

Foundations of Sequence-to-Sequence Modeling for Time Series

Vitaly Kuznetsov

Google Research, New York, NY

Zelda Mariet*

Massachusetts Institute of Technology, Cambridge, MA

vitalyk@google.com

zelda@csail.mit.edu

Abstract

The availability of large amounts of time series data, paired with the performance of deep-learning algorithms on a broad class of problems, has recently led to significant interest in the use of sequence-to-sequence models for time series forecasting. We provide the first theoretical analysis of this time series forecasting framework. We include a comparison of sequence-to-sequence modeling to classical time series models, and as such our theory can serve as a quantitative guide for practitioners choosing between different modeling methodologies.

1 Introduction

Time series analysis is a critical component of real-world applications such as climate modeling, web traffic prediction, neuroscience, as well as economics and finance. In this work, we focus on the fundamental question of time series forecasting. Specifically, we study the task of forecasting the next ℓ steps of an m -dimensional time series \mathbf{Y} , where m is assumed to be very large. For example, in climate modeling, m may correspond to the number of locations at which we collect historical observations, and more generally to the number of sources which provide us with time series.

Often, the simplest way to tackle this problem is to approach it as m separate tasks, where for each of the m dimensions we build a model to forecast the univariate time series corresponding to that dimension. Auto-regressive and state-space models [8, 4, 6, 5, 12], as well as non-parametric approaches such as RNNs [3], are often used in this setting. To account for correlations between different time series, these models have also been generalized to the multivariate case [22, 23, 34, 13, 14, 1, 2, 34, 2, 36, 30, 39, 24]. In both univariate and multivariate settings, an observation at time t is treated as a single sample point, and the model tries to capture relations between observations at times t and $t + 1$. Therefore, we refer to these models as *local*.

In contrast, an alternative methodology based on treating m univariate time series as m samples drawn from some unknown distribution has also gained popularity in recent years. In this setting, each of the m dimensions of \mathbf{Y} is treated as a separate example and a single model is learned from these m observations. This model learns to map past vectors of length $T - \ell$ to corresponding future vectors of length ℓ , where T is the length of time series. LSTMs and RNNs [15] are a popular choice of model class for this setup [9, 11, 40, 21, 20, 41].¹ Consequently, we refer to this framework as *sequence-to-sequence modeling*.

While there has been progress in understanding the generalization ability of local models [38, 26, 27, 28, 17, 18, 19, 42], to the best of our knowledge the generalization properties of sequence-to-sequence modeling have not yet been studied, raising the following natural questions:

- What is the generalization ability of sequence-to-sequence models and how is it affected by the statistical properties of the underlying stochastic processes (e.g. non-stationarity, correlations)?
- When is sequence-to-sequence modeling preferable to local modeling, and vice versa?

We provide the first generalization guarantees for time series forecasting with sequence-to-sequence models. Our results are expressed in terms of natural measures of non-stationarity and correlation strength between different time series and hence explicitly depend on the key components of the learning problem.

*Authors are in alphabetical order.

¹Sequence-to-sequence models are also among the winning solutions in the recent time series forecasting competition: <https://www.kaggle.com/c/web-traffic-time-series-forecasting>.

Table 1: Summary of local, sequence-to-sequence, and hybrid models.

LEARNING MODEL	TRAINING SET	HYPOTHESIS	EXAMPLE
UNIVAR. LOCAL	$\mathbf{Z}_i = \{(Y_{t-1-p}^{t-1}(i), Y_t(i)) : p \leq t \leq T\}$	$h_i : \mathcal{Y}^p \rightarrow \mathcal{Y}$	ARIMA
MULTIVAR. LOCAL	$\mathbf{Z} = \{(Y_{t-1-p}^{t-1}(\cdot), Y_t(\cdot)) : p \leq t \leq T\}$	$h : \mathcal{Y}^{m \times p} \rightarrow \mathcal{Y}^m$	VARMA
SEQ-TO-SEQ	$\mathbf{Z} = \{(Y_1^{T-1}(i), Y_T(i)) : 1 \leq i \leq m\}$	$h : \mathcal{Y}^T \rightarrow \mathcal{Y}$	Neural nets
HYBRID	$\mathbf{Z} = \{(Y_{t-p}^{t-1}(i), Y_t(i)) : 1 \leq i \leq m, p \leq t \leq T\}$	$h : \mathcal{Y}^p \rightarrow \mathcal{Y}$	Neural nets




Figure 1: Local and sequence-to-sequence splits of a one dimensional time series into training and test examples.

We compare our generalization bounds to guarantees for local models and identify regimes under which one methodology is superior to the other. Therefore, our theory may also serve as a quantitative guide for a practitioner choosing the right modeling approach.

The rest of the paper is organized as follows: in Section 2, we formally define sequence-to-sequence and local modeling. In Section 3, we define the key tools that we require for our analysis. Generalization bounds for sequence-to-sequence models are given in Section 4. We compare sequence-to-sequence and local models in Section 5. Section 6 concludes this paper with a study of a setup that is a hybrid of the local and sequence-to-sequence models.

2 Sequence-to-sequence modeling

We begin by providing a formal definition of sequence-to-sequence modeling. The learner receives a multi-dimensional time series \mathbf{Y} , where we denote by $Y_t(i)$ the value of the i -th time series at time t . We write $Y_a^b(i)$ to denote the sequence $(Y_a(i), Y_{a+1}(i), \dots, Y_b(i))$. Similarly, we let $Y_t = (Y_t(1), \dots, Y_t(m))$ and $Y_a^b = (Y_a, \dots, Y_b)$. Note that $\mathbf{Y} \equiv Y_1^T$. In addition, the sequence Y_1^{T-1} is of a particular importance in our analysis and we denote it by \mathbf{Y}' .

The goal of the learner is to predict Y_{T+1} . We further assume that to solve this task, our input \mathbf{Y} is partitioned into a training set of m examples $\mathbf{Z} = \{Z_1, \dots, Z_m\}$, where each $Z_i = (Y_1^{T-1}(i), Y_T(i)) \in \mathcal{Y}^T$. The learner's objective is to select a hypothesis $h : \mathcal{Y}^T \rightarrow \mathcal{Y}$ from a given hypothesis set \mathcal{H} that achieves a small generalization error:

$$\mathcal{L}(h | \mathbf{Y}) = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\mathcal{D}} [L(h(Y_1^T(i)), Y_{T+1}(i)) | \mathbf{Y}],$$

where $L: \mathcal{Y} \times \mathcal{Y} \rightarrow [0, M]$ is some bounded² loss function and \mathcal{D} is the distribution of Y_{T+1} conditioned on \mathbf{Y} . In other words, the learner seeks a hypothesis h which achieves a small expected error when forecasting all m dimensions at time $T + 1$, given the full history \mathbf{Y} .

By analogy, we write $\mathcal{L}(h | \mathbf{Y}')$ the expected error on the training set (conditioned on \mathbf{Y}'), and \mathcal{D}' the distribution of Y_T conditioned on \mathbf{Y}' .

In practice, each Z_i may start at a different, arbitrary time t_i , and may furthermore include some additional features X_i , i.e. $Z_i = (Y_{t_i}^{T-1}(i), X_i, Y_T(i))$. Additionally, we are often interested in long term forecasting, i.e. predicting $Y_{T+1}^{T+\ell}(i)$ for $\ell \geq 1$. Hence, our learner searches for a hypothesis h that maps a historical sequence $Y_{t_i}^{t_i+T}(i)$ to a future sequence $Y_{t_i+T+1}^{t_i+T+\ell}(i)$ and so we refer to this approach as *sequence-to-sequence modeling*. In fact, LSTMs and RNNs are common choice for the hypothesis space \mathcal{H} [9, 11, 40,

²Most of the results in this paper can be straightforwardly extended to unbounded case assuming Y is sub-Gaussian.

21, 20, 41]. However, feed-forward NNs have also been successfully applied in this framework [33] and in practice, our definition allows for any set of functions \mathcal{H} that map input sequences to output sequences. For instance, we can train a feed-forward NN to map $Y_1^{T-1}(i)$ to $Y_T(i)$ and at inference time use $Y_2^T(i)$ as input to obtain a forecast for $Y_T(i)$.³

We contrast sequence-to-sequence modeling to the *local modeling* approach, which consists of splitting each time series $Y(i)$ into a training set $\mathbf{Z}_i = \{Z_{i,1}, \dots, Z_{i,T}\}$, where $Z_{i,t} = (Y_{t-p}^{t-1}(i), Y_t(i))$ for some $p \in \mathbb{N}$, then learning a separate hypothesis h_i for each \mathbf{Z}_i . Each h_i models relations between observations that are close in time, which is why we refer to this framework as local modeling. As in sequence-to-sequence modeling, the goal of a local learner is to achieve a small generalization error for $h_{\text{loc}} = (h_1, \dots, h_m)$, given by:

$$\mathcal{L}(h_{\text{loc}} | \mathbf{Y}) = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\mathcal{D}} [L(h_i(Y_{T-p}^T(i)), Y_{T+1}(i)) | \mathbf{Y}].$$

In order to model correlations between different time series, it is also common to split \mathbf{Y} into one single set of multivariate examples $\mathbf{Z} = \{Z_1, \dots, Z_T\}$, where $Z_i = (Y_{t-p}^{t-1}(\cdot), Y_t(\cdot))$, and to learn a single hypothesis h that maps $\mathcal{Y}^{m \times p} \rightarrow \mathcal{Y}^m$. As mentioned earlier, we consider this approach a variant of local modeling, since h in this case again models relations between observations that are close in time.

Finally, *hybrid* or *local sequence-to-sequence* models, which interpolate between local and sequence-to-sequence approaches, have also been considered in the literature [41]. In this setting, local examples in $\mathbf{Z} = \cup_{i=1}^m \mathbf{Z}_i = \cup_{i=1}^m \{Z_{i,1}, \dots, Z_{i,T}\}$ are used to train a single sequence-to-sequence model h .

Our work focuses on the statistical properties of the sequence-to-sequence model. We provide the first generalization bounds for sequence-to-sequence and hybrid models, and compare these to similar bounds for local models. This allows us to identify regimes in which one methodology is more likely to succeed than another.

Aside from learning guarantees, there are other important considerations that may lead a practitioner to choose one approach over others. For instance, the local approach is trivially parallelizable; on the other hand, when additional features X_i are available, sequence-to-sequence modeling provides an elegant solution to the *cold start problem* in which at test time we are required to make predictions on time series for which no historical data is available.

3 Rademacher complexity, correlations and non-stationarity

In the standard supervised learning scenario, it is common to assume that training and test data are drawn *i.i.d.* from some unknown distribution. However, this assumption does not hold in general for time series, where observations at different times as well as across different series may be correlated. Furthermore, the data-generating distribution may also evolve over time.

These phenomena present a significant challenge to providing guarantees in time series forecasting. To quantify non-stationarity and correlations, we introduce the notions of mixing coefficients and discrepancy, which are defined below.

