
SutraNets: Sub-series Autoregressive Networks for Long-Sequence, Probabilistic Forecasting

Shane Bergsma Timothy Zeyl Lei Guo
Huawei Cloud, Alkaid Lab Canada
{shane.bergsma,timothy.zeyl,leiguo}@huawei.com

Abstract

We propose SutraNets, a novel method for neural probabilistic forecasting of long-sequence time series. SutraNets use an autoregressive generative model to factorize the likelihood of long sequences into products of conditional probabilities. When generating long sequences, most autoregressive approaches suffer from harmful error accumulation, as well as challenges in modeling long-distance dependencies. SutraNets treat long, univariate prediction as *multivariate* prediction over lower-frequency *sub-series*. Autoregression proceeds across time *and* across sub-series in order to ensure coherent multivariate (and, hence, high-frequency univariate) outputs. Since sub-series can be generated using fewer steps, SutraNets effectively reduce error accumulation and signal path distances. We find SutraNets to significantly improve forecasting accuracy over competitive alternatives on six real-world datasets, including when we vary the number of sub-series and scale up the depth and width of the underlying sequence models.

1 Introduction

As a cloud provider, we rely on long-range forecasts to allocate resources and plan capacity. In particular, we need to predict the *fine-grained* pattern of demand (e.g., hourly) for a large number of products, over a long time horizon (e.g., months). Since purchase decisions are large-scale (e.g., supplying regions with 100K machines), and made weeks-to-months before delivery, cloud providers have a profound need for methods that “capture the uncertainty of the future” [43]. Fine granularity, long histories, long forecast horizons, and *quantification of uncertainty* can be formalized as the long-sequence (many-step) *probabilistic* forecasting problem.

Recent work applies Transformers [82] to long-sequence forecasting. In order to circumvent the Transformer’s quadratically-scaling *attention* mechanism, long sequences necessitate trading off increased signal path for reduced computational complexity [44, 92, 87, 47]. Unfortunately, the effectiveness and even validity of these approaches have been questioned [90, 32]. Most of these approaches do not provide probabilistic outputs (rather best-guess point predictions), and all condition on fixed-size windows, typically of very short duration. Indeed, many recent papers on “long-term” forecasting [87, 86, 70, 93] condition on a maximum of 96 inputs — only 4 days of history at 1-hour granularity. Such restricted context is a serious limitation given that many time series (including in standard evaluation datasets based on electricity and traffic) have strong *weekly* seasonality.

Our production forecasting system, like many in industrial settings [65, 52, 72] and cloud services [69], is based on RNNs. The RNN’s recurrent architecture forces the compression and consolidation of long-range information, and RNNs work well on smaller datasets as they are not as “data hungry” as Transformers [16]. RNNs can theoretically condition on context of any length; in this paper, we condition RNN forecasts on length-2016 contexts without difficulties in scaling.

RNNs are often used to generate probabilistic forecasts via an autoregressive approach [21, 65, 60, 7], whereby the probability over future timesteps is factorized into a product of one-step-ahead

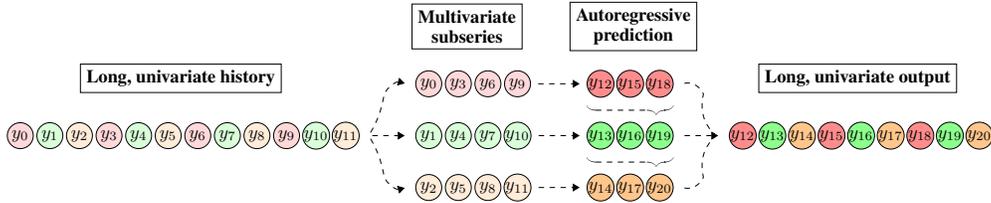


Figure 1: SutraNets: a long, univariate history is converted to lower-frequency sub-series. Multivariate autoregressive prediction, across time and sub-series, ensures coherent univariate output samples.

conditionals [6, 28]. When generating many steps, however, two key problems arise: (1) a discrepancy grows between training, where we condition on true previous values, versus inference, where we condition on sampled values (the *discrepancy* problem), and (2) long-term dependencies may be challenging to exploit (the *signal path* problem). The discrepancy problem leads to error accumulation during inference [5], but also a kind of *informational asymmetry*: because highly-informative true values are available during training, the model may ignore other useful info, e.g., seasonally-lagged inputs provided as extra features. The signal path problem, meanwhile, is especially acute for RNNs, where distant information must be preserved over many steps in the RNN hidden state.

We propose SutraNets,¹ a general method for probabilistic prediction of long time series. We focus on the application of SutraNets to RNN-based forecasting, where they offer major benefits. SutraNets transform a length- N series into K sub-series of length N/K . Sub-series forecasts are generated autoregressively, sequentially conditional on each other, enabling coherent outputs (Fig. 1). SutraNets address both the discrepancy problem (by forcing the network to take larger *generative strides*, i.e., predicting without access to immediately-preceding true values) and the signal path problem (by reducing the distance between historical and output values by a factor of K). Training of SutraNets *decomposes* over sub-series; i.e., sub-series RNNs can be trained in *parallel*, enabling a K -fold improvement in training parallelism over standard RNNs.

Experimentally, SutraNets demonstrate superior forecasting across six datasets, improving accuracy by an average of 15% over C2FAR [7] and over strong seasonality-aware baselines. We evaluate a variety of other strategies for improving long-sequence forecasting, including lags, input dropout, deeper and wider networks, and multi-rate hierarchical approaches. Through these evaluations, we gain insights into how to succeed at long-sequence generation, in time series and beyond.

2 Background and related work

Probabilistic autoregressive forecasting Autoregressive forecasting models typically combine a backbone sequence model with a technique to estimate an output probability distribution at each step. In this paper, we apply SutraNets to C2FAR [7], a state-of-the-art distribution estimator that was shown to improve over DeepAR [65], DeepAR-binned [60], SQF-RNN [21], and IQN-RNN [27]. SutraNets can serve as the sequential backbone for any of these estimators, as well as for other recent approaches such as the denoising diffusion models [62], or conditioned normalizing flows [63].

Multi-rate forecasting It is sometimes acceptable to *aggregate* high-frequency series or event data to a lower frequency prior to sequence modeling, e.g., aggregating patient events over 6-month intervals [48], or network traffic over 10 subframes [79]. When high-frequency predictions are required, lower-frequency forecasts can provide guidance. Techniques for *temporal disaggregation* convert low-frequency series to higher frequencies via related high-frequency *indicator* series [13, 67]. Athanopoulos et al. [2] simultaneously forecast at multiple frequencies and then reconcile the predictions. Such methods do not provide uncertainty estimates nor are they based on neural networks.

Scaleformer [70] aggregates series, forecasts at the lower frequency, and upsamples the lower-freq predictions in order to provide additional inputs for the original forecast model. We evaluate a similar approach, Low2HighFreq (§3), as a baseline. Other neural approaches (in forecasting and beyond) implicitly encourage learning of lower-frequency dynamics through hierarchical attention [20, 47, 16],

¹In Sanskrit, a *sūtra* is a thread or “that which like a thread runs through or holds everything together” [85]. A SutraNet prediction is woven from threads of sub-series predictions; Fig. 2 shows different stitching patterns.

pooling blocks [92, 87, 49, 93], and deeply-stacked [29, 81] and multi-rate [68, 18, 38, 73, 14, 11] recurrence layers. Unlike these approaches, SutraNets *explicitly* generate *low-frequency* sub-series conditional on other low-frequency sub-series, without directly forecasting at a high frequency at all.

