

# 000 CONTEXT PARROTING: A SIMPLE BUT TOUGH-TO-BEAT 001 BASELINE FOR FOUNDATION MODELS IN SCIENTIFIC 002 MACHINE LEARNING 003

006 **Anonymous authors**

007 Paper under double-blind review

## 011 ABSTRACT

013 Recent time-series foundation models exhibit strong abilities to predict physical  
014 systems. These abilities include zero-shot forecasting, in which a model forecasts  
015 future states of a system given only a short trajectory as context, without knowledge  
016 of the underlying physics. Here, we show that foundation models often forecast  
017 through a simple parroting strategy, and when they are not parroting they exhibit  
018 some shared failure modes such as converging to the mean. As a result, a naive  
019 context parroting model that copies directly from the context scores higher than  
020 leading time-series foundation models on predicting a diverse range of dynamical  
021 systems, including low-dimensional chaos, turbulence, coupled oscillators, and  
022 electrocardiograms, at a tiny fraction of the computational cost. We draw a parallel  
023 between context parroting and induction heads, which explains recent works show-  
024 ing that large language models can often be repurposed for time series forecasting.  
025 Our dynamical systems perspective also ties the scaling between forecast accuracy  
026 and context length to the fractal dimension of the underlying chaotic attractor,  
027 providing insight into previously observed in-context neural scaling laws. By  
028 revealing the performance gaps and failure modes of current time-series foundation  
029 models, context parroting can guide the design of future foundation models and  
030 help identify in-context learning strategies beyond parroting.

## 031 1 INTRODUCTION

033 A key test of generalization in scientific machine learning (SciML) is zero-shot forecasting: the ability  
034 to forecast future states of a new physical system based on a short context trajectory. Prior SciML  
035 approaches primarily focus on developing specialized forecasting models trained specifically on the  
036 system that needs to be predicted (Brunton et al., 2016; Weinan, 2017; Chen et al., 2018; Pathak et al.,  
037 2018; Li et al., 2020; Chen & Tao, 2021; Jordan et al., 2021; Gauthier et al., 2021; Lim & Zohren,  
038 2021; Karniadakis et al., 2021; Levine & Stuart, 2022; Mikhaeil et al., 2022; Brunton et al., 2022;  
039 Das et al., 2023; Krishnapriyan et al., 2023; Yang et al., 2024; Yu & Wang, 2024; Azizzadenesheli  
040 et al., 2024; Brenner et al., 2024a;b; Ricci et al., 2024; He et al., 2025; Cheng et al., 2025; Grigoryeva  
041 et al., 2025; Berry & Das, 2025). However, the generality of these models is limited by the amount of  
042 system-specific data available, motivating the recent development of time-series foundation models  
043 (Oreshkin et al., 2021; Garza & Mergenthaler-Canseco, 2023; Rasul et al., 2023; Jin et al., 2023; Zhou  
044 et al., 2023; Gruver et al., 2024; Dooley et al., 2024; Liu et al., 2024b; Woo et al., 2024; Ansari et al.,  
045 2024; Goswami et al., 2024; Das et al., 2024; Liang et al., 2024; Shi et al., 2025; Zhai et al., 2024; Liu  
046 et al., 2025b), which are trained on vast amounts of observed and simulated time series from diverse  
047 domains, and which can subsequently perform zero-shot forecasts for any time series—including  
048 those generated by previously-unseen dynamical systems. Interestingly, it was recently found that,  
049 when available historical data is limited, time-series foundation models outperform classical deep  
learning models in forecasting chaotic dynamical systems (Zhang & Gilpin, 2024).

050 What mechanisms do time-series foundation models use to make zero-shot forecasts, and why they are  
051 effective for dynamical systems not seen during pre-training? It was recently observed that one such  
052 foundation model, Chronos (Ansari et al., 2024), often employs an extremely simple strategy when  
053 forecasting chaotic systems (Zhang & Gilpin, 2024). The strategy, *context parroting*, scans the context  
for nearly repeating motifs and copies the part of the context following the best-matching motif as its

054 prediction (Fig. 1). This can be viewed as a kind of “in-context nearest neighbor” algorithm, which is  
 055 easy to implement during in-context computation (Garg et al., 2022). How good is context parroting  
 056 as a zero-shot forecasting strategy? By comparing it with existing foundation models, what can we  
 057 learn about current models’ strengths and limitations?

058 Here, we compare context parroting with a diverse set of competitive baselines on the challenging  
 059 task of forecasting chaotic systems. Our baselines include four state-of-the-art time-series foundation  
 060 models: Chronos and Chronos Bolt (Ansari et al., 2024), TimesFM (Das et al., 2024), Time-MoE  
 061 (Shi et al., 2025), and Moirai (Woo et al., 2024), as well as a recent foundation model specifically  
 062 designed for dynamical systems: DynaMix (Hemmer & Durstewitz, 2025). In the Appendix, we  
 063 also include two classical forecasting methods that are particularly effective in the small-data limit:  
 064 AutoARIMA (Hyndman & Athanasopoulos, 2018) and simplex projection (Sugihara & May, 1990).  
 065 The latter represents a classical nonlinear forecasting method conceptually resembling context  
 066 parroting (Appendix H). We find that parroting outperforms all baselines (including the leading  
 067 foundation models) in both zero-shot forecast accuracy and inference cost, especially for longer  
 068 context windows. Our results suggest that current time-series foundation models do not fully utilize  
 069 the information in the context data, and thus still have significant room for improvement when it  
 070 comes to SciML tasks.

071 Our main contributions are:

- 072 1. Introduce context parroting as a simple but effective baseline for zero-shot forecasting of  
 073 dynamical systems, which can guide the design of more informative benchmarks that cannot  
 074 be solved by simple repetitions and help identify forecasting strategies beyond parroting
- 075 2. Show that context parroting outperforms leading time-series foundation models in predicting  
 076 chaotic systems and reveal common failure modes of many existing foundation models,  
 077 which can guide the design of better models in the future
- 078 3. Explain the in-context neural scaling law between forecast accuracy and context length,  
 079 linking the scaling coefficient to the fractal dimension of the underlying chaotic attractor

## 081 2 RELATED WORK

082 **Foundation models for science.** Foundation models have recently been introduced for many scientific  
 083 machine-learning tasks (Miller et al., 2024), including partial differential equations (Takamoto et al.,  
 084 2022; Yang et al., 2023; Rahman et al., 2024; Subramanian et al., 2024; Herde et al., 2024; McCabe  
 085 et al., 2024; Totounferoush et al., 2025), neuroscience (Cui et al., 2024; Caro et al., 2023; McKeen  
 086 et al., 2024), and weather forecasting (Nguyen et al., 2023; Bodnar et al., 2024). However, most of  
 087 these foundation models remain a black box, and they have not yet provided interpretable strategies  
 088 for forecasting diverse physical and dynamical processes. Here, we analyze context parroting as a  
 089 simple mechanism used by time-series foundation models, noting its strengths and weaknesses as a  
 090 zero-shot forecasting strategy. This strategy, and the insights gained here, can potentially be applied  
 091 to other scientific tasks.

092 **In-context neural scaling laws.** Neural scaling laws describe the relationship between the performance  
 093 of a neural network and certain resources, such as model size, data size, or the amount of  
 094 compute (Kaplan et al., 2020; Sorscher et al., 2022; Bahri et al., 2024; Yao et al., 2024). Such scaling  
 095 laws allow practitioners to predict the performance of yet-to-be-trained models based on the available  
 096 resources and allocate them strategically to optimize compute-adjusted accuracy (Hoffmann et al.,  
 097 2022). When applying LLMs to forecast dynamical systems, Liu et al. (2024a) recently observed  
 098 an in-context neural scaling law, in which the test loss decreases with the context length following  
 099 a power law. Here, we show that this in-context neural scaling law can be reproduced when using  
 100 context parroting to predict dynamical systems, and the scaling coefficient can be linked to an invariant  
 101 property of the underlying dynamic process (the fractal dimension of the chaotic attractor). This  
 102 finding shows that neural scaling laws are intrinsically linked to invariant properties of the process  
 103 generating the data, and the theory can potentially be generalized to other models and tasks (e.g., can  
 104 we estimate the “fractal dimension” of a language from the neural scaling laws of LLMs?).

105 **In-context learning and induction heads.** Induction heads are computational circuits that naturally  
 106 emerge in simple transformers through training, and they have been hypothesized to underlie much  
 107 of the in-context learning ability of foundation models (Elhage et al., 2021; Olsson et al., 2022;

108 [Von Oswald et al., 2023](#); [Reddy, 2023](#)). In its simplest form, an induction head copies repeating tokens  
 109 in the context to make predictions. For example, when presented with a token stream  $[A][B]\dots[A]$ ,  
 110 an induction head will output  $[B]$  as the next token. Prior works train transformers on minimal  
 111 Markov chain grammars, and find that, during pretraining, models learn to identify increasingly  
 112 higher-order  $k$ -grams, with different attention heads specializing in copying, lookup, and aggregation  
 113 ([Edelman et al., 2024](#); [Chen et al., 2024a](#)). These works imply that pretraining enables models to  
 114 learn conditional distributions, allowing them to represent sequence distributions seen in the context  
 115 ([Lv et al., 2024](#); [Chen et al., 2024b](#); [Keskar et al., 2019](#); [Zekri et al., 2024](#)).

116 There is a clear parallel between context parroting and induction heads: both are essentially copy-and-  
 117 paste operations, with context parroting involving the matching of not just one but multiple contiguous  
 118 tokens. In fact, it is easy to imagine context parroting emerging naturally from combining multiple  
 119 induction heads. This parallel can potentially explain the unreasonable effectiveness of applying  
 120 language models trained on text to time-series tasks without fine-tuning or prompt engineering ([Garza](#)  
 121 & [Mergenthaler-Canseco, 2023](#); [Jin et al., 2023](#); [Zhou et al., 2023](#); [Gruver et al., 2024](#); [Liu et al.,](#)  
 122 [2024a](#)). The induction heads formed from training on natural language happen to be also effective for  
 123 predicting time series and can be easily repurposed to implement strategies such as context parroting.  
 124

### 125 3 CONTEXT PARROTING AS A ZERO-SHOT FORECASTING STRATEGY

127 **Overview of context parroting.** In this section we motivate and introduce our baseline: context  
 128 parroting. It was inspired by recent observations that Chronos often predicts chaotic systems by  
 129 copying directly from the context ([Zhang & Gilpin, 2024](#)). An example of Chronos using parroting  
 130 to forecast a partially-observed Lorenz system is shown in Fig. 1.

131 On a high level, context parroting uses the last  $D$  tokens of the context to query the remaining  
 132 context. For whatever context sequence that most closely matches the query, the subsequent tokens  
 133 in the context are copied and used as the forecast. Because the length of the motif  $D$  can be seen  
 134 as the number of delayed states in a delay embedding from the lens of Takens' embedding theorem  
 135 ([Takens, 2006](#); [Huke, 2006](#)), we also refer to  $D$  as the embedding dimension and will use the  
 136 terms embedding dimension and query length interchangeably. Interpreting  $D$  as the embedding  
 137 dimension is convenient because context parroting can be seen as a nearest neighbor algorithm in the  
 138  $D$ -dimensional delay-embedded space. During the matching process, we exclude the last  $D$  motifs  
 139 to avoid parroting too close to where the prediction starts. Framed in terms of induction heads, the  
 140 query lookup acts as a *copy* head, the nearest-neighbor match is a *selector*, and the exact repetition is  
 141 the *aggregation* operation ([Chen et al., 2024a](#)). We provide a pseudocode for context parroting in  
 142 Algorithm 1.

---

#### 143 **Algorithm 1** Context Parroting

145 **Input:** Context trajectory  $x_{1:L} = \{x_1, \dots, x_L\}$ , embedding dimension  $D$  (i.e., the length of the  
 146 motif to match), and forecast length  $H$ .

147 **Output:** Forecast trajectory  $x_{L+1:L+H} = \{x_{L+1}, \dots, x_{L+H}\}$ .

- 1: **for all** length- $D$  motif  $s: x_{s-D+1:s}$  in the context  $x_{1:L-D}$  **do**
- 2:   compute the Euclidean distance  $d_s$  between motif  $s$  and the last motif  $x_{L-D+1:L}$
- 3:   Find the best-matching motif,  $s_{opt}$ , with the smallest Euclidean distance
- 4:   Set the first  $L - s_{opt}$  predicted points to be  $x_{L+1:2L-s_{opt}} = x_{s_{opt}+1:L}$  and repeat until the  
 forecast length  $H$  is reached

---

154 **Relationship to classical nonlinear forecasting methods.** We show in Appendix H.3 that, in various  
 155 limits, context parroting is equivalent to two classical algorithms from nonlinear dynamics: the  
 156 *simplex projection* technique and the *S-map* algorithm ([Sugihara & May, 1990](#); [Sugihara, 1994](#)). Both  
 157 approaches have their foundations in Takens' embedding theorem, which states that time-delayed  
 158 low-dimensional observables derived from a nonlinear dynamical system can recover key geometric  
 159 properties of the underlying high-dimensional attractor ([Takens, 2006](#); [Huke, 2006](#)). However, unlike  
 160 context parroting, which looks for the best matching motif, simplex projection tries to identify  
 161 multiple matching motifs in the context and computes a weighted average as its forecast. This  
 can potentially make simplex projection more sensitive to the choice of the embedding dimension

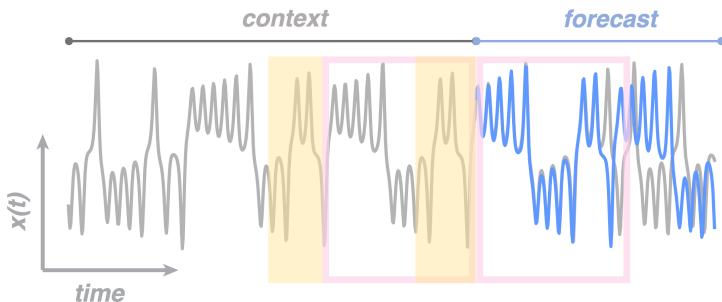


Figure 1: **Example of a foundation model forecasting chaotic dynamics with context parroting.** Here, the foundation model (Chronos) was asked to predict the  $x$  variable of the Lorenz system based on a short context trajectory with 512 data points. Blue is the prediction and gray is the ground truth. Chronos produced an accurate prediction by simply looking for a motif in the context similar to the motif immediately preceding the prediction (highlighted in yellow) and copying the evolution following the matching motif (highlighted by pink boxes). We distill this context parroting strategy into Algorithm 1 and compare it against time-series foundation models (including Chronos itself).

$D$ , limiting the method to small embedding dimensions in practice (Chang et al., 2017). Other than simplex projection and S-map, there also exists other zero-shot forecasting strategies from nonlinear dynamics. For example, the Farmer-Sidorowich method (Farmer & Sidorowich, 1987) looks at multiple nearest neighbors in the context and builds a local linear model to make forecasts. An interesting future direction is to compare these methods from nonlinear dynamics with context parroting and time-series foundation models, which might inspire new zero-shot forecasting strategies.