The final ingredient we need to analyze sequence-to-sequence learning is the notion of Rademacher complexity, which has been previously used to characterize learning in the *i.i.d.* setting [16, 29].

3.1 Rademacher complexity

Definition 1 (Rademacher complexity). Given a family of functions \mathcal{F} and a training set $\mathbf{Z} = \{Z_1, \dots, Z_m\}$, the Rademacher complexity of \mathcal{F} conditioned on \mathbf{Y}' is given by

$$\widehat{\mathfrak{R}}_{\mathbf{Z}}(\mathcal{F}) = \mathbb{E}_{\mathbf{Z}, \sigma} \left[\max_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \sigma_i f(Z_i) \middle| \mathbf{Y}' \right]$$

³As another example, a runner-up in the Kaggle forecasting competition (<https://www.kaggle.com/c/web-traffic-time-series-forecasting>) used a combination of boosted decision trees and feed-forward networks, and as such employs the sequence-to-sequence approach.

where $\sigma_1, \dots, \sigma_m$ are i.i.d. random variables uniform on $\{-1, +1\}$. The Rademacher complexity of \mathcal{F} for sample size m is given by

$$\mathfrak{R}_m(\mathcal{F}) = \mathbb{E}_{\mathbf{Y}' } \left[\widehat{\mathfrak{R}}_{\mathbf{Z}}(\mathcal{F}) \right].$$

The Rademacher complexity has been studied for a variety of function classes. For instance, for the linear hypothesis space $\mathcal{H} = \{x \rightarrow w^\top x, \|w\|_2 \leq \Lambda\}$, $\widehat{\mathfrak{R}}_{\mathbf{Z}}$ can be upper bounded by $\widehat{\mathfrak{R}}_{\mathbf{Z}}(\mathcal{H}) \leq \frac{\Lambda}{\sqrt{m}} \max_i \|Z_i\|_2$. As another example, the hypothesis class of ReLU feed-forward neural networks with d layers and weight matrices W_k such that $\prod_{k=1}^d \|W_k\|_F \leq \gamma$ verifies $\widehat{\mathfrak{R}}_{\mathbf{Z}}(\mathcal{H}) \leq \frac{2^{d-1/2}\gamma}{\sqrt{m}} \max_i \|Z_i\|_2$ [31].

3.2 Expected mixing coefficients

To measure the strength of dependency between time series, we extend the notion of β -mixing coefficients [7] to *expected β -mixing coefficients*, which turn out to be a more appropriate measure of correlation for the sequence-to-sequence modeling setup.

Definition 2 (Expected β_{s2s} coefficients). Let $i, j \in [m] \triangleq \{1, \dots, m\}$. We define

$$\beta_{s2s}(i, j) = \mathbb{E}_{\mathbf{Y}' } \left[\|P(Y_T(i)|\mathbf{Y}')P(Y_T(j)|\mathbf{Y}') - P(Y_T(i), Y_T(j)|\mathbf{Y}')\|_{TV} \right],$$

and for a subset $I \subseteq [m]$, we define

$$\beta_{s2s}(I) = \sup_{i, j \in I} \beta_{s2s}(i, j).$$

The coefficient $\beta_{s2s}(i, j)$ captures how close $Y_{T+1}(i)$ and $Y_{T+1}(j)$ are to being independent, given \mathbf{Y}' (and averaged over all realizations of \mathbf{Y}'). We further study these coefficients in Section 4, where we derive explicit upper bounds on expected β_{s2s} -mixing coefficients for various standard classes of stochastic processes, including spatio-temporal and hierarchical time series.

We also require the notion of *tangent collections*, within which series are independent from each other:

Definition 3 (Tangent collection). Given a collection of time series $C = \{Y(1), \dots, Y(c)\}$, we define the tangent collection \tilde{C} as $\{\tilde{Y}(1), \dots, \tilde{Y}(c)\}$ such that $\tilde{Y}(i)$ is drawn according to the marginal $\Pr(Y(i))$ and such that $\tilde{Y}(i)$ and $\tilde{Y}(i')$ are independent for $i \neq i'$.

The following result enables us to reduce the analysis of correlated time series in C to analysis of independent time series in \tilde{C} .

Proposition 1. Yu [38, Corollary 2.7]. Let f be a real-valued Borel measurable function such that $0 \leq f \leq 1$. Then, we have the following guarantee:

$$\left| \mathbb{E}[f(\tilde{C})] - \mathbb{E}[f(C)] \right| \leq (|C|-1)\beta,$$

where β is the total variation distance between joint distributions of C and \tilde{C} .

3.3 Discrepancy

Various notions of discrepancy have been previously used to measure the non-stationarity of the underlying stochastic processes with respect to the hypothesis set \mathcal{H} and loss function L in the analysis of local models [18, 42]. In this work, we introduce a notion of discrepancy specifically tailored to sequence-to-sequence modeling scenario.

Definition 4 (Discrepancy). Let \mathcal{D} be the distribution of Y_{T+1} conditioned on \mathbf{Y} and let \mathcal{D}' be the distribution of Y_T conditioned on \mathbf{Y}' . We define the discrepancy Δ as

$$\begin{aligned} \Delta &= \sup_{h \in \mathcal{H}} |\mathcal{L}(h | \mathbf{Y}) - \mathcal{L}(h | \mathbf{Y}')| \\ &= \frac{1}{m} \sup_{h \in \mathcal{H}} \left| \sum_{i=1}^m \mathbb{E}_{\mathcal{D}}[L(h(Y_1^T(i)), Y_{T+1}(i)) | \mathbf{Y}] - \mathbb{E}_{\mathcal{D}'}[L(h(Y_1^{T-1}(i)), Y_T(i)) | \mathbf{Y}'] \right| \end{aligned}$$

The discrepancy forms a pseudo-metric on the space of probability distributions and can be completed to a Wasserstein metric (by extending \mathcal{H} to all Lipschitz functions). This also immediately implies that the discrepancy can be further upper bounded by the l_1 -distance and by relative entropy between conditional distributions of Y_T and Y_{T+1} (via Pinsker's inequality). However, unlike these other divergences, the discrepancy takes into account both the hypothesis set and the loss function, making it a finer measure of non-stationarity.

However, the most important property of the discrepancy is that it can be upper bounded by the related notion of *symmetric discrepancy*, which can be estimated from data.

Definition 5 (Symmetric discrepancy). We define the symmetric discrepancy Δ_s as

$$\Delta_s = \frac{1}{m} \sup_{h, h' \in \mathcal{H}} \left| \sum_{i=1}^m L(h(Y_1^T(i)), h'(Y_1^T(i))) - L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) \right|.$$

Proposition 2. Let \mathcal{H} be a hypothesis space and let L be a bounded loss function which respects the triangle inequality. Let $h \in \mathcal{H}$ be any hypothesis. Then,

$$\Delta \leq \Delta_s + \mathcal{L}(h \mid \mathbf{Y}) + \mathcal{L}(h \mid \mathbf{Y}'),$$

where

$$\mathcal{L}(h \mid \mathbf{Y}') = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\mathcal{D}'} [L(h(Y_1^{T-1}(i)), Y_T(i)) \mid \mathbf{Y}'].$$

The proof of this result (as well as all other proofs in this paper) is deferred to the supplementary material.

Note that we do not require the test labels to evaluate the Δ_s , as opposed to Δ . Since Δ_s only depends on the observed data it can be computed directly from samples, making it a useful tool to assess the non-stationarity of the problem.

Another useful property of Δ_s is that, for certain classes of stochastic processes, we can provide a direct analysis of this quantity. Define the (unconditioned) $\bar{\beta}$ coefficients to be

$$\begin{aligned} \bar{\beta}(i, j) &= \|\Pr(Y_1^T(i), Y_1^T(j)) - \Pr(Y_1^T(i)) \Pr(Y_1^T(j))\|_{TV} \\ \bar{\beta}'(i, j) &= \|\Pr(Y_1^{T-1}(i), Y_1^{T-1}(j)) - \Pr(Y_1^{T-1}(i)) \Pr(Y_1^{T-1}(j))\|_{TV} \end{aligned}$$

and as before, for a subset I of $[m]$, write $\bar{\beta}(I) = \sup_{i, j \in I} \bar{\beta}(i, j)$ (and similarly for $\bar{\beta}'$).

Lemma 1. For $\bar{\beta}$ (and $\bar{\beta}'$ similarly), we have the following upper bound:

$$\bar{\beta}(i, j) \leq \beta_{s2s}(i, j) + \mathbb{E}_{\mathbf{Y}'} \left[\text{Cov} \left(\Pr(Y_T(i) \mid \mathbf{Y}'), \Pr(Y_T(j) \mid \mathbf{Y}') \right) \right]$$

Proposition 3. Let I_1, \dots, I_k be a partition of $\{1, \dots, m\}$, C_1, \dots, C_k be the corresponding partition of \mathbf{Y} and C'_1, \dots, C'_k be the corresponding partition of \mathbf{Y}' . Write $c = \min_j |C_j|$, and define the expected discrepancy

$$\Delta_e = \sup_{h, h' \in \mathcal{H}} \left[\mathbb{E}_{\mathbf{Y}} [L(h(Y_1^T), h'(Y_1^T))] - \mathbb{E}_{\mathbf{Y}'} [L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] \right].$$

Then, we have with probability $1 - \delta$,

$$\Delta_s \leq \Delta_e + \max \left(\max_j \mathfrak{R}_{|C_j|}(\tilde{C}'_j), \max_j \mathfrak{R}_{|C_j|}(\tilde{C}_j) \right) + \sqrt{\frac{1}{2c} \log \frac{2k}{\delta - \sum_j (|I_j| - 1) [\bar{\beta}(I_j) + \bar{\beta}'(I_j)]}}.$$

The expected discrepancy Δ_e can be computed analytically for many classes of stochastic processes. For example, for stationary processes, we can show that it is negligible. Similarly, for covariance-stationary⁴ processes with linear hypothesis sets and the squared loss function, the discrepancy is once again negligible. These examples justify our use of the discrepancy as a natural measure of non-stationarity. In particular, the second example highlights the fact that the discrepancy takes into account not only the distribution of the stochastic processes but also \mathcal{H} and L .

⁴Recall that a process X_1, X_2, \dots is called stationary if for any l, k, m , the distributions of (X_k, \dots, X_{k+l}) and $(X_{k+m}, \dots, X_{k+m+l})$ are the same. Covariance stationarity is a weaker condition that requires that $\mathbb{E}[X_k]$ be independent of k and that $\mathbb{E}[X_k X_m] = f(k - m)$ for some f .

Proposition 4. *If $Y(i)$ is stationary for all $1 \leq i \leq m$, and \mathcal{H} is a hypothesis space such that $h \in \mathcal{H} : \mathcal{Y}^{T-1} \rightarrow \mathcal{Y}$ (i.e. the hypotheses only consider the last $T - 1$ values of Y), then $\Delta_e = 0$.*

Proposition 5. *If \mathbf{Y} is covariance stationary for all $1 \leq i \leq m$, L is the squared loss, and \mathcal{H} is a linear hypothesis space $\{x \rightarrow w \cdot x \mid \|w\| \in \mathbb{R}^p \leq \Lambda\}$, then $\Delta_e = 0$.*

Another insightful example is the case when $\mathcal{H} = \{h\}$: then, $\Delta = 0$ even if \mathbf{Y} is non-stationary, which illustrates that learning is trivial for trivial hypothesis sets, even in non-stationary settings.