Reducing training/inference discrepancy If we know *a priori* how far to forecast, it is possible to avoid error accumulation by directly forecasting all horizons in one step [76], or via non-autoregressive decoding [20]. Such systems usually generate point predictions [92, 87, 9], although predicting certain fixed forecast quantiles is also possible [84, 20]. ProTran [77] is a probabilistic but non-autoregressive approach. Without modeling local dependencies among outputs, such models prevent generation of *coherent* samples, and preclude computing robust likelihoods for *given* output sequences — where error accumulation is not an issue (e.g., for anomaly detection, missing value interpolation, compression, etc.). SutraNets are more in-step with flexible autoregressive models like GPT3 [8].

For autoregressive models, *scheduled sampling* [5] aims to handle errors better during inference by sampling (incorrect) inputs during *training*. This approach was effective in tasks such as human motion synthesis [45], but results have been mixed in autoregressive forecasting [65, 66]. Noting that scheduled sampling produces a biased estimator, Lamb et al. [40] instead used adversarial techniques to reduce training/inference discrepancy. Wu et al. [88] applied this method to forecasting, using a GAN [26] with an autoregressive sparse transformer. However, adversarial training showed insignificant gains over the sparse transformer alone on 7-day electricity and traffic predictions.

Like scheduled sampling, *dropout* [74] is a kind of stochastic regularization that prevents overfitting by injecting noise into networks. Dropout rates of around 20% are commonly used in input layers of autoencoders and feed-forward nets [3, 74]. Dropout is also common in autoregressive models [82, 61], where, in theory, it may encourage networks to rely less on (now noisy) previous values [56].

Modeling long-term dependencies When forecasting long series, seasonally *lagged* values — historical values from, e.g., one year ago — can be used as extra features in the sequence model, acting as a kind of residual input [31]. A recent competition featured long time series (e.g., 700 days of history) [1]; the winning model found lag to be more effective than fixed attention weights [75]. Some researchers have even declared that, when it comes to attention, “lag is all you need” [19].

Learned *attention* can capture long-range dependencies [82], but is infeasible for very long sequences due to memory and time complexity quadratic in sequence length. Complexity also scales quadratically for forecasting via fully-connected layers [57, 9]. Recent work has explored sparsified or hierarchical attention in forecasting [44, 92, 47, 87]. TimesNet [86] transforms univariate time series into 2D tensors of multiple periods, which are then processed via Inception-style 2D kernels to generate point predictions. As mentioned earlier, most of these systems fail to condition on more than very short windows — a few days of hourly data. Ultimately, we are interested in *probabilistic* forecasting using *years* of fine-grained historical data. It is unclear whether lag is really all we need, or even whether seasonal-naive baselines [34] are already more effective than Transformers [32].

Unlike Transformers, RNNs can be applied to long sequences *out-of-the-box*. However, since RNN signal path is linear in sequence length, mechanisms are required to preserve long-range info, such as the special gated units of LSTMs [33]. Multi-dimensional RNNs [30] propagate distant info via row and column connections across images [78, 81]. Didolkar et al. [16] combine RNNs for long-range signals with Transformers that attend over shorter chunks. SutraNets, in contrast, do not require architectural changes *within* RNNs; indeed, they work with RNNs or Transformers. In SutraNets, low-frequency information is both *implicitly* captured in network state and *explicitly* and autoregressively communicated through generated low-frequency forecasts.

Comparison to other RNNs Table 1 compares SutraNets to prior RNNs (cf. Table 1 in [82]). PatchTST [55] groups consecutive values into input patches. While patching was proposed as a method to reduce the attentional complexity of Transformers, patching was also recently used with RNNs in SegRNN [46], where it effectively reduces the signal path by a factor of K (taking in K inputs each step — like SutraNets — but with K times fewer overall input steps). However, neither PatchTST nor SegRNN provides a mechanism to probabilistically *generate* patches, and thus generate coherent *probabilistic* outputs. Dilated LSTMs [83] and Dilated RNNs [10] with minimum dilation amounts >1 do facilitate error reduction and parallel training, but at the cost of sacrificing coherency, as the model is then “equivalent to multiple shared-weight networks, each working on partial inputs” [10, §4.4]. With such DilatedRNNs, if the time series were to spike to a high level at

Table 1: Comparison of SutraNets to prior RNNs. Like prior methods, SutraNets reduce signal path, but also improve error accumulation (via a larger generative stride), while enabling greater training parallelism (via fewer sequential operations), without sacrificing coherent *probabilistic* output.

RNN method	Meaning of K	Signal path	Generative max stride	Sequential operations	Coherent outputs
Standard RNN, e.g., LSTM [33]	N/A	$\mathcal{O}(N)$	$\mathcal{O}(1)$	$\mathcal{O}(N)$	Yes
Skips [91], lags/residuals [31]	Skip/lag amount	$\mathcal{O}(N/K)$	$\mathcal{O}(1)$	$\mathcal{O}(N)$	Yes
PatchTST [55], SegRNN [46]	Patch size/stride	$\mathcal{O}(N/K)$	N/A	$\mathcal{O}(N/K)$	N/A
DilatedRNN [10], dilations $C \dots K$	Largest dilation	$\mathcal{O}(N/K)$	$\mathcal{O}(C)$	$\mathcal{O}(N/C)$	If $C=1$
SutraNets	Num. sub-series	$\mathcal{O}(N/K)$	$\mathcal{O}(K)$	$\mathcal{O}(N/K)$	Yes

the very final conditioning input, only a subset of the outputs would be generated conditional on this spike, resulting in highly incoherent output. This can be mitigated for point predictions by adding a final *fusion layer* [10], but such an approach is not compatible with probabilistic forecasts.

Another related approach is the Subscale WaveRNN model for audio generation [35] (conditioned on input text). For Subscale WaveRNN, the motivation is not to reduce signal path nor improve error accumulation, but to have a K -fold increase in inference parallelism when generating very, very long audio sequences. Generation proceeds somewhat similar to our regular, non-alternating version (§3) but conditioning on past and some *future* values, and with a shared neural network for all sub-series.

3 SutraNets: Sub-series autoregressive networks

Let $y_t \in \mathbb{R}$ be the value of a time series at time t , and \mathbf{x}_t be a vector of time-varying features or *covariates*. Let $y_{i:j}$ denote sequence $y_i \dots y_j$. We seek a model of the distribution of N future values of y_t (the *prediction range*) given T historical values (the *conditioning range*). We can formulate this conditional distribution as an autoregressive (AR) generative model as in [28]:

$$p(y_{T+1:T+N} | y_{1:T}, \mathbf{x}_{1:T+N}) = \prod_{t=T+1}^{T+N} p(y_t | y_{1:t-1}, \mathbf{x}_{1:T+N}) \quad (1)$$

Following DeepAR [65], forecasting approaches have modeled the one-step-ahead distributions using RNNs [21, 65, 60, 7] or Transformers [44]. In these approaches, a global sequence model is trained by slicing many training series into many *windows*, i.e., conditioning+prediction ranges at different start points, and normalizing the windows using their conditioning ranges. Model parameters are fit via gradient descent, minimizing NLL of prediction-range outputs. At inference time, forecasts for a given conditioning range are generated by sampling $\hat{y}_{T+1} \dots \hat{y}_{T+N}$ sequentially, autoregressively conditioning on samples generated at previous timesteps. By repeating this procedure many times, a Monte Carlo estimate of Eq. (1) is obtained, from which desired forecast quantiles can be derived.

