos 20M | 5.99 ± 5.07 | 6.19 ± 4.85 | 6.51 ± 4.76 | 6.76 ± 4.74 | 6.94 ± 5.41 |
| Chronos 200M | 5.12 ± 5.25 | 5.49 ± 5.22 | 6.05 ± 5.30 | 6.36 ± 5.28 | 6.47 ± 5.67 |
| DynaMix | 4.75 ± 5.70 | 4.90 ± 5.65 | 5.22 ± 5.72 | 5.40 ± 5.70 | 5.51 ± 6.13 |
| $\Delta\% (\uparrow)$ | $+16.7\%$ | $+3.7\%$ | -7.9% | -13.7% | -16.0% |

499 Table 1: KL divergence between ground truth and model predictions. $\Delta\%$ denotes percentage gain
 500 of *Panda* over the best baseline. See Table 12 for per-system differences, and Appendix F.1 for
 501 implementation details.

| Average $H^2(S_{\text{Ground Truth}(L_{\text{pred}})} S_{\text{Model Prediction}(L_{\text{pred}})})$ | | | | | |
|---|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Model | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
| Panda | 0.25 ± 0.14 | 0.25 ± 0.12 | 0.25 ± 0.11 | 0.25 ± 0.11 | 0.26 ± 0.12 |
| Chronos 20M SFT | 0.29 ± 0.17 | 0.29 ± 0.16 | 0.29 ± 0.16 | 0.30 ± 0.16 | 0.30 ± 0.18 |
| Chronos 20M | 0.37 ± 0.16 | 0.36 ± 0.16 | 0.37 ± 0.16 | 0.38 ± 0.16 | 0.38 ± 0.17 |
| Chronos 200M | 0.28 ± 0.16 | 0.28 ± 0.15 | 0.29 ± 0.15 | 0.30 ± 0.15 | 0.30 ± 0.17 |
| DynaMix | 0.36 ± 0.19 | 0.34 ± 0.19 | 0.33 ± 0.19 | 0.33 ± 0.19 | 0.32 ± 0.21 |
| $\Delta\% (\uparrow)$ | $+10.7\%$ | $+10.7\%$ | $+13.8\%$ | $+16.7\%$ | $+13.3\%$ |

513 Table 2: Average per-dimension spectral Hellinger distance between ground truth and model predictions.
 514 $\Delta\%$ denotes percentage improvement of *Panda* over the next closest baseline. See Table 13 for
 515 per-system differences. We use Welch’s method for estimating the PSD.

516 Surprisingly, we observe competitive performance in the KL divergence up to $8\times$ the training
 517 prediction horizon, and an across the board advantage on spectral Hellinger distance. However, we
 518 visually confirm that *Panda* tends to regress to the mean of the context when rolling out far past the
 519 training horizon in Appendix M. We can quantify this failure mode by computing the distributional
 520 metrics on the tail forecasts in Appendix F.5 which confirms the failure of mean regression as a long
 521 term forecaster. In contrast, *Chronos* exhibits parroting for long horizons (Fig. 31, 33, 32) which
 522 serves as a decent surrogate for long term attractor geometry due to forecasting periodic orbits. We
 523 report additional distributional metrics and computed invariant quantities in Appendix F.

525 6 CONCLUSION AND FUTURE DIRECTIONS

527 Our work demonstrates the feasibility of pretrained models in discovering generalizable properties of
 528 dynamical systems, mathematical objects of intrinsic interest to the SciML and forecasting communi-
 529 ties. Our model’s emergent ability to predict higher-dimensional partial differential equations, and
 530 the scaling of its performance with the diversity of dynamical systems, show that its generalization
 531 signal stems from unique properties of dynamics relative to time series.

532 A limitation of our work stems from our focus on low-dimensional dynamical systems. We argue
 533 that low-dimensional dynamics are the building block for higher-dimensional systems like weather
 534 front or spiking neurons, because they capture essential properties like bifurcations that become more
 535 complex in extended systems. A future variant of our approach for high-dimensional dynamics could
 536 exploit the structure of coupling such as sparsity or blocks typical in these systems by allowing the
 537 channel attention layers to receive custom attention masks. Another limitation is the degradation
 538 of rollout performance from MLM pretraining. Future work will investigate the question of what
 539 pretraining task is most natural for modeling dynamical systems. We believe this is a basic question
 that necessitates further progress in SciML.

540 REFERENCES
541

542 Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers, 2020. URL
543 <https://arxiv.org/abs/2005.00928>.

544 Tosif Ahamed, Antonio C Costa, and Greg J Stephens. Capturing the continuous complexity of
545 behaviour in *caenorhabditis elegans*. *Nature Physics*, 17(2):275–283, 2021.

546 Abdul Fatir Ansari, Lorenzo Stella, Ali Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin
547 Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor,
548 et al. Chronos: Learning the language of time series. *Transactions on Machine Learning Research*,
549 2024.

550 Alexis Asseman, Tomasz Kornuta, and Ahmet Ozcan. Learning beyond simulated physics. In
551 *Modeling and Decision-making in the Spatiotemporal Domain Workshop*, 2018. URL <https://openreview.net/pdf?id=HylajWsRF7>.

552 Federico Barbero, Alex Vitvitskyi, Christos Perivolaropoulos, Razvan Pascanu, and Petar Velicković.
553 Round and round we go! what makes rotary positional encodings useful? In *The Thirteenth
554 International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=GtvuNrk58a>.

555 Sören Becker, Michal Klein, Alexander Neitz, Giambattista Parascandolo, and Niki Kilbertus.
556 Predicting ordinary differential equations with transformers. In *International conference on
557 machine learning*, pp. 1978–2002. PMLR, 2023.

558 Dhananjay Bhaskar, Daniel Sumner Magruder, Matheo Morales, Edward De Brouwer, Aarthi Venkat,
559 Frederik Wenkel, Guy Wolf, and Smita Krishnaswamy. Inferring dynamic regulatory interaction
560 graphs from time series data with perturbations. In *Learning on Graphs Conference*, pp. 22–1.
561 PMLR, 2024.

562 Manuel Brenner, Christoph Jürgen Hemmer, Zahra Monfared, and Daniel Durstewitz. Almost-linear
563 rnns yield highly interpretable symbolic codes in dynamical systems reconstruction. *Advances in
564 Neural Information Processing Systems*, 37:36829–36868, 2024.

565 Steven L Brunton, Marko Budisić, Eurika Kaiser, and J Nathan Kutz. Modern koopman theory for
566 dynamical systems. *SIAM Review*, 64(2):229–340, 2022.

567 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
568 differential equations. *Advances in neural information processing systems*, 31, 2018.

569 Wuyang Chen, Jialin Song, Pu Ren, Shashank Subramanian, Dmitry Morozov, and Michael W
570 Mahoney. Data-efficient operator learning via unsupervised pretraining and in-context learning.
571 *Advances in Neural Information Processing Systems*, 37:6213–6245, 2024.

572 Xiaoyuan Cheng, Yi He, Yiming Yang, Xiao Xue, Sibo Chen, Daniel Giles, Xiaohang Tang, and
573 Yukun Hu. Learning chaos in a linear way. In *The Thirteenth International Conference on Learning
574 Representations*, 2025.

575 Maximilian Christ, Nils Braun, Julius Neuffer, and Andreas W Kempa-Liehr. Time series feature
576 extraction on basis of scalable hypothesis tests (tsfresh—a python package). *Neurocomputing*, 307:
577 72–77, 2018.

578 Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. Power-law distributions in empirical
579 data. *SIAM Review*, 51(4):661–703, November 2009. ISSN 1095-7200. doi: 10.1137/070710111.
580 URL <http://dx.doi.org/10.1137/070710111>.

581 James P Crutchfield and Melanie Mitchell. The evolution of emergent computation. *Proceedings of
582 the National Academy of Sciences*, 92(23):10742–10746, 1995.

583 Predrag Cvitanović, Ruslan L Davidchack, and Evangelos Siminos. On the state space geometry
584 of the kuramoto–sivashinsky flow in a periodic domain. *SIAM Journal on Applied Dynamical
585 Systems*, 9(1):1–33, 2010.

594 Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for
595 time-series forecasting. In *Forty-first International Conference on Machine Learning*, 2024.
596

597 Stéphane d’Ascoli, Sören Becker, Alexander Mathis, Philippe Schwaller, and Niki Kilbertus.
598 Odeformer: Symbolic regression of dynamical systems with transformers. *arXiv preprint*
599 *arXiv:2310.05573*, 2023.

600 David A. Dickey and Wayne A. Fuller. Distribution of the estimators for autoregressive time series
601 with a unit root. *Journal of the American Statistical Association*, 74(366):427–431, 1979. ISSN
602 01621459, 1537274X. URL <http://www.jstor.org/stable/2286348>.

603

604 Reik V Donner, Yong Zou, Jonathan F Donges, Norbert Marwan, and Jürgen Kurths. Recurrence
605 networks—a novel paradigm for nonlinear time series analysis. *New Journal of Physics*, 12(3):
606 033025, 2010.

607 Mengge Du, Yuntian Chen, Zhongzheng Wang, Longfeng Nie, and Dongxiao Zhang. Large language
608 models for automatic equation discovery of nonlinear dynamics. *Physics of Fluids*, 36(9), 2024.

609

610 Ian Falconer, Georg A Gottwald, Ian Melbourne, and Kjetil Wormnes. Application of the 0-1 test for
611 chaos to experimental data. *SIAM Journal on Applied Dynamical Systems*, 6(2):395–402, 2007.

612 Martin J Falk, Finnegan D Roach, William Gilpin, and Arvind Murugan. Curiosity-driven search for
613 novel nonequilibrium behaviors. *Physical Review Research*, 6(3):033052, 2024.

614

615 Daniel J Gauthier, Erik Boltt, Aaron Griffith, and Wendson AS Barbosa. Next generation reservoir
616 computing. *Nature communications*, 12(1):1–8, 2021.

617

618 Robert Gilmore. Topological analysis of chaotic dynamical systems. *Reviews of Modern Physics*, 70
619 (4):1455, 1998.

620

621 William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling.
622 *NeurIPS*, 34, 2021.

623

624 William Gilpin. Model scale versus domain knowledge in statistical forecasting of chaotic systems.
625 *Phys. Rev. Research*, 5(4):043252, 2023.

626

627 William Gilpin. Recurrences reveal shared causal drivers of complex time series. *Physical Review X*,
628 15(1):011005, 2025.

629

630 Niclas Alexander Göring, Florian Hess, Manuel Brenner, Zahra Monfared, and Daniel Durstewitz.
631 Out-of-domain generalization in dynamical systems reconstruction. In Ruslan Salakhutdinov, Zico
632 Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp
633 (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of
634 *Proceedings of Machine Learning Research*, pp. 16071–16114. PMLR, 21–27 Jul 2024. URL
635 <https://proceedings.mlr.press/v235/goring24a.html>.

636

637 Peter Grassberger and Itamar Procaccia. Characterization of strange attractors. *Phys. Rev. Lett.*,
638 50:346–349, Jan 1983a. doi: 10.1103/PhysRevLett.50.346. URL <https://link.aps.org/doi/10.1103/PhysRevLett.50.346>.

639

640 Peter Grassberger and Itamar Procaccia. Estimation of the kolmogorov entropy from a chaotic
641 signal. *Phys. Rev. A*, 28:2591–2593, Oct 1983b. doi: 10.1103/PhysRevA.28.2591. URL <https://link.aps.org/doi/10.1103/PhysRevA.28.2591>.

642

643 Christoph Jürgen Hemmer and Daniel Durstewitz. True zero-shot inference of dynamical systems
644 preserving long-term statistics. *Advances in Neural Information Processing Systems*, 38, 2025.

645

646 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint*
647 *arXiv:1606.08415*, 2016.

648

649 Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel
650 de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *Advances in*
651 *Neural Information Processing Systems*, 37:72525–72624, 2024.

648 Florian Hess, Zahra Monfared, Manuel Brenner, and Daniel Durstewitz. Generalized teacher forcing
649 for learning chaotic dynamics. In *Proceedings of the 40th International Conference on Machine*
650 *Learning*, ICML'23. JMLR.org, 2023.

651 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott
652 Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models.
653 *arXiv preprint arXiv:2001.08361*, 2020.

654 Dmitrii Kochkov, Janni Yuval, Ian Langmore, Peter Norgaard, Jamie Smith, Griffin Mooers, Milan
655 Klöwer, James Lottes, Stephan Rasp, Peter Düben, et al. Neural general circulation models for
656 weather and climate. *Nature*, 632(8027):1060–1066, 2024.

657 Georgia Koppe, Hazem Toutounji, Peter Kirsch, Stefanie Lis, and Daniel Durstewitz. Identifying
658 nonlinear dynamical systems via generative recurrent neural networks with applications to fmri.
659 *PLoS computational biology*, 15(8):e1007263, 2019.

660 Jean Kossaifi, Nikola Kovachki, Zongyi Li, Davit Pitt, Miguel Liu-Schiaffini, Robert Joseph George,
661 Boris Bonev, Kamyar Azizzadenesheli, Julius Berner, and Anima Anandkumar. A library for
662 learning neural operators, 2024.

663 Nikola B. Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya,
664 Andrew M. Stuart, and Anima Anandkumar. Neural operator: Learning maps between function
665 spaces. *CoRR*, abs/2108.08481, 2021.

666 Akarsh Kumar, Chris Lu, Louis Kirsch, Yujin Tang, Kenneth O Stanley, Phillip Isola, and David Ha.
667 Automating the search for artificial life with foundation models. *arXiv preprint arXiv:2412.17799*,
668 2024.

669 Jose Nathan Kutz, Steven L. Brunton, Bingni W. Brunton, and Joshua L. Proctor. *Dynamic Mode*
670 *Decomposition : Data-Driven Modeling of Complex Systems*. SIAM, 2016. ISBN 9781611974492
671 1611974496. URL <http://www.dmdbook.com/>.

672 Denis Kwiatkowski, Peter CB Phillips, Peter Schmidt, and Yongcheol Shin. Testing the null
673 hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1-3):
674 159–178, 1992.

675 William La Cava, Bogdan Burlacu, Marco Virgolin, Michael Kommenda, Patryk Orzechowski,
676 Fabrício Olivetti de França, Ying Jin, and Jason H Moore. Contemporary symbolic regression
677 methods and their relative performance. *Advances in neural information processing systems*, 2021
678 (DB1):1, 2021.

679 Jonathan Lee, Annie Xie, Aldo Pacchiano, Yash Chandak, Chelsea Finn, Ofir Nachum, and Emma
680 Brunskill. In-context decision-making from supervised pretraining. In *ICML Workshop on New*
681 *Frontiers in Learning, Control, and Dynamical Systems*, 2023.

682 Zongyi Li, Miguel Liu-Schiaffini, Nikola Kovachki, Burigede Liu, Kamyar Azizzadenesheli, Kaushik
683 Bhattacharya, Andrew Stuart, and Anima Anandkumar. Learning dissipative dynamics in chaotic
684 systems. In *Proceedings of the 36th International Conference on Neural Information Processing*
685 *Systems*, pp. 16768–16781, 2022.

686 Yixin Lin, Jan Humplík, Sandy H Huang, Leonard Hasenclever, Francesco Romano, Stefano Saliceti,
687 Daniel Zheng, Jose Enrique Chen, Catarina Barros, Adrian Collister, et al. Proc4gem: Foundation
688 models for physical agency through procedural generation. *arXiv preprint arXiv:2503.08593*,
689 2025.

690 Toni JB Liu, Nicolas Boullé, Raphaël Sarfati, and Christopher J Earls. Llms learn governing principles
691 of dynamical systems, revealing an in-context neural scaling law. *arXiv preprint arXiv:2402.00795*,
692 2024.

693 Edward Lorenz. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, 20(2):130–148,
694 1963.

695 Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. DeepXDE: A deep learning library
696 for solving differential equations. *SIAM Review*, 63(1):208–228, 2021. doi: 10.1137/19M1274067.

702 Michael McCabe, Régaldo-Saint Blancard, Liam Parker, Ruben Ohana, Miles Cranmer, Alberto Bietti,
703 Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, et al. Multiple physics
704 pretraining for spatiotemporal surrogate models. *Advances in Neural Information Processing
705 Systems*, 37:119301–119335, 2024.

706 Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. Umap: Uniform manifold
707 approximation and projection. *Journal of Open Source Software*, 3(29):861, 2018.

708 Igor Mezić. Analysis of fluid flows via spectral properties of the koopman operator. *Annual review of
709 fluid mechanics*, 45(1):357–378, 2013.

710 Jonas M. Mikhæil, Zahra Monfared, and Daniel Durstewitz. On the difficulty of learning chaotic
711 dynamics with rnns. In *Proceedings of the 36th International Conference on Neural Infor-
712 mation Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN
713 9781713871088.

714 Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64
715 words: Long-term forecasting with transformers. In *The Eleventh International Conference on
716 Learning Representations*, 2022.

717 Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan,
718 Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads.
719 *arXiv preprint arXiv:2209.11895*, 2022.

720 Edward Ott. Strange attractors and chaotic motions of dynamical systems. *Rev. Mod. Phys.*, 53:
721 655–671, Oct 1981. doi: 10.1103/RevModPhys.53.655. URL <https://link.aps.org/doi/10.1103/RevModPhys.53.655>.

722 Norman H Packard, James P Crutchfield, J Doyne Farmer, and Robert S Shaw. Geometry from a
723 time series. *Physical review letters*, 45(9):712, 1980.

724 Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction
725 of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical
726 review letters*, 120(2):024102, 2018.

727 Louis M Pecora and Thomas L Carroll. Synchronization in chaotic systems. *Physical review letters*,
728 64(8):821, 1990.

729 Chris Pedersen, Laure Zanna, and Joan Bruna. Thermalizer: Stable autoregressive neural emulation
730 of spatiotemporal chaos. *arXiv preprint arXiv:2503.18731*, 2025.

731 Ya B Pesin. Characteristic lyapunov exponents and smooth ergodic theory. *Russian Mathematical
732 Surveys*, 32(4):55, 1977.

733 Ilan Price, Alvaro Sanchez-Gonzalez, Ferran Alet, Tom R Andersson, Andrew El-Kadi, Dominic
734 Masters, Timo Ewalds, Jacklynn Stott, Shakir Mohamed, Peter Battaglia, et al. Probabilistic
735 weather forecasting with machine learning. *Nature*, 637(8044):84–90, 2025.

736 Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines.
737 In J. Platt, D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural In-
738 formation Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL
739 https://proceedings.neurips.cc/paper_files/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf.

740 Chris Reinke, Mayalen Etcheverry, and Pierre-Yves Oudeyer. Intrinsically motivated discovery of di-
741 verse patterns in self-organizing systems. In *International Conference on Learning Representations*,
742 2020.

743 Michael T. Rosenstein, James J. Collins, and Carlo J. De Luca. A practical method for calculating
744 largest lyapunov exponents from small data sets. *Phys. D*, 65(1–2):117–134, May 1993. ISSN
745 0167-2789. doi: 10.1016/0167-2789(93)90009-P. URL [https://doi.org/10.1016/0167-2789\(93\)90009-P](https://doi.org/10.1016/0167-2789(93)90009-P).

756 David Ruelle and Floris Takens. On the nature of turbulence. *Communications in Mathematical*
757 *Physics*, 20(3):167 – 192, 1971.

758

759 Patrick Seifner, Kostadin Cvejoski, Antonia Körner, and Ramses J Sanchez. Zero-shot imputation with
760 foundation inference models for dynamical systems. In *The Thirteenth International Conference*
761 *on Learning Representations*, 2025.

762 Rajat Sen, Hsiang-Fu Yu, and Inderjit S Dhillon. Think globally, act locally: A deep neural network
763 approach to high-dimensional time series forecasting. *Advances in neural information processing*
764 *systems*, 32, 2019.

765

766 Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, and Ming Jin. Time-
767 moe: Billion-scale time series foundation models with mixture of experts, 2024. URL <https://arxiv.org/abs/2409.16040>.