The final example that we consider in this section is the case of non-stationary periodic time series. Remarkably, we show that the discrepancy is still negligible in this case provided that we observe all periods with equal probability.

Proposition 6. *If the $Y(i)$ are periodic with period p and the observed starting time of each $Y(i)$ is distributed uniformly at random in $[p]$, then $\Delta_e = 0$.*

4 Generalization bounds

We now present our generalization bounds for time series prediction with sequence-to-sequence models.

We write $\mathcal{F} = \{L \circ h : h \in \mathcal{H}\}$, where $f = L \circ h$ is the loss of hypothesis h given by $f(h, Z_i) = L(h(Y_1^{T-1}(i)), Y_T)$. To obtain bounds on the generalization error $\mathcal{L}(h \mid \mathbf{Y})$, we study the gap between the generalization error and the empirical error $\widehat{\mathcal{L}}(h)$ of a hypothesis h , where

$$\widehat{\mathcal{L}}(h) = \frac{1}{m} \sum_{i=1}^m f(h, Z_i).$$

That is, we aim to give a high probability bound on the supremum of the empirical process

$$\Phi(\mathbf{Y}) = \sup_h \left[\mathcal{L}(h \mid \mathbf{Y}) - \widehat{\mathcal{L}}(h) \right].$$

We take the following high-level approach: we first partition the training set \mathbf{Z} into k collections C_1, \dots, C_k such that within each collection, correlations between different time series are as weak as possible. We then analyze each collection C_j by comparing the generalization error of sequence-to-sequence learning on C_j to the sequence-to-sequence generalization error on the *tangent* collection \widetilde{C}_j .

Theorem 4.1. *Let C_1, \dots, C_k form a partition of the training input \mathbf{Z} and let I_j denote the set of indices of time series that belong to C_j . Assume that the loss function L is bounded by 1. Then, we have for any $\delta > 0$, with probability $1 - \delta$,*

$$\Phi(\mathbf{Y}) \leq \max_j \left[\widehat{\mathfrak{R}}_{\widetilde{C}_j}(\mathcal{F}) \right] + \Delta + \frac{1}{\sqrt{2 \min_j |I_j|}} \sqrt{\log \left(\frac{k}{\delta - \sum_j (|I_j| - 1) \beta_{s2s}(I_j)} \right)}.$$

Theorem 4.1 illustrates the trade-offs that are involved in sequence-to-sequence learning for time series forecasting. For example, assuming that the C_j are of the same size, if \mathcal{H} is a collection of neural networks of bounded depth and width then $\widehat{\mathfrak{R}}_{\widetilde{C}_j}(\mathcal{F}) = \mathcal{O}\left(\sqrt{kT/m}\right)$ (see Section 3). Therefore,

$$\mathcal{L}(h \mid \mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + \Delta + \mathcal{O}\left(\sqrt{\frac{kT}{m}}\right)$$

with high probability uniformly over $h \in \mathcal{H}$, provided that $\frac{m}{k} \sum_{j=1}^k \beta_{s2s}(I_j) = o(1)$. This suggests that extremely high-dimensional ($m \gg 1$) time series are beneficial for sequence-to-sequence models, whereas series with a long histories $T \gg m$ may not benefit from sequence-to-sequence learning. Note also that correlations in data reduce the effective sample size from m to m/k .

Furthermore, Theorem 4.1 indicates that balancing the complexity of the model (e.g. depth and width of a neural net) with the fit it provides to the data is critical for controlling both the discrepancy and Rademacher complexity terms. We further illustrate this bound with several concrete examples below.

4.1 Independent time series

We begin by considering the case where all dimensions of \mathbf{Y} are independent. Although this may seem a restrictive assumption, it arises in a variety of applications: in neuroscience, different dimensions may represent brain scans of different patients; in reinforcement learning, they may correspond to different trajectories of a robotic arm.

Theorem 4.2. *For a given hypothesis space \mathcal{H} and associated function family \mathcal{F} corresponding to loss function L bounded by 1, for any $\delta > 0$ with probability at least $1 - \delta$ for any $h \in \mathcal{H}$:*

$$\mathcal{L}(h|\mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + 2\mathfrak{R}_m(\mathcal{F}) + \Delta + \sqrt{\frac{\log(1/\delta)}{m}}.$$

Theorem 4.2 shows that when time series are independent, learning is not affected by correlations in the samples and can only be obstructed by the non-stationarity of the problem, captured via Δ .

4.2 Correlated time series

We now consider several concrete examples of high-dimensional time series in which different dimensions may be correlated. This setting is common in a variety of applications including stock market indicators, traffic conditions, climate observations at different locations, and energy demand.

Suppose that each $Y(i)$ is generated by the auto-regressive (AR) processes with correlated noise

$$y_{t+1}(i) = \Theta_i(y_0^t(i)) + \varepsilon_{t+1}(i) \quad (4.1)$$

where the $w_i \in \mathbb{R}^p$ are unknown parameters and the noise vectors $\varepsilon_t \in \mathbb{R}^m$ are drawn from a Gaussian distribution $\mathcal{N}(0, \Sigma)$ where, crucially, Σ is not diagonal. The following lemma is key to our analysis.

Lemma 2. *Two AR processes $Y(i), Y(j)$ generated by (4.1) such that $\sigma = \text{Cov}(Y(i), Y(j)) \leq \sigma_0 < 1$ verify $\beta_{s2s}(i, j) = \max\left(\frac{3}{2(1-\sigma_0^2)}, \frac{1}{1-2\sigma_0}\right) \sigma = \mathcal{O}(\sigma)$.*

Hierarchical time series. As our first example, we consider the case of hierarchical time series that arises in many real-world applications [37, 35]. Consider the problem of energy demand forecasting: frequently, one observes a sequence of energy demands at a variety of levels: single household, local neighborhood, city, region and country. This imposes a natural hierarchical structure on these time series.

Formally, we consider the following hierarchical scenario: a binary tree of total depth D , where time series are generated at each of the leaves. At each leaf, $Y(i)$ is given by the AR process (4.1) where we impose $\Sigma_{ij} = \left(\frac{1}{m}\right)^{d(i,j)}$ given $d(i, j)$ the length of the shortest path from either leaf to the closest common ancestor between i and j . Hence, as $d(i, j)$ increases, $Y(i)$ and $Y(j)$ grow more independent.

For the bound of Theorem 4.1 to be non-trivial, we require a partition C_1, \dots, C_k of \mathbf{Z} such that within a given C_j the time series are close to being independent. One such construction is the following: fix a depth $d \leq D$ and construct C_1, \dots, C_{2^d} such that each C_i contains exactly one time series from each sub-tree of depth $D - d$; hence, $|C_i| = 2^{D-d}$.

Lemma 2 shows that for each C_i , we have $\beta(C_i) = \mathcal{O}(m^{d-D})$. For example, setting $d = \frac{D}{2} = \frac{\log m}{2}$, it follows that for any $\delta > 0$, with probability $1 - \delta$,

$$\mathcal{L}(h|\mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + \max_j \left[\mathfrak{R}_{\tilde{C}_j}(\mathcal{F}) \right] + \Delta + \frac{1}{\sqrt{2} \sqrt[4]{m}} \sqrt{\log \left(\frac{\sqrt{m}}{\delta - \frac{m}{\mathcal{O}(\sqrt{m^{\log m}})}} \right)}.$$

Furthermore, suppose the model is a linear AR process given by $y_{t+1}(i) = w_i \cdot (y_{t-p}^t(i)) + \varepsilon_{t+1}(i)$. In that case, the underlying stochastic process is weakly stationary and by Proposition 4 our bound reduces




Figure 2: Hierarchical dependence. The collections for $d = 1$ are given by C_1 which contains the red leaves and C_2 which contains the blue right-side leaves.

to: $\mathcal{L}(h|\mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + \max_j \left[\mathfrak{R}_{\tilde{C}_j}(\mathcal{F}) \right] + \mathcal{O}\left(\frac{\sqrt{\log m}}{m^{1/4}}\right)$. By Proposition 6, similar result holds when Θ_i is periodic.

Spatio-temporal processes. Another common task is spatio-temporal forecasting, in which historical observation are available at different locations. For instance, these observations may represent temperature at different locations on the earth, as in the case of climate modeling [25, 10], or car traffic volume at different urban locations [21].

It is natural to expect the correlation between two time series to decay as the geographical distance between them increases. As a simplified example, consider that the sphere \mathbb{S}^3 is subdivided according to a geodesic grid and a time series is drawn from the center of each patch according to (4.1), also with $\Sigma_{ij} = m^{-d(i,j)}$ but this time with $d(i, j)$ equal to the (geodesic) distance between the center of two cell centers. We choose subsets C_i with the goal of minimizing the strength of dependencies between time series within each subsets. Assuming we divide the sphere into \sqrt{m} collections of approximately the same size $c = \sqrt{m}$ such that the minimal distance between two points in a set is d_0 , we obtain

$$\mathcal{L}(h | \mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + \max_j \left[\mathfrak{R}_{\tilde{C}_j}(\mathcal{F}) \right] + \Delta + \frac{1}{\sqrt{2}\sqrt[4]{m}} \sqrt{\log \left(\frac{\sqrt{m}}{\delta - \mathcal{O}(m^{1-d_0})} \right)}.$$

As in the case of hierarchical time series, Proposition 4 or Proposition 6 can be used to remove the dependence on Δ for certain families of stochastic processes.

General graphs. The reasoning for the hierarchical and spatio-temporal time series can also be extended to arbitrary graphical models. In particular, Bayes networks lend themselves well to such an analysis, with the added benefit that the structure of a Bayes net can be inferred from data.

5 Comparison to local models

This section provides comparison of learning guarantees sequence-to-sequence models with those of local models. In particular, we will compare our bounds on the generalization gap $\Phi(\mathbf{Y})$ for sequence-to-sequence models and local models, where the gap is given by

$$\Phi_{\text{loc}}(\mathbf{Y}) = \sup_{(h_1, \dots, h_m) \in H^m} \left[\mathcal{L}(h_{\text{loc}} | \mathbf{Y}) - \widehat{\mathcal{L}}(h_{\text{loc}}) \right] \quad (5.1)$$

where $\widehat{\mathcal{L}}(h_{\text{loc}})$ is the average empirical error of h_i on the sample \mathbf{Z}_i , defined as

$$\widehat{\mathcal{L}}(h_{\text{loc}}) = \frac{1}{mT} \sum_{i=1}^m \sum_{t=1}^T f(h_i, Z_{t,i})$$

where $f(h_i, Z_{t,i}) = L(h_i(Y_{t-p}^{t-1}(i)), Y_t(i))$.

