

End-to-End Modeling of Hierarchical Time Series Using Autoregressive Transformer and Conditional Normalizing Flow-based Reconciliation

Shiyu Wang*, Fan Zhou*, Yinbo Sun*

Lintao Ma, James Zhang, Yangfei Zheng, Bo Zheng, Lei Lei, Yun Hu

Ant Group

Hangzhou, China

{weiming.wsy,hanlian.zf,yinbo.syb,lintao.mlt,shijun.wang,james.z,yangfei.zyf,guangyuan.jason.ll,huyun.h}@antgroup.com

ABSTRACT

Multivariate time series forecasting with hierarchical structure is pervasive in real-world applications, demanding not only predicting each level of the hierarchy, but also reconciling all forecasts to ensure coherency, i.e., the forecasts should satisfy the hierarchical aggregation constraints. Moreover, the disparities of statistical characteristics between levels can be huge, worsened by non-Gaussian distributions and non-linear correlations. To this extent, we propose a novel end-to-end hierarchical time series forecasting model, based on conditioned normalizing flow-based autoregressive transformer reconciliation, to represent complex data distribution while simultaneously reconciling the forecasts to ensure coherency. Unlike other state-of-the-art methods, we achieve the forecasting and reconciliation simultaneously without requiring any explicit post-processing step. In addition, by harnessing the power of deep model, we do not rely on any assumption such as unbiased estimates or Gaussian distribution. Our evaluation experiments are conducted on four real-world hierarchical datasets from different industrial domains (three public ones and a dataset from the application servers of Alipay's data center) and the preliminary results demonstrate efficacy of our proposed method.

CCS CONCEPTS

• Information systems → Temporal data; • Mathematics of computing → Time series analysis; • Computing methodologies → Neural networks.

KEYWORDS

hierarchical time series, transformer reconciliation, conditioned normalizing flow

ACM Reference Format:

Shiyu Wang*, Fan Zhou*, Yinbo Sun* and Lintao Ma, James Zhang, Yangfei Zheng, Bo Zheng, Lei Lei, Yun Hu. 2018. End-to-End Modeling of Hierarchical Time Series Using Autoregressive Transformer and Conditional Normalizing Flow-based Reconciliation. In *Proceedings of Make sure to enter*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

the correct conference title from your rights confirmation email (Conference).
ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Many real-world applications [1, 5, 10, 17] involve simultaneously forecasting multiple time series that are hierarchically related via aggregation operations, e.g., department sales of multiple stores at different locations and traffic flow in hierarchical regions. These time series not only interact with each other in the hierarchy, but also imply coherency, i.e., time series at upper levels are the aggregation/summation of those at lower levels. For instance, as shown in Figure 1 (i.e., the hierarchical structure of Australian domestic tourism demand [3]), the data contains 4 levels from top to bottom including, 1 country, 7 states, 27 zones, and 82 regions, among which, different levels of forecasts have distinct goals. Specifically, bottom-level forecasts are often about specific demand to help with regional government decisions, while upper-level forecasts look at the macro perspective to assist national strategies. On top of the general statistical disparities between levels in real-world hierarchies, the entanglement of their interactions and correlations presents a great challenge to the prediction model. Therefore, how to exploit the information among hierarchical time series, while learning from non-Gaussian data distributions and non-linear correlations, becomes a critical task for improving the prediction accuracy. Furthermore, coherency constraints [25] also add more complication to the prediction model. The straight-forward methods to utilize

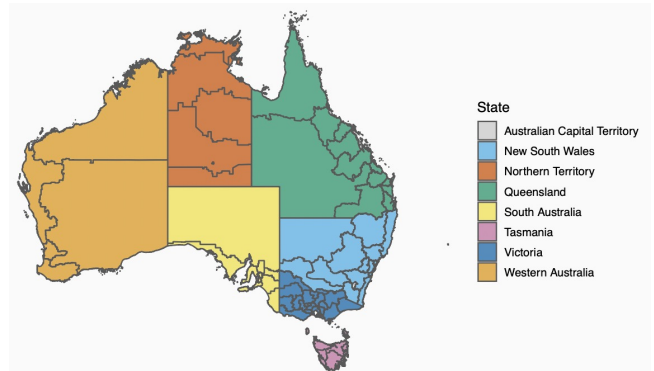


Figure 1: Hierarchical time series structures for Australian domestic tourism, aggregated by geographical locations

hierarchical structure include *bottom-up* forecasting and *top-down*

forecasting. As their names suggest, these methods make individual predictions at the bottom or upper levels and then aggregate according to the hierarchical structure. Although these methods naturally satisfy coherency, they are not able to consider statistics of all levels for prediction at the same time. Specifically, when forecasting from a single level of the aggregation structure, these methods either aggregated or disaggregated to obtain forecasts of all the rest levels. As a matter of fact, the statistical characteristics at different levels can be drastically distinct, e.g., time series of upper levels tends to be more stationary, while those at the bottom levels are often more fluctuant, for which a flexible adaptation at different levels would be ideal to better utilize the information at all levels.