768

769 Zezheng Song, Jiaxin Yuan, and Haizhao Yang. Fmint: Bridging human designed and data pretrained
770 models for differential equation foundation model. *arXiv preprint arXiv:2404.14688*, 2024.

771

772 Steven H Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry,*
773 *and engineering*. CRC press, 2018.

774

775 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: Enhanced
776 transformer with rotary position embedding, 2023. URL <https://arxiv.org/abs/2104.09864>.

777

778 Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W
779 Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Charac-
780 terizing scaling and transfer behavior. *Advances in Neural Information Processing Systems*, 36:
781 71242–71262, 2023.

782

783 Jingmin Sun, Yuxuan Liu, Zecheng Zhang, and Hayden Schaeffer. Towards a foundation model for
784 partial differential equations: Multioperator learning and extrapolation. *Physical Review E*, 111(3):
785 035304, 2025.

786

787 Floris Takens. Dynamical systems and turbulence. *Warwick, 1980*, pp. 366–381, 1981.

788

789 Tapas Tripura and Souvik Chakraborty. A foundational neural operator that continuously learns
790 without forgetting. *arXiv preprint arXiv:2310.18885*, 2023.

791

792 VP Vera-Ávila, Ricardo Sevilla-Escoboza, AA Lozano-Sánchez, RR Rivera-Durón, and Javier M
793 Buldú. Experimental datasets of networks of nonlinear oscillators: Structure and dynamics during
794 the path to synchronization. *Data in brief*, 28:105012, 2020.

795

796 Rui Wang, Robin Walters, and Rose Yu. Data augmentation vs. equivariant networks: A theory of
797 generalization on dynamics forecasting. *arXiv preprint arXiv:2206.09450*, 2022.

798

799 Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory
800 Erickson, David Held, and Chuang Gan. Robogen: Towards unleashing infinite data for automated
801 robot learning via generative simulation. In *International Conference on Machine Learning*, pp.
802 51936–51983. PMLR, 2024.

803

804 Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data–driven approximation
805 of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*,
806 25:1307–1346, 2015.

807

808 Biao Zhang and Rico Sennrich. Root mean square layer normalization. In H. Wal-
809 lach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Ad-
810 vances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.,
811 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/1e8a19426224ca89e83cef47f1e7f53b-Paper.pdf.

812

813 Xiangyuan Zhang, Weichao Mao, Haoran Qiu, and Tamer Başar. Decision transformer as a foundation
814 model for partially observable continuous control. *arXiv preprint arXiv:2404.02407*, 2024.

810 Yuanzhao Zhang and William Gilpin. Context parroting: A simple but tough-to-beat baseline for
811 foundation models in scientific machine learning. *arXiv preprint arXiv:2505.11349*, 2025a.
812
813 Yuanzhao Zhang and William Gilpin. Zero-shot forecasting of chaotic systems. In *The Thirteenth*
814 *International Conference on Learning Representations*, 2025b. URL <https://arxiv.org/abs/2409.15771>.
815
816 Martin Ziegler, Andres Felipe Posada-Moreno, Friedrich Solowjow, and Sebastian Trimpe. On foun-
817 dation models for dynamical systems from purely synthetic data. *arXiv preprint arXiv:2412.00395*,
818 2024.
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864 A GENERATION OF A NOVEL CHAOTIC SYSTEMS DATASET

865

866 A.1 SKEW-PRODUCT SYSTEMS

867

868 We algorithmically discover skew-product systems following the methodology described in Section 3.
869 Here, we present a subset of 30 of these systems, out of a total of 2×10^4 in our training set.

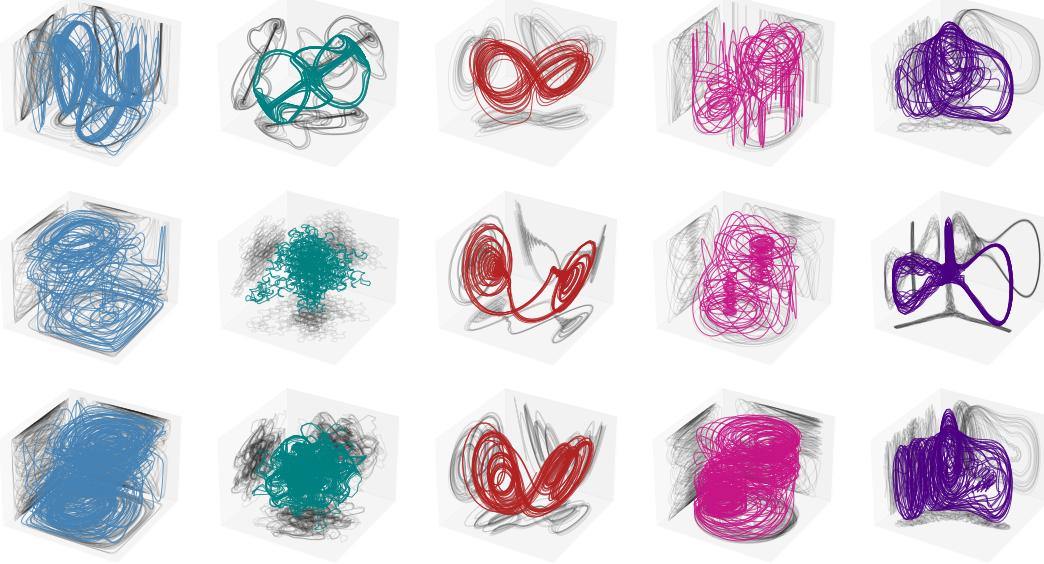


911 Figure 9: Examples of novel chaotic skew-product systems discovered via evolutionary search.
912 Shaded regions correspond to two-dimensional projections onto the corresponding axes.
913

914 Our starting point is a hand-curated, crowdsourced public dataset of 129 chaotic low-dimensional
915 dynamical systems from the nonlinear dynamics literature (Gilpin, 2021; 2023; Zhang & Gilpin,
916 2025b). Each entry comprises a set of coupled ordinary differential equations with dimensionality
917 between three and ten. The parameters and initial conditions for each system have been hand-tuned
into the chaotic regime, based on values used in previously-published studies.

918 A.2 MUTATION OF BASE SYSTEMS
919

920 We also generate new instances of the base 129 chaotic systems by perturbing the ODE parameters.
921



922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941 Figure 10: Examples of parameter perturbations of base systems. **Top row** Unperturbed original
942 systems. **Bottom rows** Parameter perturbations of the top row systems.
943

944 A.3 NUMERICAL INTEGRATION
945

946 For all ODEs, we integrate trajectories of 4096 timesteps with an integration time-span dictated
947 by a system's characteristic timescale based off the dominant modes in the power spectrum. We
948 thus call this timescale the *period* of the system and set the integration time-span to be $[0, N_p \times \varphi]$
949 where N_p is the number of periods, and φ is the "period" measured from integrating test trajectories
950 of the base system; for skew systems we take the period to be the maximum period between the
951 driver and response systems. For all experiments in the main text, we take $N_p = 40$, but use a larger
952 mixed-period dataset in our scaled up experiments (Appendix K).

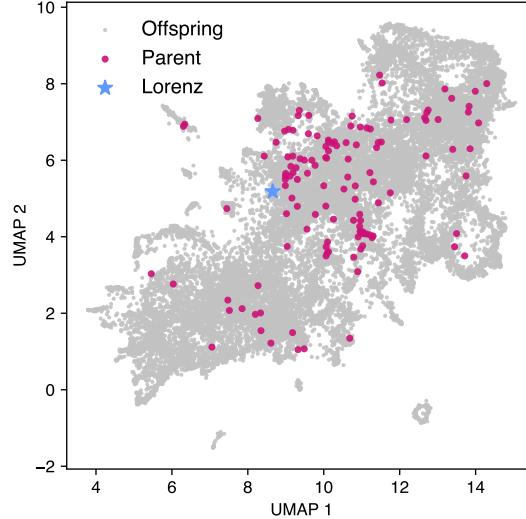
953 The numerical integration timestep is controlled via adaptive step-sizing from the *Radau* solver, a 5th
954 order implicit Runge Kutta Scheme. For high quality trajectory data, we integrate using a relative
955 tolerance 1×10^{-9} and an absolute tolerance of 1×10^{-10} . The initial conditions for discovered
956 systems are obtained by integrating a test trajectory at a lower tolerance ($\text{rtol} = 1e-6$, $\text{atol} = 1e-7$)
957 and sampling a point from the coarse trajectory which approximates starting at a point *on attractor*.

958 A.4 ATTRACTOR SELECTION
959

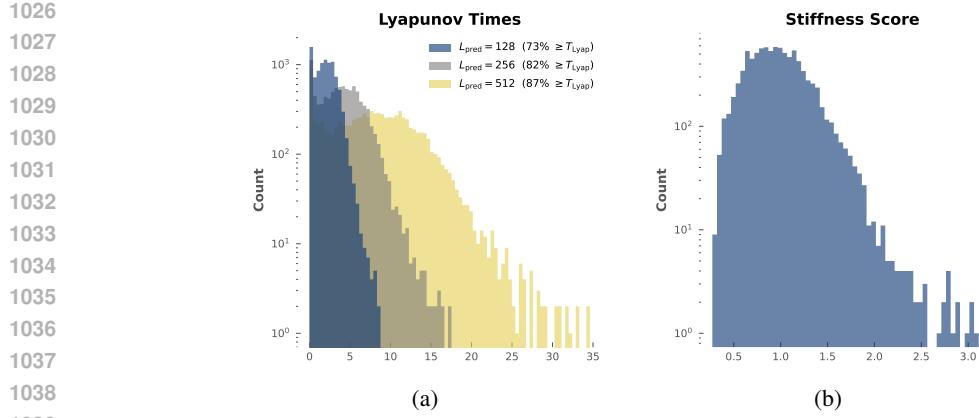
960 The only general way to identify properties about chaotic dynamical systems is to integrate them.
961 This fundamental fact makes the system discovery process described in Section 3 very expensive at
962 scale. To effectively reduce the number of incoming candidates for chaoticity selection and validation,
963 we employ callbacks during integration that will immediately kill the process and prune that system
964 candidate. Specifically, we terminate integration whenever the step size falls below 10^{-10} , any
965 bounded (non-driving dimension) coordinate exceeds 10^4 in value, and whenever the integration
966 time exceeds 5 minutes. The surviving systems will finish integration and move on to the chaoticity
967 selection phase (see the overview of our selection for chaoticity in Section 3).
968
969
970
971

972 A.5 DATASET PROPERTIES
973

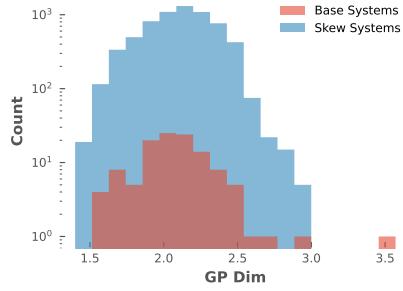
974 We verify that our integrated trajectories exhibit chaotic dynamics by measuring the average number
975 of Lyapunov times in various prediction horizons. A chaotic flow separates nearby initial conditions
976 according to $|\delta(t)| \approx \exp(\lambda_1 t)|\delta(0)|$ where δ is the time dependent separation and λ_1 is the maximum
977 Lyapunov exponent. Thus, a Lyapunov time is defined to be $T_{\text{Lyap}} := 1/\lambda_1$. Given an arbitrary time
978 series with timestep δt , the Lyapunov times per N timepoints is then $\lambda_1 \times \delta t \times N$. Since we rely on
979 an implicit integrator with adaptive step-sizing, we compute the average timestep over the integration
980 timespan and estimate the maximum Lyapunov exponent using the Rosenstein estimator (Rosenstein
981 et al., 1993) to compute the distribution of Lyapunov times per horizon length in Fig. 12a where it is
982 clear that we are predicting in the chaotic regime most of the time.

983 To ensure consistency between the founder population and offspring, we featurize all pretraining time
984 series using the same procedure as previous works reporting chaotic systems datasets (Gilpin, 2021;
985 2023). For each channel of a D -dimensional multivariate time series (4096 timepoints, 100 points per
986 dominant Fourier period) we compute a vector of 749 standard time-shift invariant time series features
987 like wavelet modes, signal power, reversion rate, etc. using the `tsfresh` library (Christ et al., 2018).
988 We average the D feature vectors for each system to produce a single channel-permutation invariant
989 feature vector for each skew-product system. We then project all 2×10^4 pretraining skew-product
990 systems into 2D using UMAP, a nonlinear embedding algorithm that preserves the local neighbors
991 of each point from the high dimensional space (distances, however, are not necessarily preserved)
992 (McInnes et al., 2018). We next featurize and embed the 135 parent systems from the founder
993 population into the same space. We observe broad dispersion of the parent systems across the child
994 population, implying the absence of mode collapse or strong distribution shift between the parents
995 and offspring (Fig. 11). We interpret this result as the absence of strong founder or inbreeding effects
996 in the offspring generation.

1015 Figure 11: A low-dimensional embedding of the 2×10^4 skew-product systems (gray), as well as
1016 the 135 founding parent systems (magenta) from which these offspring systems are evolved. The
1017 well-known chaotic Lorenz attractor is starred on the plot.

1018
1019
1020
1021 Additionally, for all trajectories in the test set, we measure the empirical stiffness score defined as
1022 $S := \log_{10}(\max_t |\Delta_t| / \mathbb{E}_t |\Delta_t|)$, where Δ_t is the finite difference (forward or backward) at time
1023 t . Fig. 12b shows that most test systems have at least an order of magnitude scaling between the
1024 largest observation jump compared to the average change *per channel*. This distribution suggests that
1025 the dataset generation algorithm generates stiff systems and reinforces the fact that the integrated
1026 trajectories exhibit non-trivial dynamics.



1040 Figure 12: Dynamical properties of systems in the test set. (a) Distribution of Lyapunov times within
1041 L_{pred} timepoints; annotated with the proportion of systems which exceed 1 Lyapunov time in the
1042 horizon. (b) Distribution of stiffness scores (log-ratio of largest delta compared to the average delta).



1055 Figure 13: Distribution of correlation
1056 dimension (Grassberger-Procaccia) of
1057 skew systems and their founder (base)
1058 systems.

1059 As shown in Fig. 13, our skew-product generation preserves the distribution of the correlation
1060 dimension, an invariant quantity used as a proxy for fractal dimension. This suggests that our dataset
1061 **does not** suffer from a "founder effect" that would kill off diversity. Table 4 further presents a
1062 comparison of the Kullback Leibler divergence between the invariant measures of the attractors, for
1063 skew systems: with the same parents; with different parents; with one parent shared; and between
1064 parent (Base) and child (Skew) systems.

| System | GP Dim (mean \pm std) |
|--------------|-------------------------|
| Base Systems | 2.09 ± 0.27 |
| Skew Systems | 2.11 ± 0.23 |

Table 3: Correlation dimension for Base (Founder) systems and Skew (Children) systems.

| Metric | mean \pm std | N_{combos} |
|---|-----------------|---------------------|
| $D_{KL}(\text{Skew} \parallel \text{Response})$ | 5.35 ± 5.81 | 10,000 |
| $D_{KL}(\text{Skew} \parallel \text{Driver})$ | 8.46 ± 6.58 | 10,000 |
| $D_{KL}(\text{Skew} \parallel \text{Non-Parent})$ | 9.01 ± 6.64 | 10,000 |
| $D_{KL}(\text{Skew Intra})$ | 3.39 ± 5.11 | 10,000 |
| $D_{KL}(\text{Skew Inter})$ | 7.07 ± 6.58 | 10,000 |
| $D_{KL}(\text{Base Intra})$ | 2.54 ± 4.13 | 6,000 |
| $D_{KL}(\text{Base Inter})$ | 8.24 ± 6.48 | 6,000 |

1074 Table 4: D_{KL} between skew systems and: 1) param parts of response; 2) param parts of driver; 3)
1075 param parts of non-parent system in the founder pool. (Skew Intra) D_{KL} between param parts of
1076 skew systems with the same parents. (Skew Inter) D_{KL} between param parts of skew systems with
1077 different parents. (Base Intra) D_{KL} between parameter perturbations of the same founder system.
1078 (Base Inter): D_{KL} between parameter perturbations of different founder systems.
1079

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095

Table 5: Model Architecture

| Parameter | Value |
|--------------------|---------|
| Context length | 512 |
| Prediction length | 128 |
| Hidden layers | 8 |
| d_{model} | 512 |
| FFN dimension | 512 |
| Attention heads | 8 |
| Activation | gelu |
| Pre-norm | True |
| Normalization | RMSNorm |
| Init std | 0.02 |

Table 6: Model Architecture (Continued)

| Parameter | Value |
|-----------------------|---------|
| Patch length / stride | 16 / 16 |
| Rope percent | 0.75 |
| Max wavelength | 500 |
| Poly features | 120 |
| Poly degrees | 2 |
| RFF count | 256 |
| RFF scale | 1.0 |

B TRAINING DETAILS

A technical difficulty of training a global multivariate model is forming batches of trajectories with mixed channel dimensions. We look to dynamical systems theory and note that it is well known that at least 3 coupled dynamical variables are necessary for a system to exhibit deterministic chaos in continuous-time (Strogatz, 2018). To this end, we fix the dimensions of each input trajectory to 3 *only during training* by randomly sampling 3 channels from each multivariate trajectory to enable efficient batching. During inference time, we process the full multivariate trajectories. For the 21.3M parameter *Panda* model we use $d_{\text{model}} = d_{\text{ffn}} = 512$, $N_{\text{heads}} = N_{\text{layers}} = 8$, $N_{\text{poly}} = 120$ with degree 2, and $N_{\text{rff}} = 256$. For the 41.5M parameter model (Appendix K), we use $d_{\text{model}} = d_{\text{ffn}} = 640$, $N_{\text{heads}} = N_{\text{layers}} = 10$, $N_{\text{poly}} = 156$ with degree 2, and $N_{\text{rff}} = 312$. And for the 71.5M parameter *Panda-72M* (Appendix K), we use $d_{\text{model}} = d_{\text{ffn}} = 768$, $N_{\text{heads}} = N_{\text{layers}} = 12$, $N_{\text{poly}} = 188$ with degree 2, and $N_{\text{rff}} = 376$. Additionally, data augmentations (Section 3) are uniformly randomly applied to 20% of the training trajectories.

We use a patch size (and patch stride) of 16. All models are trained with a context length of 512, which corresponds to 32 patches, and use a non-causal transformer encoder with 8 layers, each with $d_{\text{model}} = 512$ and 8 heads. Each attention block maps a (batch size, channels, patches, hidden) sized hidden state H via:

$$\begin{aligned} H &\leftarrow H + \text{RopeTemporalAttention} \circ \text{RMSNorm}(H) \\ H &\leftarrow H + \text{ChannelAttention} \circ \text{RMSNorm}(H)^\top \\ H &\leftarrow H + \text{FFN} \circ \text{RMSNorm}(H) \end{aligned}$$

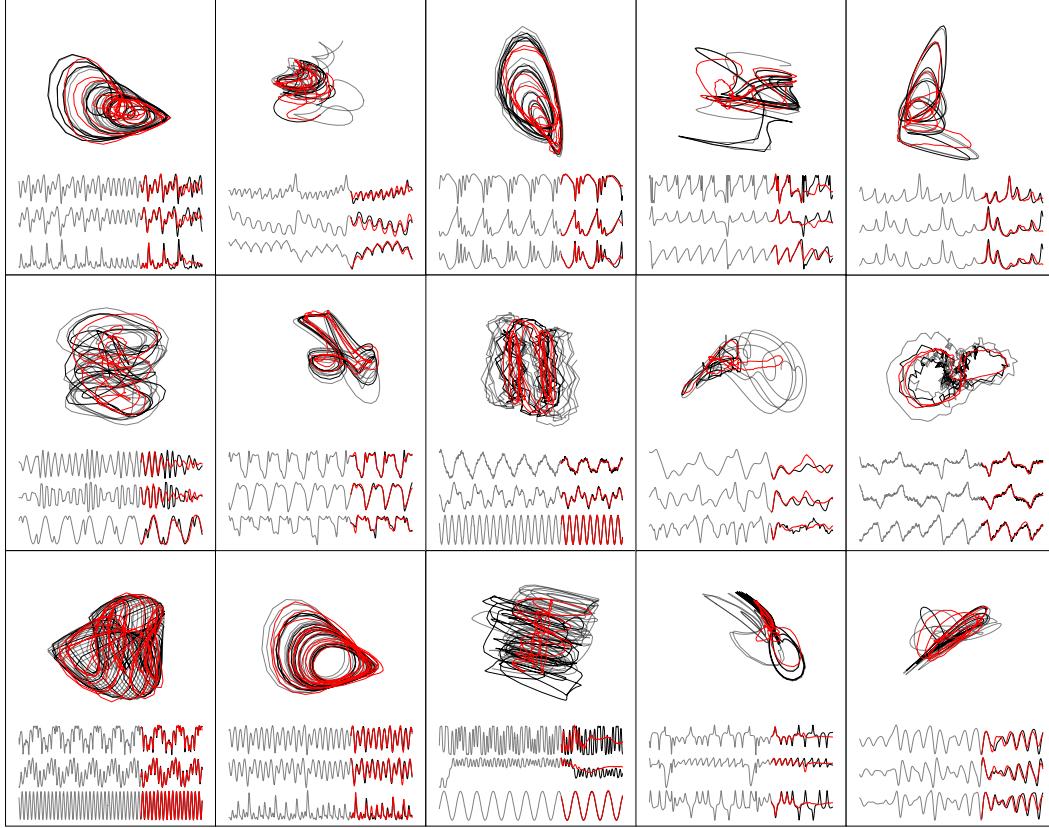
Where the transpose is applied to the channel and patch (sequence) dimension.