To give a high probability bound for this setting, we take advantage of existing results for the single local model h_i [18]. These results are given in terms of a slightly different notion of discrepancy Δ which is defined by

$$\Delta(\mathbf{Z}_i) = \sup_{h \in \mathcal{H}} \left[\mathbb{E} \left[L(h(Y_{t-p+1}^T), Y_{T+1}) \mid Y_1^T \right] - \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[L(h(Y_{t-p}^{t-1}), Y_t) \mid Y_1^{t-1} \right] \right].$$

Another required ingredient to state these results is the expected sequential covering number $\mathbb{E}_{v \sim T(\mathbb{P})}[\mathcal{N}_1(\alpha, \mathcal{F}, v)]$. For many hypothesis sets, the log of the sequential covering number admits upper bounds similar to those that were presented earlier for the Rademacher complexity. We provide some examples below and refer the interested reader to [32] for a detailed exposition.

Theorem 5.1. *For $\delta > 0$ and $\alpha > 0$, with probability at least $1 - \delta$, for any (h_1, \dots, h_m) , and any $\alpha > 0$, we have*

$$\Phi_{\text{loc}}(\mathbf{Y}) \leq \frac{1}{m} \sum_{i=1}^m \Delta(\mathbf{Z}_i) + 2\alpha + \sqrt{\frac{2}{T} \log \frac{m \max_i \mathbb{E}_{v \sim T(\mathbf{Z}_i)}[\mathcal{N}_1(\alpha, \mathcal{F}, v)]}{\delta}}.$$

Choosing $\alpha = 1/\sqrt{T}$, we can show that, for standard local models such as the linear hypothesis space $\{x \rightarrow w^\top x, w \in \mathbb{R}^p, \|w\|_2 \leq \Lambda\}$, we have

$$\sqrt{\frac{1}{T} \log \frac{2m \mathbb{E}_{v \sim T(Z)}[\mathcal{N}_1(\alpha, \mathcal{F}, v)]}{\delta}} = \mathcal{O}\left(\sqrt{\frac{\log m}{T}}\right).$$

In this case, it follows that

$$\Phi_{\text{loc}}(\mathbf{Y}) \leq \frac{1}{m} \sum_{i=1}^m \Delta(\mathbf{Z}_i) + \mathcal{O}\left(\sqrt{\frac{\log m}{T}}\right).$$

where the last term in this bound should be compared with corresponding (non-discrepancy) terms in the bound of Theorem 4.1, which, as discussed above, scales as $\mathcal{O}(\sqrt{T/m})$ for a variety of different hypothesis sets.

This suggests that when we have access to relatively few time series compared to their length ($m \ll T$), learning to predict each time series as its own independent problem may lead to a better generalization bound. On the other hand, in extremely high-dimensional settings, when we have significantly more time series than time steps ($m \gg T$), sequence-to-sequence learning may provide superior performance. We also note that we expect the performance of sequence-to-sequence models to deteriorate as the correlation between different time series increases.

A direct comparison of bounds in Theorem 4.1 and Theorem 5.1 is complicated by the fact that discrepancies that appear in these results are different. In fact, it is possible to design examples where $\frac{1}{m} \sum_{i=1}^m \Delta(\mathbf{Z}_i)$ is constant and Δ is negligible, and vice-versa.

Consider a tent function g_b such that $g_b(s) = 2bs/T$ for $s \in [0, T/2]$ and $g_b(s) = -2bs/T + 2b$ for $s \in [T/2, T]$. Let f_b be its periodic extension to the real line, and define $\mathcal{S} = \{f_b: b \in [0, 1]\}$. Suppose that we sample uniformly m times $b \in [0, 1]$ and $s \in \{0, T/2\}$, and observe time series $f_{b_i}(s_i), \dots, f_{b_i}(s_i + T)$. Then, as we have shown in Proposition 6, Δ is negligible for sequence-to-sequence models. However, unless the model class is trivial, it can be shown that $\Delta(\mathbf{Z}_i)$ is bounded away from zero for all i .

Conversely, suppose we sample uniformly m times $b \in [0, 1]$ and observe time series $f_{b_i}(0), \dots, f_{b_i}(T/2 + 1)$. Consider a set of local models that learn an offset from the previous point $\{h: x \mapsto x + c: c \in [0, 1]\}$. It can be shown that in this case $\Delta(\mathbf{Z}_i) = 0$, whereas Δ is bounded away from zero for any non-trivial class of sequence-to-sequence models.

From a practical perspective, we can simply use Δ_s and empirical estimates of $\Delta(\mathbf{Z}_i)$ to decide whether to choose sequence-to-sequence or local models.

We conclude this section with an observation that similar results can be proved for multivariate local models with the only difference that the sample complexity of the problem scales as $\mathcal{O}(\sqrt{m/T})$, and hence these models are even more prone to the curse of dimensionality.

6 Hybrid models

In this section, we discuss models that interpolate between local and sequence-to-sequence models. This hybrid approach trains a single model h on the union of *local* training sets $\mathbf{Z}_1, \dots, \mathbf{Z}_m$ used to train m models in the local approach. The bounds that we state here require the following extension of the discrepancy to Δ_t , defined as

$$\Delta_t = \frac{1}{m} \sup_{h \in \mathcal{H}} \left| \sum_{i=1}^m \mathbb{E}_{\mathcal{D}}[L(h(Y_{t-p-1}^{t-1}(i)), Y_t(i)) | Y_1^{t-1}] - \mathbb{E}_{\mathcal{D}'}[L(h(Y_{T-p}^T(i)), Y_{T+1}(i)) | \mathbf{Y}] \right|$$

Many of the properties that were discussed for the discrepancy Δ carry over to Δ_t as well. The empirical error in this case is the same as for the local models:

$$\widehat{\mathcal{L}}(h) = \frac{1}{mT} \sum_{i=1}^m \sum_{t=1}^T f(h, Z_{t,i}).$$

Observe that one straightforward way to obtain a bound for hybrid models is to apply Theorem 5.1 with $(h, \dots, h) \in \mathcal{H}^m$. Alternatively, we can apply Theorem 4.1 at every time point $t = 1, \dots, T$.

Combining these results via union bound leads to the following learning guarantee for hybrid models.

Theorem 6.1. Let C_1, \dots, C_k form a partition of the training input \mathbf{Z} and let I_j denote the set of indices of time series that belong to C_j . Assume that the loss function L is bounded by 1. Then, we have for any $\delta > 0$, with probability $1 - \delta$, for any $h \in \mathcal{H}$ and any $\alpha > 0$

$$\mathcal{L}(h \mid \mathbf{Y}) \leq \widehat{\mathcal{L}}(h) + \min(B_1, B_2),$$

where

$$B_1 = \frac{1}{T} \sum_{t=1}^T \Delta_t + \max_j \widehat{\mathfrak{R}}_{\widehat{C}_j}(\mathcal{F}) + \frac{1}{\sqrt{2 \min_j |I_j|}} \sqrt{\log \left(\frac{2Tk}{\delta - 2 \sum_j (|I_j| - 1) \beta_{s2s}(I_j)} \right)}$$

$$B_2 = \frac{1}{m} \sum_{i=1}^m \Delta(\mathbf{Z}_i) + 2\alpha + \sqrt{\frac{2}{T} \log \frac{2m \max_i \mathbb{E}_{v \sim T(\mathbf{Z}_i)}[\mathcal{N}_1(\alpha, \mathcal{F}, v)]}{\delta}}.$$

Using the same arguments for the complexity terms as in the case of sequence-to-sequence and local models, this result shows that hybrid models may be successful when $m \gg T$ or correlations between time series are strong, as well as when $T \gg m$.

Potential costs for this model are hidden in the new discrepancy term $\frac{1}{T} \sum_{t=1}^T \Delta_t$. This term leads to different bounds depending on the particular non-stationarity in the given problem. As before this trade-off can be accessed empirically using the data-dependent version of discrepancy.

Note that the above bound does not imply that hybrid models are strictly better than local models, since using m different hypotheses h_1, \dots, h_m can help us achieve a better trade-off between $\widehat{\mathcal{L}}(h)$ and B_2 , and vice versa.

7 Conclusion

Sequence-to-sequence learning has recently become relevant as large amounts of time series data are made available for analysis, and several state-of-the-art results in time series prediction use neural network models which leverage this framework.

We formally introduce sequence-to-sequence learning for time series, a framework in which a model learns to map past sequences of length T to their next values. We provide the first generalization bounds for sequence-to-sequence modeling. Our results are stated in terms of new notions of discrepancy and expected mixing coefficients. We study these new notions for several different families of stochastic processes including stationary, weakly stationary, periodic, hierarchical and spatio-temporal time series. Furthermore, we show that our discrepancy can be computed from data, making it a useful tool for practitioners to empirically assess the non-stationarity of their problem. In particular, the discrepancy can be used to determine whether the sequence-to-sequence methodology is likely to succeed based on the inherent non-stationarity of the problem.

Furthermore, compared to the local framework for time series forecasting, in which independent models for each one-dimensional time series are learned, our analysis shows that the sample complexity of sequence-to-sequence models scales as $\mathcal{O}(\sqrt{T/m})$, providing superior guarantees when the number m of time series is significantly greater than the length T of each series, provided that different series are weakly correlated. Conversely, we show that the sample complexity of local models scales as $\mathcal{O}(\sqrt{\log(m)/T})$, and should be preferred when $m \ll T$ or when time series are strongly correlated. We also study hybrid models for which learning guarantees are favorable both when $m \gg T$ and $T \gg m$, but which have a more complex trade-off in terms of discrepancy.

As a final note, note that the analysis we have carried through is easily extended to show similar results for the sequence-to-sequence scenario when the test data includes *new* series not observed during training, as is often the case in a variety of applications.

References

- [1] Marta Banbura, Domenico Giannone, and Lucrezia Reichlin. Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92, 2010.