## 4 METHODS

**Datasets.** The `dysts` dataset provides a standardized benchmark of 135 low-dimensional chaotic systems, each defined by a set of ordinary differential equations between dimensionality three and six (Gilpin, 2021). The chaotic systems are drawn from different published papers and span fields such as neuroscience, climate science, fluid dynamics, and astrophysics. Every system is annotated with its largest Lyapunov exponent  $\lambda$ , an invariant characteristic of the underlying dynamics that quantifies the rate at which small perturbations grow over time. In chaotic systems, even minor errors rapidly compound over a characteristic timescale known as the Lyapunov time, defined as  $\tau \equiv \lambda^{-1}$ . To normalize the difficulty of predicting different chaotic systems (so results from the 135 systems can be meaningfully compared and combined), we generate trajectory data with a fixed sampling rate of 30 points per Lyapunov time, and also measure the forecast performance in terms of Lyapunov times. To show the relevance of our findings to a broad class of SciML tasks, later we also go beyond low-dimensional chaotic systems and simulated data by benchmarking on real-world datasets from ECG measurements and electronic circuits.

**Models.** For time-series foundation models, we select Chronos<sub>base</sub> (200M parameters), its variant Chronos-Bolt<sub>base</sub> (205M parameters), Time-MoE<sub>large</sub> (200M parameters), TimesFM-2.0 (500M parameters) and Moirai-2.0<sub>small</sub> (11M parameters) (Das et al., 2024; Ansari et al., 2024; 2025; Shi et al., 2025; Liu et al., 2025a). All of these models are pretrained on massive amounts of real-world time series data (hundreds of billions of data points), which are often complemented by synthetic data to improve generalization. We also consider DynaMix, a foundation model pretrained on chaotic dynamical systems (Hemmer & Durstewitz, 2025). These models encompass a wide array of design choices: Time-MoE, TimesFM-2.0, and Moirai-2.0 are decoder-only architectures, Chronos is an encoder-decoder architecture, and DynaMix is an almost-linear RNN trained via teacher forcing (Brenner et al., 2024a). Chronos and Time-MoE use pointwise tokenization, DynaMix implicitly tokenizes pointwise, while TimesFM-2.0 and Moirai-2.0 use patching. Time-MoE, DynaMix, and TimesFM-2.0 by default give point forecasts, whereas Chronos, Chronos-Bolt, and Moirai-2.0 provide probabilistic forecasts with uncertainty quantification. For these models, we use the median prediction when evaluating forecast errors. An important parameter for all foundation models is the

216 maximum context length  $L_{\max}$ , which varies from 512 data points (Chronos), 1680 (Moirai-2.0),  
 217 2048 (TimesFM-2.0), 4096 (Time-MoE), to arbitrary for DynaMix due to its recurrent formulation.  
 218

219 **Pipelines.** To evaluate different models' ability to zero-shot forecast dynamical systems, we generate  
 220 a chaotic trajectory of length  $10^5$  for each of the 135 chaotic systems in `dysts`, with a granularity of  
 221 30 data points per Lyapunov time. Each trajectory is normalized to have zero mean and unit standard  
 222 deviation. For a given context length  $L$ , we randomly pick a length- $L$  segment from the chaotic  
 223 trajectory and provide it to the model as the context. The model's task is to predict the next 300 data  
 224 points (equivalent to 10 Lyapunov times) solely based on the context. We ask the model to make a  
 225 univariate forecast on each dimension independently, which is then evaluated separately for each  
 226 dimension. To obtain reliable statistics, we aggregate the results over all 135 chaotic systems, all  
 227 dimensions, and 20 random initial conditions for each system.

228 **Metrics.** In line with previous research (Hyndman & Koehler, 2006; Makridakis et al., 2022; Gilpin,  
 229 2021; 2023), we assess forecasting performance using a diverse set of complementary metrics.

230 *Symmetric Mean Absolute Percentage Error (sMAPE).*

$$232 \text{sMAPE}(\mathbf{x}, \hat{\mathbf{x}}) \equiv 2 \frac{100}{T} \sum_{t=1}^T \frac{|\mathbf{x}_t - \hat{\mathbf{x}}_t|}{|\mathbf{x}_t| + |\hat{\mathbf{x}}_t|},$$

235 where the sequence  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$  denotes the ground truth, and  $\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, \dots, \hat{\mathbf{x}}_T$  are the corresponding  
 236 predictions made by the model. To provide some context to help interpret the sMAPE value, we  
 237 note that predicting the mean of white noise would give you an sMAPE around 200.  
 238

239 *Valid Prediction Time (VPT).* This metric identifies the latest time step  $t_f$  before which the sMAPE  
 240 remains below a predefined threshold  $\epsilon$ , as described in Vlachas et al. (2020). Formally:

$$242 \text{VPT} \equiv \text{argmax}_{t_f} \{t_f | \text{sMAPE}(\mathbf{x}_t, \hat{\mathbf{x}}_t) < \epsilon, \forall t < t_f\}.$$

244 We use  $\epsilon = 30$ , consistent with prior work (Vlachas et al., 2020; Gilpin, 2023; Zhang & Gilpin, 2024).  
 245 For chaotic systems, consistently achieving a VPT over one Lyapunov time is considered impressive  
 246 (Gilpin, 2023; Zhang & Gilpin, 2024).

247 Other than sMAPE and VPT, we also show benchmark results using *Mean Square Error* (MSE) and  
 248 *Mean Absolute Error* (MAE), two other metrics commonly used in the time series literature.

250 For chaotic dynamical systems, point forecasts will inevitably fail due to the exponential rate of error  
 251 accumulation. It is thus equally important for a forecasting model to preserve the long-term statistical  
 252 and geometric properties of the chaotic attractors, such as Lyapunov exponents and the attractor  
 253 dimension. Here, we compare the structure of true and predicted attractors by calculating the KL  
 254 Divergence between their distributions.

255 *Kullback–Leibler Divergence between Attractors ( $D_{\text{stsp}}$ ).*

$$257 D_{\text{stsp}} \equiv D_{\text{KL}}(P \| Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)},$$

261 where  $P$  and  $Q$  represent the true and the predicted attractor, respectively. When estimating  $D_{\text{stsp}}$ , we  
 262 follow the methodology in Hess et al. (2023); Göring et al. (2024). Specifically, we place Gaussian  
 263 kernels at each point in the true and predicted trajectories and estimate the KL divergence between  
 264 these Gaussian mixtures using a sampling-based approximation (Hershey & Olsen, 2007).

265 In the appendix, we also measure attractor reconstruction accuracy using *Fractal Dimension* and  
 266 *Lyapunov Exponents*. The *correlation dimension* estimates the fractal dimension from a time series  
 267 by calculating the scaling of the number of attractor points that fall within a given radius of each point  
 268 (Grassberger & Procaccia, 1983). This quantity is among the few invariant quantities of a dynamical  
 269 system that can be non-parametrically estimated from a time series; however, the estimates can still  
 be unstable and brittle depending on the fitting procedure (Clauset et al., 2009).

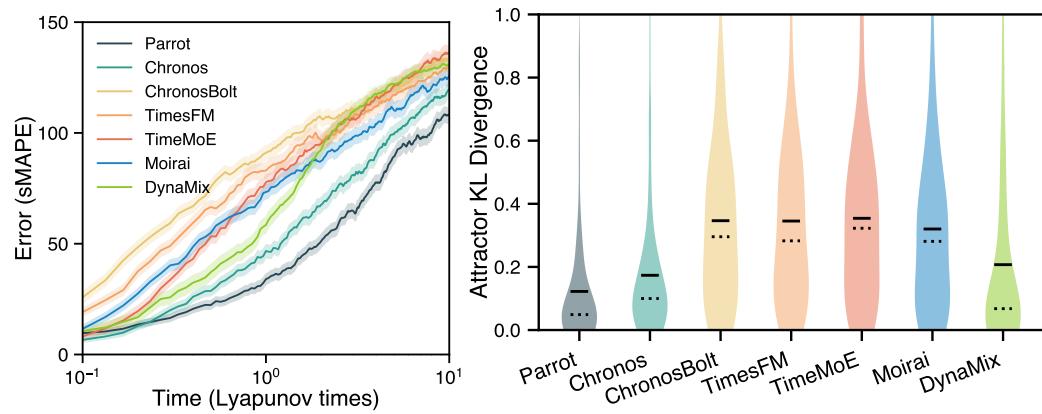
270 

## 5 RESULTS

271 

### 5.1 CONTEXT PARROTING VERSUS FOUNDATION MODELS

274 Here, we compare context parroting and foundation models in their ability to predict chaotic dynamics.  
 275 Figure 2 shows each model’s forecasting error (measured by sMAPE) as well as their accuracy in  
 276 attractor reconstruction (measured by KL Divergence). It is clear that context parroting is better than  
 277 all foundation models tested here in both metrics. In Fig. 6, we show that this remains true when  
 278 benchmarked against MSE and MAE. The results for fractal dimension accuracy are shown in Fig. 7,  
 279 where parroting and Chronos significantly outperform the rest of the foundation models. We also  
 280 explore the effects of observational noise and data granularity on the forecast performance in Tables 5  
 281 to 14. Over the wide range of noise intensity and data granularity we tested, parroting is always the  
 282 best or the second best according to all metrics (VPT, MAE, MSE, fractal dimension accuracy, and  
 283 attractor KL Divergence).



298 **Figure 2: Context parroting outperforms foundation models in zero-shot forecasting for both**  
 299 **short-term point-wise accuracy and long-term attractor reconstruction.** Left: Forecast error  
 300 of each model as a function of the forecast horizon. The context length is set to 512 for all models.  
 301 Right: KL Divergence between the predicted attractors and the true attractors (smaller is better). Solid  
 302 lines represent mean and dotted lines represent median. All results are obtained from 135 chaotic  
 303 systems in the `dysts` database, with 20 trajectories from random initial conditions for each system.

304 Among foundation models, Chronos is the best performer in predicting chaotic systems, which is  
 305 not surprising given that it utilizes parroting as a main forecasting strategy (Zhang & Gilpin, 2024).  
 306 Chronos’s tendency to context parrot arises from its distinct architecture as a language model that  
 307 implicitly quantizes time series. As a result, Chronos is trained using cross-entropy loss, which  
 308 incentivizes preservation of k-gram frequencies and encourages the generation of diverse forecast  
 309 samples consistent with the dynamical system’s underlying measure (Yu et al., 2025). In contrast,  
 310 TimeMoE and TimesFM are trained using mean squared error loss. As a result, these models  
 311 lose diversity and forecast the mean at long forecast horizons (i.e., they tend to underestimate the  
 312 oscillations). Some representative forecasts from the foundation models are shown in Fig. 5, which  
 313 shows that regressing to the mean is a common failure mode for many foundation models on chaotic  
 314 systems.

315 Moreover, as we demonstrate in Fig. 8, the inference cost of context parroting is negligible compared  
 316 to foundation models (not to mention the substantial GPU time needed to pre-train them). For  
 317 example, there is an over six orders of magnitude computational gap between Chronos and context  
 318 parroting for all context lengths. Combined with the fact that the performance of parroting is not  
 319 sensitive to the choice of the embedding dimension  $D$  (Fig. 10), these results establish context  
 320 parroting as a simple but effective baseline for zero-shot forecasting of dynamical systems.

321 Figure 3 further explores the effect of context length on forecast accuracy. We find that longer context  
 322 windows generally lead to better performance for both context parroting and Chronos. However, the  
 323 longest context length Chronos can effectively utilize is 512 data points. This limit is determined  
 by Chronos’s maximum context window chosen at pre-training. To be able to utilize longer context,

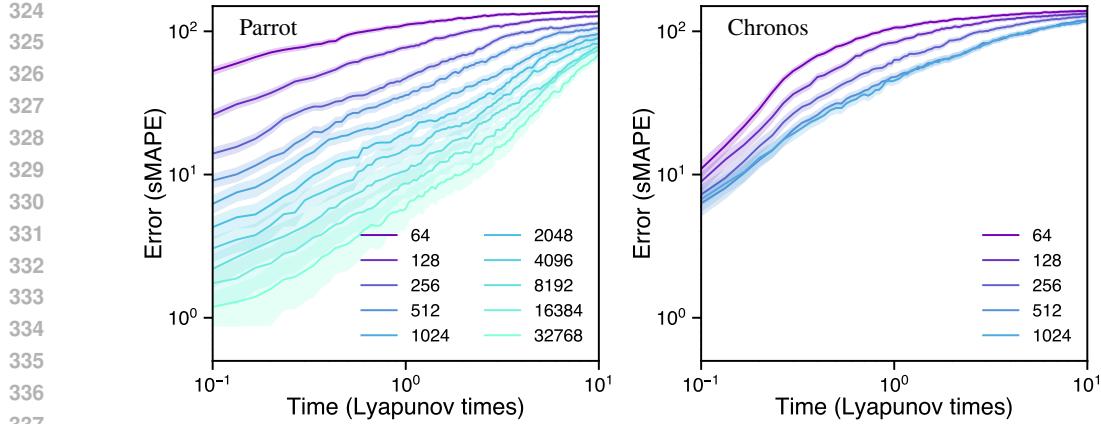


Figure 3: **Parrot can better utilize longer context data while Chronos does better for shorter contexts.** Each line represents the forecast error for a different context length. The performance of Chronos saturates once the context length exceeds its designed upper limit of 512 data points, whereas the accuracy of context parrotting keeps improving for longer context windows. Here we set the embedding dimension  $D = 10$  for the parrotting algorithm.

Chronos must be retrained from scratch with much more data and compute. Context parrotting, on the other hand, is happy to utilize context data of any length. In Fig. 9, we show VPT as a function of the context length for context parrotting and the foundation models. Again, parrotting is the clear winner for medium to long context lengths, easily reaching an average VPT of over 5 Lyapunov times when given sufficient context.

Interestingly, Chronos outperforms context parrotting on short contexts, which points to additional zero-shot learning strategies beyond parrotting employed by Chronos. This is perhaps not surprising given that at short context length, the time series becomes effectively nonstationary, which is the strength of time-series foundation models. For example, Chronos is great at continuing the local trend in the context, which can be a more effective strategy than parrotting when the length of the context is limited. Moreover, even when restricted to parrotting, the  $\sim \mathcal{O}(L^2)$  operations performed by attention heads in transformers like Chronos have, in principle, sufficient computational complexity to dynamically choose the optimal embedding dimension  $D$  for each individual time series, giving attention an advantage over parrotting algorithms with a fixed  $D$ , which have the  $\sim \mathcal{O}(DL)$  complexity of nearest-neighbor search. It would be interesting to explicitly identify the mechanisms that enable Chronos to outperform parrotting in the short-context regime.

## 5.2 IN-CONTEXT NEURAL SCALING LAW

[Liu et al. \(2024a\)](#) reported an in-context neural scaling law for LLMs applied to dynamical systems, in which the one-step forecast error decreases algebraically with context length. However, it is unclear where this scaling law came from or why LLMs trained on text can be effective for time series without fine tuning. Here, we show that context parrotting naturally gives rise to the same in-context scaling law and provides geometric insights into its origin. Given the similarity between parrotting and the induction heads implemented by LLMs ([Olsson et al., 2022](#)), the geometric explanation we develop next for context parrotting can conceivably be applied to LLMs and partially explain the observations in [Liu et al. \(2024a\)](#).