For models optimized with masked language modeling (MLM) style pretraining (masking and reconstructing intermediate patch tokens), a linear head is used to infill masked patches. For the forecasting model, a prediction head aggregates the encoder hidden states via a mean along the sequence (patch) dimension and a linear layer maps this representation to a *fixed-length* 128 forecast for **all** models. All models are trained with MSE loss and the AdamW optimizer with a maximum learning rate 1×10^{-3} on a cosine schedule with a 10% warmup. Additionally, we train with gradient norm clipping at a value of 1.0. See Tables 5, 6 for comprehensive details about model architecture.

The 20M *Panda MLM* models are trained for 200K iterations (~ 52 wallclock hours across 4 GPUs or ~ 208 GPU hours) with a batch size of 1024 with 50% of tokens randomly masked out each batch.

The forecast models are trained for 100K iterations with a batch size of 1024 and are optionally initialized with a pretrained encoder from an MLM model (Section 5.1). The 20M parameter forecasting checkpoints are trained for ~ 26 wallclock hours or ~ 104 GPU hours. The Chronos-SFT models use considerably more memory during training - permitting a batch size of 160 for 300K training iterations which required ~ 48 wallclock hours or ~ 192 GPU hours.

1132
1133

1134 **C FORECASTS**
1135
1136
1137

1164 Figure 14: Examples of zero-shot forecasts ($L_{\text{pred}} = 256$) on held-out chaotic dynamical systems.
1165
1166
1167

1168 For additional forecasts, see Appendix L (Figs. 28, 29, and 30). All forecasts plotted are with
1169 prediction length $L_{\text{pred}} = 256$.
1170
1171
1172

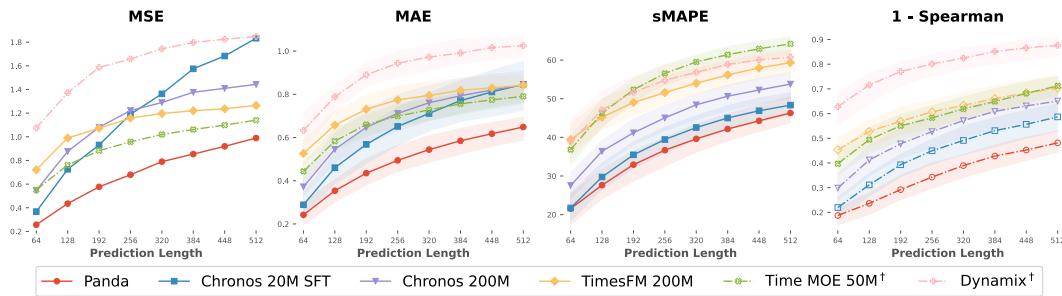
1173 **D ADDITIONAL FORECAST METRICS**
1174
1175

1176 Table 7 shows statistical significance testing of *Panda* metrics against other baselines. Note that we
1177 do not report results for TimeMoE due to the presence of NaNs, and instead test against a 200M
1178 Chronos baseline in greedy decoding and probabilistic mode. *Panda* clearly achieves lower error
1179 across the board; the gap closes with Chronos 20M SFT but still remains statistically significant.
1180
1181
1182
1183
1184
1185

1186 In Fig. 2 we presented the sMAPE and MAE comparison for *Panda* versus our baseline models. We
1187 now present more zero-shot forecast metrics, but using the probabilistic forecasting mode for Chronos
1188 and Chronos-SFT. When finetuning *Chronos 20M* on our dataset (i.e. *Chronos 20M SFT*), we used
1189 the default top-k and top-p (nucleus sampling) and temperature settings. We use these same settings,
1190 $\text{top-}k = 50$, $\text{top-}p = 1.0$, $\text{temperature} = 1.0$ for the Chronos probabilistic forecasting, aggregating
1191 our metrics over 10 sample forecasts per context window per system.
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2287
2288
2289
2289
2290
2291
2292
2293
2294
2295
2296
2297
2297
2298
2299
2299
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2328
2329
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2338
2339
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2348
2349
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2358
2359
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2368
2369
2369
2370
2371
2372
2373
2374
2375
2376
2377
2377
2378
2378
2379
2379
2380
2381
2382
2383
2384
2385
2386
2386
2387
2387
2388
2388
2389
2389
2390
2391
2392
2393
2394
2395
2396
2396
2397
2397
2398
2398
2399
2399
2400
2401
2402
2403
2404
2405
2406
2406
2407
2407
2408
2408
2409
2409
2410
2411
2412
2413
2414
2415
2416
2416
2417
2417
2418
2418
2419
2419
2420
2421
2422
2423
2424
2425
2426
2426
2427
2427
2428
2428
2429
2429
2430
2431
2432
2433
2434
2435
2436
2436
2437
2437
2438
2438
2439
2439
2440
2441
2442
2443
2444
2445
2446
2446
2447
2447
2448
2448
2449
2449
2450
2451
2452
2453
2454
2455
2456
2456
2457
2457
2458
2458
2459
2459
2460
2461
2462
2463
2464
2465
2466
2466
2467
2467
2468
2468
2469
2469
2470
2471
2472
2473
2474
2475
2475
2476
2476
2477
2477
2478
2478
2479
2479
2480
2481
2482
2483
2484
2485
2485
2486
2486
2487
2487
2488
2488
2489
2489
2490
2491
2492
2493
2494
2495
2495
2496
2496
2497
2497
2498
2498
2499
2499
2500
2501
2502
2503
2504
2505
2505
2506
2506
2507
2507
2508
2508
2509
2509
2510
2511
2512
2513
2514
2515
2515
2516
2516
2517
2517
2518
2518
2519
2519
2520
2521
2522
2523
2524
2525
2525
2526
2526
2527
2527
2528
2528
2529
2529
2530
2531
2532
2533
2534
2535
2535
2536
2536
2537
2537
2538
2538
2539
2539
2540
2541
2542
2543
2544
2545
2545
2546
2546
2547
2547
2548
2548
2549
2549
2550
2551
2552
2553
2554
2555
2555
2556
2556
2557
2557
2558
2558
2559
2559
2560
2561
2562
2563
2564
2565
2565
2566
2566
2567
2567
2568
2568
2569
2569
2570
2571
2572
2573
2574
2575
2575
2576
2576
2577
2577
2578
2578
2579
2579
2580
2581
2582
2583
2584
2585
2585
2586
2586
2587
2587
2588
2588
2589
2589
2590
2591
2592
2593
2594
2595
2595
2596
2596
2597
2597
2598
2598
2599
2599
2600
2601
2602
2603
2604
2605
2605
2606
2606
2607
2607
2608
2608
2609
2609
2610
2611
2612
2613
2614
2615
2615
2616
2616
2617
2617
2618
2618
2619
2619
2620
2621
2622
2623
2624
2625
2625
2626
2626
2627
2627
2628
2628
2629
2629
2630
2631
2632
2633
2634
2635
2635
2636
2636
2637
2637
2638
2638
2639
2639
2640
2641
2642
2643
2644
2645
2645
2646
2646
2647
2647
2648
2648
2649
2649
2650
2651
2652
2653
2654
2655
2655
2656
2656
2657
2657
2658
2658
2659
2659
2660
2661
2662
2663
2664
2665
2665
2666
2666
2667
2667
2668
2668
2669
2669
2670
2671
2672
2673
2674
2675
2675
2676
2676
2677
2677
2678
2678
2679
2679
2680
2681
2682
2683
2684
2685
2685
2686
2686
2687
2687
2688
2688
2689
2689
2690
2691
2692
2693
2694
2695
2695
2696
2696
2697
2697
2698
2698
2699
2699
2700
2701
2702
2703
2704
2705
2705
2706
2706
2707
2707
2708
2708
2709
2709
2710
2711
2712
2713
2714
2715
2715
2716
2716
2717
2717
2718
2718
2719
2719
2720
2721
2722
2723
2724
2725
2725
2726
2726
2727
2727
2728
2728
2729
2729
2730
2731
2732
2733
2734
2735
2735
2736
2736
2737
2737
2738
2738
2739
2739
2740
2741
2742
2743
2744
2745
2745
2746
2746
2747
2747
2748
2748
2749
2749
2750
2751
2752
2753
2754
2755
2755
2756
2756
2757
2757
2758
2758
2759
2759
2760
2761
2762
2763
2764
2765
2765
2766
2766
2767
2767
2768
2768
2769
2769
2770
2771
2772
2773
2774
2775
2775
2776
2776
2777
2777
2778
2778
2779
2779
2780
2781
2782
2783
2784
2785
2785
2786
2786
2787
2787
2788
2788
2789
2789
2790
2791
2792
2793
2794
2795
2795
2796
2796
2797
2797
2798
2798
2799
2799
2800
2801
2802
2803
2804
2805
2805
2806
2806
2807
2807
2808
2808
2809
2809
2810
2811
2812
2813
2814
2815
2815
2816
2816
2817
2817
2818
2818
2819
2819
2820
2821
2822
2823
2824
2825
2825
2826
2826
2827
2827
2828
2828
2829
2829
2830
2831
2832
2833
2834
2835
2835
2836
2836
2837
2837
2838
2838
2839
2839
2840
2841
2842
2843
2844
2845
2845
2846
2846
2847
2847
2848
2848
2849
2849
2850
2851
2852
2853
2854
2855
2855
2856
2856
2857
2857
2858
2858
2859
2859
2860
2861
2862
2863
2864
2865
2865
2866
2866
2867
2867
2868
2868
2869
2869
2870
2871
2872
2873
2874
2875
2875
2876
2876
2877
2877
2878
2878
2879
2879
2880
2881
2882
2883
2884
2885
2885
2886
2886
2887
2887
2888
2888
2889
2889
2890
2891
2892
2893
2894
2895
2895
2896
2896
2897
2897
2898
2898
2899
2899
2900
2901
2902
2903
2904
2905
2905
2906
2906
2907
2907
2908
2908
2909
2909
2910
2911
2912
2913
2914
2915
2915
2916
2916
2917
2917
2918
2918
2919
2919
2920
2921
2922
2923
2924
2925
2925
2926
2926
2927
2927
2928
2928
2929
2929
2930
2931
2932
2933
2934
2935
2935
2936
2936
2937
2937
2938
2938
2939
2939
2940
2941
2942
2943
2944
2945
2945
2946
2946
2947
2947
2948
2948
2949
2949
2950
2951
2952
2953
2954
2955
2955
2956
2956
2957
2957
2958
2958
2959
2959
2960
2961
2962
2963
2964
2965
2965
2966
2966
2967
2967
2968
2968
2969
2969
2970
2971
2972
2973
2974
2975
2975
2

1188 Table 7: Wilcoxon Signed Ranked test for Panda errors vs. baseline errors (Holm–Sidák adjusted
 1189 *p*-values)

| Model | Prediction Horizon | MSE | | MAE | | sMAPE | |
|----------------------|--------------------|------------------------|-----------|------------------------|-----------|------------------------|-----------|
| | | <i>p</i> -value | statistic | <i>p</i> -value | statistic | <i>p</i> -value | statistic |
| Chronos 20M SFT | $L = 128$ | 5.96×10^{-48} | 13 484 | 9.01×10^{-51} | 12 153 | 6.35×10^{-3} | 49 421 |
| | $L = 512$ | 8.24×10^{-50} | 12 275 | 3.07×10^{-54} | 10 432 | 3.34×10^{-2} | 51 252 |
| Chronos 200M | $L = 128$ | 5.84×10^{-49} | 12 778 | 6.88×10^{-57} | 9182 | 1.20×10^{-28} | 23 770 |
| | $L = 512$ | 5.33×10^{-48} | 13 233 | 1.45×10^{-56} | 9323 | 1.36×10^{-7} | 41 310 |
| Chronos 20M SFT Prob | $L = 128$ | 6.93×10^{-41} | 16 851 | 3.09×10^{-33} | 21 198 | 1.30×10^{-5} | 44 464 |
| | $L = 512$ | 3.83×10^{-41} | 16 513 | 1.18×10^{-35} | 19 657 | 3.26×10^{-1} | 54 730 |
| Chronos 200M Prob | $L = 128$ | 2.91×10^{-49} | 12 577 | 1.22×10^{-56} | 9290 | 5.69×10^{-31} | 22 353 |
| | $L = 512$ | 2.37×10^{-38} | 18 024 | 1.83×10^{-49} | 12 541 | 7.73×10^{-11} | 37 624 |
| TimesFM 200M | $L = 128$ | 4.05×10^{-55} | 9863 | 1.67×10^{-69} | 3863 | 4.89×10^{-77} | 972 |
| | $L = 512$ | 2.50×10^{-38} | 18 103 | 1.59×10^{-59} | 7949 | 2.66×10^{-70} | 3588 |
| Time MOE 50M | $L = 128$ | 1.03×10^{-41} | 14 499 | 9.66×10^{-58} | 7320 | 1.68×10^{-76} | 170 |
| | $L = 512$ | 2.58×10^{-29} | 21 152 | 3.14×10^{-52} | 9639 | 3.13×10^{-75} | 589 |
| DynaMix | $L = 128$ | 1.46×10^{-23} | 26 724 | 3.03×10^{-35} | 19 399 | 1.86×10^{-40} | 16 533 |
| | $L = 512$ | 4.54×10^{-8} | 39 970 | 7.00×10^{-22} | 27 602 | 1.83×10^{-27} | 23 592 |



1218 Figure 15: Zero-shot forecast metrics for our baselines, using *probabilistic* (10 samples) forecasts for
 1219 the *Chronos* models. [†]Dash-dotted lines indicate presence of NaNs for some systems (4 - 12% of
 1220 systems for Spearman).

| Time per Forecast (s) | |
|-----------------------|-------------------------------------|
| Model | Time (mean \pm std) |
| Panda | 0.031 ± 0.001 |
| TimeMOE 50M | 0.336 ± 0.060 |
| DynaMix | 0.526 ± 0.016 |
| TimesFM 200M | 0.605 ± 0.032 |
| Chronos 20M | 1.880 ± 0.041 |
| Chronos 200M | 4.233 ± 0.121 |

1230 Table 8: Inference time per forecast ($L_{pred} = 512$), computed over $N = 1000$ calls to each model,
 1231 on a single H100 GPU. Each model call uses context length 512 timesteps, from our multivariate
 1232 data, which has variable number of channels (at least 3). The univariate models (Chronos, TimesFM,
 1233 TimeMOE) treat the channels as batch dimension, for each call.

1234
 1235 We also provide median forecast metrics with IQR for the metrics in Fig. 2 over multiple prediction
 1236 horizons for the best baselines in Tables 9, 10.

| sMAPE Median [P25, P75] | | | |
|-------------------------------|--------------------------|--------------------------|--------------------------|
| Model | $L_{pred} = 128$ | $L_{pred} = 256$ | $L_{pred} = 512$ |
| Panda | 27.6 [18.5, 39.3] | 36.7 [26.2, 47.6] | 46.3 [37.0, 57.0] |
| Chronos 20M SFT | 30.3 [21.9, 40.0] | 40.1 [30.3, 48.3] | 48.8 [37.3, 56.8] |
| Chronos 200M | 36.0 [26.4, 44.6] | 44.6 [34.6, 52.9] | 53.5 [42.8, 60.8] |
| Chronos 20M SFT Probabilistic | 29.7 [21.3, 40.7] | 39.4 [29.3, 48.7] | 48.3 [37.8, 57.1] |
| Chronos 200M Probabilistic | 36.4 [26.7, 44.7] | 45.0 [34.2, 53.1] | 53.8 [42.8, 60.6] |
| DynaMix | 47.1 [37.9, 57.9] | 54.7 [46.8, 62.6] | 60.8 [53.9, 65.8] |

Table 9: Median sMAPE and interquartile range [P25, P75].

| MAE Median [P25, P75] | | | |
|-------------------------------|--------------------------|--------------------------|--------------------------|
| Model | $L_{pred} = 128$ | $L_{pred} = 256$ | $L_{pred} = 512$ |
| Panda | 0.35 [0.22, 0.54] | 0.49 [0.35, 0.70] | 0.65 [0.48, 0.84] |
| Chronos 20M SFT | 0.56 [0.34, 0.80] | 0.85 [0.52, 1.26] | 1.25 [0.72, 2.14] |
| Chronos 200M | 0.61 [0.41, 0.80] | 0.86 [0.58, 1.18] | 1.07 [0.75, 1.83] |
| Chronos 20M SFT Probabilistic | 0.46 [0.29, 0.69] | 0.65 [0.43, 0.93] | 0.85 [0.57, 1.34] |
| Chronos 200M Probabilistic | 0.55 [0.38, 0.73] | 0.71 [0.51, 0.91] | 0.84 [0.63, 1.12] |
| DynaMix | 0.79 [0.60, 0.94] | 0.94 [0.76, 1.06] | 1.02 [0.88, 1.15] |

Table 10: Median MAE and interquartile range [P25, P75].

D.0.1 ADDITIONAL METRICS FOR MODEL ABLATIONS

In Fig. 3 we presented a sMAPE comparison for several key ablations of our model. Here, we provide additional zero-shot forecast metrics for these ablations, supporting our conclusion that our dynamics embedding with polynomial features (PolyEmbed) is best for long-horizon forecasting via rollouts.