- [2] Sumanta Basu and George Michailidis. Regularized estimation in sparse high-dimensional time series models. *Ann. Statist.*, 43(4):1535–1567, 08 2015.
- [3] Filippo Maria Bianchi, Enrico Maiorino, Michael C. Kampffmeyer, Antonello Rizzi, and Robert Jenssen. *Recurrent Neural Networks for Short-Term Load Forecasting - An Overview and Comparative Analysis*. Springer Briefs in Computer Science. Springer, 2017.
- [4] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307 – 327, 1986.
- [5] George Edward Pelham Box and Gwilym Jenkins. *Time Series Analysis, Forecasting and Control*. Holden-Day, Incorporated, 1990.
- [6] Peter J Brockwell and Richard A Davis. *Time Series: Theory and Methods*. Springer-Verlag New York, Inc., New York, NY, USA, 1986.
- [7] P. Doukhan. *Mixing: Properties and Examples*. Lecture notes in statistics. Springer, 1994.
- [8] Robert Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
- [9] Valentin Flunkert, David Salinas, and Jan Gasthaus. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *Arxiv:1704.04110*, 2017.
- [10] Mahsa Ghafarianzadeh and Claire Monteleoni. Climate prediction via matrix completion. In *Late-Breaking Developments in the Field of Artificial Intelligence*, volume WS-13-17 of *AAAI Workshops*. AAAI, 2013.
- [11] Hardik Goel, Igor Melnyk, and Arindam Banerjee. R2N2: Residual recurrent neural networks for multivariate time series forecasting. *arXiv:1709.03159*, 2017.
- [12] James Douglas Hamilton. *Time series analysis*. Princeton Univ. Press, Princeton, NJ, 1994.
- [13] Fang Han, Huanran Lu, and Han Liu. A direct estimation of high dimensional stationary vector autoregressions. *Journal of Machine Learning Research*, 16:3115–3150, 2015.
- [14] Fang Han, Sheng Xu, and Han Liu. Rate optimal estimation of high dimensional time series. Technical report, Technical Report, Johns Hopkins University, 2015.
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667.
- [16] V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Ann. Statist.*, 30(1):1–50, 02 2002.
- [17] Vitaly Kuznetsov and Mehryar Mohri. *Generalization Bounds for Time Series Prediction with Non-stationary Processes*, pages 260–274. Springer International Publishing, Cham, 2014.
- [18] Vitaly Kuznetsov and Mehryar Mohri. Learning theory and algorithms for forecasting non-stationary time series. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 541–549. Curran Associates, Inc., 2015.
- [19] Vitaly Kuznetsov and Mehryar Mohri. Time series prediction and online learning. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 1190–1213, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR.
- [20] Nikolay Laptev, Jason Yosinski, Li Erran Li, and Slawek Smyl. Time-series extreme event forecasting with neural networks at uber. In *ICML Workshop*, 2017.
- [21] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv:1707.01926*, 2017.
- [22] Helmut Lütkepohl. Chapter 6 Forecasting with VARMA models. In G. Elliott, C.W.J. Granger, and A. Timmermann, editors, *Handbook of Economic Forecasting*, volume 1, pages 287 – 325. Elsevier, 2006.
- [23] Helmut Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer Publishing Company, Incorporated, 2007.
- [24] Yisheng Lv, Yanjie Duan, Wenwen Kang, Zhengxi Li, and Fei-Yue Wang. Traffic flow prediction with big data: A deep learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 16:865–873, 2015.

- [25] Scott McQuade and Claire Monteleoni. Global climate model tracking using geospatial neighborhoods. In Jörg Hoffmann and Bart Selman, editors, *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22–26, 2012, Toronto, Ontario, Canada*. AAAI Press, 2012.
- [26] Ron Meir and Lisa Hellerstein. Nonparametric time series prediction through adaptive model selection. In *Machine Learning*, pages 5–34, 2000.
- [27] Mehryar Mohri and Afshin Rostamizadeh. Rademacher complexity bounds for non-i.i.d. processes. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1097–1104. Curran Associates, Inc., 2009.
- [28] Mehryar Mohri and Afshin Rostamizadeh. Stability bounds for stationary ϕ -mixing and β -mixing processes. *J. Mach. Learn. Res.*, 11:789–814, March 2010. ISSN 1532-4435.
- [29] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2012.
- [30] Sahand Negahban and Martin J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Ann. Statist.*, 39(2):1069–1097, 04 2011.
- [31] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 1376–1401, Paris, France, 03–06 Jul 2015. PMLR.
- [32] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Sequential complexities and uniform martingale laws of large numbers. *Probability Theory and Related Fields*, 161(1-2):111–153, 2015.
- [33] Pablo Romeu, Francisco Zamora-Martínez, Paloma Botella-Rocamora, and Juan Pardo. *Time-Series Forecasting of Indoor Temperature Using Pre-trained Deep Neural Networks*, pages 451–458. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [34] Song Song and Peter J. Bickel. Large vector auto regressions. *arXiv:1106.3915*, 2011.
- [35] Souhaib Ben Taieb, James W. Taylor, and Rob J. Hyndman. Coherent probabilistic forecasts for hierarchical time series. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3348–3357, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [36] D. Malioutov W. Sun. Time series forecasting with shared seasonality patterns using non-negative matrix factorization. In *The 29th Annual Conference on Neural Information Processing Systems (NIPS)*. Time Series Workshop, 2015.
- [37] Shanika L Wickramasuriya, George Athanasopoulos, and Rob J Hyndman. Forecasting hierarchical and grouped time series through trace minimization. Technical Report 15/15, Monash University, Department of Econometrics and Business Statistics, 2015.
- [38] Bin Yu. Rates of convergence for empirical processes of stationary mixing sequences. *Ann. Probab.*, 22(1): 94–116, 01 1994.
- [39] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. Temporal regularized matrix factorization for high-dimensional time series prediction. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 847–855. Curran Associates, Inc., 2016.
- [40] Rose Yu, Stephan Zheng, Anima Anandkumar, and Yisong Yue. Long-term forecasting using tensor-train RNNs. *Arxiv:1711.00073*, 2017.
- [41] Lingxue Zhu and Nikolay Laptev. Deep and confident prediction for time series at Uber. *arXiv:1709.01907*, 2017.
- [42] Alexander Zimin and Christoph H. Lampert. Learning theory for conditional risk minimization. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20–22 April 2017, Fort Lauderdale, FL, USA*, pages 213–222, 2017.

A Discrepancy analysis

Proposition 2. *Let \mathcal{H} be a hypothesis space and let L be a bounded loss function which respects the triangle inequality. Let $h' \in \mathcal{H}$. Then,*

$$\Delta \leq \Delta_s + \mathcal{L}(h \mid \mathbf{Y}) + \mathcal{L}(h \mid \mathbf{Y}')$$

Proof. Let $h, h' \in \mathcal{H}$. For ease of notation, we write

$$\Delta_s(h, h', \mathbf{Y}') = \frac{1}{m} \sum_i L(h(Y_1^T(i)), h'(Y_1^T(i))) - \frac{1}{m} \sum_i L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))).$$

Applying the triangle inequality to L ,

$$\begin{aligned} \mathcal{L}(h \mid \mathbf{Y}) &= \frac{1}{m} \sum_i \mathbb{E}[L(h(Y_1^T(i)), Y_{T+1}(i)) \mid \mathbf{Y}] \\ &\leq \frac{1}{m} \sum_i L(h(Y_1^T(i)), h'(Y_1^T(i))) + \frac{1}{m} \sum_i \mathbb{E}[L(h'(Y_1^T(i)), Y_{T+1}(i)) \mid \mathbf{Y}] \\ &= \frac{1}{m} \sum_i L(h(Y_1^T(i)), h'(Y_1^T(i))) + \mathcal{L}(h' \mid \mathbf{Y}). \end{aligned}$$

Then, by definition of $\Delta_s(h, h', \mathbf{Y}')$, we have

$$\begin{aligned} \mathcal{L}(h \mid \mathbf{Y}) &\leq \frac{1}{m} \sum_i L(h(Y_1^T(i)), h'(Y_1^T(i))) - \frac{1}{m} \sum_i L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) \\ &\quad + \frac{1}{m} \sum_i L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) + \mathcal{L}(h' \mid \mathbf{Y}) \\ &\leq \Delta_s(h, h', \mathbf{Y}') + \frac{1}{m} \sum_i L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) + \mathcal{L}(h' \mid \mathbf{Y}). \end{aligned}$$

By an application of the triangle inequality to L ,

$$\begin{aligned} \mathcal{L}(h, \mathcal{D}) &\leq \Delta_s(h, h', \mathbf{Y}') + \mathcal{L}(h' \mid \mathbf{Y}) + \frac{1}{m} \sum_i \mathbb{E}[L(h(Y_1^{T-1}(i)), Y_T(i)) \mid \mathbf{Y}'] \\ &\quad + \frac{1}{m} \sum_i \mathbb{E}[L(h'(Y_1^{T-1}(i)), Y_T(i)) \mid \mathbf{Y}'] \\ &= \Delta_s(h, h', \mathbf{Y}') + \mathcal{L}(h' \mid \mathbf{Y}) + \mathcal{L}(h \mid \mathbf{Y}') + \mathcal{L}(h' \mid \mathbf{Y}'). \end{aligned}$$

Finally, we obtain

$$\mathcal{L}(h \mid \mathbf{Y}) - \mathcal{L}(h \mid \mathbf{Y}') \leq \Delta_s(h, h', \mathbf{Y}') + \mathcal{L}(h' \mid \mathbf{Y}) + \mathcal{L}(h' \mid \mathbf{Y}')$$

and the result announced in the theorem follows by taking the supremum over \mathcal{H} on both sides. \square

Proposition 3. *Let I_1, \dots, I_k be a partition of $\{1, \dots, m\}$, and C_1, \dots, C_k be the corresponding partition of $\tilde{\mathbf{Y}}$. Write $c = \min_j |C_j|$. Then we have with probability $1 - \delta$,*

$$\Delta_s \leq \Delta_e + \max \left(\max_j \mathfrak{R}_{|C_j|}(\tilde{C}'_j), \max_j \mathfrak{R}_{|C_j|}(\tilde{I}_j) \right) + \sqrt{\frac{1}{2c} \log \frac{2k}{\delta - \sum_j (|I_j| - 1)[\tilde{\beta}(I_j) + \tilde{\beta}'(I_j)]}}.$$

Proof. By definition of Δ_s ,

$$\begin{aligned}\Delta_s &= \sup_{h, h' \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \left[L(h(Y_1^T(i)), h'(Y_1^T(i))) - L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) \right] \\ &\leq \sup_{h, h' \in \mathcal{H}} \left[\frac{1}{m} \sum_{i=1}^m L(h(Y_1^T(i)), h'(Y_1^T(i))) - \mathbb{E}_Y[L(h(Y_1^T), h'(Y_1^T))] \right] \\ &\quad + \sup_{h, h' \in \mathcal{H}} \left[\mathbb{E}_Y[L(h(Y_1^T), h'(Y_1^T))] - \mathbb{E}_Y[L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] \right] \\ &\quad + \sup_{h, h' \in \mathcal{H}} \left[\mathbb{E}_Y[L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] - \frac{1}{m} \sum_{i=1}^m L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) \right]\end{aligned}$$

by sub-additivity of the supremum. Now, define

$$\phi(\mathbf{Y}) \triangleq \sup_{h, h' \in \mathcal{H}} \left[\frac{1}{m} \sum_{i=1}^m L(h(Y_1^T(i)), h'(Y_1^T(i))) - \mathbb{E}_Y[L(h(Y_1^T), h'(Y_1^T))] \right]$$

$$\psi(\mathbf{Y}') \triangleq \sup_{h, h' \in \mathcal{H}} \left[\mathbb{E}_Y[L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] - \frac{1}{m} \sum_{i=1}^m L(h(Y_1^{T-1}(i)), h'(Y_1^{T-1}(i))) \right].$$

By definition of Δ_e , we have from the previous inequality

$$\Delta_s \leq \Delta_e + \phi(\mathbf{Y}_1^T) + \psi(\mathbf{Y}_1^{T-1}).$$

We now proceed to give a high-probability bound for ϕ ; the same reasoning will yield a bound for ψ . By sub-additivity of the max,

$$\begin{aligned}\phi(\mathbf{Y}) &\leq \sum_j \frac{|C_j|}{m} \sup_{h \in \mathcal{H}} \left[\mathbb{E}_Y[f(h, Y_1^T)] - \frac{1}{|C_j|} \sum_{Y \in C_j} f(h, Y_1^T) \right] \\ &\leq \sum_j \frac{|C_j|}{m} \phi(C_j)\end{aligned}$$

and so by union bound, for $\epsilon > 0$

$$\Pr(\phi(\mathbf{Y}) > \epsilon) \leq \sum_j \Pr(\phi(C_j) > \epsilon).$$

Let $\epsilon > \max_j \mathbb{E}[\phi(\tilde{C}_j)]$ and set $\epsilon_j = \epsilon - \mathbb{E}[\phi(\tilde{C}_j)]$.