The left panel in Fig. 4 shows the power law relation between one-step forecast error (measured by sMAPE) and context length for the parrotting method. Longer context lengths improve predictions because more context data allows the algorithm to find better matching motifs, and a closely-matched motif allows the parroted sequence to shadow the ground truth for longer. The overlap between the matching motifs can be measured by their Euclidean distance. For length- $D$  motifs, this is equivalent to embedding the context trajectory in a  $D$ -dimensional delay-embedded space (i.e., mapping  $x_s$  to  $x_{s-D+1:s}$ ) and finding the distance between the embedded last context point  $x_{L-D+1:L}$  and its nearest embedded neighbor. The right panel in Fig. 4 shows the improving overlap explicitly, with the

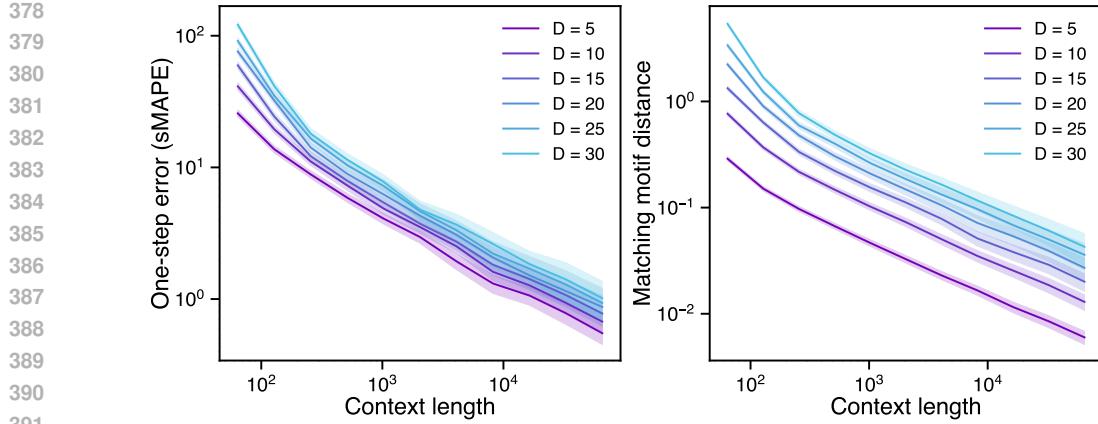


Figure 4: **Scaling laws with context length.** Left: One-step forecast error versus context length. The scaling follows a power law for all embedding dimensions  $D$  considered. Smaller  $D$  is more accurate here because of the one-step forecast error. Larger  $D$  can be more accurate for longer forecasting horizons. Right: Euclidean distance between the last context motif  $x_{L-D+1:L}$  and its closest match, as a function of context length. Again, the scaling follows a power law for all  $D$ . The forecast accuracy is directly tied to the motif distance: Smaller distances translate into better predictions. As the number of context data points is increased, it becomes more and more likely that a context motif will land in the vicinity of the last context motif, and the fractal dimension of the chaotic attractor determines the rate of approach. In principle, infinite context length should give context parrotting infinite accuracy for any deterministic system with a well-defined attractor.

distance between the matching motifs decreasing algebraically with context length. It is easy to see that matching motif distance should map linearly to the expected one-step forecast error, so the power law for the matching motif distance implies the power law for one-step forecast error, as confirmed by Fig. 4. Although we focus on deterministic ODEs in Fig. 4, we note that the same power law scaling is expected to hold for discrete maps and for systems subject to weak noise. Explaining the power law in the case of stochastic systems, such as Markov chains with randomly generated transition matrices (Liu et al., 2024a), is a promising direction for future research.

Why then does the matching motif distance follow a power law with context length,  $\ell \propto L^{-\alpha}$ ? We can link the scaling coefficient  $\alpha$  to the fractal dimension of the chaotic attractor. The fractal dimension of an attractor, as estimated by the correlation dimension  $d_{\text{cor}}$ , is defined as

$$d_{\text{cor}} \equiv \lim_{\epsilon, \epsilon' \rightarrow 0^+} \frac{\ln \left[ \frac{C(\epsilon)}{C(\epsilon')} \right]}{\ln \left( \frac{\epsilon}{\epsilon'} \right)},$$

where  $C(\epsilon)$  is the number of point pairs in the attractor that are within a given radius  $\epsilon$  from each other. In other words, if we plot  $C(\epsilon)$  against  $\epsilon$  on a log-log plot,  $d_{\text{cor}}$  would be the slope of the plotted line. Due to the ergodic property of chaotic attractors, the context trajectory can be seen as a random sample of the attractor. Longer context trajectory contains more samples, and the expected distance between two context points in a delay-embedded space decreases with context length as  $L^{-1/d_{\text{cor}}}$ . For example, for a two-dimensional attractor, the distance between two random points on the attractor will decrease as  $1/\sqrt{L}$ . Fractal dimension thus measures the speed at which the minimum embedding distance between points on an attractor can be reduced by including more samples, and higher dimensionality requires more points to reduce the distance to the same extent. Mathematically, we thus expect  $\alpha = 1/d_{\text{cor}}$ . A similar scaling law has been derived for the Farmer-Sidorowich forecasting method from the nonlinear dynamics community (Farmer & Sidorowich, 1987). Despite the theoretical correspondence, however, numerically the fractal dimension is challenging to accurately estimate from finite time series, due to instabilities in estimating the scaling coefficient  $\alpha$  as the data resolution and fitting conditions vary (Clauset et al., 2009; Grassberger & Procaccia, 1983). Nonetheless, in Fig. 11, we observe relatively strong correlation between  $d_{\text{cor}}$  and  $1/\alpha$ , supporting our theoretical argument above.

432 5.3 SCIML TASKS BEYOND LOW-DIMENSIONAL CHAOTIC SYSTEMS  
433

434 So far we focused on low-dimensional chaotic systems from the `dysts` dataset, which enabled  
435 systematic comparison between different forecasting models with standardized benchmarks. Here, we  
436 show that parrotting also outperforms foundation models on a broader class of SciML tasks, including  
437 real-world datasets of current scientific interest. Our datasets are: (1) the von Karman vortex street at  
438 Reynolds number  $Re = 900$ , a standard problem in fluid dynamics representing a flow exhibiting  
439 intermittency. We generated time series corresponding to the top PCA modes, in order to capture  
440 global structure; (2) electrocardiogram recordings (via the QT Database in PhysioNet); (3) 28 coupled  
441 electronic circuits measured experimentally from [Vera-Ávila et al. \(2020\)](#); and (4) 23 Kuramoto  
442 oscillators coupled through frustrated and nonreciprocal interactions, recently studied in [León & Pazó \(2025\)](#).  
443 These are all high-dimensional systems, two generated from simulations and two measured  
444 in the real world. For the metrics, we use MAE and MSE to measure pointwise forecast accuracy,  
445 and KL Divergence to measure the accuracy in attractor reconstruction. The results are summarized  
446 below. Parrotting is the only model that ranks in the top three for all tasks and all metrics. Other  
447 metrics, such as valid prediction time and fractal dimension accuracy, give similar rankings.  
448

449 Table 1: Performance comparison (**MAE @ 50 steps**, mean  $\pm$  standard deviation) of forecasting  
450 models across SciML tasks. **Bold** = best, *italic* = second and third best.

Task	Parrot	DynaMix	Chronos	Chronos Bolt	TimesFM	TimeMoE	Moirai
Turbulence	<i>0.403</i> $\pm$ 0.210	0.505 $\pm$ 0.247	0.431 $\pm$ 0.237	0.567 $\pm$ 0.247	0.510 $\pm$ 0.174	<i>0.394</i> $\pm$ 0.172	<b>0.382</b> $\pm$ <b>0.189</b>
ECG	<b>0.624</b> $\pm$ <b>0.315</b>	0.777 $\pm$ 0.241	0.873 $\pm$ 0.422	0.752 $\pm$ 0.279	<i>0.723</i> $\pm$ 0.259	0.799 $\pm$ 0.158	0.684 $\pm$ 0.237
Circuit	<b>0.083</b> $\pm$ <b>0.050</b>	0.425 $\pm$ 0.172	<i>0.111</i> $\pm$ 0.065	0.349 $\pm$ 0.120	<i>0.196</i> $\pm$ 0.090	0.206 $\pm$ 0.102	0.213 $\pm$ 0.093
Kuramoto	<b>0.004</b> $\pm$ <b>0.001</b>	0.076 $\pm$ 0.002	0.072 $\pm$ 0.029	0.961 $\pm$ 0.084	0.624 $\pm$ 0.061	<i>0.070</i> $\pm$ 0.011	<b>0.004</b> $\pm$ <b>0.001</b>

451  
452 Table 2: Performance comparison (**MSE @ 50 steps**) of forecasting models across SciML tasks.  
453 **Bold** = best, *italic* = second and third best.

Task	Parrot	DynaMix	Chronos	Chronos Bolt	TimesFM	TimeMoE	Moirai
Turbulence	<i>0.322</i> $\pm$ 0.333	0.490 $\pm$ 0.4530	0.380 $\pm$ 0.408	0.531 $\pm$ 0.447	0.403 $\pm$ 0.262	<b>0.278</b> $\pm$ <b>0.268</b>	<b>0.278</b> $\pm$ <b>0.267</b>
ECG	<i>0.916</i> $\pm$ 0.630	1.063 $\pm$ 0.488	1.461 $\pm$ 1.097	0.950 $\pm$ 0.581	0.940 $\pm$ 0.530	<i>0.893</i> $\pm$ 0.287	<b>0.851</b> $\pm$ <b>0.488</b>
Circuit	<b>0.012</b> $\pm$ <b>0.016</b>	0.297 $\pm$ 0.294	<i>0.024</i> $\pm$ 0.030	0.181 $\pm$ 0.122	<i>0.065</i> $\pm$ 0.056	0.076 $\pm$ 0.080	0.075 $\pm$ 0.060
Kuramoto	<b>0.001</b> $\pm$ <b>0.002</b>	0.006 $\pm$ 0.001	0.009 $\pm$ 0.007	1.296 $\pm$ 0.188	0.512 $\pm$ 0.096	<i>0.008</i> $\pm$ 0.002	<b>0.001</b> $\pm$ <b>0.001</b>

454  
455 Table 3: Performance comparison (**KL Divergence between predicted and true attractors**) of  
456 forecasting models across SciML tasks. **Bold** = best, *italic* = second and third best.

Task	Parrot	DynaMix	Chronos	Chronos Bolt	TimesFM	TimeMoE	Moirai
Turbulence	<i>0.028</i> $\pm$ 0.044	<b>0.005</b> $\pm$ <b>0.008</b>	0.041 $\pm$ 0.046	0.048 $\pm$ 0.058	0.111 $\pm$ 0.072	0.070 $\pm$ 0.058	<i>0.030</i> $\pm$ 0.041
ECG	<b>0.065</b> $\pm$ <b>0.089</b>	<i>0.099</i> $\pm$ 0.104	0.403 $\pm$ 0.367	0.253 $\pm$ 0.185	0.220 $\pm$ 0.153	<i>0.188</i> $\pm$ 0.094	0.276 $\pm$ 0.311
Circuit	<i>0.572</i> $\pm$ 0.082	2.940 $\pm$ 0.528	<i>0.630</i> $\pm$ 0.118	1.710 $\pm$ 0.255	<b>0.383</b> $\pm$ <b>0.087</b>	0.816 $\pm$ 0.200	0.848 $\pm$ 0.155
Kuramoto	<b>0.001</b> $\pm$ <b>0.001</b>	1.010 $\pm$ 0.150	0.537 $\pm$ 0.087	3.116 $\pm$ 0.202	4.489 $\pm$ 0.363	<i>0.076</i> $\pm$ 0.040	<i>0.010</i> $\pm$ 0.011

457 6 CONCLUSION AND FUTURE DIRECTIONS  
458

459 We find that a simple forecast strategy—context parrotting—outperforms leading foundation models  
460 on dynamical systems forecasting, a critical task in scientific machine learning. This surprising  
461 finding exposes a limitation of current time-series foundation models as general-purpose time-series  
462 forecasters and highlights the need to further scale them or to fine-tune them for specific domains. It  
463 also suggests that accurately measuring the performance of foundation models can be difficult for  
464 scientific machine learning tasks, because strategies like parrotting can effectively game both short-  
465 and long-term accuracy metrics.

466 Finding a simple but effective baseline for a challenging task can encourage rethinking of the status  
467 quo, motivating the development of better model architectures ([Arora et al., 2017](#)). For example,

486 context parroting formalizes an explicit baseline to compare against in the time-series domain and can  
 487 help discover beyond-parroting strategies. Identifying in-context learning strategies beyond parroting  
 488 can spur the development of next-generation foundation models and contribute to the debate on  
 489 whether (or to what extent) large language models are stochastic parrots (Bender et al., 2021; Mitchell  
 490 & Krakauer, 2023; Arora & Goyal, 2023; McCoy et al., 2024).

491 An interesting future direction is to generalize context parroting to deal with non-stationary time series  
 492 while keeping the simplicity and efficiency of the method. Context parroting assumes the existence  
 493 of a stationary underlying measure; for an ergodic deterministic system this implies that conditional  
 494 probabilities of timepoints are stationary up to any order (Appendix H). However, newer foundation  
 495 models readily handle simple nonstationarity like baseline drift, implying that a modified parroting  
 496 strategy may be possible in-context (Das et al., 2024). A promising avenue involves combining  
 497 parroting with a probabilistic model, such as a Gaussian Process, to account for nonstationary trends.  
 498 Additionally, the diversity of long-term forecasts can be improved by allowing stochastic selection  
 499 among multiple candidate forecasts (e.g., from different matching motifs). Once generalized, the  
 500 non-stationary parroting method can replace Naive and Seasonal Naive to serve as a more informative  
 501 baseline for the zero-shot forecasting of general time series (weather, traffic, finance, etc.).

502 Finally, we want to emphasize that we are not proposing to replace time-series foundation models  
 503 with context parroting. Instead, the value of parroting is as a simple baseline that can reveal the gaps  
 504 in current foundation models and guide the design of new ones. When foundation models under  
 505 perform relative to context parroting, it reveals that they haven't learned to fully utilize the context  
 506 data. For example, a common failure mode we observed across a range of leading foundation models  
 507 (TimesFM, TimeMoE, Chronos Bolt) is that they tend to underestimate oscillations in the dynamics  
 508 and the predictions often quickly converge to the mean (Fig. 5). Being aware of context parroting as  
 509 a baseline also guides the design of more informative benchmarks in the future. For example, for  
 510 dynamical systems in `dysts`, one can generate test context data that are impossible to parrot by  
 511 using only initial conditions that lead to trajectories without nearly repeating motifs.

## 512 7 REPRODUCIBILITY STATEMENT

513 A Python implementation of the context parroting algorithm and the benchmarks are available at  
 514 <https://anonymous.4open.science/r/parroting-4D26>.

## 518 REFERENCES

519 Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen,  
 520 Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al.  
 521 Chronos: Learning the language of time series. *arXiv:2403.07815*, 2024.