To address these problems, the reconciliation method, revising the predictions for coherency after forecasting all of the series individually, has become popular in the recent research [8]. These works usually follow the two-stage approach: First, independently forecast all the series (i.e., generate incoherent base forecasts); Second, reconcile the base forecast by applying forecast combination using a bottom-up procedure. In other words, the base forecasts are adjusted so that they become coherent. MinT [28] is one of the representative methods, which reconciles the base forecasts via the optimal combination with minimum variance among all unbiased revised forecasts.

Most previous works, reconciling forecasts of all levels to ensure coherency, face the following challenges: **(i)** The base forecast is obtained independently without any shared information from other time series. **(ii)** State-of-the-art methods rely on strong statistical assumptions, such as, unbiased forecasts and Gaussian noises, but the data distributions in real-world are mostly non-Gaussian/non-linear across the hierarchy, which calls for a method that can transform data into Gaussian space where tractable methods can be applied. **(iii)** The two-stage approaches reconcile the base forecast without any regard to the learned model parameters, and cannot fully utilize the power of deep parametric models, resulting in the lack of information sharing between the process of prediction and reconciliation. **(iv)** Most methods only focus on generating point estimates, but the probabilistic forecasts are often necessary in practice to facilitate the subsequent decision-making processes.

In this paper, we present a novel end-to-end approach that tackles forecasting and reconciliation simultaneously for hierarchical time series, without requiring any explicit post-processing step, by combining the recently popular *autoregressive transformer* [11] and *conditioned normalizing flow* [15] to generate the coherent probabilistic forecasts with the state-of-the-art performance.

Specifically, we **first** obtain the base forecast via the autoregressive transformer, modeling the multivariate time series of all-levels. Using encoder-decoder transformer structure, which has been successful in recent advances in the multivariate time series forecasting [30], we achieve the information fusion of all levels in hierarchy via globally shared parameters, while benefiting from the representation power of the autoregressive transformer model. Transformer models have also shown superior performance in capturing long-range dependency than RNN-based models. **Second**, we reconcile the base forecasts into coherent forecasts via conditioned normalizing flow (CNF) [15] with bottom-up aggregation matrix. Since we need to model complex statistical properties in hierarchical data for our probabilistic forecasting, normalizing flow (NF), a proven

powerful density approximator, becomes our natural choice. By extending NF to the conditional case, we can incorporate base forecasts from all levels into CNF as additional condition for the latent space, leveraging the information available across all levels for reconciliation, while modeling the non-Gaussian distributions and non-linear correlations in the hierarchy to obtain probabilistic forecasts. **Finally**, throughout the whole process, we combined forecasting and reconciliation simultaneously for end-to-end training, while ensuring the continuity and differentiability at all steps. Moreover, our framework can accommodate different loss functions besides log-likelihood, and by sampling from the forecast distribution, sufficient statistics can also be obtained via the empirical distribution to facilitate more complicated optimization objectives.

Our Contributions. We summarize our contributions as follows:

- A multivariate autoregressive transformer architecture, modeling each time series simultaneously via globally shared parameters
- A novel reconciliation method via CNF with bottom-up aggregation matrix to integrate information of all level in the hierarchy for coherent probabilistic forecast without relying on any statistical assumption
- An end-to-end learning framework of hierarchical time series achieving forecasting and reconciliation simultaneously, without requiring any explicit post-processing step
- Extensive experiments on real-world hierarchical datasets from various industrial domains and real-life deployment for traffic forecasting on application servers from Alipay’s data center.

In the following content, we introduce the involved background and review related work in Section 2 and Section 3. We then present our proposed method and describe the procedure of training and inference in Section 4. Finally in Section 5, we analyse the non-Gaussian/non-linear properties of real-world hierarchical data and demonstrate the advantages of our approach through experiments and ablation study.