SutraNets consider each univariate window to be comprised of an ordered collection of K lower-frequency sub-series, each $1/K$ of the original length (Fig. 1). The k th sub-series is obtained by selecting every K th value in the original window, each sub-series starting at a unique offset in $1 \dots K$.² Let y^k be the k th sub-series. Let $\tilde{N} = N/K$ and $\tilde{T} = T/K$ be prediction and conditioning lengths corresponding to each sub-series, and let $<k$ and $>k$ denote indices from $1 \dots K$ that are less than and greater than k . Through another application of the chain rule, now across sub-series, Eq. (1) becomes:

$$p(y_{T+1:T+N} | y_{1:T}, \mathbf{x}_{1:T+N}) = \prod_{\tilde{t}=\tilde{T}+1}^{\tilde{T}+\tilde{N}} \prod_{k=1}^K p(y_{\tilde{t}}^k | y_{1:\tilde{t}}^{<k}, y_{1:\tilde{t}-1}^k, y_{1:\tilde{t}-1}^{>k}, \mathbf{x}_{1:T+N}) \quad (2)$$

where \tilde{t} iterates over timesteps in each sub-series. Log-likelihoods of sub-series prediction ranges can be summed to compute overall NLL, which is minimized during SutraNet training. Rather than modeling these conditionals using a single RNN, SutraNets use a separate RNN for each sub-series.

²Since the k th sub-series is assigned after slicing training/testing windows, it will vary in terms of which offset it corresponds to in the *original* time series, e.g., it may start at 1pm, 2pm, 3pm, etc., in different windows.

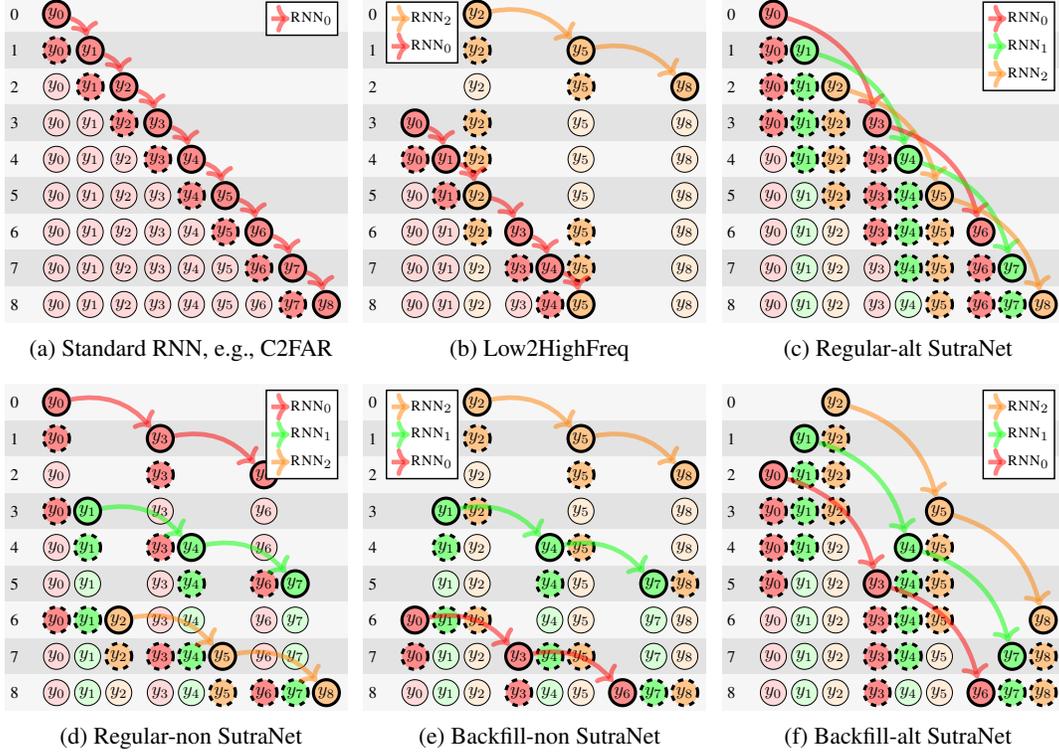


Figure 2: Ordering of generative steps in SutraNets. One output value (**bold border**) is generated in each row, conditional on (1) feature nodes (**dashed border**) *in that row*, and on (2) state from previous steps (\Rightarrow arrow connections). SutraNet RNNs (c-f) use features generated in earlier steps by both themselves and other RNNs; these values can even be from future timesteps (backfill cases).

Each RNN has its own parameters and hidden state, but SutraNets are trained and tuned collectively as a single network. The k th RNN determines the k th conditional probability at each timestep:

$$p(y_{T+1:T+N} | y_{1:T}, \mathbf{x}_{1:T+N}) = \prod_{\tilde{t}=\tilde{T}+1}^{\tilde{T}+\tilde{N}} \prod_{k=1}^K p(y_{\tilde{t}}^k | \theta_{\tilde{t}}^k = f(h_{\tilde{t}}^k)) \quad (3)$$

where $h_{\tilde{t}}^k = \text{rnn}^k(h_{\tilde{t}-1}^k, y_{\tilde{t}}^{<k}, y_{\tilde{t}-1}^k, y_{\tilde{t}-1}^{>k}, \mathbf{x}_{\tilde{t}}^k)$ and $f(\cdot)$ is a function mapping the k th RNN state to parameters of a parametric output distribution, $p(y_{\tilde{t}}^k | \theta_{\tilde{t}}^k)$.³ So while prior models condition y_t only on y_{t-1} , SutraNets autoregressively condition the generation of $y_{\tilde{t}}^k$ on:

1. $y_{\tilde{t}}^{<k}$: *Current* values (at \tilde{t}) of sub-series with indices $< k$ (always)
2. $y_{\tilde{t}-1}^k$: The *previous* value (at $\tilde{t} - 1$) of the current sub-series, k (always)
3. $y_{\tilde{t}-1}^{>k}$: *Previous* values (at $\tilde{t} - 1$) of sub-series with indices $> k$ (optional)

Conditioning $y_{\tilde{t}}^k$ on $y_{\tilde{t}}^{<k}$ makes the model autoregressive across sub-series; this ensures sub-series are generated conditional on each other, and thus (recombined) samples of the original high-frequency series will be coherent. Conditioning on $y_{\tilde{t}-1}^{>k}$ is optional; when using $y_{\tilde{t}-1}^{>k}$, SutraNets must generate in an *alternating* manner, generating one value for each sub-series at each timestep, \tilde{t} (Fig. 2c). Without $y_{\tilde{t}-1}^{>k}$, SutraNets can generate the complete prediction range of the k th sub-series before generating any predictions for $> k$ ones (Fig. 2d); in other words, we can exchange ordering of the products in Eq. (2), remove $y_{\tilde{t}-1}^{>k}$ from the equation, and proceed in a *non-alternating* manner.

³Note we can input any subset of the covariates to the RNN at each timestep. Here we include those specific to the given timestep of the given sub-series, $\mathbf{x}_{\tilde{t}}^k$, as in practice most covariates are either static across the prediction range (e.g., a product ID), or provide timestep-specific info (e.g., the hour-of-the-day) [65].

SutraNets must also specify how the sub-series index, k , relates to the offset in the original series. Figs. 2c and 2d illustrate *regular* ordering, where sub-series k begins at offset k in the original series. An alternative is *backfill* ordering, where sub-series k starts at offset $K-k+1$ (Figs. 2e, 2f). We must also specify the number of sub-series, K . We pick K so it divides into the primary seasonal period (e.g., $K=6$ for 24-hour seasons), as each sub-series then exhibits seasonality; signal processing techniques can reveal seasonality when it is not known *a priori*. Sub-series with seasonality may be more predictable, which will be especially useful for non-alternating models, as these models generate initial sub-series while conditioning on only $1/K$ of historical values. Success for these models depends on whether there is *redundancy* in the series, or whether values from other sub-series are sufficiently important that non-alternating generation creates a *missing information* problem.