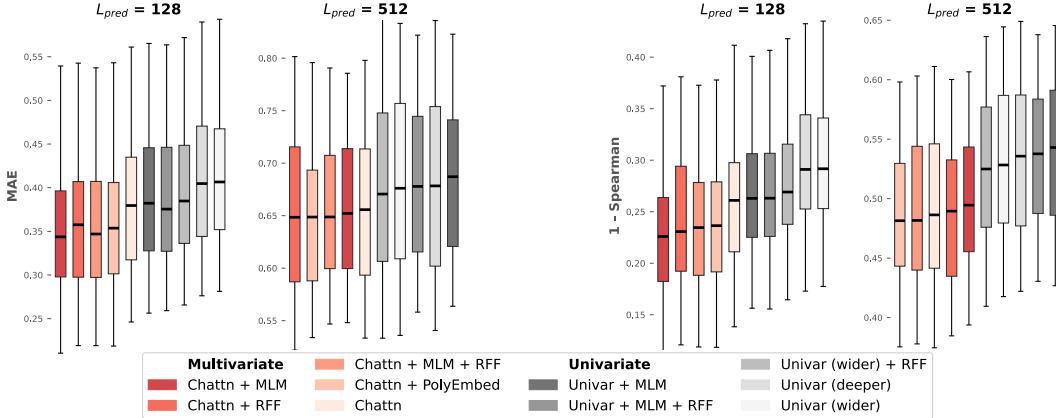
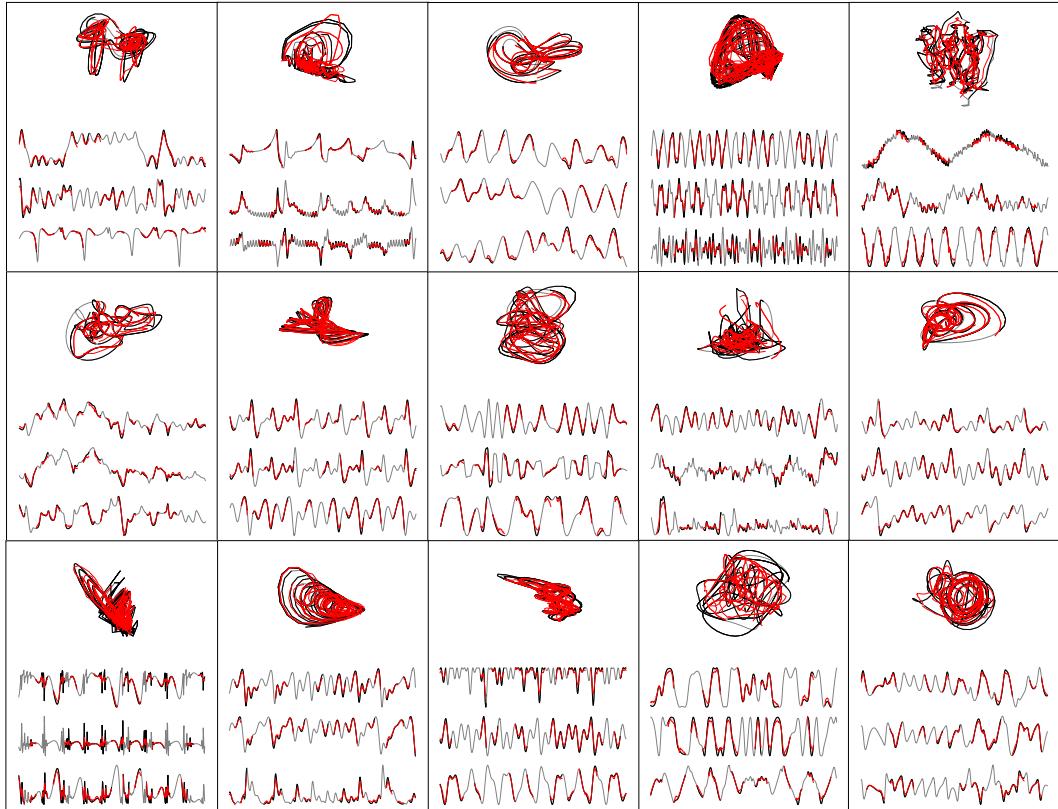


Figure 16: Zero-shot forecast metrics for our ablation experiments.

1296 **E MLM COMPLETIONS**
1297

1325 Figure 17: Examples of zero-shot completions on held-out chaotic dynamical systems. Each completion plotted was with a context length of 512 time points, with half the patches (patch length 16) 1326 randomly masked out in a channel-inconsistent manner. These plots show *Panda MLM*, our 20M 1327 parameter checkpoint, completing the masked-out trajectories i.e. 256 time points. See Appendix N 1328 for more examples from *Panda MLM* and from our scaled-up model *Panda MLM-66M*. 1329

1330 We present examples of *Panda MLM* completions on our held-out test set in Fig. 17. For more 1331 examples of zero-shot completions, see Appendix N. 1332

1333 For the completions task, we randomly mask out half of patches for each coordinate dimension 1334 separately i.e. channel-independent masking. We trained *Panda MLM* with patch length 16 and 1335 context length 512, so each context window has 32 patches on the time axis. But we can generate 1336 completions with any context length. We refer to the masked-out portions of the trajectory as the 1337 erasures. We seek to measure how the model learns the cross-channel coupling relationships and 1338 statistical dependencies. 1339

1340 In future work, we hope to investigate more sophisticated masking strategies, such as masking out 1341 contiguous blocks of patches and investigating channel-dependent masking, which is closer to a 1342 forecasting task. Recall from our discussion of Fig. 3 that MLM pretraining reduces performance on 1343 autoregressive rollout (c.f. Fig. 16). Determining the optimal MLM pretraining objective for long 1344 horizon forecasts on autoregressive rollout remains an intriguing area to investigate. 1345

1346 To quantify the performance of our MLM checkpoint on the completion task, we compute the 1347 correlation dimension (Fig. 18) of completions versus ground truth trajectories using the Grassberger- 1348 Procaccia algorithm (Grassberger & Procaccia, 1983a;b). This algorithm was developed to quantify 1349 the strangeness (Lorenz, 1963; Ott, 1981; Ruelle & Takens, 1971) of chaotic attractors via a 1350 computable metric related to the fractal (Hausdorff) dimension and information entropy. For the result in 1351 Fig. 18, we take the entire length 4096 trajectory for each of our 9.3×10^3 held-out systems and we 1352 randomly mask out (erase) half of the patches (patch length 16) in a channel-inconsistent manner. 1353

1350 Let $\{x_i\}_{i=1}^T \subset \mathbb{R}^D$ be a time series of T points in D dimensions. First, we compute pairwise Euclidean
 1351 distances (excluding $i = j$):
 1352

$$1353 \quad \mathcal{R} := \{r_{ij} = \|x_i - x_j\|_2 \mid 1 \leq i, j \leq T, i \neq j\}$$

1355 Next, we select the scaling region. Let $r_{(5\%)}$ and
 1356 $r_{(50\%)}$ denote the empirical 5th and 50th percentiles
 1357 of \mathcal{R} . Then truncate to:
 1358

$$1359 \quad \mathcal{R}^* = \{r \in \mathcal{R} \mid r_{(5\%)} < r < r_{(50\%)}\}$$

1360 Now denote $n := |\mathcal{R}^*|$ and $r_{\min} := \min_{r \in \mathcal{R}^*} r$.
 1361

1363 Following Clauset, Shalizi, and Newman (Clauzet
 1364 et al., 2009), we identify a power law fit using
 1365 maximum likelihood estimation (MLE). Assume for
 1366 $r \geq r_{\min}$ that the distances follow $p(r) = Z r^{-\alpha}$,
 1367 where Z is the normalizing constant. Then,

$$1368 \quad \hat{\alpha} = 1 + \frac{n}{\sum_{r \in \mathcal{R}^*} \ln\left(\frac{r}{r_{\min}}\right)}$$

1372 In the Grassberger–Procaccia method one examines
 1373 a correlation integral with unbiased estimator:

$$1374 \quad \mathcal{C}(r) = \frac{2}{T(T-1)} \sum_{i < j} H(r - \|x_i - x_j\|_2), \quad \mathcal{C}(r) \sim r^{D_2} \quad (r \rightarrow 0),$$

1377 so that $D_2 = \frac{d \ln \mathcal{C}(r)}{d \ln r}$. Fitting $\mathcal{C}(r) \propto r^{D_2}$ is equivalent to fitting the distribution of pairwise distances
 1378 to a power law, yielding $D_2 \approx \hat{\alpha}$ as the estimated correlation dimension.
 1379

1380 *Panda MLM*, with $20M$ parameters, shows promise in recovering the correlation dimension, a
 1381 statistical invariant of the attractor, even when given much longer context (length 4096) than seen
 1382 during training (recall the context length for training was 512), and with half of the timesteps masked
 1383 out (in patches) per dimension. We also trained a scaled-up checkpoint, *Panda MLM-66M* (with
 1384 details in Appendix K), which demonstrates improved performance (Fig. 18).

1385 In 11, we present the comparison against interpolation baselines. For polynomial inter-
 1386 polation, we use `numpy.polyfit` to fit a polynomial to the unmasked timesteps,
 1387 and evaluate with `numpy.polyval` at masked positions. For linear interpolation,
 1388 we use `scipy.interpolate.interp1d`, with extrapolation for timesteps out-
 1389 side the range of known values. For the piecewise cubic spline baseline, we use
 1390 `scipy.interpolate.make_interp_spline` with $k = 3$ (cubic spline).

1391 Comparison with Baselines for Completions Task

| 1392 Method | $L_{\text{context}} = 4096$, with 50% Erasure (in patches) |
|---------------------------------------|---|
| 1393 Panda MLM-66M | 0.91 |
| 1394 Panda MLM | 0.78 |
| 1395 Piecewise Cubic Spline | 0.71 |
| 1396 Linear Interpolation | 0.61 |
| 1397 Polynomial Interpolation (Deg 3) | 0.21 |

1399 Table 11: Coefficient of Determination (R^2) between the correlation dimension (via Grassberger-
 1400 Procaccia) computed on the completions versus the full length 4096 ground truth trajectories. For
 1401 each of our 9347 held-out test systems, we average across 8 random seeds, which determine the
 1402 timestep masks for each trajectory (in patches of length 16). See Appendix N for further discussion.
 1403

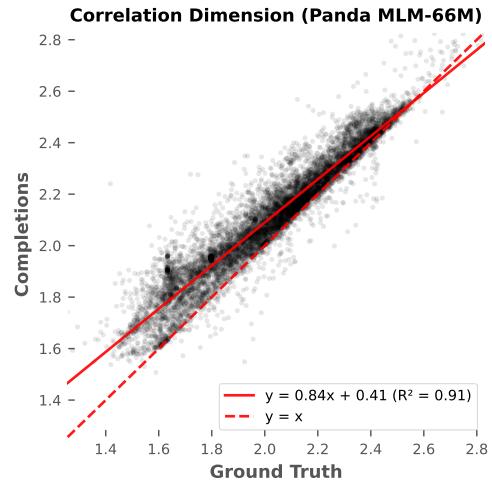


Figure 18: Correlation dimension comparison on held-out systems. Computed for ground truth and completions from *Panda MLM-66M* using the Grassberger-Procaccia method, averaged over 8 independent masks for each trajectory, using context length 4096 with half the patches (patch length 16) randomly masked out in a channel-inconsistent manner.

1404
1405

F MORE DISTRIBUTIONAL METRICS AND INVARIANT QUANTITIES

1406
1407

F.1 IMPLEMENTATION DETAILS: KL DIVERGENCE VIA GMMs

1408
1409
1410
1411
1412

Algorithm 1 presents our implementation of the Kullback Leibler divergence between ground truth and model predictions. This is the implementation we use for our main evaluations (Tables 1 and 12), although in Subsection F.3 we also present results using an alternative implementation found in the literature. In particular, we construct GMMs by fitting Gaussians to each point, with local scale parameter determined by the simplex neighbors algorithm.

1413

Algorithm 1 KL Divergence Estimation via Gaussian Mixture Models1414
1415
1416

Require: Ground Truth $\mathbf{X} = \{\mathbf{x}_t\}_{t=1}^T$, Generated Predictions $\mathbf{Y} = \{\mathbf{y}_t\}_{t=1}^T$, Number of Monte Carlo Samples n_s , small $\varepsilon > 0$ and $\text{tol} > 0$

1417

```

1: Function ESTIMATEKLDIVERGENCE( $\mathbf{X}, \mathbf{Y}, n_s, \varepsilon$ )
2:   // Step 1: Local bandwidth (scale) estimation
3:    $\sigma^X \leftarrow \text{SIMPLEXNEIGHBORS}(\mathbf{X}, k = 10)$ 
4:    $\sigma^Y \leftarrow \text{SIMPLEXNEIGHBORS}(\mathbf{Y}, k = 10)$ 
5:   // Step 2: Construct Gaussian Mixture Models
6:    $p \leftarrow \text{GAUSSIANMIXTURE}(\text{means} = \mathbf{X}, \text{covariances} = \sigma^X)$ 
7:    $q \leftarrow \text{GAUSSIANMIXTURE}(\text{means} = \mathbf{Y}, \text{covariances} = \sigma^Y)$ 
8:   // Step 3: Monte Carlo KLD Estimate
9:    $\{\mathbf{z}_i\}_{i=1}^{n_s} \leftarrow p.\text{SAMPLE}(n_s)$ 
10:  for  $i = 1$  to  $n_s$  do
11:     $p_i \leftarrow p(\mathbf{z}_i)$ 
12:     $q_i \leftarrow q(\mathbf{z}_i)$ 
13:     $q_i \leftarrow \max(q_i, \varepsilon)$ 
14:     $r_i \leftarrow \log(p_i/q_i)$ 
15:  end for
16:  return  $\widehat{\text{KLD}} = \frac{1}{n_s} \sum_{i=1}^{n_s} r_i$ 

17: Function SIMPLEXNEIGHBORS( $\mathbf{Z}, k$ )
18: Let  $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^n$  with  $\mathbf{z}_i \in \mathbb{R}^d$ 
19: Build a  $(k+1)$ -nearest neighbor search structure on  $\mathbf{Z}$  (e.g., using Euclidean distance)
20: for  $i = 1$  to  $n$  do
21:   Query  $(k+1)$  nearest neighbors of  $\mathbf{z}_i$ , including itself
22:   Discard the self-neighbor to obtain neighbors  $\{\mathbf{z}_{i,j}\}_{j=1}^k$  with distances  $d_{i,j}$ 
23:    $\sigma_i \leftarrow \text{FINDSIGMA}((d_{i,1}, \dots, d_{i,k}), k)$            ▷ Estimate local scale parameter
24:   Let  $\rho_i \leftarrow \min_j d_{i,j}$ 
25: end for
26: return  $\sigma = (\sigma_i)_{i=1}^n$ 

27: Function FINDSIGMA( $\mathbf{d}, k$ )           ▷  $\mathbf{d} = (d_1, \dots, d_k)$  are distances to  $k$  nearest neighbors
28:  $\rho \leftarrow \min_j d_j$ 
29: Define  $\Delta_j \leftarrow \max(d_j - \rho, 0)$  for  $j = 1, \dots, k$            ▷ ReLU on shifted distances
30: Define objective
31: 
$$\phi(\sigma) = \left( \sum_{j=1}^k \exp(-\Delta_j/(\sigma + \text{tol})) - \log_2(k) \right)^2, \quad \sigma > 0$$

32: Minimize  $\phi(\sigma)$  using 1D optimization (e.g., root-finding in  $\log \sigma$  with initial guess  $\rho$ ).
33: Let  $\sigma^*$  be the resulting positive solution
34: return  $\sigma^*$ 

```

1455
1456
1457

1458 F.2 PER-SYSTEM DIFFERENCES IN DISTRIBUTIONAL METRICS
1459

1460 We report mean \pm std. dev. of per-system differences across all test systems, averaged over 5 context
1461 windows for prediction horizons $L_{\text{pred}} \in \{512, 1024, 2048, 3072\}$.
1462

1463 Per-system Difference in $D_{KL}(\text{Ground Truth}(L_{\text{pred}}) \parallel \text{Model Prediction}(L_{\text{pred}}))$ between Baselines
1464

| Comparison | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
|--------------------------------|-----------------------------------|-----------------------------------|------------------------------------|------------------------------------|------------------------------------|
| Chronos 20M SFT – Panda | 0.79 \pm 4.60 | 0.36 \pm 5.03 | -0.01 \pm 5.14 | -0.22 \pm 5.09 | -0.33 \pm 5.84 |
| DynaMix – Panda | 0.81 \pm 5.12 | 0.18 \pm 5.58 | -0.42 \pm 5.59 | -0.75 \pm 5.52 | -0.88 \pm 6.07 |
| Chronos 20M – Chronos 20M SFT | 1.27 \pm 4.64 | 1.09 \pm 5.39 | 0.89 \pm 5.50 | 0.84 \pm 5.53 | 0.89 \pm 6.62 |
| Chronos 200M – Chronos 20M SFT | 0.39 \pm 4.35 | 0.41 \pm 5.36 | 0.43 \pm 5.61 | 0.44 \pm 5.73 | 0.42 \pm 6.59 |

1469 Table 12: The (mean \pm std) of per-system diff. in KL divergence between models, a fine-grained view
1470 of Table 1. *DynaMix* and *Chronos 20M SFT* outperform *Panda* on very long prediction horizons.
1471

1472

1473 Per-System Difference in Average $H^2(S_{\text{Ground Truth}(L_{\text{pred}})} \parallel S_{\text{Model Prediction}(L_{\text{pred}})})$ between Baselines
1474

| Comparison | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
|--------------------------------|-------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Chronos 20M SFT – Panda | 0.04 \pm 0.19 | 0.04 \pm 0.18 | 0.04 \pm 0.18 | 0.04 \pm 0.18 | 0.04 \pm 0.20 |
| DynaMix – Panda | 0.11 \pm 0.22 | 0.09 \pm 0.22 | 0.08 \pm 0.21 | 0.07 \pm 0.21 | 0.07 \pm 0.24 |
| Chronos 20M – Chronos 20M SFT | 0.08 \pm 0.20 | 0.08 \pm 0.20 | 0.08 \pm 0.20 | 0.08 \pm 0.20 | 0.08 \pm 0.23 |
| Chronos 200M – Chronos 20M SFT | 0.00 \pm 0.20 | 0.00 \pm 0.20 | 0.00 \pm 0.20 | 0.00 \pm 0.20 | 0.00 \pm 0.23 |

1480 Table 13: The (mean \pm std) of per-system differences in average spectral Hellinger distance between
1481 models, a fine-grained view of Table 2 showing that *Panda* outperforms the baselines.
1482

1483

1484 F.3 AN ALTERNATIVE KL DIVERGENCE IMPLEMENTATION (GEOMETRIC MISALIGNMENT)
1485

1486 In addition to our GMM-based KL divergence implementation (Tables 1 and 12), we also use the
1487 implementation of (Hemmer & Durstewitz, 2025) based on geometric misalignment.
1488

1489 Per-System Difference in $D_{KL}(\text{Ground Truth}(L_{\text{pred}}) \parallel \text{Model Prediction}(L_{\text{pred}}))$ via Geometric Misalignment
1490

| Model | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
|------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Panda | 2.82 \pm 2.67 | 3.29 \pm 2.79 | 3.88 \pm 2.85 | 4.26 \pm 2.88 | 4.44 \pm 3.14 |
| Chronos 20M SFT | 2.52 \pm 2.63 | 2.81 \pm 2.94 | 3.09 \pm 3.16 | 3.25 \pm 3.28 | 3.34 \pm 3.72 |
| Chronos 20M | 4.33 \pm 3.20 | 4.67 \pm 3.53 | 5.03 \pm 3.76 | 5.24 \pm 3.88 | 5.37 \pm 4.40 |
| Chronos 200M | 2.96 \pm 2.86 | 3.19 \pm 3.15 | 3.47 \pm 3.37 | 3.64 \pm 3.49 | 3.73 \pm 4.01 |
| DynaMix | 3.06 \pm 4.07 | 3.15 \pm 4.42 | 3.24 \pm 4.68 | 3.30 \pm 4.81 | 3.37 \pm 5.40 |
| $\Delta\% (\uparrow)$ | -11.9% | -17.1% | -25.6% | -31.1% | -32.9% |

1498 Table 14: KL divergence between the ground truth and model predictions. $\Delta\%$ denotes percentage
1499 gain of *Panda* over the best baseline. See Table 15 for per-system differences. Here, we use the
1500 implementation of (Hemmer & Durstewitz, 2025).
1501

1502

1503 Per-system Difference in $D_{KL}(\text{Ground Truth}(L_{\text{pred}}) \parallel \text{Model Prediction}(L_{\text{pred}}))$ between Baselines
1504

| Comparison | $L_{\text{pred}} = 512$ | $L_{\text{pred}} = 1024$ | $L_{\text{pred}} = 2048$ | $L_{\text{pred}} = 3072$ | $L_{\text{pred}} = 3584$ |
|--------------------------------|-------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Chronos 20M SFT – Panda | -0.30 \pm 3.50 | -0.49 \pm 3.75 | -0.79 \pm 3.95 | -1.00 \pm 4.01 | -1.09 \pm 4.54 |
| DynaMix – Panda | 0.23 \pm 4.71 | -0.14 \pm 5.06 | -0.64 \pm 5.25 | -0.96 \pm 5.38 | -1.07 \pm 6.01 |
| Chronos 20M – Chronos 20M SFT | 1.82 \pm 3.60 | 1.86 \pm 3.93 | 1.94 \pm 4.25 | 1.99 \pm 4.33 | 2.03 \pm 5.14 |
| Chronos 200M – Chronos 20M SFT | 0.44 \pm 3.33 | 0.39 \pm 3.66 | 0.38 \pm 4.01 | 0.39 \pm 4.10 | 0.38 \pm 4.88 |

1509 Table 15: The (mean \pm std) of per-system diff. in KL divergence between models, a fine-grained view
1510 of Table 14. *DynaMix* and *Chronos 20M SFT* outperform *Panda* on very long prediction horizons.
1511

1512
1513

F.4 VISUALIZATION OF METRICS DISTRIBUTION ACROSS ALL TEST SYSTEMS

1514
1515
1516
1517

As seen in Figure 19, *Dynamix* has a failure mode for a large group of systems for which it performs badly in KL. However, we emphasize that a direct comparison is not fair because *Dynamix* is much smaller, trained on fewer systems, cannot handle systems with > 3 dimensions, and likely has train data leakage on our test set.