2 BACKGROUND

2.1 Hierarchical Time Series and Reconciliation

The hierarchical time series can be expressed as a tree structure (see Figure 2) with linear aggregation constraints, represented by aggregation matrix $S \in \mathbb{R}^{n \times m}$ (m is number of bottom-level nodes, n is total number of nodes). Each node in the hierarchy represents one time series, to be predicted over a time horizon. Given a time horizon $t \in \{1, 2, \dots, T\}$, we use $y_{t,i} \in \mathbb{R}$ to denote the values of a multivariate hierarchical time series, where $i \in \{1, 2, \dots, n\}$ is the index of the individual univariate time series. Here we assume that the index i of the individual time series abides by the level-order traversal of the hierarchical tree, going from left to right at each level, and we use $x_{t,i}$ to denote time-varying covariate vectors associated with each univariate time series i at time step t .

In our tree hierarchy, the time series of leaf nodes are called the *bottom-level* series $b_t \in \mathbb{R}^m$, and those of the remaining nodes are termed the *upper-level* series $u_t \in \mathbb{R}^r$. Obviously, the total number of nodes $n = r + m$, and $y_t = [u_t, b_t]^T \in \mathbb{R}^n$ contains observations at time t for all levels, which satisfy (by an aggregation matrix




Figure 2: Example of hierarchical time series structure for $n=7$ time series with $m=4$ bottom-level series and $r=3$ upper-level series

$S \in \{0, 1\}^{n \times m}$:

$$y_t = [u_t, b_t]^T \Leftrightarrow y_t = S b_t \Leftrightarrow \begin{bmatrix} u_t \\ b_t \end{bmatrix} = S b_t, \quad (1)$$

for every time step t . For example, in Figure 2, the total number of series in the hierarchy is $n = r + m = 3 + 4$, i.e., the number of series at the bottom-level is $m = 4$ and the number of upper-level series is $r = 3$. For every time step t , $u_t = [y_{A,t}, y_{B,t}, y_{C,t}] \in \mathbb{R}^3$ and $b_t = [y_{D,t}, y_{E,t}, y_{F,t}, y_{G,t}] \in \mathbb{R}^4$, the $y_{A,t} = y_{B,t} + y_{C,t} = y_{D,t} + y_{E,t} + y_{F,t} + y_{G,t}$, and the aggregation matrix $S \in \{0, 1\}^{7 \times 4}$.

Reconciliation of Hierarchical Forecasting. For given multivariate hierarchical time series, we first consider ignoring the aggregation constraints and forecasting all time series separately, which are called *base forecasts* denoted by $\hat{y}_t \in \mathbb{R}^n$, where h is the forecast horizon. Then all forecasting approaches for hierarchical structures can be represented as

$$\tilde{y}_t = S P \hat{y}_t, \quad (2)$$

where $P \in \mathbb{R}^{m \times n}$ is a matrix that projects the base forecasts (of dimension n) into the bottom-forecast (of dimension m), which is then summed up by the aggregation matrix $S \in \{0, 1\}^{n \times m}$ using the aggregation structure to produce a set of *coherent forecasts* $\tilde{y}_t \in \mathbb{R}^n$, which satisfy the aggregation constraints. Equation 2 can be easily extended to h -period-ahead forecast

$$\tilde{y}_{t+h} = S P y_{t+h}.$$

Typically, the base and coherent forecasts can be linked by Equation 2 above, which is also called **reconciliation process**. A main effort of the state-of-the-art methods is to improve the reconciliation process, and in Section 3, we will review the related work in this domain.

2.2 Autoregressive Transformer

Recently, the encoder-decoder transformer structure has been highly successful in advancing the research on multivariate time series forecasting, enabled by its multi-head self-attention mechanism to capture both long- and short-term dependencies in time series data [16, 30]. Meanwhile, combined with classical autoregressive models for time series, we can extend the transformer structure to an autoregressive deep learning model using causal masking, which

preserves the autoregressive property by utilizing a mask that reflects the causal direction of the progressing time, i.e., masking out data from future [11].

Specifically, given \mathcal{D} , defined as a batch of time series $Y = [y_1, \dots, y_T] \in \mathbb{R}^{T \times D}$, the transformer takes in the above sequence Y , and then transforms this into dimension l to obtain the query matrix Q , the key matrix K and the value matrix V as follows:

$$Q = Y W_l^Q \quad K = Y W_l^K \quad V = Y W_l^V, \quad (3)$$

where $Q \in \mathbb{R}^{dl \times dh}$, $K \in \mathbb{R}^{dl \times dh}$, and $V \in \mathbb{R}^{dl \times dh}$ (we denote that dh is the length of time steps), and the W_l^Q, W_l^K, W_l^V are learnable parameters. After these linear transformation, the scaled dot-product attention computes the sequence of vector outputs via:

$$S = \text{Attention}(Q, K, V) = \text{softmax} \left(\frac{Q K^T}{\sqrt{d_K}} \cdot M \right) V, \quad (4)$$

where the mask matrix M can be applied to filter out right-ward attention (or future information leakage) by setting its upper-triangular elements to $-\infty$ and normalization factor d_K is the dimension of the W_h^K matrix. Finally, all outputs S are concatenated and linearly projected again into the next layer.