Reducing training/inference discrepancy, signal path Non-alternating and backfill approaches force the network to predict K steps ahead, without access to true previous values. Non-alternating also generates far horizons in fewer steps (e.g., for the $k=1$ sub-series), reducing error accumulation. In contrast, Regular-alt generates in the same order, and conditional on the same values, as standard RNN models (Fig. 2a), and so should suffer the same effects from training/inference discrepancy.

All SutraNets reduce RNN signal path by a factor of K , as can be seen by tracing paths between features and outputs in Fig. 2. Note Regular-alt has low signal path but standard discrepancy (as noted above). The relative effectiveness of Regular-alt can thereby provide a diagnostic for whether discrepancy or signal path affects a given forecasting task (§4).

Form of output distribution and input encoding Unlike sampling text from a softmax, forecasting requires a parametric output distribution, $p(y_t^k | \theta_t^k)$ (Eq. (3)), that can account for mixed discrete and continuous outputs, potentially of unbounded dynamic range. C2FAR [7] proposes such an output distribution, via an autoregressive generative model over a hierarchical coarse-to-fine discretization of time series amplitudes (with Pareto-distributed tails). We adopt C2FAR-LSTMs for our sub-series sequence models. We use a 3-level C2F model with 12 bins at each level, discretizing normalized values to a precision of 12^3 but requiring only 12×3 output dimensions at each timestep. Previous values from each sub-series and covariate features from other sub-series are also C2F-encoded and provided as inputs at each step. Here the efficiency of C2FAR proves helpful; at each step we can encode K covariate values (i.e., from K other sub-series) using only $12 \times 3 \times K$ input dimensions, as opposed to, e.g., the $12^3 \times K$ dimensions that would be required with equivalent flat binnings [60].

Low2HighFreq Instead of generating all lower-frequency sub-series, we experiment with generating a single low-freq sub-series and conditioning on it within a high-freq model. This Low2HighFreq approach (Fig. 2b) is similar to multi-rate approaches (§2), except here low-freq samples provide not only features, but hard constraints on every corresponding K th value in the high-freq output. During inference, when we use Monte Carlo sampling to estimate $p(y_{T+1:T+N} | \dots)$, we can actually *sample* high-freq values *conditional* on low-freq ones. As such, techniques such as rejection sampling [53] are applicable; we experimented with importance sampling via likelihood weighting (fixing the low-freq values and sampling the rest), but ultimately found the likelihood weights had such high variance that it was more effective to use uniform weights and accept inconsistent estimation.

Complexity Let H be the number of RNN hidden units and K the number of sub-series. The complexity-per-timestep of each SutraNet RNN scales with the sum of the RNN’s recurrence, $H \times H$, and projection of (at most) K inputs (from other sub-series) to the RNN hidden state, $K \times H$. Even though SutraNets have K separate RNNs, only one runs at each original timestep (i.e., each SutraNet RNN runs $1/K$ of total timesteps). In consequence, SutraNet complexity-per-timestep scales with $H \times H + K \times H$, whereas standard RNNs scale $H \times H$. In practice, $H \gg K$, and we find empirically SutraNets also scale $H \times H$, resulting in similar training/inference speeds and memory requirements compared to standard models, given the same individual LSTM depth and width (§4).

One key advantage of SutraNets is that training *decomposes* over sub-series. Sub-series RNNs do not condition on hidden states from other sub-series RNNs, rather only on *generated* values. During training, when we have access to the true values of other sub-series, all inputs are known in advance and sub-series RNNs can be trained in *parallel*. The only sequential action is evolution of each RNN’s hidden state — over $1/K$ fewer steps than standard or multidimensional RNNs [78, 81]. In this way, SutraNet-RNNs provide middle-ground between RNNs and standard Transformers (with no sequential computation in training). Inference in SutraNets, like all AR models, is sequential [36].

Application to Transformers Rather than using K RNNs to parameterize the conditional probabilities as in Eq. (3), we could instead use K autoregressive Transformers. Compared to a standard Transformer, backfill and non-alternating sub-series Transformers would have the advantage of reducing *error accumulation* (by increasing the generative stride), but not of *signal path* (already $\mathcal{O}(1)$ for Transformers [82]). Moreover, at each timestep, a sub-series Transformer could attend to essentially all prior values, limited only by the generative order. In a naive implementation, SutraNet Transformers would therefore have an asymptotic complexity of $\mathcal{O}(N^2)$ — the same as standard Transformers. However, we can attain $\mathcal{O}(N^2/K)$ complexity by restricting each sub-series Transformer to only attend to values from its own sub-series, plus a few proximal values from other sub-series (similar to strided or banded self-attention [12, 8]). Although asymptotically larger than the $\mathcal{O}(N(\log N)^2)$ of LogSparse attention [44], it merits empirical investigation to determine which approach achieves the most favorable accuracy vs. complexity trade-off.

Limitations and broader impact SutraNets are more complex than standard sequence models in a few ways. First, they require more decisions, such as specifying the number of sub-series. While superior accuracy can be obtained when making these choices heuristically, we evaluate the extent to which tuning can achieve further gains (§4). Secondly, having K separate RNNs results in roughly K times as many total parameters. Although, as noted above, this does not translate into requiring more memory or training/inference time, SutraNet models are larger to store and transmit.

As a flexible predictor of both long and short horizons (and without needing *a priori* specification of confidence quantiles), SutraNets could also serve as the basis for a general-purpose large forecasting model (LFM), along the lines of LLMs like GPT3 [8]. Like modern LLMs, LFM could potentially be used with little task-specific data and no fine-tuning, and may have similarly-large environmental and financial costs [4]. As with LLMs, new techniques will be needed to detect [64, 71, 54], document [51, 22], and mitigate [25, 58] unsafe usage and other failure cases.

4 Experiments

Training. We implemented SutraNets in PyTorch [59]. We use Adam [37] (default $\beta_1=0.9$, $\beta_2=0.999$, $\epsilon=10^{-8}$), early stopping, and a decaying learning rate ($\gamma=0.99$). The default RNN is a 1-layer LSTM with 64 hidden units; when using >1 layers (Table 3), inter-layer dropout [50] with $p_{drop}=0.001$ is applied. For every model/dataset pair, we tune weight decay and initial learning rate over a 4×4 grid.

Systems. We evaluate the following systems. Naïve and SeasNaïve are non-probabilistic, untrained baselines, while the C2FAR variants follow the standard RNN ordering (Fig. 2a):

Naïve: At each horizon, repeats last-observed historical value [34].

SeasNaïve: At each horizon, repeat previous value at same *season* as horizon [34], e.g., same hour-of-day (SeasNaïve1d) or hour-of-week (SeasNaïve1w).

C2FAR: LSTM-based forecasting model [7], configured as described above (§3).

C2FAR+lags: C2FAR, but with extra features for previous hour-of-day values (C2F-encoded).

C2FAR+dropout: C2FAR, but with $p_{drop}=0.2$ (and re-scaling [74]) applied to inputs during training.

Low2HighFreq: High-frequency forecast generated conditional on low-frequency values (§3).

SutraNets: Variations Regular-alt, Regular-non, Backfill-alt, Backfill-non (§3).