523 Abdul Fatir Ansari, Oleksandr Shchur, Jaris Küken, Andreas Auer, Boran Han, Pedro Mercado,  
 524 Syama Sundar Rangapuram, Huibin Shen, Lorenzo Stella, Xiyuan Zhang, et al. Chronos-2: From  
 525 univariate to universal forecasting. *arXiv preprint arXiv:2510.15821*, 2025.

526 Sanjeev Arora and Anirudh Goyal. A theory for emergence of complex skills in language models.  
 527 *arXiv:2307.15936*, 2023.

529 Sanjeev Arora, Yingyu Liang, and Tengyu Ma. A simple but tough-to-beat baseline for sentence  
 530 embeddings. In *International conference on learning representations*, 2017.

532 Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and  
 533 Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nat. Rev.  
 534 Phys.*, pp. 1–9, 2024.

535 Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural  
 536 scaling laws. *Proc. Natl. Acad. Sci. U.S.A.*, 121(27):e2311878121, 2024.

538 Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the  
 539 dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM  
 conference on fairness, accountability, and transparency*, pp. 610–623, 2021.

540 Tyrus Berry and Sudhansu Das. Limits of learning dynamical systems. *SIAM Review*, 67(1):  
 541 107–137, 2025.  
 542

543 Cristian Bodnar, Wessel P Bruinsma, Ana Lucic, Megan Stanley, Johannes Brandstetter, Patrick  
 544 Garvan, Maik Riechert, Jonathan Weyn, Haiyu Dong, Anna Vaughan, et al. Aurora: A foundation  
 545 model of the atmosphere. *arXiv:2405.13063*, 2024.

546 Richard C Bradley. Basic properties of strong mixing conditions. a survey and some open questions.  
 547 *Probability Surveys*, 2:107–144, 2005.  
 548

549 Manuel Brenner, Florian Hess, Jonas M Mikhail, Leonard F Bereska, Zahra Monfared, Po-Chen  
 550 Kuo, and Daniel Durstewitz. Tractable dendritic rnns for reconstructing nonlinear dynamical  
 551 systems. In *International conference on machine learning*, pp. 2292–2320. Pmlr, 2022.

552 Manuel Brenner, Christoph Jürgen Hemmer, Zahra Monfared, and Daniel Durstewitz. Almost-linear  
 553 rnns yield highly interpretable symbolic codes in dynamical systems reconstruction. *Advances in  
 554 Neural Information Processing Systems*, 37:36829–36868, 2024a.  
 555

556 Manuel Brenner, Elias Weber, Georgia Koppe, and Daniel Durstewitz. Learning interpretable  
 557 hierarchical dynamical systems models from time series data. *arXiv preprint arXiv:2410.04814*,  
 558 2024b.

559 Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data  
 560 by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. U.S.A.*, 113(15):  
 561 3932–3937, 2016.  
 562

563 Steven L Brunton, Marko Budisić, Eurika Kaiser, and J Nathan Kutz. Modern Koopman theory for  
 564 dynamical systems. *SIAM Rev.*, 64(2):229–340, 2022.

565 Josue Ortega Caro, Antonio H de O Fonseca, Christopher Averill, Syed A Rizvi, Matteo Rosati,  
 566 James L Cross, Prateek Mittal, Emanuele Zappala, Daniel Levine, Rahul M Dhodapkar, et al.  
 567 Brainlm: A foundation model for brain activity recordings. *bioRxiv*, pp. 2023–09, 2023.  
 568

569 Chun-Wei Chang, Masayuki Ushio, and Chih-hao Hsieh. Empirical dynamic modeling for beginners.  
 570 *Ecological research*, 32:785–796, 2017.

571 Renyi Chen and Molei Tao. Data-driven prediction of general hamiltonian dynamics via learning  
 572 exactly-symplectic maps. In *International Conference on Machine Learning*, pp. 1717–1727.  
 573 PMLR, 2021.

574 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary  
 575 differential equations. *NeurIPS*, 31, 2018.

577 Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. Unveiling induction heads: Provable  
 578 training dynamics and feature learning in transformers. In *The Thirty-eighth Annual Conference  
 579 on Neural Information Processing Systems*, 2024a.

580 Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. Parallel structures in pre-training  
 581 data yield in-context learning. In *62nd Annual Meeting of the Association for Computational  
 582 Linguistics, ACL 2024*, pp. 8582–8592. Association for Computational Linguistics (ACL), 2024b.

584 Xiaoyuan Cheng, Yi He, Yiming Yang, Xiao Xue, Sibo Chen, Daniel Giles, Xiaohang Tang, and  
 585 Yukun Hu. Learning chaos in a linear way. *arXiv preprint arXiv:2503.14702*, 2025.

587 Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical  
 588 data. *SIAM review*, 51(4):661–703, 2009.

589 Wenhui Cui, Woojae Jeong, Philipp Thölke, Takfarinas Medani, Karim Jerbi, Anand A Joshi, and  
 590 Richard M Leahy. Neuro-gpt: Towards a foundation model for eeg. In *2024 IEEE International  
 591 Symposium on Biomedical Imaging (ISBI)*, pp. 1–5. IEEE, 2024.

593 Abhimanyu Das, Weihao Kong, Andrew Leach, Shaan Mathur, Rajat Sen, and Rose Yu. Long-term  
 forecasting with tide: Time-series dense encoder. *arXiv:2304.08424*, 2023.

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 Samuel Dooley, Gurnoor Singh Khurana, Chirag Mohapatra, Siddartha V Naidu, and Colin White.  
 598 Forecastpfn: Synthetically-trained zero-shot forecasting. *Advances in Neural Information Processing  
 599 Systems*, 36, 2024.

600 Ezra Edelman, Nikolaos Tsilivis, Benjamin Edelman, Eran Malach, and Surbhi Goel. The evolution  
 601 of statistical induction heads: In-context learning markov chains. *Advances in Neural Information  
 602 Processing Systems*, 37:64273–64311, 2024.

603 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda  
 604 Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer  
 605 circuits. *Transformer Circuits Thread*, 1(1):12, 2021.

606

607 Jianqing Fan and Qiwei Yao. *Nonlinear time series: nonparametric and parametric methods*. Springer  
 608 Science & Business Media, 2008.

609 J Doyne Farmer and John J Sidorowich. Predicting chaotic time series. *Physical review letters*, 59(8):  
 610 845, 1987.

611

612 Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn  
 613 in-context? a case study of simple function classes. *Advances in neural information processing  
 614 systems*, 35:30583–30598, 2022.

615 Azul Garza and Max Mergenthaler-Canseco. Timegpt-1. *arXiv:2310.03589*, 2023.

616

617 Daniel J Gauthier, Erik Bollt, Aaron Griffith, and Wendson AS Barbosa. Next generation reservoir  
 618 computing. *Nat. Commun.*, 12:5564, 2021.

619

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

622

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

625

626 Niclas Göring, Florian Hess, Manuel Brenner, Zahra Monfared, and Daniel Durstewitz. Out-of-  
 627 domain generalization in dynamical systems reconstruction. *arXiv:2402.18377*, 2024.

628

629 Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski.  
 630 Moment: A family of open time-series foundation models. *arXiv:2402.03885*, 2024.

631

632 Peter Grassberger and Itamar Procaccia. Measuring the strangeness of strange attractors. *Physica D:  
 633 nonlinear phenomena*, 9(1-2):189–208, 1983.

634

635 Lyudmila Grigoryeva, Hannah Lim Jing Ting, and Juan-Pablo Ortega. Infinite-dimensional next-  
 636 generation reservoir computing. *Physical Review E*, 111(3):035305, 2025.

637

638 Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew G Wilson. Large language models are zero-shot  
 639 time series forecasters. *Advances in Neural Information Processing Systems*, 36, 2024.

640

641 Yi He, Yiming Yang, Xiaoyuan Cheng, Hai Wang, Xiao Xue, Boli Chen, and Yukun Hu. Chaos meets  
 642 attention: Transformers for large-scale dynamical prediction. *arXiv preprint arXiv:2504.20858*,  
 643 2025.

644

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

647

648 Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Em-  
 649 manuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes.  
 650 *arXiv:2405.19101*, 2024.

651

652 John R Hershey and Peder A Olsen. Approximating the kullback leibler divergence between  
 653 gaussian mixture models. In *2007 IEEE International Conference on Acoustics, Speech and Signal  
 654 Processing-ICASSP'07*, volume 4, pp. IV–317. IEEE, 2007.

648 Florian Hess, Zahra Monfared, Manuel Brenner, and Daniel Durstewitz. Generalized teacher forcing  
 649 for learning chaotic dynamics. In *International Conference on Machine Learning*, pp. 13017–  
 650 13049. PMLR, 2023.

651

652 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza  
 653 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al.  
 654 Training compute-optimal large language models. *arXiv:2203.15556*, 2022.

655 Jeremy P Huke. Embedding nonlinear dynamical systems: A guide to takens' theorem. 2006.

656

657 Rob J Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018.

658

659 Rob J Hyndman and Anne B Koehler. Another look at measures of forecast accuracy. *International  
 660 journal of forecasting*, 22(4):679–688, 2006.

661

662 Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yux-  
 663 uan Liang, Yuan-Fang Li, Shirui Pan, et al. Time-LLM: Time series forecasting by reprogramming  
 664 large language models. *arXiv:2310.01728*, 2023.

665

666 Ian D Jordan, Piotr Aleksander Sokół, and Il Memming Park. Gated recurrent units viewed through  
 667 the lens of continuous time dynamical systems. *Frontiers in computational neuroscience*, 15:  
 668 678158, 2021.

669

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

673

674 George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang.  
 675 Physics-informed machine learning. *Nat. Rev. Phys.*, 3(6):422–440, 2021.

676

677 Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher.  
 678 Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint  
 679 arXiv:1909.05858*, 2019.

680

681 Aditi S Krishnapriyan, Alejandro F Queiruga, N Benjamin Erichson, and Michael W Mahoney.  
 682 Learning continuous models for continuous physics. *Communications Physics*, 6(1):319, 2023.

683

684 Iván León and Diego Pazó. Dynamics and chaotic properties of the fully disordered kuramoto model.  
 685 *Chaos*, 35(7), 2025.

686

687 Matthew Levine and Andrew Stuart. A framework for machine learning of model error in dynamical  
 688 systems. *Commun. Am. Math. Soc.*, 2(07):283–344, 2022.

689

690 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew  
 691 Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations.  
 692 *arXiv:2010.08895*, 2020.

693

694 Yuxuan Liang, Haomin Wen, Yuqi Nie, Yushan Jiang, Ming Jin, Dongjin Song, Shirui Pan, and Qing-  
 695 song Wen. Foundation models for time series analysis: A tutorial and survey. *arXiv:2403.14735*,  
 696 2024.

697

698 Bryan Lim and Stefan Zohren. Time-series forecasting with deep learning: a survey. *Philosophical  
 699 Transactions of the Royal Society A*, 379(2194):20200209, 2021.

700

701 Chenghao Liu, Taha Aksu, Juncheng Liu, Xu Liu, Hanshu Yan, Quang Pham, Doyen Sahoo, Caiming  
 702 Xiong, Silvio Savarese, and Junnan Li. Moirai 2.0: When less is more for time series forecasting.  
 703 *arXiv:2511.11698*, 2025a.

704

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

707

708 Yong Liu, Guo Qin, Xiangdong Huang, Jianmin Wang, and Mingsheng Long. Autotimes: Autore-  
 709 gressive time series forecasters via large language models. *arXiv:2402.02370*, 2024b.

702 Yong Liu, Guo Qin, Zhiyuan Shi, Zhi Chen, Caiyin Yang, Xiangdong Huang, Jianmin Wang,  
 703 and Mingsheng Long. Sundial: A family of highly capable time series foundation models.  
 704 *arXiv:2502.00816*, 2025b.

705 Ang Lv, Ruobing Xie, Xingwu Sun, Zhanhui Kang, and Rui Yan. Language models" grok" to copy.  
 706 *arXiv preprint arXiv:2409.09281*, 2024.

708 Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. The m5 competition:  
 709 Background, organization, and implementation. *International Journal of Forecasting*, 38(4):  
 710 1325–1336, 2022.

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

715 R Thomas McCoy, Shunyu Yao, Dan Friedman, Mathew D Hardy, and Thomas L Griffiths. Embers  
 716 of autoregression show how large language models are shaped by the problem they are trained to  
 717 solve. *Proc. Natl. Acad. Sci. U.S.A.*, 121(41):e2322420121, 2024.

719 Kaden McKeen, Laura Oliva, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. Ecg-fm:  
 720 An open electrocardiogram foundation model. *arXiv preprint arXiv:2408.05178*, 2024.

721 Jonas Mikhaeil, Zahra Monfared, and Daniel Durstewitz. On the difficulty of learning chaotic  
 722 dynamics with rnns. *Advances in Neural Information Processing Systems*, 35:11297–11312, 2022.

724 John A Miller, Mohammed Aldosari, Farah Saeed, Nasid Habib Barna, Subas Rana, I Budak Arpinar,  
 725 and Ninghao Liu. A survey of deep learning and foundation models for time series forecasting.  
 726 *arXiv:2401.13912*, 2024.

727 Melanie Mitchell and David C Krakauer. The debate over understanding in ai's large language  
 728 models. *Proc. Natl. Acad. Sci. U.S.A.*, 120(13):e2215907120, 2023.

730 Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, and Aditya Grover. Climax:  
 731 A foundation model for weather and climate. *arXiv:2301.10343*, 2023.

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

735 Boris N Oreshkin, Dmitri Carpow, Nicolas Chapados, and Yoshua Bengio. Meta-learning framework  
 736 with applications to zero-shot time-series forecasting. In *Proceedings of the AAAI conference on  
 737 artificial intelligence*, volume 35, pp. 9242–9250, 2021.

739 Jaideep Pathak, Brian Hunt, Michelle Girvan, Zhixin Lu, and Edward Ott. Model-free prediction of  
 740 large spatiotemporally chaotic systems from data: A reservoir computing approach. *Phys. Rev.  
 741 Lett.*, 120(2):024102, 2018.

742 Md Ashiqur Rahman, Robert Joseph George, Mogab Elleithy, Daniel Leibovici, Zongyi Li, Boris  
 743 Bonev, Colin White, Julius Berner, Raymond A Yeh, Jean Kossaifi, et al. Pretraining codomain  
 744 attention neural operators for solving multiphysics pdes. *arXiv:2403.12553*, 2024.

746 Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Arian Khorasani, George Adamopoulos,  
 747 Rishika Bhagwatkar, Marin Bilos, Hena Ghonia, Nadhir Vincent Hassen, Anderson Schneider,  
 748 et al. Lag-llama: Towards foundation models for time series forecasting. *arXiv:2310.08278*, 2023.

749 Gautam Reddy. The mechanistic basis of data dependence and abrupt learning in an in-context  
 750 classification task. *arXiv:2312.03002*, 2023.

751 Matthew Ricci, Guy Pelc, Zoe Piran, Noa Moriel, and Mor Nitzan. Trendy: Temporal regression of  
 752 effective non-linear dynamics. *arXiv preprint arXiv:2412.03496*, 2024.