Transformer is commonly used in an encoder-decoder architecture, where some warm-up time series as context are passed through the encoder to train the decoder and autoregressively obtain predictions.

In this work, we employ the multivariate autoregressive transformer architecture, to model each time series simultaneously via globally shared parameters. By doing so, we achieve the information fusion of all levels in hierarchy to generate the base forecasts. In fact, in our case, *base forecasts* from the base model do not directly correspond to un-reconciled forecasts for the base series, but rather represent predictions of an unobserved latent states, which can be used for subsequent density estimations.

2.3 Density Estimation Via Normalizing Flow

Normalizing flows (NF) [20], which learn a distribution by transforming the data to samples from a tractable distribution where both sampling and density estimation can be efficient and exact, have been proven to be powerful density approximators. The change of variables formula (Equation 5 below) empowers the computation of exact likelihood [14], which is in contrast to other powerful density estimator such as Variational Autoencoders or Generative Adversarial Networks. Impressive estimation results, especially in the field of nonlinear high-dimensional data generation, have led to great popularity of flow-based deep models [12].

NF are invertible neural networks that typically transform isotropic Gaussians to characterize a more complex data distribution. They map from \mathbb{R}^D to \mathbb{R}^D such that densities p_Y on the input space $Y \in \mathbb{R}^D$ are transformed into some tractable distribution p_Z (e.g., an isotropic Gaussian) on space $Z \in \mathbb{R}^D$. This mapping function, $f: Y \rightarrow Z$ and inverse mapping function, $f^{-1}: Z \rightarrow Y$ are composed of a sequence of bijections or invertible functions, and we can express the target distribution densities $p_Y(y)$ by

$$p_Y(y) = p_Z(z) \left| \det \left(\frac{\partial f(y)}{\partial y} \right) \right|, \quad (5)$$

where $\partial f(y)/\partial y$ is the Jacobian of f at y . NF have the property that the inverse $y = f^{-1}(z)$ is easy to evaluate and computing the Jacobian determinant takes $O(D)$ time.

For mapping functions f , the bijection introduced by RealNVP architecture (the coupling layer) [6] satisfies the above properties, leaving part of its inputs unchanged, while transforming the other part via functions of the un-transformed variables (with superscript denoting the coordinate indices)

$$\begin{cases} y^{1:d} = z^{1:d} \\ y^{d+1:D} = z^{d+1:D} \odot \exp(s(z)^{1:d} + t(z^{1:d})) \end{cases}, \quad (6)$$

where \odot is an element wise product, $s(\cdot)$ is a scaling and $t(\cdot)$ is a translation function from $\mathbb{R}^D \mapsto \mathbb{R}^{D-d}$, using neural networks. To model a nonlinear density map $f(x)$, a number of coupling layers, mapping $\mathcal{Y} \mapsto \mathcal{Y}_\infty \mapsto \dots \mapsto \mathcal{Y}_{K-\infty} \mapsto \mathcal{Y}_K \mapsto \mathcal{Z}$, are composed together with unchanged dimensions.

Conditional Normalizing Flow. Inspired by the conditional extension of NF [15], with the conditional distribution $p_Y(\mathbf{y}|\mathbf{h})$, we realize that the scaling and translation function approximators do not need to be invertible [15], which means we can make the transformation dependant on condition $\mathbf{h} \in \mathbb{R}^H$. Implementing $p_Y(\mathbf{y}|\mathbf{h})$ on \mathbf{h} is straight-forward: we concatenate \mathbf{h} to both the inputs of the scaling and translation function approximators of the coupling layers, i.e., $s(\text{concat}(z^{1:d}, \mathbf{h}))$ and $t(\text{concat}(z^{1:d}, \mathbf{h}))$, which are modified to map $\mathbb{R}^{d+H} \mapsto \mathbb{R}^{D-d}$.