Datasets. We evaluate on the following datasets, where we note the number of SutraNet sub-series, K , as well as lengths of conditioning, C , and prediction, P , ranges:

azure: $K=6$, $C=2016$, $P=288$: 5-minutely usage of 2085 VM flavors, groupings (by tenant, etc.) in Azure cloud. Originally from [15], aggregated here instead to 5-minute intervals.

elec: $K=6$, $C=168$, $P=168$: Hourly electricity usage of 321 customers [17], version from [39].

traff: $K=6$, $C=168$, $P=168$: Hourly occupancy for 862 car lanes [17], version from [89].

wiki: $K=7$, $C=91$, $P=91$: Daily hit-count for 9535 pages, first used in [21], version from [27].

mnist: $K=4$, $C=392$, $P=392$: row-wise sequential pixel values for 70,000 28×28 digits [42].

mnist $^\pi$: $K=4$, $C=392$, $P=392$: *mnist* sequences, but with a fixed, random permutation applied.

Table 2: ND%, wQL% across datasets, best results in **bold**. For *mnist/mnist π* , previous row value used for SeasNaive1d. Poor results for C2FAR relative to Regular-alt indicate *signal path problem* (green cells). Poor results for both C2FAR and Regular-alt indicate *discrepancy problem* (pink). Non-alternating weakness indicates *missing info problem* (yellow). Backfill-alt reduces signal path and discrepancy without missing info, therefore offering best or near-best solution on each dataset.

	<i>azure</i>	<i>elec</i>	<i>traff</i>	<i>wiki</i>	<i>mnist</i>	<i>mnistπ</i>
Naïve	3.2, 3.2	43.0, 43.0	75.5, 75.5	44.0, 44.0	100.0, 100.0	117.9, 117.9
SeasNaive1d	2.8, 2.8	10.5, 10.5	31.1, 31.1	44.0, 44.0	145.2, 145.2	163.5, 163.5
SeasNaive1w	5.2, 5.2	11.1, 11.1	17.5, 17.5	41.8, 41.8	N/A	N/A
C2FAR	3.2, 2.5	10.6, 8.4	19.3, 16.0	31.1, 27.2	67.9, 52.3	100.0, 95.7
C2FAR+lags	2.7, 2.2	10.4, 8.3	18.0, 15.0	31.2, 27.2	68.1, 52.8	99.3, 89.9
C2FAR+dropout	3.2, 2.6	10.9, 8.7	18.5, 15.3	31.3, 27.0	68.6, 53.3	99.7, 90.4
Low2HighFreq	2.6, 2.1	10.3, 8.2	20.2, 16.6	30.9, 27.2	62.4, 47.6	98.1, 88.3
Regular-alt	3.3, 2.6	9.9, 7.9	15.5, 13.0	30.4, 26.4	67.9, 52.4	48.7, 38.3
Regular-non	2.6, 2.1	9.7, 7.8	15.6, 13.1	31.2, 27.3	59.4, 45.4	90.1, 69.7
Backfill-alt	2.5, 2.0	9.3, 7.4	15.3, 12.8	30.1, 26.1	64.4, 49.7	72.0, 54.6
Backfill-non	2.3, 1.9	9.3, 7.4	15.7, 13.1	30.5, 26.8	58.9, 44.9	78.2, 58.3

Sequential (binarized) MNIST has been used for NLL [40] and classification [41] evaluations. Le et al. [41] used a fixed permutation “to make the task even harder.” Here, we forecast a distribution over *real-valued* pixel intensities in the second half of an image, based on values in the first. Unlike the other datasets, train/dev/test splits are not based on time, rather we use the standard test digits [42].

Metrics. We evaluate the median of our forecast distribution by computing *normalized deviation* (ND) from true values. We evaluate the full distribution using *weighted quantile loss* (wQL), evaluated at forecast quantiles $\{0.1, 0.2 \dots 0.9\}$, as in [60, 7]. For non-probabilistic Naïve and SeasNaïve baselines, note wQL reduces to ND, similarly to how CRPS reduces to absolute error [24, §4.2].

Results. We summarize the most important experimental findings as follows:

Finding 1: *Backfill-alt and Backfill-non improve over C2FAR and other baselines on every dataset, with both obtaining an average relative reduction in ND of around 15% over the best non-SutraNet.*

Table 2 has the main results; note the differences in ND% are very stable: we repeated *grid search tuning and testing* multiple times with different random seeds and found *SutraNets superior to standard models across all repeats* (supplemental §B.2). Backfill-non does best when the training/inference discrepancy problem predominates (*azure & mnist*, Table 2). Adverse effects of discrepancy seem to be magnified by high data *redundancy*; in these cases, non-alternating generation reduces discrepancy while having sufficient info to forecast accurately. We also hypothesize that Regular-non performs poorer because it generates less-*coherent* output: by construction, it only conditions on past values, even when subsequent values have already been generated (e.g., row 7 in Fig. 2d). In contrast, when generating a given point (e.g., row 7 in Fig. 2e), Backfill-non has always seen the previous and future proximal values, which were input at the previous and current timesteps, respectively. When signal path problems predominate (*traff & mnist π*), alternating generation, with access to more info, performs better. On *mnist π* , differences between regular and backfill depend solely on the random permutation; the superiority of Regular-alt here motivates tuning over multiple orderings of sub-series, perhaps averaging over them, as is done in some prior non-sequential models [23, 80].

Accuracy gains seem to derive from solving both discrepancy *and* signal path problems together. When discrepancy predominates, Regular-alt performs poorly, even though signal path is improved. For the opposite case, where error accumulation alone is improved, we implemented an RNN that processes values in backfill order, but otherwise acts as a standard RNN. This model performs much poorer than SutraNet-based Backfill-alt (supplemental §B.4). The lack of a holistic solution may also explain the inconsistency of lags and dropout, which sometimes help, sometimes hurt compared to vanilla C2FAR. Low2HighFreq works best on datasets where discrepancy predominates, while performing much poorer when signal path is important. Note that weekly-seasonal baselines have not previously been evaluated on these datasets. Observing in particular the relatively strong accuracy of SeasNaive1w on *traff*, we suggest they be included in future evaluations.

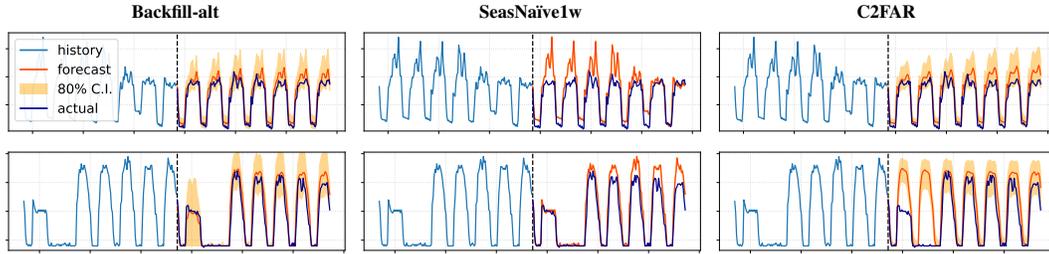


Figure 3: Forecasts and 80% intervals for *elec* series. Backfill-alt performs well with a changepoint (row 1, where SeasNaive1w fails) and with weekly seasonality (row 2, where C2FAR fails).