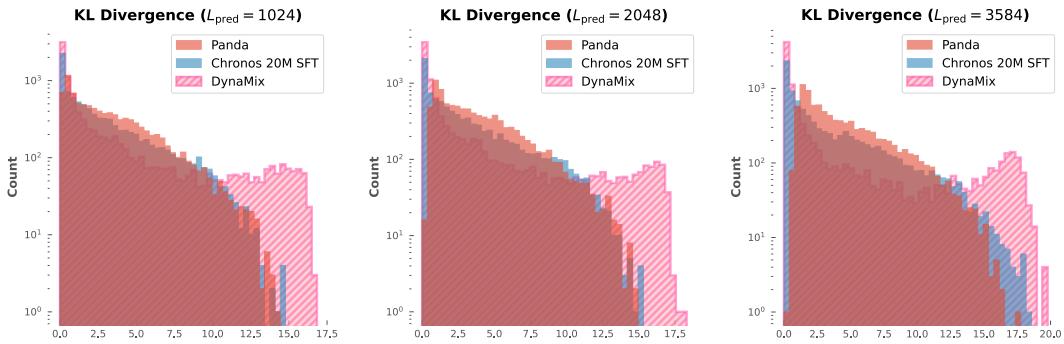
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
15291530
1531
1532

Figure 19: KL divergence (via geometric misalignment) between ground truth (L_{pred}) and model predictions (L_{pred}). See Table 14 for aggregate values. Note that $L_{\text{pred}} = 3584$ is 28x the prediction length used in training *Panda*.

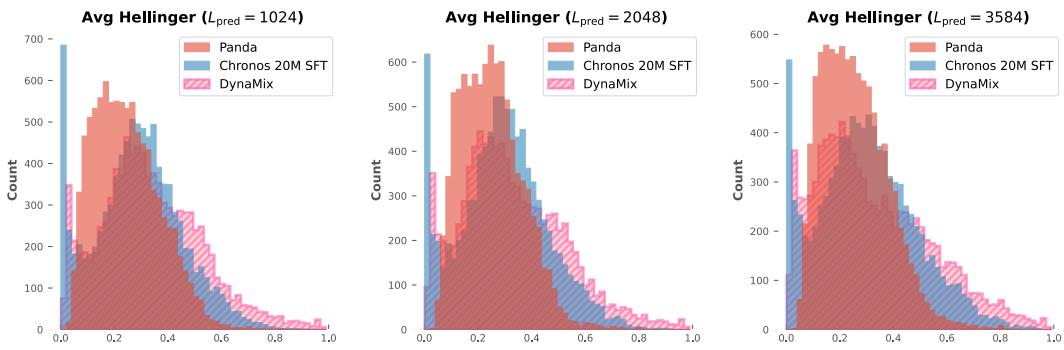
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
15441545
1546
1547

Figure 20: Average spectral Hellinger distance, between the power spectra of ground truth (L_{pred}) and model predictions (L_{pred}). See Table 2 for aggregate values.

1548
1549
1550
1551
1552
1553
1554
1555

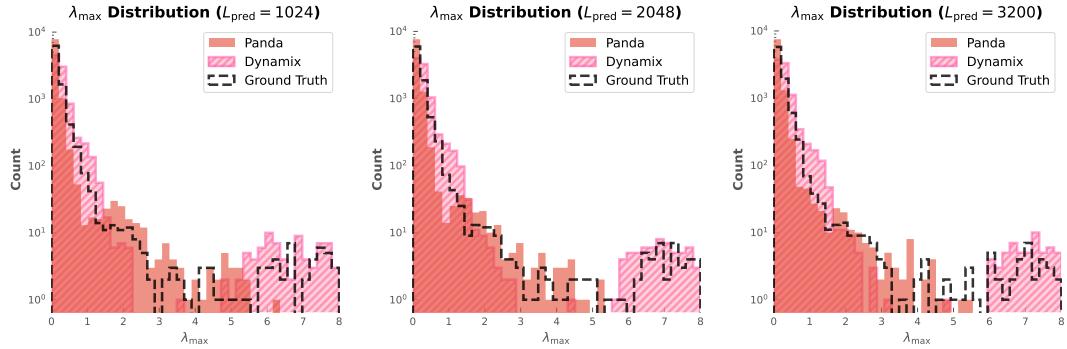
F.5 QUANTIFYING MEAN REGRESSION ON VERY LONG HORIZONS

Mean regression is a common failure mode for TSFMs on very long prediction horizons. To quantify this failure mode, we compute the distributional metrics at $L_{\text{pred}} = 3584$, which is the longest possible horizon for evaluation, since our dataset contains trajectories of 4096 timepoints. However, we *cut off* the first $N_{\text{cutoff}} = 1536$ timepoints of model predictions and ground truth, to compute the metrics on the last 2048 timepoints - solidly within the mean regression regime.

1556
1557
1558
1559
1560
1561
1562
1563
1564
1565

| Metrics on $t_{\text{pred}} = [1536, 3584]$ (Cut Off First 1536 Timepoints) | | |
|---|--|-----------------------------|
| Model | KL Divergence (Geometric Misalignment) | Spectral Hellinger Distance |
| Panda | 15.25 ± 2.46 | 0.49 ± 0.11 |
| Chronos 20M SFT | 7.00 ± 5.63 | 0.34 ± 0.18 |
| Chronos 20M | 9.68 ± 6.12 | 0.48 ± 0.18 |
| Chronos 200M | 7.37 ± 6.05 | 0.36 ± 0.19 |
| DynaMix | 3.50 ± 5.44 | 0.35 ± 0.22 |

Table 16: Metrics between ground truth and model predictions *after cutting off* the first 1536 timepoints of $L_{\text{pred}} = 3584$ (keeping only the last 2048). We present (mean \pm std) across test systems.

1566 F.6 MAXIMUM LYAPUNOV EXPONENT
1567
1568
1569

1581 Figure 21: Distributional comparison of the maximum Lyapunov exponents estimated from the
1582 ground truth (L_{pred}) versus estimates from the model predictions (L_{pred}) of *Panda* and of *DynaMix*.
1583 Note that the y-axis is on a log scale.

1584
1585
1586 We compute the maximum Lyapunov exponents for long prediction horizons, using the data-driven
1587 Rosenstein estimator. In Fig. 21, we compare the distribution of estimated (λ_{max}) for *Panda* versus
1588 that of *DynaMix* at prediction lengths $L_{\text{pred}} = 1024, 2048, 3200$. Note that *DynaMix* was trained
1589 pointwise autoregressively for a prediction horizon of 550 points. Despite *Panda* being trained
1590 non-autoregressively for a $78 \times$ shorter prediction horizon, it is capable of producing forecasts which
1591 maintain the characteristic Lyapunov exponent out to $25 \times$ the prediction horizon it was trained on.

1592 We do observe that *Panda* struggles to capture systems with $\lambda_{\text{max}} > 6$ in Fig. 21. This is likely due
1593 to the failure mode of mean regression over long enough prediction horizons.

1594
1595
1596 G IMPLICIT SPATIO-TEMPORAL COUPLING
1597
1598
1599

1600 Temporal attention and channel attention layers independently mix information along the patch and
1601 channel dimensions. For a system like the Lorenz attractor with coupled phase coordinates $[x, y, z]$,
1602 we would ideally want information to mix across space *and* time. We will show that by composing
1603 temporal and channel attention in sequence, *Panda* implicitly performs spatio-temporal coupling.

1604 Let W_Q, W_K, W_V denote the learned projections for temporal attention and $\bar{W}_Q, \bar{W}_K, \bar{W}_V$ for
1605 channel attention. For simplicity, we will focus on the linear attention setting without the row-wise
1606 softmax. Let $P \in \mathbb{R}^{N \times C \times d_{\text{model}}}$ be a stack of N , d_{model} -dimensional patch embeddings with C
1607 channels, and $p_i^{(c)} \in \mathbb{R}^{d_{\text{model}}}$ an individual patch embedding for patch i and channel c . The linear
1608 attention output is $(PW_QW_K^\top P)PW_V$. In vector form,

1610
1611
$$(\mathbf{TA}) : \phi_i^{(\cdot)} = \sum_{j=1}^T \left\langle W_Q^\top p_i^{(\cdot)}, W_K^\top p_j^{(\cdot)} \right\rangle W_V^\top p_j^{(\cdot)} = \sum_{j=1}^T \left\langle p_i^{(\cdot)}, A_{\mathbf{TA}}, p_j^{(\cdot)} \right\rangle W_V^\top p_j^{(\cdot)} \quad (3)$$
1612
1613

1614
1615
1616
1617
$$(\mathbf{CA}) : \bar{\phi}_i^{(k)} = \sum_{\ell=1}^c \left\langle \bar{W}_Q^\top \phi_i^{(k)}, \bar{W}_K^\top \phi_i^{(\ell)} \right\rangle \bar{W}_V^\top \phi_i^{(\ell)} = \sum_{\ell=1}^c \underbrace{\left\langle \phi_i^{(k)}, A_{\mathbf{CA}}, \phi_i^{(\ell)} \right\rangle}_{M_i^{k\ell}} \bar{W}_V^\top \phi_i^{(\ell)} \quad (4)$$
1618
1619

1620 Where **TA** denotes temporal attention and **CA** channel attention, and $A_{\text{TA}} := W_Q W_K^\top$ and $A_{\text{CA}} :=$
 1621 $\overline{W}_Q \overline{W}_K^\top$. Looking at an element of the 3-tensor $M_i^{k\ell}$ we see that:
 1622

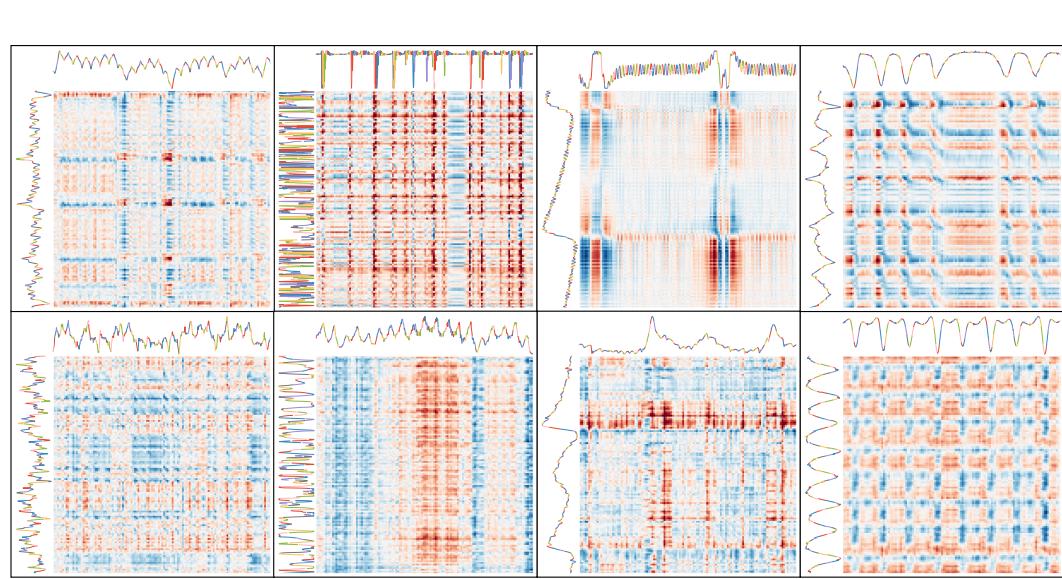
$$1623 M_{k\ell} = \left\langle \sum_{j=1}^T \left\langle \mathbf{p}_i^{(k)}, A_{\text{TA}} \mathbf{p}_j^{(k)} \right\rangle W_V^\top \mathbf{p}_j^{(k)}, A_{\text{CA}} \sum_{j'=1}^T \left\langle \mathbf{p}_i^{(\ell)}, A_{\text{TA}} \mathbf{p}_{j'}^{(\ell)} \right\rangle W_V^\top \mathbf{p}_{j'}^{(\ell)} \right\rangle \quad (5)$$

$$1626 = \sum_{j,j'=1}^T \left\langle \left\langle \mathbf{p}_i^{(k)}, A_{\text{TA}} \mathbf{p}_j^{(k)} \right\rangle W_V^\top \mathbf{p}_j^{(k)}, \left\langle \mathbf{p}_i^{(\ell)}, A_{\text{TA}} \mathbf{p}_{j'}^{(\ell)} \right\rangle A_{\text{CA}} W_V^\top \mathbf{p}_{j'}^{(\ell)} \right\rangle \quad (6)$$

$$1629 = \sum_{j,j'=1}^T \left\langle \mathbf{p}_i^{(k)}, A_{\text{TA}} \mathbf{p}_j^{(k)} \right\rangle \left\langle \mathbf{p}_i^{(\ell)}, A_{\text{TA}} \mathbf{p}_{j'}^{(\ell)} \right\rangle \left\langle \mathbf{p}_j^{(k)}, \underbrace{\left(W_V A_{\text{CA}} W_V^\top \right)}_{\tilde{\mathcal{A}}_{\text{CA}}} \mathbf{p}_{j'}^{(\ell)} \right\rangle \quad (7)$$

1633 Where $\tilde{\mathcal{A}}_{\text{CA}}$ prescribes how patches from different channels attend to each other. In matrix form,
 1634

$$1635 M_i^{k\ell} = (\mathbf{p}_{\text{TA}}^{(k)})^\top \underbrace{\begin{bmatrix} \left\langle \mathbf{p}_1^{(k)}, \tilde{\mathcal{A}}_{\text{CA}} \mathbf{p}_1^{(\ell)} \right\rangle & \dots & \left\langle \mathbf{p}_1^{(k)}, \tilde{\mathcal{A}}_{\text{CA}} \mathbf{p}_T^{(\ell)} \right\rangle \\ \vdots & \ddots & \vdots \\ \left\langle \mathbf{p}_T^{(k)}, \tilde{\mathcal{A}}_{\text{CA}} \mathbf{p}_1^{(\ell)} \right\rangle & \dots & \left\langle \mathbf{p}_T^{(k)}, \tilde{\mathcal{A}}_{\text{CA}} \mathbf{p}_T^{(\ell)} \right\rangle \end{bmatrix}}_{\text{Cross-Channel Mixing Map}} \mathbf{p}_{\text{TA}}^{(\ell)}, \quad \mathbf{p}_{\text{TA}}^{(k)} := \begin{bmatrix} \left\langle \mathbf{p}_i^{(k)}, A_{\text{TA}} \mathbf{p}_1^{(k)} \right\rangle \\ \vdots \\ \left\langle \mathbf{p}_i^{(k)}, A_{\text{TA}} \mathbf{p}_T^{(k)} \right\rangle \end{bmatrix}$$



1659 Figure 22: Cross-channel mixing maps across patches for different channels from different held-out
 1660 systems. Each mixing map is max-scaled to the range $[-1, 1]$.
 1661

1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673

1674 H FORECASTS ON PDES

1675
 1676 Other than foundation models, we include a Fourier neural operator (FNO) baseline that was trained
 1677 using the `neuraloperator` framework (Kossaifi et al., 2024; Kovachki et al., 2021), and a
 1678 DeepONet baseline trained with the `deepxde` framework (Lu et al., 2021). Unless otherwise
 1679 specified, experiment parameters follow the default values in these libraries. For both operator
 1680 learning baselines, we tune the parameters for each PDE and use the best training checkpoint
 1681 according to the validation loss. Both operator learning baselines are trained (on a single MI100X
 1682 GPU) for one-step-ahead prediction on length 512 context windows and rolled out for 512 prediction
 1683 steps for each evaluation window in Fig. 6c.

1684 For the Kuramoto-Shivashinsky (KS) PDE, we integrate the equations pseudospectrally with 64
 1685 Fourier modes and the spatial length parameter $L = 100$. We use an explicit eighth-order Dormand-
 1686 Prince scheme (DOP853) to integrate the discretized PDE with a relative and absolute tolerance of
 1687 $1e-8$ from $t = 0$ to $t = 100$ save the trajectory at 4096 uniformly spaced timepoints. We sample 40
 1688 initial conditions from $u_i \sim \mathcal{N}(0, \varepsilon^2 \mathbb{I}_{64 \times 64})$ where we choose $\varepsilon = 0.1$ and use the length 512 context
 1689 window starting at the 1024-th timepoint for training and the following 512 for prediction/rollout for
 1690 each sample to produce the error bars in Fig. 6c. See Tables 17 for comprehensive details.

1691 For the Von-Karman vortex street (VKVS) data, we use 4600 timepoints of velocity field data in
 1692 the domain $\Omega = [0, 2] \times [0, 1]$ on a 256×128 grid simulated via a Lattice Boltzmann simu-
 1693 lation at a Reynolds number of 450. We then compute the vorticity field via second-order finite
 1694 difference and reduce the dimensionality by keeping the top 512 principal components. For eval-
 1695 uation in Fig. 6c, we train on length 512 training context windows starting at the time indices
 1696 $\{0, 1024, 2048, 3072\}$ and cross-validate on the length 512 prediction windows starting at the time
 1697 indices $\{512, 1536, 2560, 3584\}$ (avoiding train-set leakage) to produce the error bars in Fig. 6c. See
 1698 Tables 18 for comprehensive details on the operator learning baselines.

| 1700 Component | 1701 Specification | 1702 Component | 1703 Specification |
|-----------------|-------------------------------|----------------|---------------------------|
| Model | Fourier Neural Operator (FNO) | Model | DeepONet |
| Modes | 256 | Branch Net | $[128, 256 \times 6]$ |
| Hidden Channels | 256 | Trunk Net | $[1, 256 \times 6]$ |
| Layers | 6 | Activation | tanh |
| Activation | GELU | Initializer | He normal |
| Optimizer | AdamW | Optimizer | AdamW |
| Learning Rate | 1×10^{-3} | Learning Rate | 1×10^{-3} |
| LR Scheduler | Cosine decay | LR Scheduler | Cosine decay |
| Epochs | 5000 | Iterations | 2×10^6 |
| Batch Size | 512 | Batch Size | 512 |
| Loss Function | L_2 | Metric | Mean relative L^2 error |

(a) FNO configuration
(b) DeepONet configuration

1712 Table 17: Kuramoto-Shivashinsky PDE Operator Learning Configurations.
 1713

| 1715 Component | 1716 Specification | 1717 Component | 1718 Specification |
|-----------------|-------------------------------|----------------|---------------------------|
| Model | Fourier Neural Operator (FNO) | Model | DeepONet |
| Modes | 512 | Branch Net | $[512, 512 \times 5]$ |
| Hidden Channels | 256 | Trunk Net | $[1, 512 \times 5]$ |
| Layers | 5 | Activation | tanh |
| Activation | GELU | Initializer | He normal |
| Optimizer | AdamW | Optimizer | AdamW |
| Learning Rate | 1×10^{-3} | Learning Rate | 1×10^{-3} |
| LR Scheduler | Cosine decay | LR Scheduler | Cosine decay |
| Epochs | 5000 | Iterations | 1×10^6 |
| Batch Size | 512 | Batch Size | 512 |
| Loss Function | L_2 | Metric | Mean relative L^2 error |

(a) FNO configuration
(b) DeepONet configuration

1726 Table 18: Von-Karman PDE Operator Learning Configurations.
 1727

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
Compared to the foundation models, the operator learning baselines under-perform mostly since they are limited to a context and prediction length of 1 for one-step-ahead prediction in contrast to the much larger context and prediction lengths of foundation models. We do not claim that foundation models are superior operator learning methods, but merely aim to provide a baseline for the PDE problems. The dash-dotted lines in Fig. 6 indicate that these methods are not directly comparable.