Define for time series $Y(i), Y(j)$ the mixing coefficient

$$\bar{\beta}(i, j) = \|\Pr(Y_1^T(i), Y_1^T(j)) - \Pr(Y_1^T(i)) \Pr(Y_1^T(j))\|_{TV}$$

where we also extend the usual notation to $\bar{\beta}(C_j)$.

$$\begin{aligned}\Pr(\phi(C_j) > \epsilon) &= \Pr\left(\phi(C_j) - \mathbb{E}[\phi(\tilde{C}_j)] > \epsilon_j\right) \\ &\stackrel{(a)}{\leq} \Pr\left(\phi(\tilde{C}_j) - \mathbb{E}[\phi(\tilde{C}_j)] > \epsilon_j\right) + (|I_j| - 1)\bar{\beta}(I_j) \\ &\stackrel{(b)}{\leq} e^{-2c\epsilon_j^2} + (|I_j| - 1)\bar{\beta}(I_j),\end{aligned}$$

where (a) follows by applying Prop. 1 to the indicator function of the event $\Pr(\phi(C_j) - \mathbb{E}[\phi(\tilde{C}_j)] \geq \epsilon)$, and (b) is a direct application of McDiarmid's inequality to $\phi(\tilde{C}_j) - \mathbb{E}[\phi(\tilde{C}_j)]$.

Hence, by summing over j we obtain

$$\Pr(\phi(\mathbf{Y}) > \epsilon) \leq ke^{-2 \min_j |C_j| (\epsilon - \max_j \mathbb{E}[\phi(\tilde{C}_j)])^2} + \sum_j (|I_j| - 1) \bar{\beta}(I_j)$$

and similarly

$$\Pr(\psi(\mathbf{Y}') > \epsilon) \leq ke^{-2 \min_j |C'_j| (\epsilon - \max_j \mathbb{E}[\psi(\tilde{C}'_j)])^2} + \sum_j (|I'_j| - 1) \bar{\beta}'(I'_j),$$

which finally yields

$$\begin{aligned} \Pr(\Delta_s - \Delta_e > \epsilon) &\leq \Pr(\phi(\mathbf{Y}) > \epsilon) + \Pr(\psi(\mathbf{Y}') > \epsilon) \\ &\leq 2k \exp(-2c(\epsilon - \max_j \mathbb{E}[\phi(\tilde{C}_j)], \max_j \mathbb{E}[\psi(\tilde{C}'_j)]))^2 + \sum_j (|I_j| - 1) [\bar{\beta}(I_j) + \bar{\beta}'(I'_j)], \end{aligned}$$

where we recall that we write $c = \min_j |C_j|$. We invert the previous equation by setting

$$\epsilon = \max(\max_j \mathbb{E}[\phi(\tilde{C}_j)], \max_j \mathbb{E}[\psi(\tilde{C}'_j)]) + \sqrt{\frac{1}{2c} \log \frac{2k}{\delta - \sum_j (|I_j| - 1) [\bar{\beta}(I_j) + \bar{\beta}'(I'_j)]}},$$

yielding with probability $1 - \delta$,

$$\Delta_s \leq \Delta_e + \max(\max_j \mathbb{E}[\phi(\tilde{C}_j)], \max_j \mathbb{E}[\psi(\tilde{C}'_j)]) + \sqrt{\frac{1}{2c} \log \frac{2k}{\delta - \sum_j (|I_j| - 1) [\bar{\beta}(I_j) + \bar{\beta}'(I'_j)]}}.$$

We now bound $\mathbb{E}[\phi(\tilde{C}_j)]$ by $\mathfrak{R}_{|C_j|}(\tilde{C}_j)$. A similar argument yields the bound for ψ . By definition, we have

$$\begin{aligned} \mathbb{E}[\phi(\tilde{C}_j)] &= \mathbb{E} \left[\sup_{h \in \mathcal{H}} \frac{1}{|C_j|} \sum_{Z \in \tilde{C}_j} f(h, Y_1^T(i)) - \mathbb{E}_Y[f(h, Y_1^T)] \right] \\ &= \frac{1}{|C_j|} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \sum_{Z \in \tilde{C}_j} \underbrace{f(h, Y_1^T(i)) - \mathbb{E}_Y[f(h, Y_1^T)]}_{g(h, Y_1^T(i))} \right] \\ &= \frac{1}{|C_j|} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \sum_{Z \in \tilde{C}_j} g(h, Y_1^T(i)) \right] \end{aligned}$$

Standard symmetrization arguments as those used for the proof of the famous result by [16], which hold also when data is drawn independently but not identically at random, yield

$$\mathbb{E}[\phi(\tilde{C}_j)] \leq \mathfrak{R}_{|C_j|}(\tilde{C}_j).$$

The same argument yields for ψ

$$\mathbb{E}[\psi(\tilde{C}'_j)] \leq \mathfrak{R}_{|C'_j|}(\tilde{C}'_j).$$

To conclude our proof, it only remains to prove the bound

$$\bar{\beta}(i, j) \leq \beta_{s2s}(i, j) + \mathbb{E}_{\mathbf{Y}'} \left[\text{Cov} \left(\Pr(Y_T(i) | \mathbf{Y}'), \Pr(Y_T(j) | \mathbf{Y}') \right) \right]$$

Let $Y(i), Y(j)$ be two time series, and write $X_i = \mathbb{E}[\Pr(Y_1^T(i)) \mid \mathbf{Y}']$. Then the following bound holds

$$\begin{aligned}\bar{\beta}(i, j) &= \|\Pr(Y_1^T(i), Y_1^T(j)) - \Pr(Y_1^T(i)) \Pr(Y_1^T(j))\|_{TV} \\ &= \|\mathbb{E}[\Pr(Y_1^T(i), Y_1^T(j)) \mid \mathbf{Y}'] - \mathbb{E}[X_i] \mathbb{E}[X_j]\|_{TV} \\ &= \|\mathbb{E}[\Pr(Y_1^T(i), Y_1^T(j)) \mid \mathbf{Y}_1^{T-1}] - \mathbb{E}[X_i, X_j] - \mathbb{E}[\text{Cov}(X_i, X_j)]\|_{TV} \\ &\leq \beta_{s2s}(i, j) + \mathbb{E}_{\mathbf{Y}'}[\text{Cov}(X_i, X_j)],\end{aligned}$$

which is the desired inequality. \square

We now show two useful lemmas for various specific cases of time series and hypothesis spaces.

Proposition 4. *If $Y(i)$ is stationary for all $1 \leq i \leq m$, and \mathcal{H} is a hypothesis space such that $h \in \mathcal{H} : \mathcal{Y}^{T-1} \rightarrow \mathcal{Y}$ (i.e. the hypotheses only consider the last $T - 1$ values of Y), then $\Delta_e = 0$.*

Proof. Let $h, h' \in \mathcal{H}$. For stationary $Y(i)$, we have $\Pr(Y_1^T(i)) = \Pr(Y_2^T(i))$, and so

$$\mathbb{E}[L(h(Y_2^T), h'(Y_2^T))] - \mathbb{E}[L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] = 0$$

and so taking the supremum over h, h' yields the desired result. \square

Proposition 5. *If $Y(i)$ is covariance stationary for all $1 \leq i \leq m$, L is the squared loss, and \mathcal{H} is a linear hypothesis space $\{x \rightarrow w \cdot x \mid \|w\| \in \mathbb{R}^p \leq \Lambda\}$, then $\Delta_e = 0$.*

Proof. Recall that a time series Y is covariance stationary if $\mathbb{E}_Y[Y_t]$ does not depend on t and $\mathbb{E}_Y[Y_t Y_s] = f(t - s)$ for some function f .

Let now $(h, h') \in \mathcal{H} \equiv (w, w') \in \mathbb{R}^p$. We write $\Sigma = \Sigma_2^T(Y) = \Sigma_1^T(Y)$ the covariance matrix of Y where the equality follows from covariance stationarity. Without loss of generality, we consider $p = T - 1$. Then,

$$\begin{aligned}\mathbb{E}[L(h(Y_2^T), h'(Y_2^T))] - \mathbb{E}[L(h(Y_1^{T-1}), h'(Y_1^{T-1}))] \\ = \mathbb{E}[\langle (w - w')^\top \Sigma_2^T(Y) (w - w') \rangle] - \mathbb{E}[\langle (w - w')^\top \Sigma_1^{T-1}(Y) (w - w') \rangle] \\ = 0.\end{aligned}$$

Taking the supremum over h, h' yields the desired result. \square

Proposition 6. *If the $Y(i)$ are periodic of period p and the observed starting time of each $Y(i)$ is distributed uniformly at random in $[p]$, then $\Delta_e = 0$.*

Proof. This proof is similar to the stationary case: indeed, we can write $\Pr(Y_1^{T-1}(i)) = \frac{1}{p} \Pr(Y(i))$ due to the uniform distribution on starting times. Then, by the same reasoning, we have also

$$\Pr(Y_2^T(i)) = \frac{1}{p} \Pr(Y(i)) = \Pr(Y_1^{T-1}(i)),$$

from which the result follows. \square

B Generalization bounds

Theorem 4.1. *Let \mathcal{H} be a hypothesis space, and $h \in \mathcal{H}$. Let C_1, \dots, C_k form a partition of the training input \mathbf{Y}_1^T , and consider that the loss function L is bounded by 1. Then, we have for $\delta > 0$, with probability $1 - \delta$,*

$$\Phi_{s2s}(h) \leq \Delta + \max_j \left[\mathfrak{R}_{|C_j|}(\tilde{C}_j \mid \mathbf{Y}) \right] + \frac{1}{\sqrt{2 \min_j |I_j|}} \sqrt{\log \left(\frac{k}{\delta - \sum_j (|I_j| - 1) \beta_{s2s}(I_j)} \right)}.$$

For ease of notation, we write

$$\begin{aligned}\phi(\mathbf{Y}) &= \sup_{h \in \mathcal{H}} \mathcal{L}(h \mid \mathbf{Y}') - \widehat{\mathcal{L}}(h, \mathbf{Y}) \\ &= \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \mathbb{E}[f(h, Y_1^T(i)) \mid \mathbf{Y}'] - \frac{1}{m} \sum_{i=1}^m f(h, Y_1^T(i)).\end{aligned}$$

We begin by proving the following lemma.