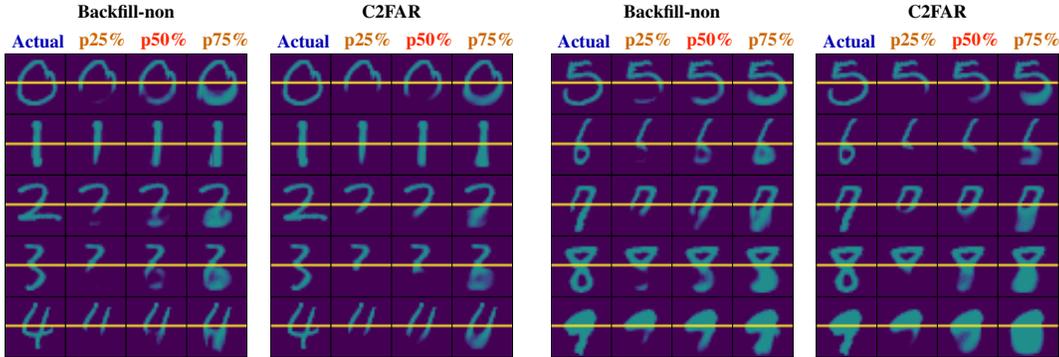


Figure 4: Percentiles of forecast distributions for first occurrence of 0 . . . 9 in *mnist* test data. Pixels before yellow line (in row-order) are history. Backfill-non more confident in bottom part of digits.

Finding 2: *SutraNets are able to effectively model both long and short-term dependencies.*

On *elec* data, where SeasNaive1d is more effective than longer-term SeasNaive1w (Table 2), modeling long-term dependencies remains important (Fig. 3). Backfill-alt is able to capture both short-term changes and long-term seasonality, the latter of which is neglected by C2FAR. In plots of the *mnist* forecasts (Fig. 4), we see that compared to C2FAR, Backfill-non generates more non-zero values in the bottom parts of digits 0, 2, 5, etc., e.g., at p50%, implying the networks have more certainty or *confidence* (i.e., tighter prediction intervals) on the digit being generated. Also, Table 2 results indicate that discrepancy is the predominant problem on *mnist*; this may help explain why autoregressive models like PixelRNN [81] score very high in likelihood-based evaluation (where there is no sampling), but produce less-realistic samples compared to GANs. SutraNets may therefore provide a useful tool for reducing training/inference discrepancy in a variety of image generators.

Finding 3: *SutraNets are more accurate than C2FAR across all forecast horizons.*

Autoregressive models are normally thought to accumulate and compound [56] errors “along the generated sequence” [5]. We therefore expect the accuracy gap between SutraNets and C2FAR to grow across horizons. Instead, on *elec* and *traff*, we already see a large gap over the first 24 hours (supplemental §B.3). Improvements in both informational asymmetry (§1) and signal path evidently enable gains across all steps. Low2HighFreq, in contrast, is only accurate at horizons corresponding to its low-freq forecast, while Backfill-alt is slightly better at horizons corresponding to its first sub-series. This suggests errors do “accumulate” to a small extent across sub-series autoregression.

Finding 4: *Deeper models enable improved long-sequence forecasting, for SutraNets and C2FAR.*

In particular, going from 1 to 2 layers greatly improves accuracy on *traff* (Table 3), a dataset where signal path is the predominant problem. Smaller improvements are also seen on *elec*. Since both LSTM stacking and use of SutraNets increase complexity, a natural question is which approach offers the best computational tradeoff for a desired level of accuracy. Fig. 5 shows that for a given number of parameters — and especially for a given inference time — SutraNets are able to obtain better

Table 3: ND% decreases for SutraNets & C2FAR as number of RNN layers/hidden units increases.

Dataset	<i>elec</i>				<i>traff</i>				
	#layers	1	2	3	4	1	2	3	4
#hidden	64	128	256	256	64	128	256	256	
C2FAR	10.6	10.2	9.9	9.7	19.3	14.7	14.8	14.4	
Regular-alt	9.9	9.3	9.0	9.2	15.5	14.3	14.3	14.2	
Backfill-alt	9.3	9.3	9.0	9.3	15.3	14.2	14.1	14.0	
Backfill-non	9.3	9.1	9.5	9.3	15.7	14.4	14.4	14.4	

Table 4: ND% lowest when #sub-series divides into seasonal period (24), but alternating models less sensitive on *elec* (cf. C2FAR with 10.6% on *elec*, 19.3% on *traff*.)

Dataset	<i>elec</i>			<i>traff</i>		
	#subseries	6	7	12	6	7
Regular-alt	9.9	10.0	10.0	15.5	17.2	15.0
Backfill-alt	9.3	9.4	9.4	15.3	16.9	14.9
Backfill-non	9.3	10.0	9.5	15.7	17.3	15.4

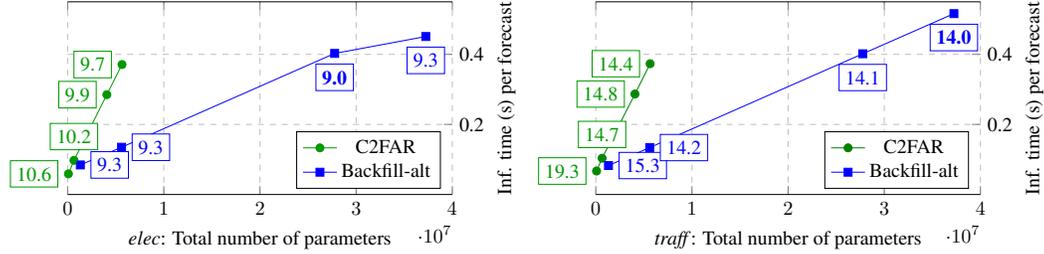


Figure 5: Total parameters (x-axis) and average inference time per forecast (y-axis, lower better) for C2FAR and Backfill-alt systems of Table 3, for *elec* (left) and *traff* (right). ND% labeled at each point in boxes. Backfill-alt enables more parameters and better ND% at equivalent speeds.

ND% (plots of memory usage and of training time are similar, see supplemental). SutraNet speed and memory usage scale with the complexity of each RNN rather than the number of sub-series RNNs (§3); they can consequently make use of $K \times$ more parameters without $K \times$ the cost.

Finding 5: *Backfill-alt is robust across different configurations, including when we vary K .*

We also experimented with varying the number of sub-series. For all $K > 1$, SutraNets work much better than standard C2FAR (Table 4). Interestingly, $K=12$ led to the highest accuracy on *traff*, where the signal path problem predominates. When K does not divide evenly into the seasonal period, non-alternating models have a missing info problem (e.g., they miss predictive preceding values at the same hour-of-day). Both Backfill-alt and Regular-alt can access these values through covariate features (cf. Fig. 2f), and hence perform relatively better when $K=7$ on *elec*. Backfill-alt is relatively weaker at $K=7$ when there is less redundancy (*traff*), although still superior to standard C2FAR.

See the supplemental for further experimental details and results. Code for SutraNets is available at https://github.com/huaweicloud/c2far_forecasting/wiki/SutraNets.

5 Conclusion

Probabilistic forecasting of long time series is an under-studied problem, with most recent work focusing on point prediction. We presented SutraNets, a novel autoregressive approach to probabilistic forecasting of long-sequence time series. SutraNets convert a univariate series into a K -dimensional multivariate series, each dimension comprising a distinct every- K th-value sub-series of the original sequence. The SutraNet autoregressive model generates each sub-series conditional on both its own prior values, and on other sub-series, ensuring coherent output samples. From these samples, an estimate of the full joint probability of future values is made. SutraNets leverage C2FAR under-the-hood, allowing for efficient representation of time series amplitudes, and low-overhead input encoding of covariate sub-series. SutraNets provide a holistic solution to challenges in long-sequence sampling, reducing training/inference discrepancy (by generating conditional on longer-term information) and signal path distances (by a factor of K). Experimentally, SutraNets achieve state-of-the-art results on a variety of long-sequence forecasting data sets, while demonstrating similar speed and memory requirements as standard sequence models, even when using $K \times$ more parameters.