I COMPUTING AND HARDWARE REQUIREMENTS

All training runs were conducted on 4x AMD MI100X GPUs, each with 32 GB of HBM2 memory. Inference was performed on a single AMD MI100X GPU.

J EFFECT OF PATCH LENGTH

To investigate the effect of patch length on our model’s performance, we conduct an ablation study in which we train a version of our model with various patch lengths. To isolate the effect of patch size, we **remove** our dynamics embedding for these ablations. This is because each patch gets embedded to dimension d_{model} , making the dynamics embedding incomparable between models using different patch lengths. Keeping a fixed compute budget, we also must halve the batch size every time we halve the patch length, as a tradeoff patch length $\propto 1/\text{batch size}$ exists between the two quantities: half the patch length implies twice as many patches, all embedded to dimension d_{model} . In the tables below, $\Delta\%$ denotes percentage improvement of the best ablation over the next closest.

| sMAPE Median [P25, P75] | | | |
|-------------------------|--------------------------|--------------------------|--------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Patch 4 | 26.6 [17.5, 37.4] | 36.2 [26.5, 47.5] | 48.1 [39.0, 58.4] |
| Patch 8 | 28.6 [19.6, 40.2] | 37.7 [27.7, 48.8] | 48.2 [39.9, 58.7] |
| Patch 12 | 28.9 [19.4, 40.6] | 37.8 [27.9, 49.5] | 47.9 [39.6, 59.0] |
| Patch 16 | 29.1 [19.7, 41.1] | 37.6 [27.8, 49.4] | 47.7 [38.5, 58.7] |
| Patch 24 | 30.1 [20.6, 41.4] | 37.9 [27.9, 49.5] | 47.0 [38.7, 58.3] |
| Patch 32 | 30.1 [20.3, 42.2] | 37.7 [28.2, 49.3] | 46.6 [38.1, 57.4] |
| $\Delta\% (\uparrow)$ | +7.0% | +3.7% | +0.9% |

Table 19: Median sMAPE and interquartile range [P25, P75] for various patch lengths.

| MAE Median [P25, P75] | | | |
|-----------------------|-----------------------------|-----------------------------|-----------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Patch 4 | 0.321 [0.198, 0.509] | 0.472 [0.329, 0.685] | 0.642 [0.481, 0.824] |
| Patch 8 | 0.359 [0.224, 0.543] | 0.498 [0.359, 0.710] | 0.668 [0.502, 0.860] |
| Patch 12 | 0.364 [0.234, 0.559] | 0.501 [0.366, 0.711] | 0.662 [0.490, 0.859] |
| Patch 16 | 0.380 [0.246, 0.561] | 0.507 [0.365, 0.724] | 0.656 [0.486, 0.846] |
| Patch 24 | 0.382 [0.245, 0.578] | 0.519 [0.375, 0.729] | 0.668 [0.495, 0.857] |
| Patch 32 | 0.382 [0.248, 0.584] | 0.519 [0.376, 0.722] | 0.664 [0.492, 0.852] |
| $\Delta\% (\uparrow)$ | +10.7% | +5.2% | +2.1% |

Table 20: Median MAE and interquartile range [P25, P75] for various patch lengths.

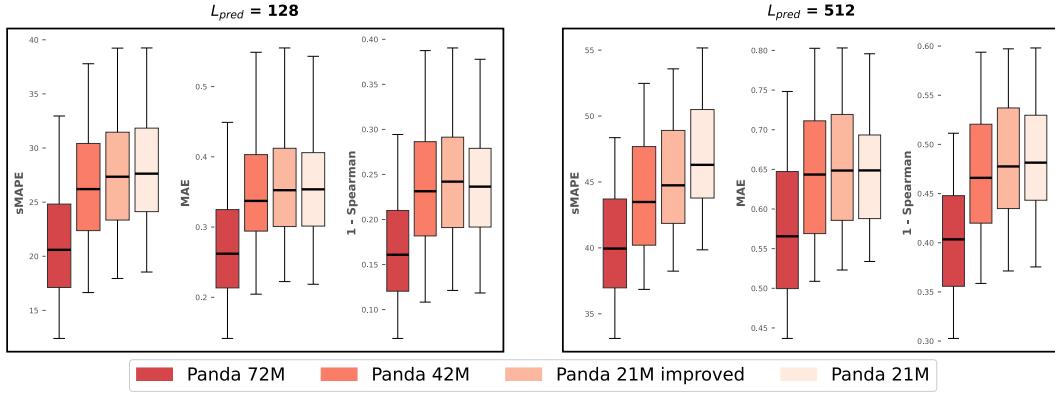
| 1 - ρ (Spearman distance) Median [P25, P75] | | | |
|--|-----------------------------|-----------------------------|-----------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Patch 4 | 0.219 [0.117, 0.349] | 0.322 [0.197, 0.464] | 0.480 [0.326, 0.612] |
| Patch 8 | 0.244 [0.138, 0.376] | 0.350 [0.229, 0.485] | 0.495 [0.350, 0.637] |
| Patch 12 | 0.246 [0.139, 0.384] | 0.357 [0.225, 0.507] | 0.496 [0.338, 0.640] |
| Patch 16 | 0.261 [0.138, 0.412] | 0.360 [0.221, 0.504] | 0.486 [0.341, 0.644] |
| Patch 24 | 0.264 [0.150, 0.408] | 0.361 [0.235, 0.512] | 0.492 [0.348, 0.643] |
| Patch 32 | 0.268 [0.150, 0.413] | 0.364 [0.230, 0.513] | 0.485 [0.342, 0.644] |
| $\Delta\% (\uparrow)$ | +10.2% | +8.0% | +1.0% |

Table 21: Median $1 - \rho$ and interquartile range [P25, P75] for various patch lengths.

1782 K SCALING UP 1783

1784 We scale up our model parameters and training to investigate the improvement in performance. Our
1785 21M, 42M, and 72M parameter models have values of (n_{heads} , n_{layers} , d_{model}) set to (8, 8, 512), (10,
1786 10, 640), and (12, 12, 768) respectively. For the scaled-up training, we had (N_{iters} , batch size per
1787 device, number of GPUs) set to (400K, 1024, 4), (400K, 512, 6), and (800K, 384, 6) respectively.

1788 We also scaled up our MLM checkpoint to create *Panda MLM-66M* with (n_{heads} , n_{layers} , d_{model}) set to
1789 (12, 12, 768), trained for 800K iterations, with batch size 192, and on 6 GPUs.
1790



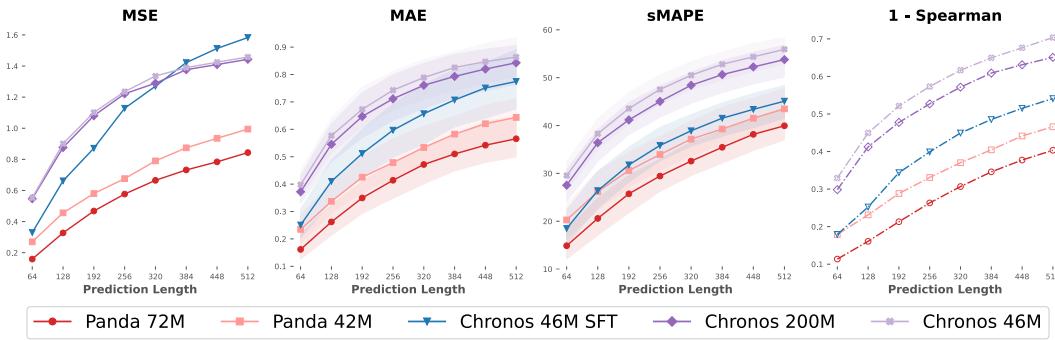
1804 Figure 23: Zero-shot metrics for scaled-up checkpoints of *Panda* with increasing number of parameters.
1805 Here, Panda 21M is our original model presented in the main body, for Panda 21M *improved*,
1806 we trained for 400K iterations (compared to the 100K iterations for our original Panda 21M),
1807 and on an improved dataset with $\approx 20\%$ more systems, which we also use for the Panda 42M training.
1808 For Panda 72M, we trained on a larger version of our improved dataset with 8 initial conditions per
1809 system and with mixed periods. For presentation, bars show a semi-interquartile range (40th to 60th
1810 percentile); for numerical values of medians and interquartile ranges, see Tables 22, 23, 24.

1811 K.1 SCALED-UP BASELINES 1812

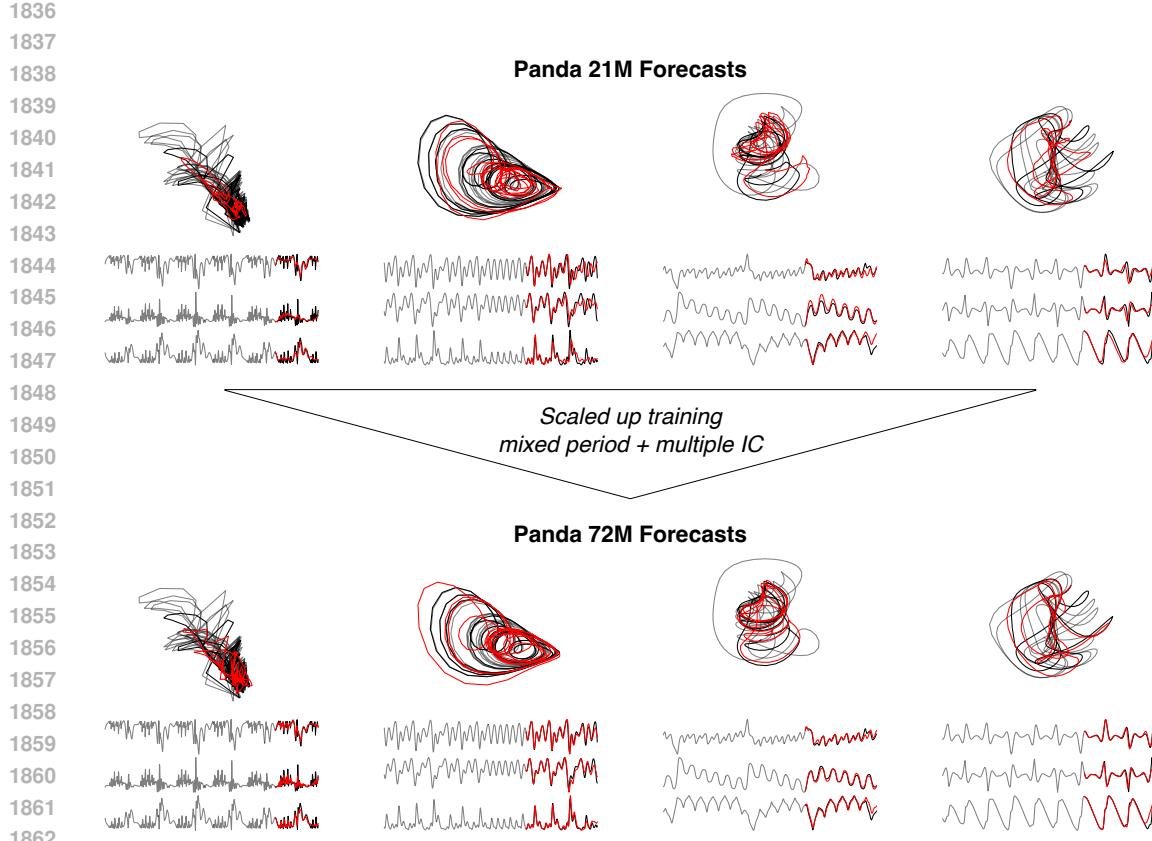
1814 We also scaled up the model parameters for Chronos as well as the training for the Chronos SFT
1815 baseline. Hardware limitations prevented us from fine-tuning Chronos 200M and larger model classes.
1816 We observe that our model continues to beat the baselines.

1817 For our scaled-up training of Chronos 46M SFT, we used ($N_{\text{iters}} = 400K$, batch size per device =
1818 100, number of GPUs = 6).

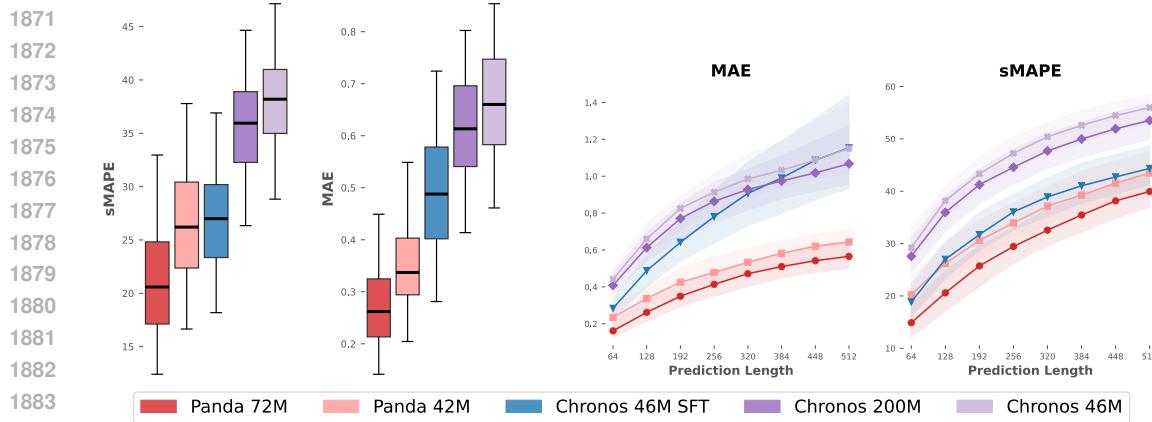
1819 For Panda 72M and for Panda MLM-66M, we trained on a larger dataset with 8 initial conditions per
1820 system and with mixed periods.
1821



1833 Figure 26: Zero-shot forecast metrics for scaled-up baselines, using *probabilistic* forecasts for
1834 Chronos and Chronos SFT. Dash-dotted lines indicate presence of NaNs for some systems (4% of
1835 systems for Spearman).



1863 Figure 24: Comparison of sample zero-shot forecasts between the Panda 21M model (8 heads, 8
1864 layers), and the Panda 72M model (12 heads, 12 layers), with the latter trained on a larger dataset
1865 with 8 initial conditions and mixed periods. As reflected in the metrics of Fig. 23, the scaled-up
1866 model forecasts appear to decrease error and capture higher-frequency details.



1885 Figure 25: Zero-shot forecast metrics for scaled-up baselines, using deterministic forecasts for
1886 Chronos and Chronos SFT. For Panda 72M, we trained on a larger dataset with 8 initial conditions per
1887 system and with mixed periods. Spearman correlation is not shown because of the high proportion of
1888 NaNs for the Chronos deterministic forecasts, which we attribute to mean regression.

1890 We present the metrics shown in Fig. 25 and Fig. 26 in tabular form in Tables 22, 23, 24.
1891

| sMAPE Median [P25, P75] | | | |
|-------------------------------|--------------------------|--------------------------|--------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Panda 72M | 20.6 [12.4, 32.9] | 29.4 [20.0, 41.4] | 39.9 [30.9, 50.6] |
| Panda 42M | 26.2 [16.6, 37.8] | 33.9 [24.4, 45.7] | 43.5 [35.0, 54.9] |
| Chronos 46M SFT | 27.0 [18.2, 36.9] | 36.0 [26.2, 45.5] | 44.4 [32.7, 54.3] |
| Chronos 200M | 36.0 [26.3, 44.7] | 44.6 [34.6, 52.9] | 53.5 [42.7, 60.8] |
| Chronos 46M | 38.2 [28.8, 47.1] | 47.2 [38.6, 54.6] | 56.0 [46.9, 61.8] |
| Chronos 46M SFT Probabilistic | 26.4 [18.0, 37.3] | 35.8 [26.0, 46.0] | 45.1 [33.8, 54.1] |
| Chronos 200M Probabilistic | 36.4 [26.7, 44.7] | 45.0 [34.2, 53.1] | 53.8 [42.8, 60.6] |
| Chronos 46M Probabilistic | 38.3 [28.6, 46.6] | 47.5 [37.9, 54.7] | 55.9 [47.2, 61.6] |

1902 Table 22: Median sMAPE and interquartile range [P25, P75] for scaled-up models.
1903

| MAE Median [P25, P75] | | | |
|-------------------------------|--------------------------|--------------------------|--------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Panda 72M | 0.26 [0.14, 0.45] | 0.41 [0.26, 0.60] | 0.57 [0.41, 0.78] |
| Panda 42M | 0.34 [0.20, 0.55] | 0.48 [0.33, 0.71] | 0.64 [0.46, 0.85] |
| Chronos 46M SFT | 0.49 [0.28, 0.72] | 0.78 [0.45, 1.17] | 1.15 [0.61, 1.95] |
| Chronos 200M | 0.61 [0.41, 0.80] | 0.86 [0.58, 1.18] | 1.07 [0.75, 1.83] |
| Chronos 46M | 0.66 [0.46, 0.85] | 0.91 [0.63, 1.23] | 1.15 [0.81, 1.95] |
| Chronos 46M SFT Probabilistic | 0.41 [0.25, 0.64] | 0.60 [0.37, 0.89] | 0.77 [0.50, 1.21] |
| Chronos 200M Probabilistic | 0.55 [0.38, 0.73] | 0.71 [0.51, 0.91] | 0.84 [0.63, 1.12] |
| Chronos 46M Probabilistic | 0.58 [0.41, 0.75] | 0.74 [0.54, 0.93] | 0.86 [0.69, 1.17] |

1916 Table 23: Median MAE and interquartile range [P25, P75] for scaled-up models.
1917

| 1 - ρ (Spearman distance) Median [P25, P75] | | | |
|--|--------------------------|--------------------------|--------------------------|
| Model | $L_{\text{pred}} = 128$ | $L_{\text{pred}} = 256$ | $L_{\text{pred}} = 512$ |
| Panda 72M | 0.16 [0.07, 0.29] | 0.26 [0.14, 0.40] | 0.40 [0.26, 0.56] |
| Panda 42M | 0.23 [0.11, 0.39] | 0.33 [0.20, 0.48] | 0.47 [0.32, 0.62] |
| Chronos 46M SFT Probabilistic | 0.25 [0.14, 0.40] | 0.40 [0.25, 0.54] | 0.54 [0.36, 0.68] |
| Chronos 200M Probabilistic | 0.41 [0.25, 0.55] | 0.53 [0.35, 0.67] | 0.65 [0.48, 0.79] |
| Chronos 46M Probabilistic | 0.45 [0.28, 0.60] | 0.57 [0.40, 0.71] | 0.70 [0.55, 0.82] |

1926 Table 24: Median $1 - \rho$ and interquartile range [P25, P75] for scaled-up models.
1927

1929 K.2 DATASET WITH MULTIPLE INITIAL CONDITIONS AND MIXED PERIODS

1930 For our scaled-up training, we used larger dataset of multiple initial conditions and mixed periods.
1931 We present a sample of this dataset in Fig. 27. We vary the number of periods (on Fourier timescale),
1932 from 20 to 100 to produce multiple periods, and carry out the numerical integration with up to 16
1933 different initial conditions (although we only use 8 initial conditions per system for training, due to
1934 compute budget restrictions). We integrate the same set of 2×10^4 systems used in our training set.
1935 The scaled-up training thus allows us to assess the effect of varying the timescales present in our
1936 training data.
1937