Lemma 3. *Let $\bar{\mathbf{Y}}$ be equal to \mathbf{Y} on all time series except for the last, where we have $\bar{Y}(m) = Y(m)$ at all times except for time $t = T$. Then*

$$|\phi(\mathbf{Y}) - \phi(\bar{\mathbf{Y}})| \leq \frac{1}{m}$$

Proof. Fix $h^* \in \mathcal{H}$. Then,

$$\begin{aligned}\mathcal{L}(h^* \mid \mathbf{Y}') - \widehat{\mathcal{L}}(h^*, \mathbf{Y}) &- \sup_{h \in \mathcal{H}} [\mathcal{L}(h \mid \bar{\mathbf{Y}}') - \widehat{\mathcal{L}}(h, \bar{\mathbf{Y}})] \\ &\leq \mathcal{L}(h^* \mid \mathbf{Y}') - \widehat{\mathcal{L}}(h^*, \mathbf{Y}) - [\mathcal{L}(h^* \mid \bar{\mathbf{Y}}') - \widehat{\mathcal{L}}(h^*, \bar{\mathbf{Y}})] \\ &\stackrel{(a)}{\leq} \widehat{\mathcal{L}}(h^*, \bar{\mathbf{Y}}) - \widehat{\mathcal{L}}(h^*, \mathbf{Y}) \\ &\leq \frac{1}{m} [f(h^*, \bar{Y}_1^T(m)) - f(h^*, Y_1^T(m))] \leq \frac{1}{m}.\end{aligned}$$

where (a) follows from the fact that $\mathbf{Y}' = \bar{\mathbf{Y}}'$ and the last inequality follows from the fact that f is bounded by 1.

By taking the supremum over h^* , the previous calculations show that $\phi(\mathbf{Y}) - \phi(\bar{\mathbf{Y}}) \leq 1/m$; by symmetry, we obtain $\phi(\bar{\mathbf{Y}}) - \phi(\mathbf{Y}) \leq 1/m$ which proves the lemma. \square

We now prove the main theorem.

Proof. Observe that the following bounds holds

$$\begin{aligned}\Phi_{s2s}(\mathbf{Y}) &= \mathcal{L}(h \mid \mathbf{Y}) - \widehat{\mathcal{L}}(h, \mathbf{Y}) \\ &\leq \sup_{h \in \mathcal{H}} [\mathcal{L}(h \mid \mathbf{Y}) - \mathcal{L}(h \mid \mathbf{Y}')] + \sup_{h \in \mathcal{H}} [\mathcal{L}(h \mid \mathbf{Y}') - \widehat{\mathcal{L}}(h, \mathbf{Y})].\end{aligned}$$

and so

$$\Phi_{s2s}(\mathbf{Y}) - \Delta \leq \underbrace{\sup_{h \in \mathcal{H}} \mathcal{L}(h, \mid \mathbf{Y}') - \widehat{\mathcal{L}}(h, \mathbf{Y})}_{\phi(\mathbf{Y})}.$$

Define $M = \max_j \mathbb{E}[\phi(\tilde{C}_j) \mid \bar{\mathbf{Y}}']$. Then,

$$\Pr\left(\Phi_{s2s}(\mathbf{Y}) - \Delta - M > \epsilon \mid \mathbf{Y}'\right) \leq \Pr(\phi(\mathbf{Y}) - M > \epsilon \mid \mathbf{Y}'). \quad (\text{B.1})$$

By sub-additivity of the supremum, we have

$$\phi(\mathbf{Y}) - M \leq \sum_j \frac{|C_j|}{m} \sup_{h \in \mathcal{H}} [\mathcal{L}(h \mid \mathbf{Y}) - \widehat{\mathcal{L}}(h, C_j) - M]$$

and so by union bound,

$$\Pr(\phi(\mathbf{Y}) - M \geq \epsilon \mid \mathbf{Y}') \leq \sum_j \Pr(\phi(C_j) - M \geq \epsilon \mid \mathbf{Y}').$$

By definition of M ,

$$\begin{aligned}
\Pr\left(\phi(C_j) - M \geq \epsilon \mid \mathbf{Y}'\right) &\leq \Pr\left(\phi(C_j) - \mathbb{E}[\phi(\tilde{C}_j) \mid \tilde{\mathbf{Y}}'] \geq \epsilon \mid \mathbf{Y}'\right) \\
&\stackrel{(a)}{\leq} \Pr\left(\phi(\tilde{C}_j) - \mathbb{E}[\phi(\tilde{C}_j) \mid \tilde{\mathbf{Y}}'] \geq \epsilon \mid \mathbf{Y}'\right) + (|I_j| - 1)\beta_{s2s}(I_j \mid \mathbf{Y}') \\
&\stackrel{(b)}{\leq} e^{-2|C_j|\epsilon^2} + (|I_j| - 1)\beta_{s2s}(I_j \mid \mathbf{Y}').
\end{aligned}$$

where (a) follows by applying Prop. 1 to the indicator function of the event $\Pr(\phi(C_j) - \mathbb{E}[\phi(\tilde{C}_j) \mid \tilde{\mathbf{Y}}'] \geq \epsilon)$, and (b) is a direct application of McDiarmid's inequality, following Lemma 3. The notation $\beta_{s2s}(I_j \mid \mathbf{Y}')$ indicates the total variation distance between the joint distributions of C_j and \tilde{C}_j conditioned on \mathbf{Y}' . In particular, we have $\mathbb{E}_{\mathbf{Y}'}\beta_{s2s}(C_j \mid \mathbf{Y}') = \beta_{s2s}(C_j)$.

Finally, taking the expectation of the previous term over all possible \mathbf{Y}' values and summing over j , we obtain

$$\Pr(\mathcal{L}(h \mid \mathbf{Y}) - \hat{\mathcal{L}}(h, \mathbf{Y}) - \mathbb{E}_{\tilde{C}_j'}[\phi(\tilde{C}_j') \mid \tilde{\mathbf{Y}}] \geq \epsilon) \leq \sum_j e^{-2|C_j|\epsilon^2} + \sum_j (|I_j| - 1)\beta_{s2s}(I_j).$$

Combining this bound with (B.1), we obtain

$$\begin{aligned}
\Pr\left(\Phi_{s2s}(\mathbf{Y}) - \Delta - M > \epsilon\right) &\leq \sum_j e^{-2|C_j|\epsilon^2} + \sum_j (|I_j| - 1)\beta_{s2s}(I_j) \\
&\leq ke^{-2\min_j |C_j|\epsilon^2} + \sum_j (|I_j| - 1)\beta_{s2s}(I_j)
\end{aligned}$$

We invert the previous equation by choosing $\delta > \sum_j (|I_j| - 1)\beta_{s2s}(I_j)$ and setting

$$\epsilon = \sqrt{\frac{\log \frac{k}{\delta - \sum_j (|I_j| - 1)\beta_{s2s}(I_j)}}{2\min_j |I_j|}},$$

which yields that with probability $1 - \delta$, we have

$$\Phi_{s2s}(Z) \leq M + \Delta + \sqrt{\frac{\log \left(\frac{k}{\delta - \sum_j (|I_j| - 1)\beta_{s2s}(I_j)} \right)}{2\min_j |I_j|}}.$$

To conclude our proof, it remains to show that

$$M \leq \mathfrak{R}_{|C_j|}(\tilde{C}_j \mid \tilde{\mathbf{Y}}').$$

$$\begin{aligned}
\mathbb{E}[\phi(\tilde{C}_j) \mid \tilde{\mathbf{Y}}'] &= \mathbb{E}\left[\sup_{h \in \mathcal{H}} \mathcal{L}(h \mid \tilde{\mathbf{Y}}') - \frac{1}{|C_j|} \sum_{i=1}^m f(h, \tilde{Y}_1^T(i)) \mid \tilde{\mathbf{Y}}'\right] \\
&= \frac{1}{|C_j|} \mathbb{E}\left[\sup_{h \in \mathcal{H}} \sum_{\tilde{Y}_1^T \in \tilde{C}_j} \mathbb{E}[f(h, \tilde{Y}_1^T) \mid \tilde{\mathbf{Y}}'] - f(h, \tilde{Y}_1^T(i)) \mid \tilde{\mathbf{Y}}'\right] \\
&\leq \frac{1}{|C_j|} \mathbb{E}\left[\sup_{h \in \mathcal{H}} \sum_{\tilde{Y}_1^T \in \tilde{C}_j} g(h, \tilde{Y}_1^T(i)) \mid \tilde{\mathbf{Y}}'\right]
\end{aligned}$$

where we've defined

$$g(h, \tilde{Y}_1^T(i)) \triangleq \mathbb{E}[f(h, \tilde{Y}_1^T(i)) \mid \tilde{\mathbf{Y}}'] - f(h, \tilde{Y}_1^T(i)).$$

Similar arguments to those used at the end of Appendix A yield the desired result, which concludes the proof of Theorem 4.1. \square

C Generalization bounds for local models

Theorem 5.1. Let $h = (h_1, \dots, h_m)$ where each h_i is a hypothesis learned via a local method to predict the univariate time series Z_i . For $\delta > 0$ and any $\alpha > 0$, we have w.p. with $1 - \delta$

$$\Phi_{loc}(\mathbf{Z}) \leq \frac{1}{m} \sum_i \Delta(Y(i)) + 2\alpha + \sqrt{\frac{2}{T} \log \frac{m \max_i (\mathbb{E}_{v \sim T(Y(i))} [\mathcal{N}_1(\alpha, \mathcal{F}, v)])}{\delta}}$$

Proof. Write

$$\Phi(Y_1^T(i)) = \sup_{h \in \mathcal{H}} \mathbb{E}[f(h, Y_1^{T+1}) \mid Y_1^T] - \frac{1}{T} \sum_{t=1}^T f(h, Y_t^{t+T}(i)).$$

By [18, Theorem 1], we have that for $\epsilon > 0$, and $1 \leq i \leq m$,

$$\Pr(\Phi(Y_1^T(i)) - \Delta(Y(i)) > \epsilon) \leq \mathbb{E}_{v \sim T(p)}[\mathcal{N}_1(\alpha, \mathcal{F}, v)] \times \exp\left(-\frac{T(\epsilon - 2\alpha)^2}{2}\right).$$

By union bound,

$$\Pr\left(\frac{1}{m} \sum_i \Phi(Y_1^T(i)) - \Delta(Y(i)) > \epsilon\right) \leq m \max_i (\mathbb{E}_{v \sim T(Y(i))}[\mathcal{N}_1(\alpha, \mathcal{F}, v)]) \times \exp\left(-\frac{T(\epsilon - 2\alpha)^2}{2}\right)$$

We invert the previous equation by letting

$$\epsilon = 2\alpha + \sqrt{\frac{2}{T} \log \frac{m \max_i (\mathbb{E}_{v \sim T(Y(i))}[\mathcal{N}_1(\alpha, \mathcal{F}, v)])}{\delta}}.$$

which yields the desired result. \square

D Analysis of expected mixing coefficients

Lemma 2. Two AR processes $Y(i), Y(j)$ generated by (4.1) such that $\sigma = \text{Cov}(Y(i), Y(j)) \leq \sigma_0 < 1$ verify $\beta_{s2s}(i, j) = \max\left(\frac{3}{2(1-\sigma_0^2)}, \frac{1}{1-2\sigma_0}\right) \sigma$.