754 Michael T Rosenstein, James J Collins, and Carlo J De Luca. A practical method for calculating  
 755 largest lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65(1-2):  
 117–134, 1993.

756 Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, and Ming Jin. Time-  
 757 moe: Billion-scale time series foundation models with mixture of experts. In *The Thirteenth*  
 758 *International Conference on Learning Representations*, 2025.

759

760 Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. Beyond neural  
 761 scaling laws: beating power law scaling via data pruning. *Advances in Neural Information*  
 762 *Processing Systems*, 35:19523–19536, 2022.

763 Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W  
 764 Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Charac-  
 765 terizing scaling and transfer behavior. *Advances in Neural Information Processing Systems*, 36,  
 766 2024.

767

768 George Sugihara. Nonlinear forecasting for the classification of natural time series. *Philosophical*  
 769 *Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 348  
 770 (1688):477–495, 1994.

771

772 George Sugihara and Robert M May. Nonlinear forecasting as a way of distinguishing chaos from  
 measurement error in time series. *Nature*, 344(6268):734–741, 1990.

773

774 Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk  
 775 Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning.  
 776 *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.

777

778 Floris Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence*,  
 779 *Warwick 1980: proceedings of a symposium held at the University of Warwick 1979/80*, pp.  
 366–381. Springer, 2006.

780

781 Kunio Takezawa. *Introduction to nonparametric regression*. John Wiley & Sons, 2005.

782

783 Amin Totounferoush, Serge Kotchourko, Michael W Mahoney, and Steffen Staab. Paving the way  
 784 for scientific foundation models: enhancing generalization and robustness in pdes with constraint-  
 aware pre-training. *arXiv preprint arXiv:2503.19081*, 2025.

785

786 Yao-Hung Hubert Tsai, Shaojie Bai, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhut-  
 787 dinov. Transformer dissection: An unified understanding for transformer’s attention via the lens of  
 788 kernel. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*,  
 2019.

789

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

793

794 Pantelis R Vlachas, Jaideep Pathak, Brian R Hunt, Themistoklis P Sapsis, Michelle Girvan, Edward  
 795 Ott, and Petros Koumoutsakos. Backpropagation algorithms and reservoir computing in recurrent  
 796 neural networks for the forecasting of complex spatiotemporal dynamics. *Neural. Netw.*, 126:  
 191–217, 2020.

797

798 Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev,  
 799 Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In  
 800 *International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.

801

802 P Walters. An introduction to ergodic theory. *Springer-Verlag*, 1982.

803

804 E Weinan. A proposal on machine learning via dynamical systems. *Commun. Math. Stat.*, 1(5):1–11,  
 2017.

805

806 Gerald Woo, Chenghao Liu, Akshat Kumar, Caiming Xiong, Silvio Savarese, and Doyen Sahoo.  
 807 Unified training of universal time series forecasting transformers. *arXiv:2402.02592*, 2024.

808

809 Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning with data  
 810 prompts for differential equation problems. *Proc. Natl. Acad. Sci. U.S.A.*, 120(39):e2310142120,  
 2023.

810 Lu Yang, Xiuwen Sun, Boumediene Hamzi, Houman Owhadi, and Naiming Xie. Learning dynamical  
 811 systems from data: A simple cross-validation perspective, part v: Sparse kernel flows for 132  
 812 chaotic dynamical systems. *Physica D: Nonlinear Phenomena*, 460:134070, 2024.

813

814 Qingren Yao, Chao-Han Huck Yang, Renhe Jiang, Yuxuan Liang, Ming Jin, and Shirui Pan. Towards  
 815 neural scaling laws for time series foundation models. *arXiv:2410.12360*, 2024.

816

817 Annan Yu, Danielle C Maddix, Boran Han, Xiyuan Zhang, Abdul Fatir Ansari, Oleksandr Shchur,  
 818 Christos Faloutsos, Andrew Gordon Wilson, Michael W Mahoney, and Yuyang Wang. Under-  
 819 standing the implicit biases of design choices for time series foundation models. *arXiv preprint*  
 820 *arXiv:2510.19236*, 2025.

821

822 Rose Yu and Rui Wang. Learning dynamical systems from data: An introduction to physics-guided  
 823 deep learning. *Proc. Natl. Acad. Sci. U.S.A.*, 121(27):e2311808121, 2024.

824

825 Oussama Zekri, Ambroise Odonnat, Abdelhakim Benechehab, Linus Bleistein, Nicolas Boullé, and  
 826 Ievgen Redko. Large language models as markov chains. *arXiv preprint arXiv:2410.02724*, 2024.

827

828 Zheng-Meng Zhai, Jun-Yin Huang, Benjamin D Stern, and Ying-Cheng Lai. Reconstructing dynamics  
 829 from sparse observations with no training on target system. *arXiv:2410.21222*, 2024.

830

831 Yuanzhao Zhang and William Gilpin. Zero-shot forecasting of chaotic systems. *arXiv:2409.15771*,  
 832 2024.

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

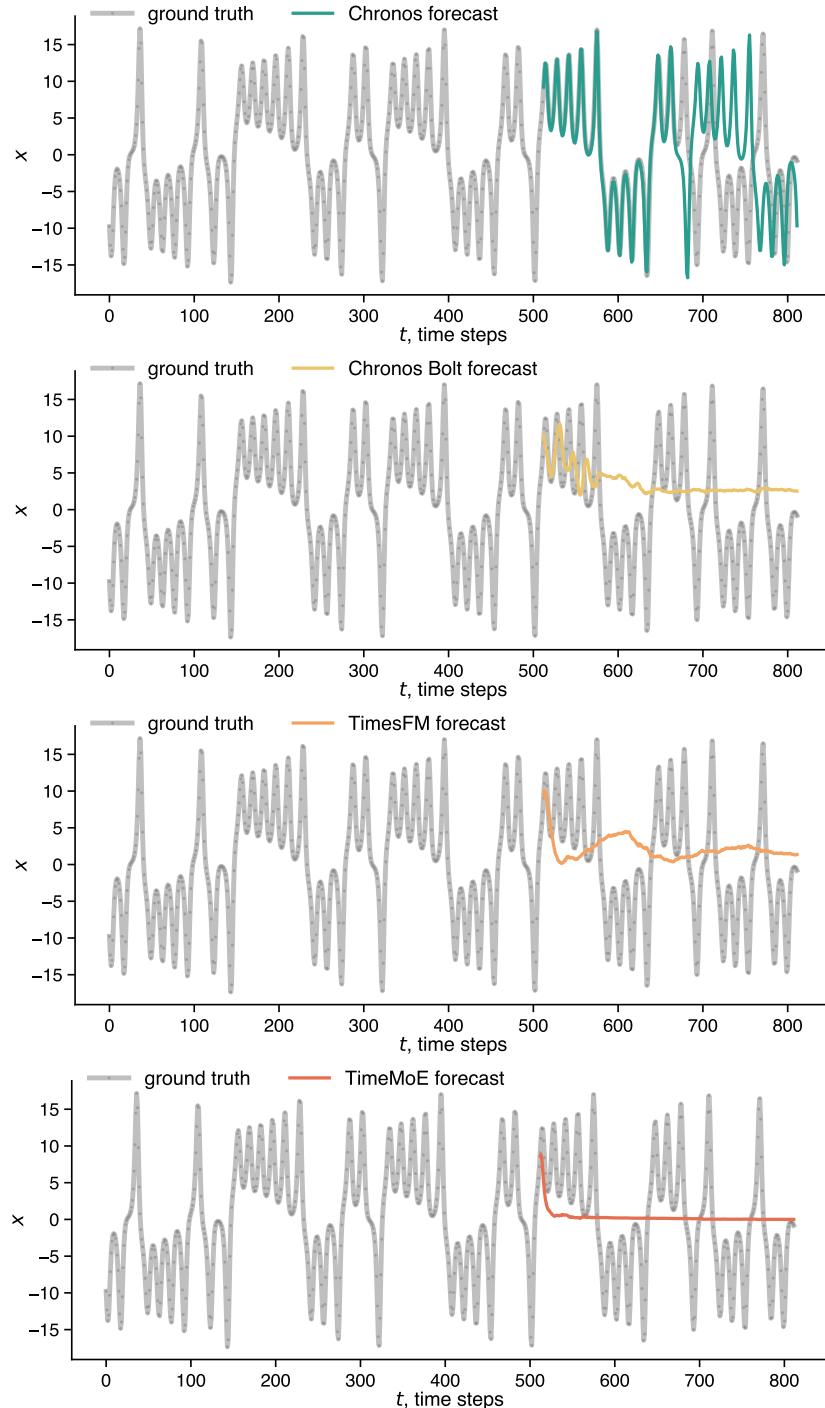
860

861

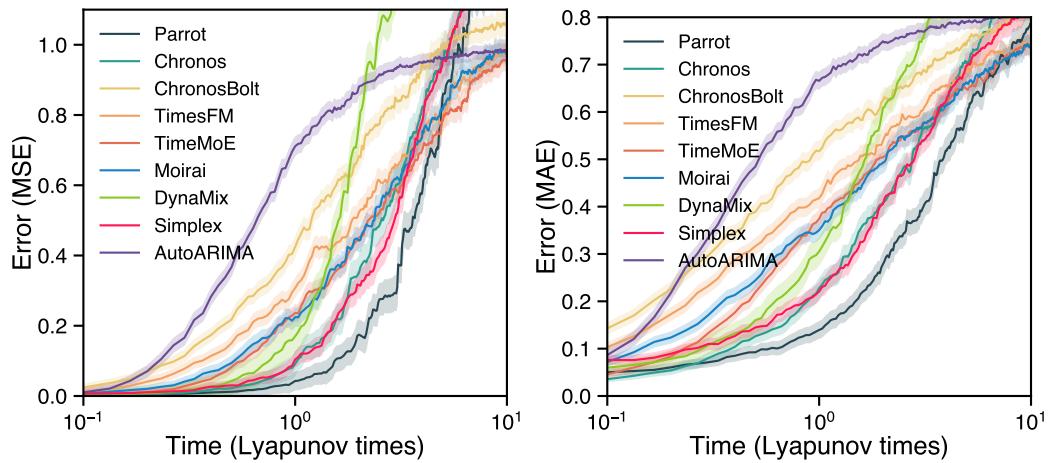
862

863

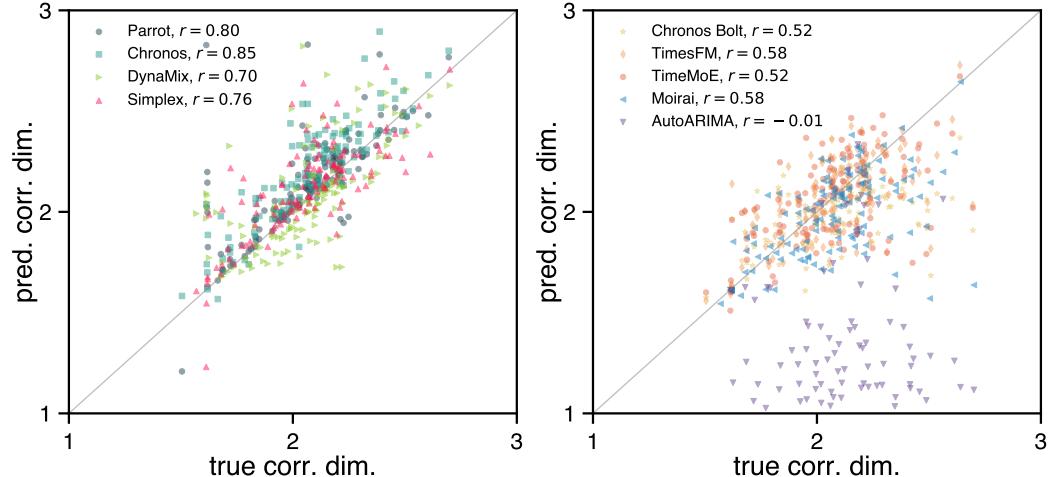
## A SAMPLE PREDICTIONS FROM FOUNDATION MODELS



**Figure 5: Example forecasts on a chaotic system from foundation models reveal common failure modes.** This is the same task as presented in Fig. 1 (predicting the  $x$  variable of the Lorenz system based on a short context trajectory with 512 data points). Chronos does extremely well with a parrotting strategy. The other models perform comparatively poorly and all exhibit a tendency to underestimate the oscillations (e.g., by quickly converging towards the mean). This is a general trend that we consistently observe across different chaotic systems and initial conditions.

918  
919  
920  
921 B BENCHMARKING WITH OTHER METRICS  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936

937 **Figure 6: Context parrotting outperforms foundation models in zero-shot forecasting.** Same  
938 setup as in Fig. 2, but with the forecast error measured by MSE (left) and MAE (right). On top of the  
939 foundation models, we also include two classical forecasting methods in the comparison: simplex  
940 projection (Sugihara & May, 1990) and AutoARIMA (Hyndman & Athanasopoulos, 2018).



959 **Figure 7: Fractal dimension accuracy for parrotting and foundation models.** Each point represents  
960 the predicted fractal dimension of a chaotic attractor by a model (median of 20 predictions from  
961 random initial conditions). The accuracy is measured by the Spearman correlation  $r$  between the 135  
962 predicted fractal dimensions and the true fractal dimensions.

963  
964  
965 C EFFECTS OF EMBEDDING DIMENSION  $D$   
966  
967

968 Fig. 10 investigates how the choice of the embedding dimension  $D$  affects the performance of  
969 context parrotting. Overall, the valid prediction time stays consistent over a wide range of embedding  
970 dimension  $D$ . For short context windows, there is a slight advantage to small  $D$ . For long context  
971 windows, larger embedding dimensions are marginally better. This observation suggests potential  
improvements in the future that choose  $D$  adaptively based on factors such as context length.

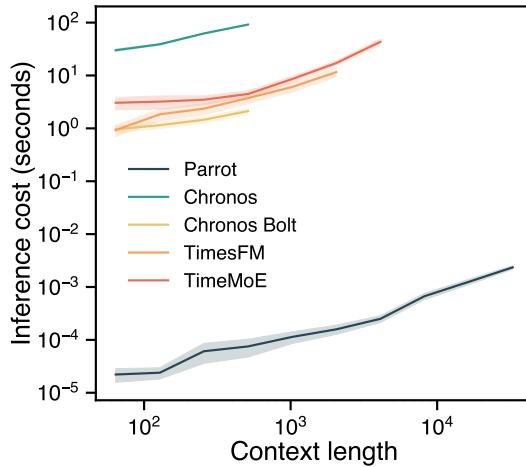


Figure 8: **Context parrotting is computationally more efficient than foundation models.** There is generally a gap of five or six orders of magnitude between context parrotting and foundation models. For each foundation model, context lengths from 64 to the maximum context window are considered. All inferences are performed with CPUs, and the forecast horizon is fixed to 300 steps.