References

- [1] Oren Anava and Vitaly Kuznetsov. 2017. Web Traffic Time Series Forecasting. <https://www.kaggle.com/c/web-traffic-time-series-forecasting>
- [2] George Athanasopoulos, Rob J Hyndman, Nikolaos Kourentzes, and Fotios Petropoulos. 2017. Forecasting with temporal hierarchies. *European Journal of Operational Research* 262, 1 (2017), 60–74.
- [3] Pierre Baldi and Peter J Sadowski. 2013. Understanding Dropout. *Advances in neural information processing systems* 26 (2013).
- [4] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [5] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. *Advances in Neural Information Processing Systems* 28 (2015).
- [6] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *Journal of Machine Learning Research* 3 (2003), 1137–1155.
- [7] Shane Bergsma, Tim Zeyl, Javad Rahimipour Anaraki, and Lei Guo. 2022. C2FAR: Coarse-to-Fine Autoregressive Networks for Precise Probabilistic Forecasting. In *Advances in Neural Information Processing Systems*, Vol. 35.
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- [9] Cristian Challu, Kin G Olivares, Boris N Oreshkin, Federico Garza, Max Mergenthaler-Canseco, and Artur Dubrawski. 2022. N-HiTS: Neural Hierarchical Interpolation for Time Series Forecasting. *arXiv preprint arXiv:2201.12886* (2022).
- [10] Shiyu Chang, Yang Zhang, Wei Han, Mo Yu, Xiaoxiao Guo, Wei Tan, Xiaodong Cui, Michael Witbrock, Mark A Hasegawa-Johnson, and Thomas S Huang. 2017. Dilated recurrent neural networks. *Advances in Neural Information Processing Systems* 30 (2017).
- [11] Zipeng Chen, Qianli Ma, and Zhenxi Lin. 2021. Time-Aware Multi-Scale RNNs for Time Series Modeling.. In *International Joint Conference On Artificial Intelligence*. 2285–2291.
- [12] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509* (2019).
- [13] Gregory C Chow and An-loh Lin. 1971. Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics* (1971), 372–375.
- [14] Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. 2017. Hierarchical Multiscale Recurrent Neural Networks. *International Conference on Learning Representations* (2017).
- [15] Eli Cortez, Anand Bonde, Alexandre Muzio, Mark Russinovich, Marcus Fontoura, and Ricardo Bianchini. 2017. Resource Central: Understanding and Predicting Workloads for Improved Resource Management in Large Cloud Platforms. In *Symposium on Operating Systems Principles*. 153–167.
- [16] Aniket Didolkar, Kshitij Gupta, Anirudh Goyal, Nitesh Bharadwaj Gundavarapu, Alex M Lamb, Nan Rosemary Ke, and Yoshua Bengio. 2022. Temporal Latent Bottleneck: Synthesis of Fast and Slow Processing Mechanisms in Sequence Learning. *Advances in Neural Information Processing Systems* 35 (2022), 10505–10520.
- [17] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>

- [18] Salah El Hihi and Yoshua Bengio. 1995. Hierarchical Recurrent Neural Networks for Long-Term Dependencies. *Advances in Neural Information Processing Systems* 8 (1995).
- [19] Christos Faloutsos, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, and Yuyang Wang. 2019. Forecasting Big Time Series: Theory and Practice. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 3209–3210.
- [20] Chenyou Fan, Yuze Zhang, Yi Pan, Xiaoyue Li, Chi Zhang, Rong Yuan, Di Wu, Wensheng Wang, Jian Pei, and Heng Huang. 2019. Multi-Horizon Time Series Forecasting with Temporal Attention Learning. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2527–2535.
- [21] Jan Gasthaus, Konstantinos Benidis, Yuyang Wang, Syama Sundar Rangapuram, David Salinas, Valentin Flunkert, and Tim Januschowski. 2019. Probabilistic Forecasting with Spline Quantile Function RNNs. In *International Conference on Artificial Intelligence and Statistics*. 1901–1910.
- [22] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. *Commun. ACM* 64, 12 (2021), 86–92.
- [23] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. 2015. MADE: Masked Autoencoder for Distribution Estimation. In *International Conference on Machine Learning*. 881–889.
- [24] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- [25] Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862* (2019).
- [26] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. *Advances in Neural Information Processing Systems* (2014), 2672–2680.
- [27] Adèle Gouttes, Kashif Rasul, Mateusz Koren, Johannes Stephan, and Tofigh Naghibi. 2021. Probabilistic Time Series Forecasting with Implicit Quantile Networks. *arXiv preprint arXiv:2107.03743* (2021).
- [28] Alex Graves. 2013. Generating Sequences with Recurrent Neural Networks. *arXiv preprint arXiv:1308.0850* (2013).
- [29] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech Recognition with Deep Recurrent Neural Networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. 6645–6649.
- [30] Alex Graves and Jürgen Schmidhuber. 2008. Offline handwriting recognition with multidimensional recurrent neural networks. *Advances in Neural Information Processing Systems* 21 (2008).
- [31] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [32] Hansika Hewamalage, Klaus Ackermann, and Christoph Bergmeir. 2023. Forecast evaluation for data scientists: common pitfalls and best practices. *Data Mining and Knowledge Discovery* 37, 2 (2023), 788–832.
- [33] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [34] Rob J Hyndman and George Athanasopoulos. 2018. *Forecasting: Principles and Practice*. OTexts.

- [35] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aaron van den Oord, Sander Dieleman, and Koray Kavukcuoglu. 2018. Efficient Neural Audio Synthesis. In *International Conference on Machine Learning*. 2410–2419.
- [36] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*. 5156–5165.
- [37] Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [38] Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. 2014. A clockwork RNN. In *International conference on machine learning*. 1863–1871.
- [39] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. 2018. Modeling Long-and Short-Term Temporal Patterns with Deep Neural Networks. In *ACM SIGIR Conference on Research & Development in Information Retrieval*. 95–104.
- [40] Alex M Lamb, Anirudh Goyal Alias Parth Goyal, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. 2016. Professor forcing: A new algorithm for training recurrent networks. *Advances in Neural Information Processing Systems* 29 (2016).
- [41] Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. 2015. A Simple Way to Initialize Recurrent Networks of Rectified Linear Units. *arXiv preprint arXiv:1504.00941* (2015).
- [42] Yann LeCun, Corinna Cortes, and Christopher JC Burges. 1998. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>
- [43] Ji You Li, Jiachi Zhang, Wenchao Zhou, Yuhang Liu, Shuai Zhang, Zhuoming Xue, Ding Xu, Hua Fan, Fangyuan Zhou, and Feifei Li. 2023. Eigen: End-to-end Resource Optimization for Large-Scale Databases on the Cloud. *Proceedings of the VLDB Endowment* 16, 12 (2023), 3795–3807.
- [44] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhua Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. Enhancing the Locality and Breaking the Memory Bottleneck of Transformer on Time Series Forecasting. *Advances in Neural Information Processing Systems* 32 (2019).
- [45] Zimo Li, Yi Zhou, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. 2017. Auto-conditioned recurrent networks for extended complex human motion synthesis. *arXiv preprint arXiv:1707.05363* (2017).
- [46] Shengsheng Lin, Weiwei Lin, Wentai Wu, Feiyu Zhao, Ruichao Mo, and Haotong Zhang. 2023. SegRNN: Segment Recurrent Neural Network for Long-Term Time Series Forecasting. *arXiv preprint arXiv:2308.11200* (2023).
- [47] Shizhan Liu, Hang Yu, Cong Liao, Jianguo Li, Weiyao Lin, Alex X Liu, and Schahram Dustdar. 2021. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*.
- [48] Xing Han Lu, Aihua Liu, Shih-Chieh Fuh, Yi Lian, Liming Guo, Yi Yang, Ariane Marelli, and Yue Li. 2021. Recurrent disease progression networks for modelling risk trajectory of heart failure. *PLoS one* 16, 1 (2021).
- [49] Kiran Madhusudhanan, Johannes Burchert, Nghia Duong-Trung, Stefan Born, and Lars Schmidt-Thieme. 2021. Yformer: U-net inspired transformer architecture for far horizon time series forecasting. *arXiv preprint arXiv:2110.08255* (2021).
- [50] Gábor Melis, Chris Dyer, and Phil Blunsom. 2017. On the State of the Art of Evaluation in Neural Language Models. *arXiv preprint arXiv:1707.05589* (2017).
- [51] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Conference on Fairness, Accountability, and Transparency*. 220–229.