Proof. For simplicity, we write $U = Y(i)$ and $V = Y(j)$.

Write

$$\begin{aligned} \beta &= \|P(U_T | \mathbf{Y}') P(V_T | \mathbf{Y}') - P(U_T, V_T | \mathbf{Y}')\|_{TV} \\ &= \sup_{u, v} |P(U_T = u) P(V_T = v) - P(U_T = u, V_T = v)| \\ &= \sup_{u, v} |P(U_T = u \mid U_0^{T-1}) P(V_T = v \mid v_0^{T-1}) - P(U_T = u, V_T = v \mid v_0^{T-1}, u_0^{T-1})| \\ &= \sup_{u, v} \left| \left[P(u, v \mid U_0^{T-1}, V_0^{T-1}) + f(\sigma, \delta, \epsilon) \right] - P(u, v \mid U_0^{T-1}, V_0^{T-1}) \right| \end{aligned}$$

where we've written $\delta = u - \Theta_i(U_0^{T-1})$ (and ϵ similarly for v), and we've defined

$$\begin{aligned} f(\sigma, \delta, \epsilon) &= P(u \mid U_0^{T-1}) P(v \mid V_0^{T-1}) - P(u, v \mid U_0^{T-1}, V_0^{T-1}) \\ &= e^{-\frac{1}{2}(\delta^2 + \epsilon^2)} - \frac{1}{1 - \sigma^2} e^{-\frac{1}{2} \frac{1}{1 - \sigma^2} (\delta^2 + \epsilon^2 - 2\sigma\epsilon\delta)}. \end{aligned}$$

Assuming we can bound $f(\sigma, \delta, \epsilon)$ by a function $g(\sigma)$ independent of δ, ϵ , we can then derive a bound on β .

Let $x = \sqrt{\delta^2 + \epsilon^2}$ be a measure of how far the AR process noises lie from their mean $\mu = 0$. Using the inequality

$$|\delta\epsilon| \leq \delta^2 + \epsilon^2,$$

we proceed to bound $|f(\sigma, \delta, \epsilon)|$ by bounding f and $-f$.

$$\begin{aligned} f(\sigma, \delta, \epsilon) &\leq e^{-\frac{1}{2}(\delta^2 + \epsilon^2)} - e^{-\frac{1}{2}\frac{1}{1-\sigma^2}(\delta^2 + \epsilon^2 + 2\sigma|\delta\epsilon|)} \\ &\leq e^{-\frac{1}{2}x^2} - e^{-\frac{1}{2}\frac{1}{1-\sigma^2}(1+2\sigma)x^2} \\ &\leq e^{-\frac{1}{2}x^2} \left(1 - e^{-\frac{1}{2}\frac{2\sigma + \sigma^2}{1-\sigma^2}x^2}\right) \end{aligned}$$

Using the inequality $1 - x \leq e^{-x}$, it then follows that

$$\begin{aligned} f(\sigma, \delta, \epsilon) &\leq e^{-\frac{1}{2}x^2} \left(1 - \left(1 - \frac{1}{2}\frac{2\sigma + \sigma^2}{1-\sigma^2}x^2\right)\right) \\ &\leq \frac{1}{2}\frac{3}{1-\sigma^2}\sigma x^2 e^{-\frac{1}{2}x^2} \\ &\stackrel{(a)}{\leq} \frac{3}{e(1-\sigma^2)}\sigma \end{aligned} \tag{D.1}$$

where inequality (a) follows from the fact that $y \rightarrow ye^{-y}$ is bounded by $1/e$.

Similarly, we now bound $-f$:

$$\begin{aligned} -f(\sigma, \delta, \epsilon) &\leq \frac{1}{1-\sigma^2} e^{-\frac{1}{2}\frac{1}{1-\sigma^2}(\delta^2 + \epsilon^2 - 2\sigma|\delta\epsilon|)} - e^{-\frac{1}{2}(\delta^2 + \epsilon^2)} \\ &\leq \frac{1}{1-\sigma^2} e^{-\frac{1}{2}\frac{1-2\sigma}{1-\sigma^2}x^2} - e^{-\frac{1}{2}x^2} \\ &\leq \frac{1}{1-\sigma^2} e^{-\frac{1}{2}(1-2\sigma)x^2} - e^{-\frac{1}{2}x^2}. \end{aligned}$$

One shows easily that this last function reaches its maximum for $x_0^2 = \frac{1}{\sigma} \log\left(\frac{1-\sigma^2}{1-2\sigma}\right)$, at which point it verifies

$$-f(\sigma, x_0) = \frac{2\sigma}{1-2\sigma} e^{-\frac{1}{2\sigma} \log\left(\frac{1-\sigma^2}{1-2\sigma}\right)} \leq \frac{2\sigma}{1-2\sigma} \tag{D.2}$$

Putting (D.1) and (D.2) together, we obtain

$$\begin{aligned} |f(\sigma, \delta, \epsilon)| &\leq \sigma \max\left(\frac{3}{e(1-\sigma^2)}, \frac{1}{1-2\sigma}\right) \\ &\leq \max\left(\frac{3}{2(1-\sigma_0^2)}, \frac{1}{1-2\sigma_0}\right) \sigma \end{aligned}$$

Taking the expectation over all possible realizations of \mathbf{Y}' yields the desired result. \square

E Generalization bounds for hybrid models

Recall that \mathbf{Y} contains $m' = mT$ examples, which we denote $Y_{t-p}^t(i)$ for $1 \leq i \leq m$ and $1 \leq t \leq T$ (when $t - p < 0$, we truncate the time series appropriately). We define

$$\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}) = \frac{1}{m} \sum_{i=1}^m \mathbb{E}[L(h(Y_{T-p+1}^T(i)), Y_{T+1}(i)) | \mathbf{Y}]$$

$$\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}') = \frac{1}{m} \sum_{i=1}^m \frac{1}{T} \sum_{t=1}^T \mathbb{E}[L(h(Y_{t-p}^{t-1}(i)), Y_t(i)) | \mathbf{Y}']$$

$$\widehat{\mathcal{L}}_{\text{hyb}}(h) = \frac{1}{m} \sum_{i=1}^m \frac{1}{T} \sum_{t=1}^T L(h(Y_{t-p}^{t-1}(i)), Y_t(i))$$

where we note that here \mathbf{Y}' indicates each of the mT training samples excluding their last time point.

Proof. Observe that the following chain of inequalities holds

$$\begin{aligned} \Phi_{\text{hyb}}(\mathbf{Y}) &= \sup_{h \in \mathcal{H}} \mathcal{L}_{\text{hyb}}(h | \mathbf{Y}) - \widehat{\mathcal{L}}_{\text{hyb}}(h) \\ &\leq \sup_{h \in \mathcal{H}} \left[\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}) - \mathcal{L}_{\text{hyb}}(h | \mathbf{Y}') \right] + \sup_{h \in \mathcal{H}} \left[\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}') - \widehat{\mathcal{L}}_{\text{hyb}}(h, \mathbf{Y}) \right]. \\ &\leq \frac{1}{T} \sum_{t=1}^T \sup_{h \in \mathcal{H}} \left[\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}) - \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\mathcal{D}'} [L(h(Y_{t-p}^{t-1}(i)), Y_t(i)) | \mathbf{Y}'] \right] + \sup_{h \in \mathcal{H}} \left[\mathcal{L}_{\text{hyb}}(h | \mathbf{Y}') - \widehat{\mathcal{L}}_{\text{hyb}}(h, \mathbf{Y}) \right]. \end{aligned}$$

and so

$$\Phi_{\text{hyb}}(\mathbf{Y}) - \frac{1}{T} \sum_t \Delta_t \leq \underbrace{\sup_{h \in \mathcal{H}} \mathcal{L}_{\text{hyb}}(h, | \mathbf{Y}') - \widehat{\mathcal{L}}_{\text{hyb}}(h, \mathbf{Y})}_{\phi(\mathbf{Y})}.$$

Then, following the exact same reasoning as above for $\Phi_{s_{2s}}$ shows that for $\delta > 0$, we have with probability $1 - \delta/2$

$$\Phi_{\text{hyb}}(\mathbf{Y}) \leq \underbrace{\max_j \widehat{\mathfrak{R}}_{\tilde{C}_j}(\mathcal{F}) + \frac{1}{T} \sum_t \Delta_t + \sqrt{\frac{\log \left(\frac{2k}{\delta - \sum_j (|I_j| - 1)\beta(I_j)} \right)}}{2 \min_j |I_j|}}_{B_1}$$

However, upper bounding Φ_{hyb} can also be approached using the same techniques as Kuznetsov and Mohri [18], which we now describe. Let $\alpha > 0$. For a given h , computing $\mathcal{L}_{\text{hyb}}(h, \mathbf{Y})$ is similar in expectation to running h on each of the m time series, yielding for each time series $Y_{T-p+1}^T(i)$ the bound

$$\begin{aligned} \mathbb{E}[L(h(Y_{T-p+1}^T(i)), Y_{T+1}(i)) | \mathbf{Y}] &\leq \frac{1}{T} \sum_{t=1}^T L(h(Y_{t-p}^{t-1}(i)), Y_t(i)) + \Delta(\mathbf{Y}_i) + 2\alpha \\ &\quad + \sqrt{\frac{2}{T} \log \frac{\max_i \mathbb{E}_{v \sim T(\mathbf{Y}_i)} [\mathcal{N}_i(\alpha, \mathcal{F}, v)]}{\delta}} \end{aligned}$$

and so by union bound, as above, we obtain with probability $1 - \delta/2$

$$\Phi_{\text{hyb}}(\mathbf{Y}) \leq \underbrace{\frac{1}{m} \sum \Delta(\mathbf{Y}_i) + 2\alpha + \sqrt{\frac{2}{T} \log \frac{2m \max_i \mathbb{E}_{v \sim T(\mathbf{Y}_i)} [\mathcal{N}_i(\alpha, \mathcal{F}, v)]}{\delta}}}_{B_2}$$

We conclude by a final union bound on the event $\{\Phi_{\text{hyb}}(\mathbf{Y}) \geq B_1 \cup \Phi_{\text{hyb}}(\mathbf{Y}) \geq B_2\}$, we obtain with probability $1 - \delta$,

$$\Phi_{\text{hyb}}(\mathbf{Y}) \leq \min(B_1, B_2)$$

□