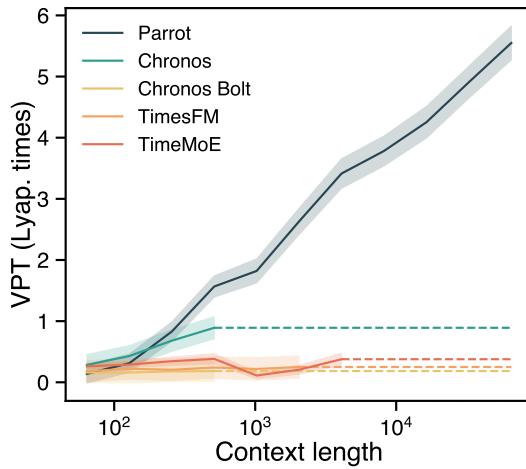


Figure 9: **Valid prediction time as a function of the context length.** Dashed lines indicate context lengths exceeding the maximum context window of the model. Results are averaged over 135 chaotic systems in the `dysts` database, with 20 trajectories from random initial conditions for each system.

## D PREDICTING SCALING COEFFICIENT WITH THE ATTRACTOR DIMENSION

Figure 11 shows the relation between the scaling coefficient of the in-context neural scaling law and the correlation dimension of the chaotic attractor.

## E EFFECT OF CONTEXT LENGTH ON INVARIANT PROPERTIES

In Figure 12, we explore how the accuracy of four representative zero-shot models changes as the context length varies over an order of magnitude. We observe that all models exhibit monotonic scaling, consistent with additional samples from the attractor enabled by larger context acting to improve the stability of reconstructed attractors.

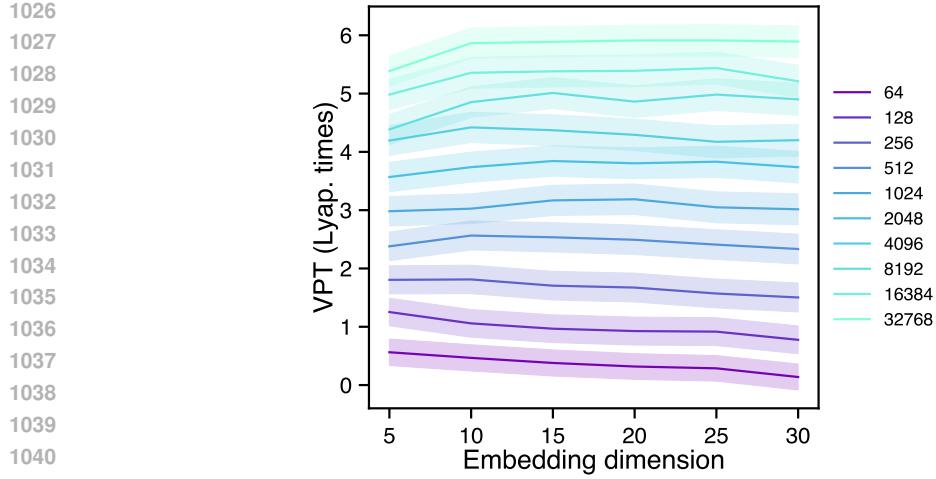


Figure 10: **Effect of the embedding dimension  $D$  on the forecast accuracy of context parroting.** Each curve represents a different context length. Results are averaged over 135 chaotic systems in the dysts database, with 100 trajectories from random initial conditions for each system.

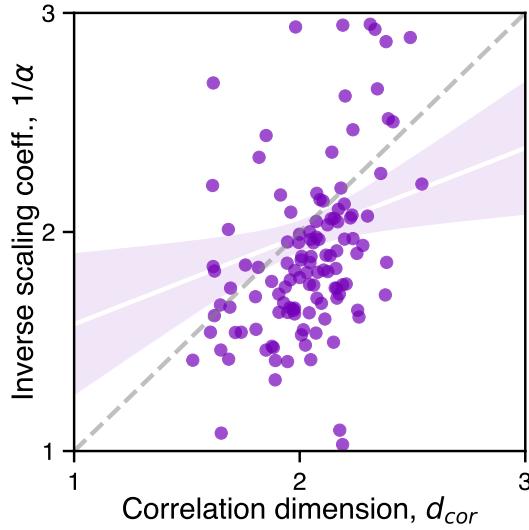


Figure 11: **Inverse scaling coefficient  $1/\alpha$  versus fractal dimension  $d_{\text{cor}}$  of the chaotic attractor.** We estimate  $\alpha$  from the  $\ell$  vs  $L$  plot obtained with context parroting at  $D = 10$ , where  $L$  ranges from  $2^6$  to  $2^{16}$ . This is done separately for each chaotic system. Thus, each dot represents one of the 135 chaotic systems included in our benchmark. The Spearman rank-order correlation coefficient is  $0.51 \pm 0.08$  (bootstrapped standard error), and a linear regression with standard error range is underscored (shaded interval). The dashed line corresponds to an exact  $1 : 1$  scaling.

## F EFFECT OF FORECAST HORIZON ON INVARIANT PROPERTIES

We next test the performance of parroting for long forecast horizons. We fix the context length  $L = 512$  and then generate forecasts of length  $H = 10000 - 512 = 9488$  (equivalent to over 316 Lyapunov Times). Table 4 shows the results of generating forecasts using the best-performing models from our shorter-horizon experiments. For each model, we evaluate its global accuracy by calculating (1) the correlation between the fractal dimension of the long forecast, and an estimate generated from the ground truth; (2) the correlation between the largest Lyapunov exponent of the long forecast, and an estimate generated from the ground truth; and (3) the attractor KL-divergence between the long forecast and ground truth (Grassberger & Procaccia, 1983; Rosenstein et al., 1993; Hess et al.,

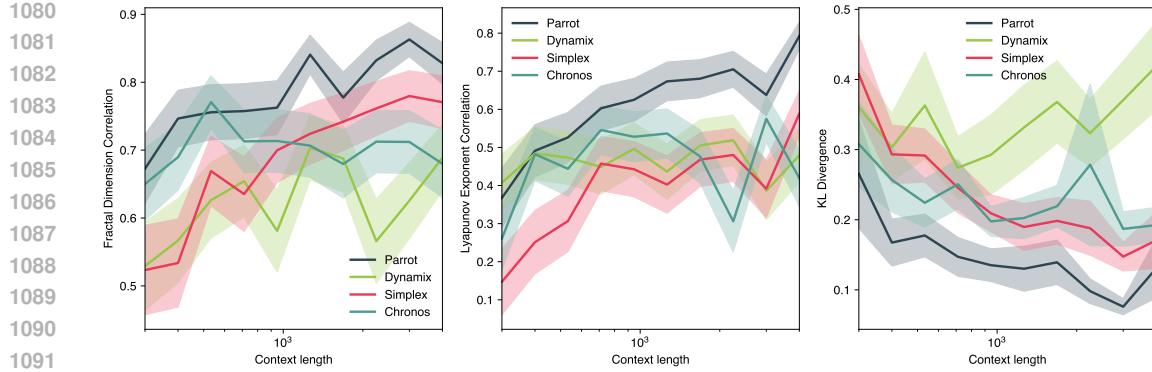


Figure 12: **Dependence of invariant properties on context length.** (Left) The Spearman correlation between the fractal dimension of the true trajectory and predicted trajectory as the context length varies. (Middle) The Spearman correlation between the estimated maximum Lyapunov exponent of the true trajectory and predicted trajectory as the context length varies. (Right) The estimated KL divergence between Gaussian mixtures fit to the true trajectory and predicted trajectory, as the context length varies. Note that Chronos has a finite context length, and so its performance saturates due to architectural constraints, and not intrinsic limitations. For this plot, forecasts are generated for 300 timepoints past the context. For the same figure with longer forecasts, see Figure 13.

2023). We find that context parrotting and DynaMix both perform well, indicating that DynaMix is a general-purpose time-series foundation model with a unique ability to capture the long-term climate of chaotic systems (Hemmer & Durstewitz, 2025).

Metric	Parrot	Chronos	Dynamix	Simplex
Attractor KL Divergence	<b><math>0.412 \pm 0.141</math></b>	$0.679 \pm 0.101$	$0.508 \pm 0.147$	$0.546 \pm 0.140$
Fractal Dimension Correlation	<b><math>0.723 \pm 0.042</math></b>	$0.120 \pm 0.118$	$0.521 \pm 0.057$	$0.341 \pm 0.072$
Largest Lyapunov Correlation	$0.343 \pm 0.018$	$0.269 \pm 0.114$	<b><math>0.466 \pm 0.071</math></b>	$0.343 \pm 0.085$

Table 4: KL Divergence and correlation of invariant properties between predicted and true attractors for different models for long forecast horizons. Error bars are standard deviation across all attractors for the KL Divergence, and uncertainty bounds based on the p-value for correlations. Bold = best, italic = second best.

## G EFFECTS OF NOISE AND SAMPLING RATE

We add Gaussian noise to normalized chaotic trajectories and repeat the original experiments (a noise level of 0.1 translates to 10% perturbation on each data point on average). The results are consistent across different orders of magnitude in noise, and parrotting is consistently the best or the second best in all experiments.

Noise level	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
$10^{-3}$	<b><math>2.17 \pm 0.19</math></b>	$1.68 \pm 0.18$	$0.79 \pm 0.18$	$1.07 \pm 0.19$	$0.92 \pm 0.15$
$10^{-2}$	<b><math>2.10 \pm 0.18</math></b>	$1.65 \pm 0.18$	$0.79 \pm 0.18$	$1.05 \pm 0.19$	$0.92 \pm 0.16$
$10^{-1}$	<b><math>1.04 \pm 0.14</math></b>	$0.89 \pm 0.15$	$0.71 \pm 0.17$	$0.89 \pm 0.17$	$0.66 \pm 0.11$

Table 5: **Valid prediction time** across noise levels (higher is better). Bold = best, italic = second best. Shading highlights best (dark) and second best (light).

In the main text, we set an intermediate granularity of 30 points per Lyapunov time. Below we compare it with results obtained for granularities of 10 points per Lyapunov time and 50 points per

Noise level	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
$10^{-3}$	<b>0.233<math>\pm</math>0.221</b>	0.297 $\pm$ 0.243	0.491 $\pm$ 0.223	0.440 $\pm$ 0.228	0.377 $\pm$ 0.211
$10^{-2}$	<b>0.235<math>\pm</math>0.220</b>	0.311 $\pm$ 0.245	0.492 $\pm$ 0.223	0.441 $\pm$ 0.228	0.383 $\pm$ 0.211
$10^{-1}$	<b>0.286<math>\pm</math>0.209</b>	0.366 $\pm$ 0.235	0.509 $\pm$ 0.219	0.455 $\pm$ 0.224	0.415 $\pm$ 0.216

Table 6: **MAE @ 1 Lyapunov Time** across noise levels (lower is better). **Bold** = best, *italic* = second best. Shading highlights best (dark) and second best (light).

Noise level	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
$10^{-3}$	<b>0.183<math>\pm</math>0.339</b>	0.268 $\pm$ 0.367	0.473 $\pm$ 0.314	0.394 $\pm$ 0.307	0.315 $\pm$ 0.273
$10^{-2}$	<b>0.185<math>\pm</math>0.340</b>	0.282 $\pm$ 0.361	0.474 $\pm$ 0.314	0.394 $\pm$ 0.306	0.318 $\pm$ 0.274
$10^{-1}$	<b>0.220<math>\pm</math>0.346</b>	0.328 $\pm$ 0.373	0.489 $\pm$ 0.314	0.407 $\pm$ 0.302	0.349 $\pm$ 0.285

Table 7: **MSE @ 1 Lyapunov Time** across noise levels (lower is better). **Bold** = best, *italic* = second best. Shading highlights best (dark) and second best (light).

Noise level	Parrot	Chronos	Chronos	Bolt	TimesFM	TimeMoE
$10^{-3}$	0.73	<b>0.85</b>		0.52	0.58	0.52
$10^{-2}$	0.63	<b>0.77</b>		0.51	0.55	0.45
$10^{-1}$	<b>0.59</b>	0.57		0.37	0.49	0.16

Table 8: **Fractal dimension accuracy** (Spearman correlation) across noise levels (higher is better). **Bold** = best, *italic* = second best. Shading highlights best (dark) and second best (light).

Noise level	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
$10^{-3}$	<b>0.113±0.205</b>	0.173±0.209	0.346±0.297	0.345±0.298	0.354±0.290
$10^{-2}$	<b>0.115±0.201</b>	0.189±0.244	0.344±0.292	0.356±0.298	0.344±0.285
$10^{-1}$	<b>0.141±0.207</b>	0.218±0.263	0.382±0.314	0.389±0.306	0.433±0.306

Table 9: **KL Divergence** between predicted and true attractors across noise levels (lower is better). **Bold** = best, *italic* = second best. Shading highlights best (dark) and second best (light).

Lyapunov time. Granularity does not strongly affect the results or relative model ranking. Parroting is either the best or the second best in all experiments. This makes sense, as we would expect changing granularity to have a similar effect as rescaling of time (although with bigger or smaller gaps between data points). For example, if we use finer granularity by a factor of 2, then we would need to double the context length to get the same lookback window.

Granularity	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
10	<b>4.70±0.57</b>	3.93±0.59	1.55±0.50	1.92±0.54	1.43±0.26
30	<b>2.15±0.19</b>	1.68±0.18	0.79±0.18	1.07±0.19	0.92±0.15
50	<b>1.41±0.11</b>	1.12±0.11	0.55±0.10	0.79±0.11	0.54±0.05

Table 10: **Valid prediction time** across different granularities (higher is better). Bold = best, italic = second best.

## H THEORETICAL PROPERTIES OF CONTEXT PARROTING

## H.1 OVERVIEW

**Mathematical Formulation.** Context parroting corresponds to a continuous 1-nearest-neighbor search over sequences of length  $D$  in the context of length  $L$ . It thus corresponds to a limit of a

Granularity	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
10	<b>0.219±0.204</b>	<i>0.316±0.256</i>	0.567±0.226	0.481±0.218	0.414±0.232
30	<b>0.233±0.221</b>	<i>0.297±0.243</i>	0.491±0.223	0.440±0.228	0.377±0.211
50	<b>0.270±0.226</b>	<i>0.329±0.241</i>	0.527±0.224	0.448±0.235	0.412±0.193

Table 11: **MAE @ 1 Lyapunov Time** across different granularities (lower is better). Bold = best, italic = second best.

Granularity	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
10	<b>0.163±0.311</b>	<i>0.291±0.364</i>	0.565±0.314	0.429±0.293	0.349±0.300
30	<b>0.164±0.295</b>	<i>0.268±0.367</i>	0.473±0.314	0.394±0.307	0.315±0.273
50	<b>0.224±0.347</b>	<i>0.310±0.377</i>	0.536±0.341	0.426±0.331	0.352±0.272

Table 12: **MSE @ 1 Lyapunov Time** across different granularities (lower is better). Bold = best, italic = second best.

Granularity	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
10	<b>0.87</b>	0.82	0.34	0.39	0.36
30	0.80	<b>0.85</b>	0.52	0.58	0.52
50	<b>0.89</b>	0.86	0.41	0.60	0.56

Table 13: **Fractal dimension accuracy** (Spearman correlation) across different granularities (higher is better). Bold = best, italic = second best.