- [52] Srayanta Mukherjee, Devashish Shankar, Atin Ghosh, Nilam Tathawadekar, Pramod Kompalli, Sunita Sarawagi, and Krishnendu Chaudhury. 2018. AR-MDN: Associative and Recurrent Mixture Density Networks for eRetail Demand Forecasting. *arXiv preprint arXiv:1803.03800* (2018).
- [53] Kevin P Murphy. 2012. *Machine Learning: A Probabilistic Perspective*. MIT press.
- [54] Moin Nadeem, Anna Bethke, and Siva Reddy. 2020. StereoSet: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456* (2020).
- [55] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730* (2022).
- [56] James O’Neill and Danushka Bollegala. 2018. Analysing Dropout and Compounding Errors in Neural Language Models. *arXiv preprint arXiv:1811.00998* (2018).
- [57] Boris N Oreshkin, Dmitri Carпов, Nicolas Chapados, and Yoshua Bengio. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437* (2019).
- [58] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [59] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Tejani Alykhan, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems* 32 (2019).
- [60] Stephan Rabanser, Tim Januschowski, Valentin Flunkert, David Salinas, and Jan Gasthaus. 2020. The Effectiveness of Discretization in Forecasting: An Empirical Study on Neural Time Series Models. *arXiv preprint arXiv:2005.10111* (2020).
- [61] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- [62] Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. 2021. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*. PMLR, 8857–8868.
- [63] Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs Bergmann, and Roland Vollgraf. 2020. Multi-Variate Probabilistic Time Series Forecasting via Conditioned Normalizing Flows. *arXiv preprint arXiv:2002.06103* (2020).
- [64] Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. *arXiv preprint arXiv:1804.09301* (2018).
- [65] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. 2020. DeepAR: Probabilistic Forecasting with Autoregressive Recurrent Networks. *International Journal of Forecasting* 36, 3 (2020), 1181–1191.
- [66] Matteo Sangiorgio and Fabio Dercole. 2020. Robustness of LSTM neural networks for multi-step forecasting of chaotic time series. *Chaos, Solitons & Fractals* 139 (2020), 110045.
- [67] Christoph Sax and Peter Steiner. 2013. Temporal disaggregation of time series. *The R Journal* 5, 2 (2013), 80–87.
- [68] Jürgen Schmidhuber. 1992. Learning Complex, Extended Sequences Using the Principle of History Compression. *Neural Computation* 4, 2 (1992), 234–242.

- [69] Amazon Web Services. 2023. DeepAR+ Algorithm - Amazon Forecast. <https://docs.aws.amazon.com/forecast/latest/dg/aws-forecast-recipe-deeparplus.html> Accessed: 2023-05-01.
- [70] Amin Shabani, Amir Abdi, Lili Meng, and Tristan Sylvain. 2022. Scaleformer: Iterative Multi-Scale Refining Transformers for Time Series Forecasting. *arXiv preprint arXiv:2206.04038* (2022).
- [71] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326* (2019).
- [72] Slawek Smyl. 2020. A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting* 36, 1 (2020), 75–85.
- [73] Alessandro Sordoni, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A Hierarchical Recurrent Encoder-Decoder for Generative Context-Aware Query Suggestion. In *Conference on Information and Knowledge Management*. 553–562.
- [74] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [75] Arthur Suilin. 2017. Kaggle Web Traffic Time Series Forecasting: How it works. https://github.com/Arturus/kaggle-web-traffic/blob/master/how_it_works.md
- [76] Souhaib Ben Taieb and Amir F Atiya. 2015. A Bias and Variance Analysis for Multistep-Ahead Time Series Forecasting. *IEEE Transactions on Neural Networks and Learning Systems* 27, 1 (2015), 62–76.
- [77] Binh Tang and David S Matteson. 2021. Probabilistic Transformer for Time Series Analysis. *Advances in Neural Information Processing Systems* 34 (2021), 23592–23608.
- [78] Lucas Theis and Matthias Bethge. 2015. Generative image modeling using spatial LSTMs. *Advances in Neural Information Processing Systems* 28 (2015).
- [79] Hoang Duy Trinh, Lorenza Giupponi, and Paolo Dini. 2018. Mobile traffic prediction from raw data using LSTM networks. In *2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. 1827–1832.
- [80] Benigno Uribe, Marc-Alexandre Côté, Karol Gregor, Iain Murray, and Hugo Larochelle. 2016. Neural Autoregressive Distribution Estimation. *Journal of Machine Learning Research* 17, 1 (2016), 7184–7220.
- [81] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel Recurrent Neural Networks. In *International Conference on Machine Learning*. 1747–1756.
- [82] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. *Advances in Neural Information Processing Systems* 30 (2017).
- [83] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. 2017. Feudal networks for hierarchical reinforcement learning. In *International Conference on Machine Learning*. PMLR, 3540–3549.
- [84] Ruofeng Wen, Kari Torkkola, Balakrishnan Narayanaswamy, and Dhruv Madeka. 2017. A Multi-Horizon Quantile Recurrent Forecaster. *arXiv preprint arXiv:1711.11053* (2017).
- [85] Monier Williams. 1963. *Sanskrit-English Dictionary*. Oxford: Clarendon Press.
- [86] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2022. TimesNet: Temporal 2D-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186* (2022).

- [87] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. *Advances in Neural Information Processing Systems* 34 (2021).
- [88] Sifan Wu, Xi Xiao, Qianggang Ding, Peilin Zhao, Ying Wei, and Junzhou Huang. 2020. Adversarial Sparse Transformer for Time Series Forecasting. *Advances in Neural Information Processing Systems* 33 (2020), 17105–17115.
- [89] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. 2016. Temporal Regularized Matrix Factorization for High-Dimensional Time Series Prediction. *Advances in Neural Information Processing Systems* 29 (2016).
- [90] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2022. Are Transformers effective for time series forecasting? *arXiv preprint arXiv:2205.13504* (2022).
- [91] Saizheng Zhang, Yuhuai Wu, Tong Che, Zhouhan Lin, Roland Memisevic, Russ R Salakhutdinov, and Yoshua Bengio. 2016. Architectural complexity measures of recurrent neural networks. *Advances in Neural Information Processing Systems* 29 (2016).
- [92] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. In *AAAI Conference on Artificial Intelligence*. 11106–11115.
- [93] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International Conference on Machine Learning*. 27268–27286.