1938
1939
1940
1941
1942
1943

1944

1945

1946

1947

1948

1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997

1998

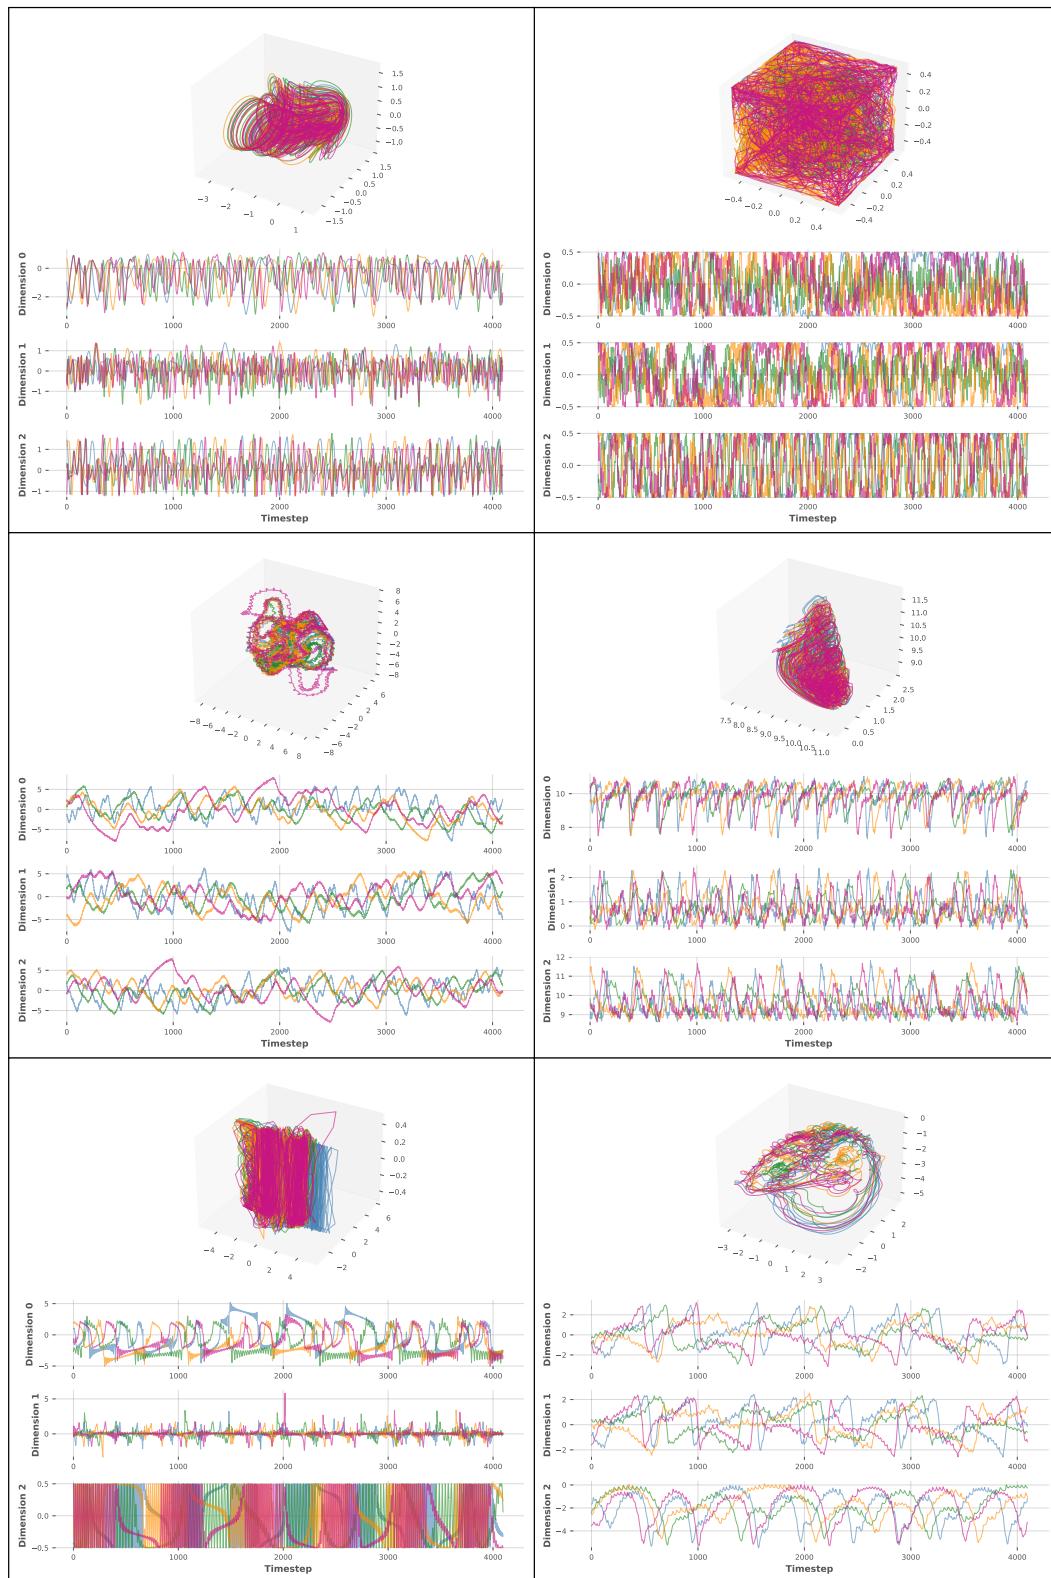


Figure 27: Examples of systems from our mixed period multi-IC training dataset. Each subplot shows multiple (4) initial conditions for a single system (integrated with different timescale).

1998
1999

L ADDITIONAL FORECASTS

2000
2001
2002
2003
2004

In Appendix C, we presented a sample of forecasts from *Panda* on our held-out test set. Here, we provide more forecasts. As done previously, we keep the prediction length fixed at $L_{\text{pred}} = 256$ for consistency and clearer visibility. Our model was trained with $L_{\text{pred}} = 128$, so these forecasts include an autoregressive rollout. Fig. 30 presents more forecasts, and Appendix Section M presents failure modes and comparison against baseline models.

2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032

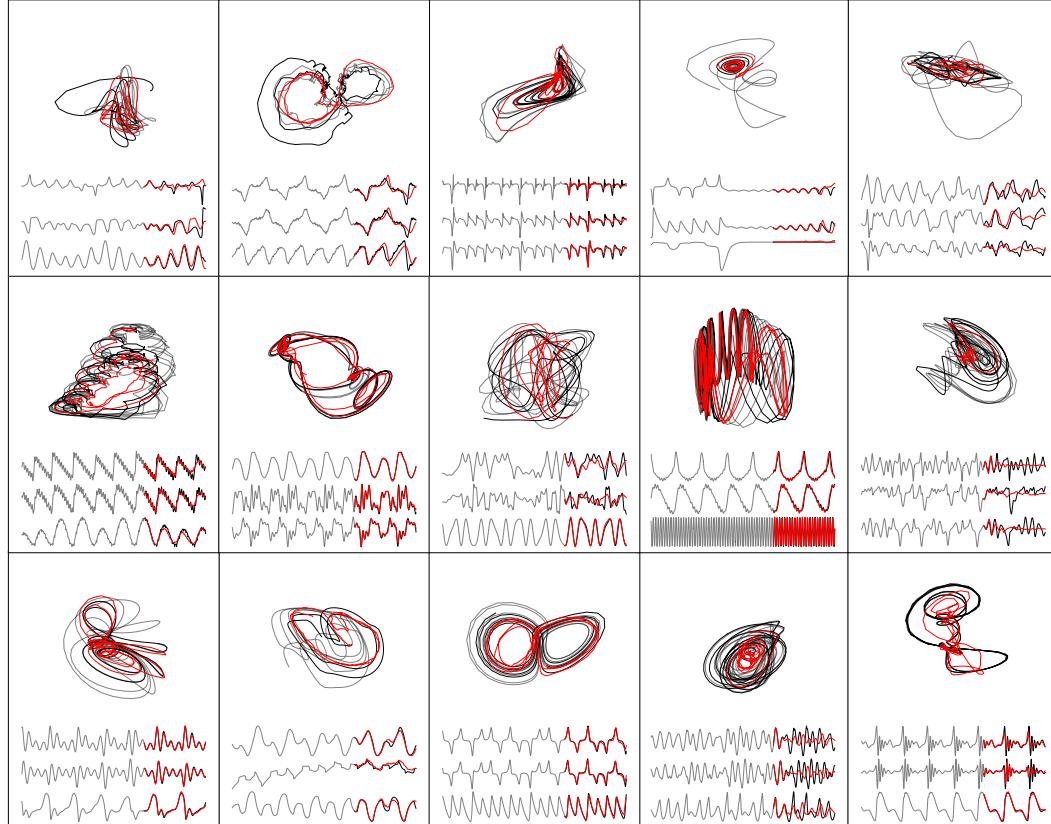


Figure 28: Examples of zero-shot forecasts ($L_{\text{pred}} = 256$) on held-out chaotic dynamical systems.

2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

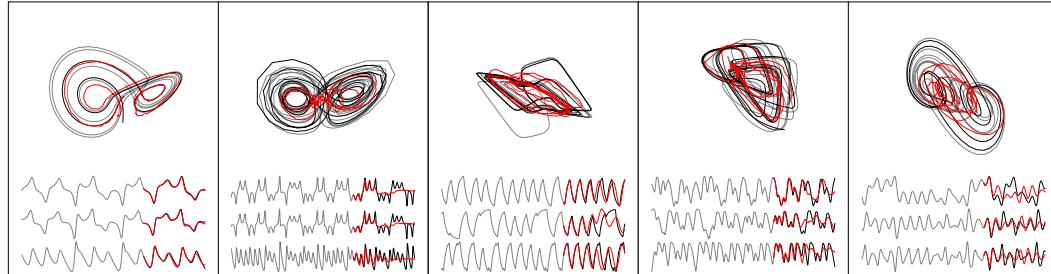


Figure 29: Examples of zero-shot forecasts ($L_{\text{pred}} = 256$) on held-out base systems (from the founder pool, parents of the skew-product systems).

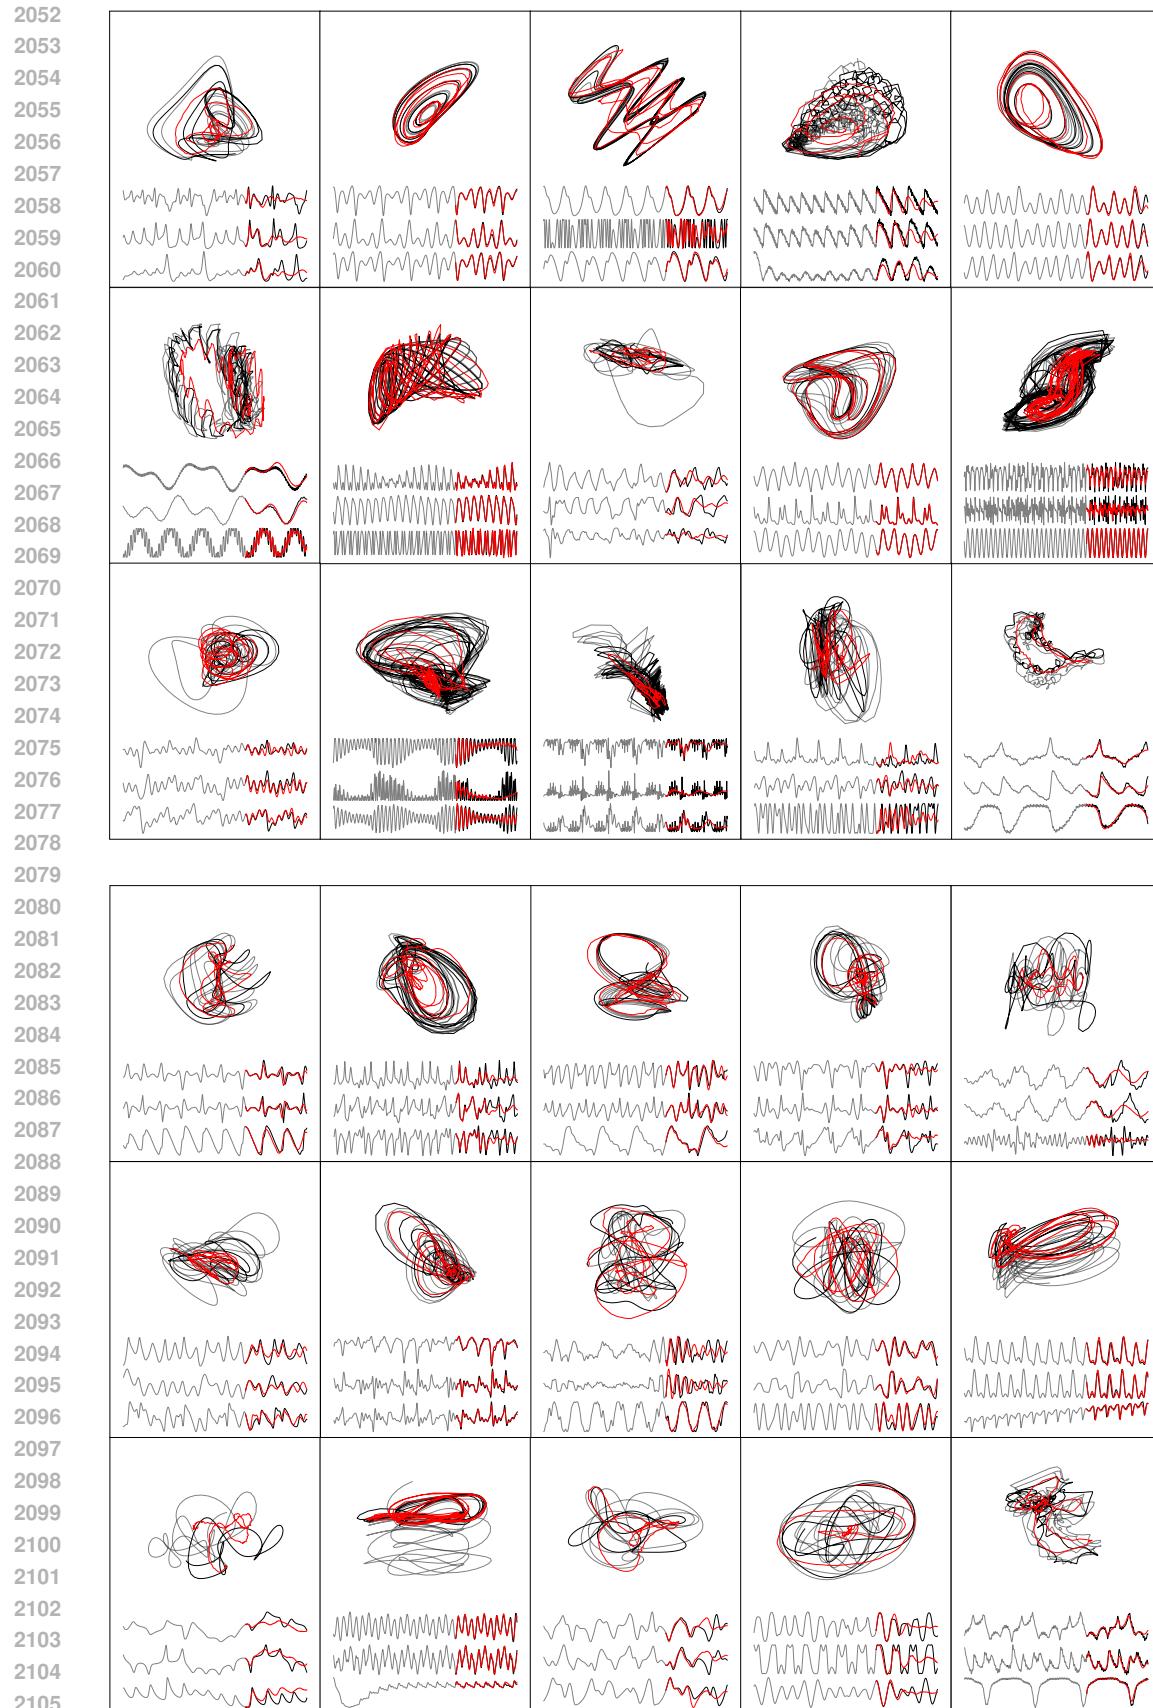
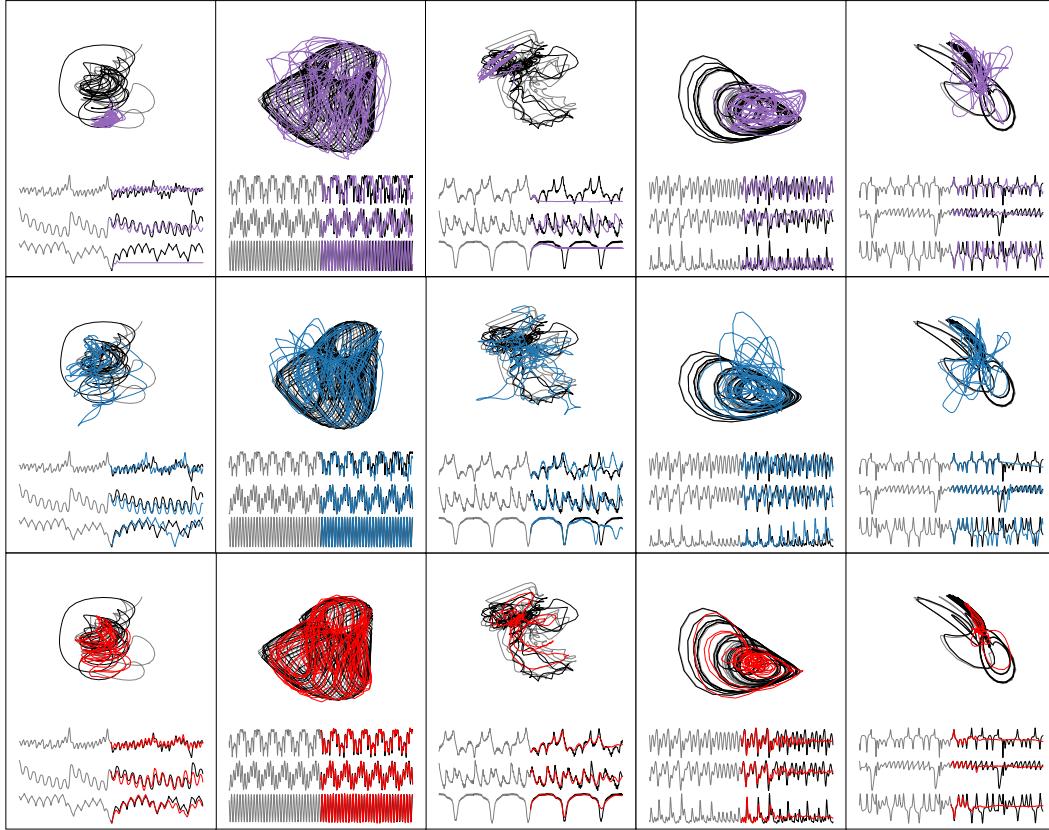


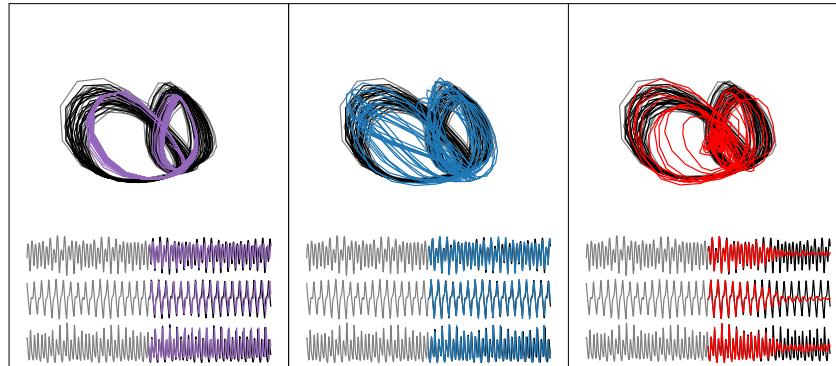
Figure 30: Examples of zero-shot forecasts ($L_{\text{pred}} = 256$) on held-out chaotic dynamical systems.