Granularity	Parrot	Chronos	Chronos Bolt	TimesFM	TimeMoE
10	<b>0.087±0.137</b>	<i>0.127±0.173</i>	0.573±0.307	0.444±0.326	0.467±0.368
30	<b>0.122±0.194</b>	<i>0.173±0.209</i>	0.346±0.297	0.345±0.298	0.354±0.290
50	<b>0.137±0.207</b>	<i>0.230±0.256</i>	0.406±0.305	0.361±0.323	0.370±0.338

Table 14: **KL Divergence** between predicted and true attractors across different granularities (lower is better). Bold = best, italic = second best.

Nadaraya–Watson model of the time series,

$$\hat{p}(\mathbf{y} \mid \mathbf{q}) = \frac{\sum_{j=D}^{L-H} K_\sigma(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) K_\sigma(\mathbf{y}, \mathbf{x}_{j+1:j+H})}{\sum_{j=D}^{L-H} K_\sigma(\mathbf{q}, \mathbf{x}_{j-(D-1):j})}, \quad (1)$$

where the query  $\mathbf{q}$  represents the length- $D$  motif immediately preceding the end of the context window.  $\mathbf{y}$  represents a length- $H$  forecast of subsequent values. The forecast sequence  $\mathbf{y}$  has probability  $\hat{p}$  conditioned on the query. The symmetric kernel  $K_\sigma(\mathbf{u}, \mathbf{v}) = \sigma^{-d} K((\mathbf{u} - \mathbf{v})/h\sigma)$  has bandwidth  $\sigma$  in dimension  $d = D \cdot \dim(x_t)$ . Assuming mean-squared error as a distance function in sequence space, we use a Gaussian kernel

$$K_\sigma(\mathbf{u}, \mathbf{v}) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{\|\mathbf{u} - \mathbf{v}\|^2}{2\sigma^2}\right)$$

We set the second kernel on  $\mathbf{y}$  in Eq. 1 to a delta function, in order to output predictions that exactly match sequences from the context, rather than nearby sequences in a least-squares sense. We write the conditional mean predictor

$$\hat{y}(\mathbf{q}) = \sum_{j=D}^{L-H} w(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) \mathbf{x}_{j+1:j+H}, \quad w(\mathbf{q}, \mathbf{z}) \equiv \frac{K_\sigma(\mathbf{q}, \mathbf{z})}{\sum_{j=D}^{L-H} K_\sigma(\mathbf{q}, \mathbf{x}_{j-(D-1):j})}. \quad (2)$$

Context parrotting corresponds to the 1-nearest-neighbor limit  $\sigma \rightarrow 0$ .

1242 **Context parroting preserves attractor properties at long context lengths.** In Appendix H.4, we  
 1243 derive the following proposition,

$$1245 \lim_{L \rightarrow \infty} \mathbb{E}_p[F(\mathbf{y})|\mathbf{q}] = \mathbb{E}_\mu[F(\mathbf{x})]$$

1246 where  $L$  is the context length for an Nadaraya–Watson estimator  $p$ ,  $F(\mathbf{y})$  is an estimate from a  
 1247 forecast sequence  $\mathbf{y}$  of a property  $F$  of an ergodic dynamical system, which has an invariant value  
 1248  $\mathbb{E}_\mu[F(\mathbf{x})]$  when calculated over the full attractor with underlying measure  $\mu$ . The query  $\mathbf{q}$  is an  
 1249 arbitrary sequence of consecutive timepoints from the dynamical system. This proposition states that,  
 1250 when the context is sufficiently long, context parroting of an ergodic system preserves invariant values  
 1251 of the underlying dynamics. Context parroting thus represents an effective baseline for dynamical  
 1252 systems forecasting, because, in the limit of long context, it will preserve global properties like  
 1253 conditional distributions of values, Lyapunov exponents, or entropy production rates.

## 1254 H.2 DISCRETE-TOKEN PARROTING

1255 For fully-discrete tokens, a  $D^{th}$  order Markov chain fit to the context has the form

$$1258 p(\mathbf{y}|\mathbf{q}) = \frac{\#\{j : (x_{j-(D-1):j} = \mathbf{q}) \wedge (x_{j+1:j+H} = \mathbf{y})}{\sum_{\mathbf{y}'} \#\{i : (x_{j-(D-1):j} = \mathbf{q}) \wedge (x_{j+1:j+H} = \mathbf{y}')\}} \quad (3)$$

1259 where the overall context has length  $L$ , and the Markov chain conditions the  $H < L$  future tokens on  
 1260 the  $D < L$  preceding tokens. The index  $j$  runs over all contiguous sequences of length  $D + H$  in  
 1261 the context,  $j \in \{D - 1, D, \dots, L - H - 2, L - H - 1\}$ . The vector  $\mathbf{q} \in \mathbb{R}^D$  represents the query,  
 1262 and the vector  $\mathbf{y} \in \mathbb{R}^H$  represents the prediction in response to this query. Eq. 3 simply counts the  
 1263 number of token sequences of length  $D + H$  that start with a given sequence of  $D$  query tokens.  
 1264 A maximum-likelihood estimator derived from this model always samples the highest-likelihood  
 1265 sequence  $\mathbf{y}$ ,

$$1266 \hat{\mathbf{y}}_{\text{MLE}}(\mathbf{q}) = \operatorname{argmax}_{\mathbf{y}} \log p(\mathbf{y}|\mathbf{q})$$

1267 However, this estimator may be unstable due to the appearance of queries  $\mathbf{q}$  not seen in the context,  
 1268 motivating the use of *token smoothing*, in which Eq. 3 is replaced by the distribution

$$1271 p(\mathbf{y}|\mathbf{q}) = \frac{\#\{j : (x_{j-(D-1):j} = \mathbf{q}) \wedge (x_{j+1:j+H} = \mathbf{y}) + \alpha}{\sum_{\mathbf{y}'} (\#\{i : (x_{j-(D-1):j} = \mathbf{q}) \wedge (x_{j+1:j+H} = \mathbf{y}') + \alpha)} \quad (4)$$

1272 with increasing values of the parameter  $\alpha$  causing predictions to converge to a uniform sample  
 1273 over possible predictions  $\mathbf{y}$ . The parameter value  $\alpha = 0$  reduces to no smoothing, while  $\alpha = 0.5$   
 1274 corresponds to the Jeffreys prior and  $\alpha = 1$  corresponds to Laplace’s rule of succession.

## 1275 H.3 CONTINUOUS-TOKEN PARROTING

1276 A more general time series model treats tokens as continuous-valued. Some time series foundation  
 1277 models like Chronos use binning to discretize time series values, allowing the direct use of discrete-  
 1278 token architectures (Ansari et al., 2024). However, many time series models assume effective  
 1279 continuity in token values, and we favor a continuous formulation in order to highlight connections to  
 1280 dynamical systems theory.

1281 To model continuous-valued tokens directly, we replace the discrete count in §H.2 with a kernel-  
 1282 weighted estimate over all past subsequences. Let  $\{\mathbf{x}_t\}$  denote a univariate or multivariate time series.  
 1283 For context length  $L$  and prediction horizon  $H$ , the Nadaraya–Watson estimate of the conditional  
 1284 density is

$$1285 \hat{p}(\mathbf{y} | \mathbf{q}) = \frac{\sum_{j=D}^{L-H} K_h(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) K_h(\mathbf{y}, \mathbf{x}_{j+1:j+H})}{\sum_{j=D}^{L-H} K_h(\mathbf{q}, \mathbf{x}_{j-(D-1):j})}, \quad (5)$$

1286 where  $K_h(\mathbf{u}, \mathbf{v}) = h^{-d} K((\mathbf{u} - \mathbf{v})/h)$  is a symmetric kernel with bandwidth  $h$  in dimension  
 1287  $d = D \cdot \dim(\mathbf{x}_t)$  for the first kernel, and  $d = H \cdot \dim(\mathbf{x}_t)$  for the second kernel. Assuming  
 1288 mean-squared error as a distance function in sequence space, we use a Gaussian kernel

$$1295 K_h(\mathbf{u}, \mathbf{v}) = \frac{1}{(2\pi h^2)^{d/2}} \exp\left(-\frac{\|\mathbf{u} - \mathbf{v}\|^2}{2h^2}\right)$$

1296 In practice, we drop the second kernel on  $\mathbf{y}$  in Eq. 5 in order to output a prediction that exactly  
 1297 matches sequences from the context, rather than nearby sequences in a least-squares sense. We thus  
 1298 write the conditional mean predictor  
 1299

$$1300 \hat{\mathbf{y}}(\mathbf{q}) = \sum_{j=D}^{L-H} w(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) \mathbf{x}_{j+1:j+H} \quad (6)$$

$$1301$$

$$1302$$

1303 where we have isolated a term corresponding to the weight of each sequence,  
 1304

$$1305 w(\mathbf{q}, \mathbf{z}) \equiv \frac{K_h(\mathbf{q}, \mathbf{z})}{\sum_{j=D}^{L-H} K_h(\mathbf{q}, \mathbf{x}_{j-(D-1):j})}.$$

$$1306$$

$$1307$$

1308 **Nearest-neighbor and global average limits.** The bandwidth  $h$  plays the role of a smoothing  
 1309 parameter (analogous to  $\alpha$  in Eq. 4). As  $h \rightarrow 0$  the scheme approximates a single-nearest neighbor  
 1310 parrot, while as  $h \rightarrow \infty$  it converges to a global average over all sequences.  
 1311

1312 **Connection to attention.** If one takes

$$1313 K(\mathbf{u}, \mathbf{v}) = \exp(\mathbf{u}^\top \mathbf{v} / \tau),$$

$$1314$$

1315 then Eq. 6 recovers a simplified form of softmax-attention, with the temperature hyperparameter  $\tau$   
 1316 controlling smoothness. In this view, the continuous parrotting scheme is a kernel-regression analogue  
 1317 of discrete  $k$ -gram smoothing (Tsai et al., 2019).

1318  **$k$ -nearest-neighbor limit.** We define a set  $\text{Top}_k$  corresponding to a subset of the possible values of  
 1319 the index  $j \in \{D, D+1, \dots, L-H-1, L-H\}$ . The  $k$  elements of  $\text{Top}_k$  correspond to the indices  
 1320  $j$  that produce the  $k$  largest values of  $w(\mathbf{q}, \mathbf{x}_{j-(D-1):j})$  across all values of  $j$ . We compute a simple  
 1321 average of these  $k$  closest matches  
 1322

$$1323 \hat{\mathbf{y}}(\mathbf{q}) = \frac{1}{k} \sum_{j \in \text{Top}_k} w(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) \mathbf{x}_{j+1:j+H} \quad (7)$$

$$1324$$

$$1325$$

1326 yielding a  $k$ -nearest-neighbors parrotting scheme. As  $k$  increases, this estimator interpolates between  
 1327 exact parrotting ( $k=1$ ) and global average ( $k \rightarrow L$ ).  
 1328

1329 **Simplex projection.** Simplex projection, a classical forecasting method in nonlinear dynamics,  
 1330 corresponds to the special case  $H=1$  (single step prediction),  $k=D+1$  in Eq. 7. The condition  
 1331  $k=D+1$  represents the minimal number of affinely independent neighbors needed to triangulate a  
 1332 point in a  $D$ -dimensional space (Sugihara & May, 1990).

1333 In simplex projection, the query  $\mathbf{q}$  is interpreted as a time-delay embedding of the time series  
 1334 observable  $\mathbf{x}$ . Takens' theorem argues that, under mild conditions, a finite number of time delay  
 1335 embeddings of an observable drawn from a deterministic ergodic system will be diffeomorphic  
 1336 (smoothly mappables) to the full-state dynamics (Takens, 2006). Because simplex projection uses  
 1337 only neighbor identities, and not absolute distances, to weight context points, a delay embedding is  
 1338 sufficient to calculate the appropriate weights.  
 1339

1340 **S-map forecasts.** Another common nonlinear forecasting technique retains all terms in the sum  
 1341 Eq. 6, but instead performs a nonlinear weighting of the form  
 1342

$$1343 K_\theta(\mathbf{u}, \mathbf{v}) = \exp(-\theta \|\mathbf{u} - \mathbf{v}\| / \bar{d})$$

$$1344$$

1345 where the scale parameter  $\bar{d}$  is determined by the distribution of distances among queries and points in  
 1346 the context. In practice, this parameter is often set to the mean pairwise distance among all sequences  
 1347 of length  $D$  in the context. The optimal value of the hyperparameter  $\theta$  increases as the underlying  
 1348 dynamics become more strongly nonlinear Sugihara (1994). We note that, in the classical formulation  
 1349 of the S-map, a locally-linear model is fit based on all sequences of length  $D+H$  seen in the context,  
 while here we use the Nadaraya–Watson estimator in order to emphasize connections to modern  
 kernel regression.

1350 H.4 INVARIANTS OF MOTION  
13511352 For ergodic dynamical systems in continuous time, there exists a natural measure  $\mu(\mathbf{x})$  such that, for  
1353 certain observables  $F(\mathbf{x})$ , the following condition almost surely holds,  
1354

1355 
$$\mathbb{E}_\mu[F] \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T F(\mathbf{x}) = \int F(\mathbf{x}) d\mu(\mathbf{x}) = \text{constant}$$
  
1356

1357 where the second equality arises from the Birkhoff ergodic theorem (Walters, 1982).  
13581359 We use the following convention for expectation values of sequences and single tokens; the expectation  
1360  $\mathbb{E}_\mu[\mathbf{x}_{t:t+T}]$  refers to the expected value of the sequence  $\mathbf{x}_{t:t+T}$  given pointwise measure  $\mu$ . We note  
1361 that, for deterministic dynamical systems, once a given point is sampled on the attractor with measure  
1362  $\mu(\mathbf{x}_t)$ , subsequent points have delta function conditional probability on the first point. Thus, we use  
1363 the convention  $\mu(\mathbf{x}_t) = \mu(\mathbf{x}_{t:t+T})$  and we use the measure to refer to both the probability of a given  
1364 timepoint, or a sequence of arbitrary length originating from that timepoint.  
1365**Proposition.** Under appropriate kernel conditions,

1366 
$$\lim_{L \rightarrow \infty} \mathbb{E}_p[F(\mathbf{y})|\mathbf{q}] = \mathbb{E}_\mu[F(\mathbf{x})]$$
  
1367

1368 where  $L$  is the context length for a Nadaraya–Watson estimator  $p$ ,  $F(\mathbf{y})$  is an estimate on a sequence  
1369  $\mathbf{y}$  of an invariant property of an ergodic dynamical system with measure  $\mu$ , and  $\mathbf{q}$  is an arbitrary  
1370 sequence of consecutive timepoints from the dynamical system. This proposition states that, when  
1371 the context is sufficiently long, a Nadaraya–Watson estimator of an ergodic system preserves the  
1372 invariant values of the underlying dynamics.  
1373**Derivation.** We start with the definition of the dynamical average,

1375 
$$\mathbb{E}_\mu[F] = \int F(\mathbf{x}) d\mu(\mathbf{x})$$
  
1376

1377 Inserting Eq. 5 into this expression,  
1378

1379 
$$\mathbb{E}_\mu[F(\mathbf{y})|\mathbf{q}] = \frac{\sum_{j=D}^{L-H} K_h(\mathbf{q}, \mathbf{x}_{j-(D-1):j}) \int F(\mathbf{y}) K_h(\mathbf{y}, \mathbf{x}_{j+1:j+H}) d\mu(\mathbf{y})}{\sum_{j=D}^{L-H} K_h(\mathbf{q}, \mathbf{x}_{j-(D-1):j})},$$
  
1380  
1381

1382 We multiply both the numerator and denominator by  $1/L$  and take the limit  $L \rightarrow \infty$ , in order to  
1383 convert the summations to expectations,  
1384

1385 
$$\lim_{L \rightarrow \infty} \mathbb{E}_\mu[F(\mathbf{y})|\mathbf{q}] = \frac{\mathbb{E}_\mu[K_h(\mathbf{q}, \mathbf{x}_{\leftarrow}) \int F(\mathbf{y}) K_h(\mathbf{y}, \mathbf{x}_{\rightarrow}) d\mu(\mathbf{y})]}{\mathbb{E}_\mu[K_h(\mathbf{q}, \mathbf{x}_{\leftarrow})]},$$
  
1386

1387 where  $\mathbf{x}_{\leftarrow}$  represents the first  $D$  points of random lookback window of length  $D + H$  sampled  
1388 from the underlying dynamical system, while  $\mathbf{x}_{\rightarrow}$  denotes the next  $H$  timepoints. In practice, this  
1389 corresponds to a time series of  $D + H$  points generated by simulating the dynamics starting at a point  
1390 on the attractor randomly-sampled according to the measure  $\mu$ .  
13911392 If we take the limit  $h \rightarrow 0$  (exact matching), then the kernel  $K_h$  becomes a delta function, yielding  
1393

1394 
$$\lim_{h \rightarrow 0} \lim_{L \rightarrow \infty} \mathbb{E}_\mu[F(\mathbf{y})|\mathbf{q}] = \mathbb{E}_\mu[F(\mathbf{x}_{\rightarrow})|\mathbf{x}_{\leftarrow} = \mathbf{q}]$$
  
1395

1396 If the measure  $\mu$  is ergodic, then the conditional expectation of an invariant  $F$  given any query  $\mathbf{q}$  is  
1397 simply its unconditional expectation,  
1398

1399 
$$\lim_{h \rightarrow 0} \lim_{L \rightarrow \infty} \mathbb{E}_\mu[F(\mathbf{y})|\mathbf{q}] = \mathbb{E}_\mu[F(\mathbf{x})]$$
  
1400

1401 H.5 SCALING LAWS LIMITING PREDICTION OF STOCHASTIC SYSTEMS  
14021403 For a stochastic time series  $\mathbf{x}_{1:T}$  with autocorrelation given by

1404 
$$|\text{Corr}(\mathbf{x}_t, \mathbf{x}_{t+\tau})| \leq C e^{-\alpha\tau}, \quad \alpha > 0$$
  
1405

1404 with  $C$  representing a proportionality constant, the expected mean squared error of a forecast scales  
 1405 as

$$\mathbb{E}[\|\hat{\mathbf{y}} - \mathbf{y}\|^2] \sim e^{-\alpha L}, \quad L \rightarrow \infty.$$

1406 Thus, under exponential decay of correlations (mixing), the amount of information about future states  
 1407 in a length- $L$  context window saturates exponentially quickly Bradley (2005). Thus, forecasts derived  
 1408 from increasingly large context windows converge exponentially quickly to optimal conditional  
 1409 forecasts under the invariant measure  $\mu$ .

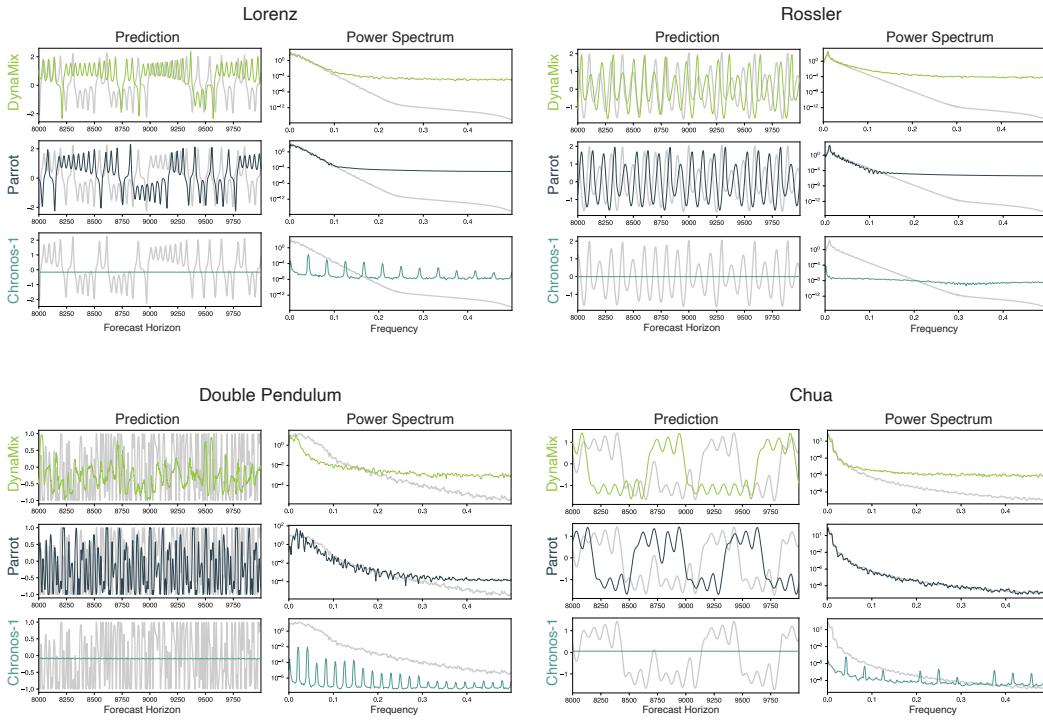
1410 Under standard smoothness conditions Fan & Yao (2008); Takezawa (2005), the forecast error also  
 1411 exhibits a standard bias-variance tradeoff of the form

$$\mathbb{E}[\|\hat{\mathbf{y}} - \mathbf{y}\|^2] = \mathcal{O}(h^4) + \mathcal{O}\left(\frac{1}{Ch^{L+H}}\right).$$

1412 The optimal width of the kernel thus scales as,

$$h_{\text{opt}} \sim C^{-1/(4+L+H)}$$

## I STABILITY OF FORECASTS AT LONG PREDICTION HORIZONS



1447 **Figure 13: Properties of forecast models at long forecast horizons.** Predictions of four distinct  
 1448 chaotic systems using various forecast models. Forecasts are generated for 10,000 points past the  
 1449 end of the context, and the last 2000 timepoints are shown. The power spectrum is estimated using  
 1450 Welch's method on the last 5000 timepoints of a 10,000 timepoint prediction.

1451  
 1452  
 1453  
 1454  
 1455  
 1456  
 1457

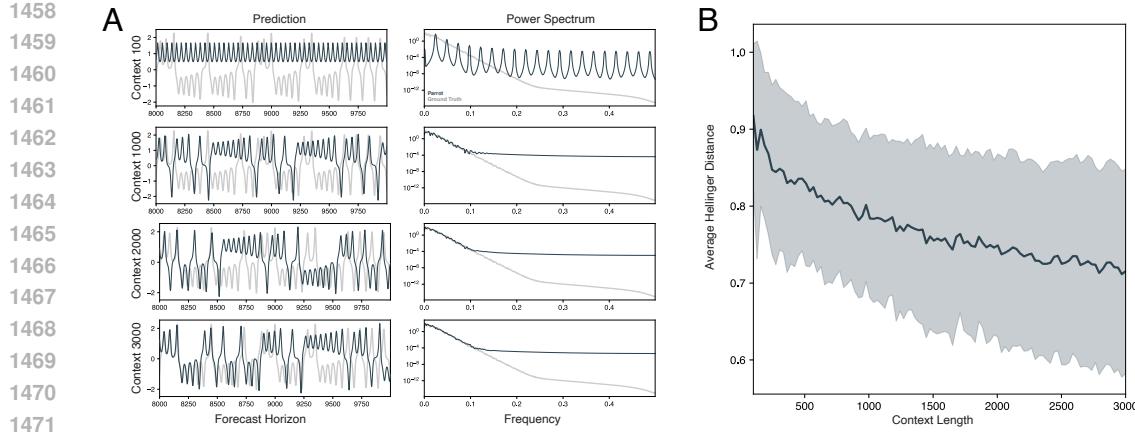


Figure 14: **Attractor properties improve with context length.** (A) Predictions of a single chaotic system, the Lorenz attractor, by the parrotting model as the context length increases. Forecasts are generated for 10,000 points past the end of the context, and the last 2000 timepoints are shown. The power spectrum is estimated using Welch’s method on the last 5000 timepoints of a 10,000 timepoint prediction. (B) The average Hellinger distance between the true and predicted power spectrum as a function of context length, averaged over 129 distinct chaotic systems (including the Lorenz attractor). Error bars correspond to standard deviations. The averaged Hellinger distance is introduced as a long-term metric for chaotic systems in [Mikhaeil et al. \(2022\)](#) and [Brenner et al. \(2022\)](#).

Metric	DynaMix	Parrot
Average Hellinger Distance	$0.595 \pm 0.166$	$0.591 \pm 0.198$
KL Divergence	$0.624 \pm 0.128$	$0.469 \pm 0.162$

Table 15: **Global attractor fidelity metrics calculated at long forecast horizons.** For these experiments, the context length is 2000 timepoints, and the prediction horizon is 10,000 timepoints. Values correspond to mean and standard errors across 129 distinct chaotic systems.

Invariant Property	DynaMix	Parrot
Largest Lyapunov Exponent	$0.278 \pm 0.100$	$0.328 \pm 0.097$
Fractal Dimension	$0.441 \pm 0.073$	$0.832 \pm 0.028$

Table 16: **Correlation of forecast properties with invariant properties.** For these experiments, the context length is 2000 timepoints, and the prediction horizon is 10,000 timepoints. Values correspond to Pearson correlations and standard errors (Fisher’s transformation) across 129 distinct chaotic systems.

## J STABILITY OF DISTRIBUTIONAL METRICS AGAINST HYPERPARAMETER CHOICES

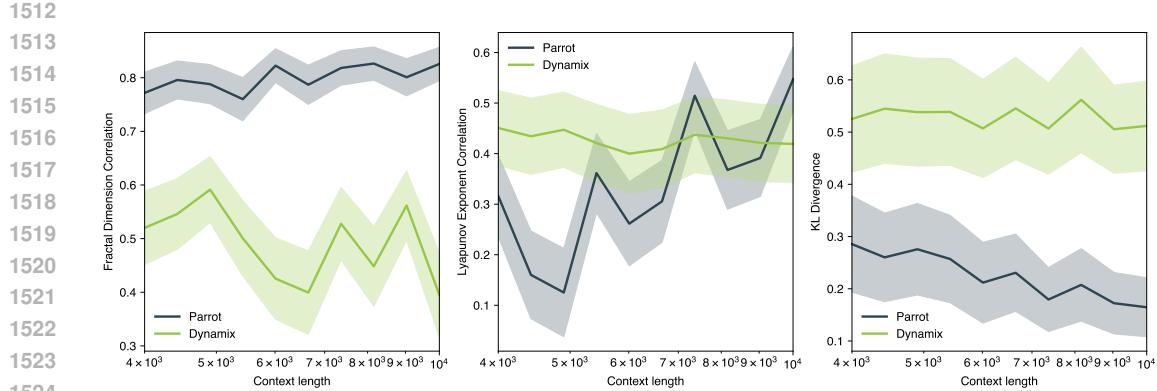


Figure 15: **Properties of forecast models at long forecast horizons as context length varies.** Three measures of long-term forecast properties at very long rollouts (10,000 timepoints) as the context length varies. For the fractal dimension and Lyapunov exponents, values correspond to Pearson correlations and standard errors (Fisher’s transformation) across 129 distinct chaotic systems. For the KL Divergence, values correspond to mean and standard errors across 129 distinct chaotic systems.

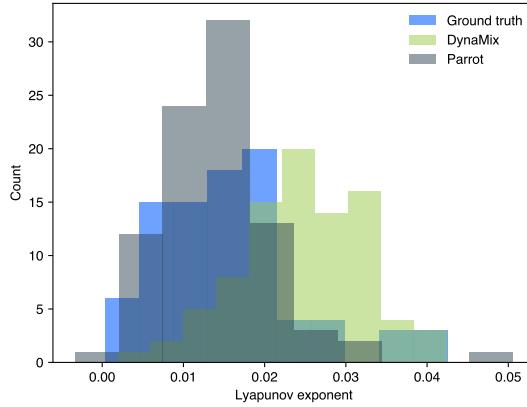


Figure 16: **Histograms of Lyapunov Exponents.** We estimate the Lyapunov exponents from the ground-truth time series, as well as from long rollouts from parrotting and DynaMix. These rollouts are generated with a context length of 2000 and a prediction horizon of 10000, and correspond to estimates from all distinct chaotic systems in dysts.

Kernel Radius	DynaMix	Parrot
0.01	$6.921 \pm 0.594$	$4.194 \pm 0.474$
0.1	$2.786 \pm 0.361$	$1.589 \pm 0.281$
1.0	$0.624 \pm 0.128$	$0.469 \pm 0.162$
$10.0^\dagger$	$-0.098 \pm 0.076$	$-0.075 \pm 0.078$

Table 17: **Variation of KL divergence at long forecast horizons, for different values of the Gaussian kernel.** For these experiments, the context length is 2000 timepoints, and the prediction horizon is 10,000 timepoints. Values correspond to mean and standard errors across 129 distinct chaotic systems.  $^\dagger$ Note: sampling-based KL divergence calculations can fluctuate below zero when the center of the estimate is close to zero. To avoid introducing directional bias, we do not enforce non-negativity with clipping or renormalization, and so these values should be interpreted as being near zero.

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

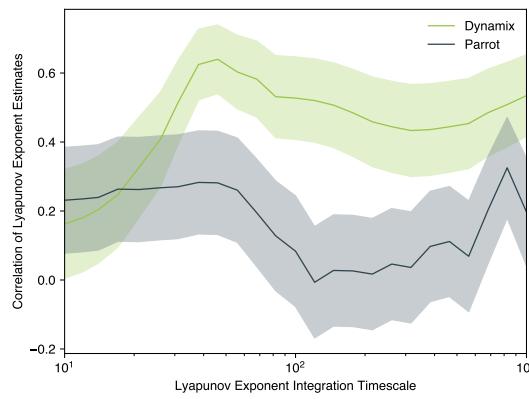


Figure 17: **Dependence of Lyapunov exponent accuracy on estimation algorithm integration time.** We estimate the Lyapunov exponents from the ground-truth time series, as well as from long rollouts from parrotting and DynaMix, for varying integration timescales in the Rosenstein algorithm (Rosenstein et al., 1993). These rollouts are generated with a context length of 2000 and a prediction horizon of 10000, and correspond to estimates from all distinct chaotic systems in dysts.

1602  
 1603  
 1604  
 1605  
 1606  
 1607  
 1608  
 1609  
 1610  
 1611  
 1612  
 1613  
 1614  
 1615  
 1616  
 1617  
 1618  
 1619