2106 M COMPARISON WITH BASELINE MODEL FORECASTS AND FAILURE MODES 2107

2108 We compare long-term ($L_{\text{pred}} = 512$) forecasts between *Panda* and the *Chronos SFT* and *Chronos*
2109 baselines. The following plots highlight some failure modes of each model, and also emphasize the
2110 advantage of our multivariate approach. Clearly, a univariate model can do well on a single channel
2111 (dimension) but fail to respect the attractor geometry. The coupling between channels encodes
2112 important information.



2141 Figure 31: Long-term zero-shot forecasts ($L_{\text{pred}} = 512$) on held-out chaotic dynamical systems.
2142 Comparison between *Panda* (Red), *Chronos SFT* (Blue), and *Chronos* (Purple).



2143
2144 Figure 32: Comparison ($L_{\text{pred}} = 512$) between *Panda* (Red), *Chronos SFT* (Blue), and *Chronos*
2145 (Purple). An illustrative example of a held-out system where *Chronos* appears to parrot (limit cycle),
2146 *Chronos SFT* does not respect the attractor geometry, and *Panda* mean regresses.
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

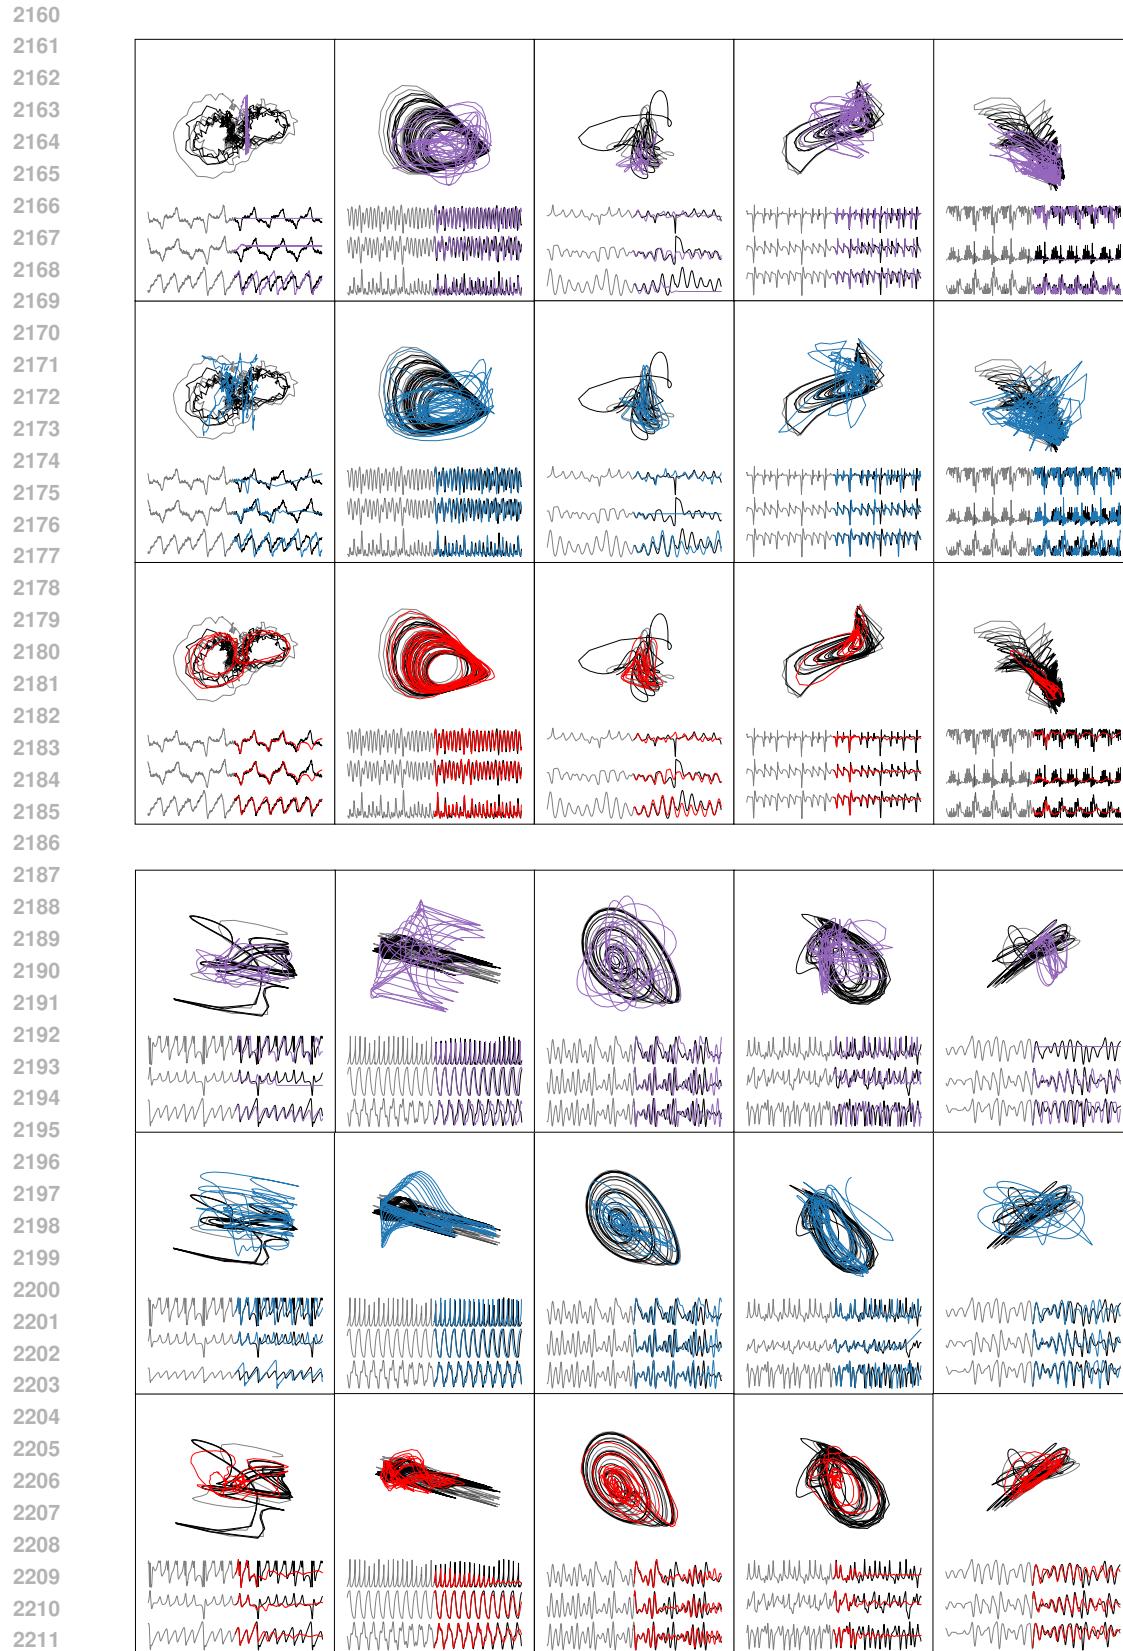


Figure 33: Long-term zero-shot forecasts ($L_{\text{pred}} = 512$) on held-out chaotic dynamical systems. Comparison between *Panda* (Red), *Chronos SFT* (Blue), and *Chronos* (Purple).

2214

N ADDITIONAL COMPLETIONS

2215

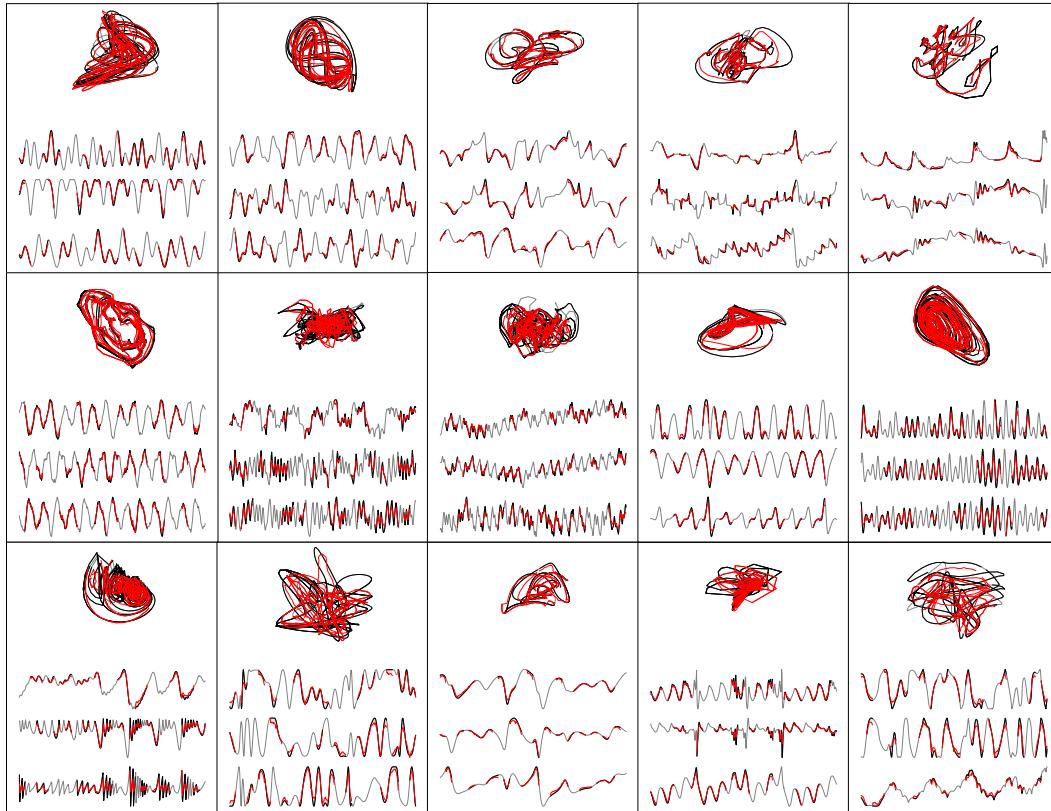
2216 In Appendix E, we presented a sample of completions from our *Panda MLM* checkpoint on our
2217 held-out test set. Here, we provide more completions from *Panda MLM* (Fig. 34) and long-context
2218 completions from our scaled-up checkpoint *Panda MLM-66M* (Fig. 36). We also provide qualitative
2219 comparison between *Panda MLM* completions and piecewise cubic spline interpolation (Fig. 35) to
2220 further demonstrate the advantage of our method.

2248

Figure 34: Examples of zero-shot completions on held-out chaotic dynamical systems. Each com-
2249 pletion plotted was with a context length of 512 time points, with half the patches (patch length 16)
2250 randomly masked out in a channel-inconsistent manner. These plots show *Panda MLM*, our 20M
2251 parameter checkpoint, completing the masked-out trajectories i.e. 256 time points.

2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267

2268

2269

2270

2271

2272

2273

2274

2275

2276

2277

2278

2279

2280

Panda MLM

2281

2282

2283

2284

2285

2286

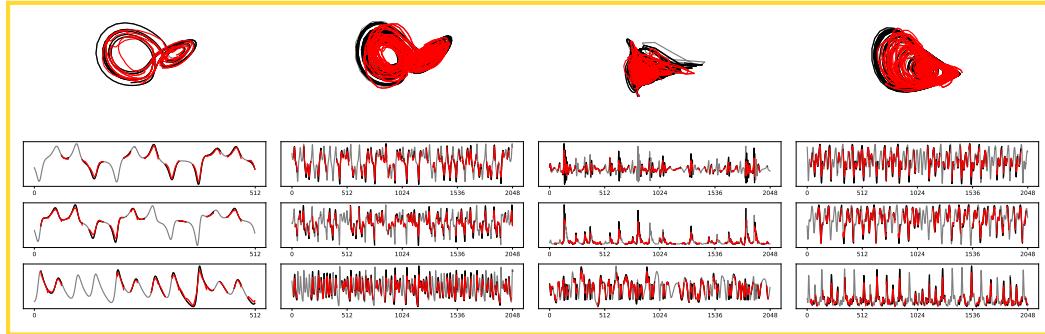
2287

2288

2289

2290

2291



2292

Piecewise Cubic Spline

2293

2294

2295

2296

2297

2298

2299

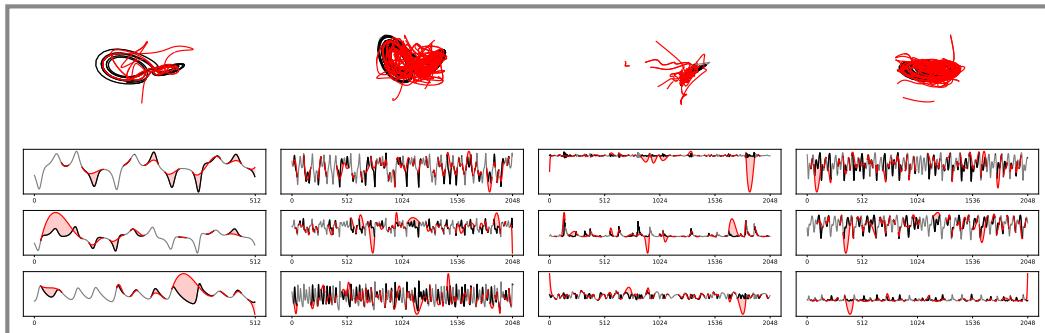
2300

2301

2302

2303

2304



2305

Figure 35: Qualitative comparison between completions generated by *Panda MLM* (with 20M parameters) and by piecewise cubic splines. First (leftmost) panel provides an example with context length 512 for clearer presentation; all other panels show context length 2048. Shaded red regions show the difference from the ground truth. Piecewise cubic spline interpolation is the most successful naive baseline, and although it achieves near competitive performance on preserving the correlation dimension (Table 11), it is not competitive with respect to pointwise error or attractor geometry.

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321

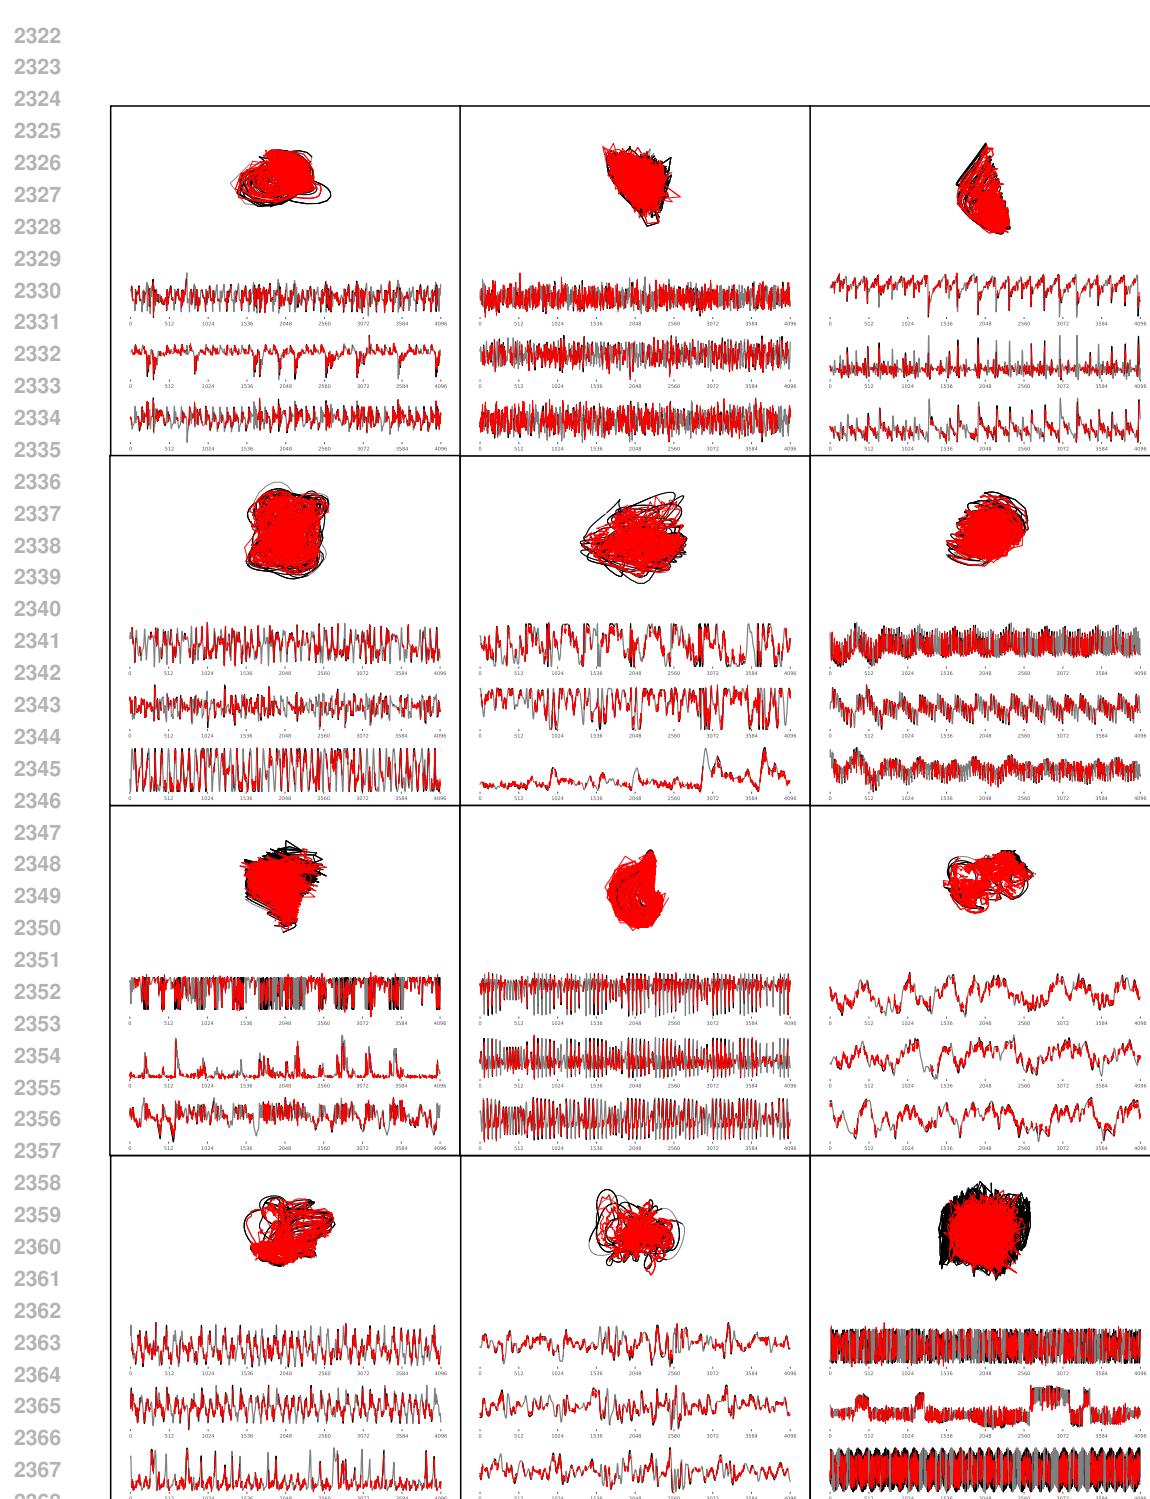


Figure 36: Examples of zero-shot completions on held-out chaotic dynamical systems. Each completion plotted was with a context length of 4096 time points, with half the patches (patch length 16) randomly masked out in a channel-inconsistent manner. These plots show *Panda MLM-66M*, completing the masked-out trajectories i.e. 2048 time points, despite only being trained on context length 512.