In our work, we can incorporate base forecasts(in Section 2.2, from the outputs of Autoregressive Transformer) from all levels into CNF as additional condition in the latent space, leveraging the information available across all levels for reconciliation and while modeling the non-Gaussian distributions and non-linear correlations in the hierarchy to obtain the coherent probabilistic forecasts. In Section 4, we will detail this process.

3 RELATED WORK

Existing hierarchical time series forecasting methods mainly follow the two-stage approach: (i) Obtain each h -period-ahead base forecasts $\hat{\mathbf{y}}_{T+h}$ independently; (ii) Reconcile the base forecasts by the reconciliation process (Equation 2) to obtain the coherent forecasts $\tilde{\mathbf{y}}_{T+h}$.

This approach has the following advantages: (1) The forecasts are coherent by construction; (2) The combination of forecasts from all levels is applied via the projection matrix \mathbf{P} , where information from all levels of hierarchy is incorporated simultaneously. The main work of the current state-of-the-art method is to improve the reconciliation process, which will be reviewed in this section.

Wickramasuriya *et al.* proposed **Mint** [28] to optimally combine the base forecasts. Specifically, assuming the base forecasts $\hat{\mathbf{y}}_{T+h}$ are unbiased, Mint computes the projection matrix $\mathbf{P} = (\mathbf{S}^T \mathbf{W}_h^{-1} \mathbf{S})^{-1} (\mathbf{S}^T \mathbf{W}_h^{-1})$ giving the minimum variance unbiased revised forecasts, i.e., minimizing $\text{tr}[\mathbf{S} \mathbf{P} \mathbf{W}_h \mathbf{P}^T \mathbf{S}^T]$ with constraint $\mathbf{S} \mathbf{P} \mathbf{S}^T = \mathbf{S}$, where \mathbf{W}_h is the covariance matrix of the h -period-ahead forecast errors $\hat{\mathbf{e}}_{T+h} = \mathbf{y}_{T+h} - \hat{\mathbf{y}}_{T+h}$ (\mathbf{S} and \mathbf{P} are matrices defined in eq. (2)). However, the covariance of errors \mathbf{W}_h is hard to obtain for general h and the strong assumption of unbiased base forecasts is normally unrealistic.

The unbiased assumption is relaxed in **Regularized Regression for Hierarchical Forecasting** [2], which also follows the two-stage approach and seeks the revised forecasts with the trade-off between bias and variance of the forecast by solving an empirical risk minimization (ERM) problem.

Probabilistic Method for Hierarchical Forecasting [19] also employs the two-stage approach, but in contrast to the above methods, it considers forecasting probability distributions rather than just the means (point forecasts). This probabilistic method starts by generating independent forecasts of the conditional marginal distributions, followed by samplings from the above distributions as the base forecasts, which are then reconciled using Equation 2. In this approach, existing reconciliation methods for point forecasts can be extended to a probabilistic setting. But the reconciliation is only applied to the samples rather than the forecast distribution, which creates another level and uncertainty with unsalable computation complexity.

One typical work in end-to-end modeling of hierarchical time series is **Hier-E2E** [22], where base forecasts are obtained using DeepVAR [23] with diagonal Gaussian distribution, followed by reconciliation using a closed-form formulation of an optimization problem, i.e., minimizing the errors of base forecast $\hat{\mathbf{y}}$ and reconciling forecast $\tilde{\mathbf{y}}$ subject to coherent constraints of hierarchical structure. The reconciliation process is as follows:

$$\begin{aligned} \tilde{\mathbf{y}}_t &= \mathbf{M} \hat{\mathbf{y}}_t \\ \mathbf{M} &:= \mathbf{I} - \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1} \mathbf{A}, \end{aligned}$$

where \mathbf{M} is a fixed matrix and \mathbf{A} is a structure matrix from the upper half of the aggregated matrix. In contrast to the previous work, the model no longer has a relationship with the predicted value \mathbf{y} . The matrix \mathbf{M} is time-invariant and can be computed offline, prior to the training. This reconciliation approach is essentially a fine-tuning of the base forecasts under the coherence constraint, and the learned model parameters are not used to revise the base forecast in the reconciliation stage.

Considering the pros and cons of the above models, we combine the advantages and propose a novel end-to-end model that can not only obtain the coherent probabilistic forecasts, but can also achieve the reconciliation in the forecast distributions (rather than just on samples), ensuring that the reconciliation is related the predicted value \mathbf{y} via integrating information from all levels, to improve the overall performance.

4 METHODOLOGY

In this section, we detail our novel approach, which combines the multivariate autoregressive transformer and CNF for coherent probabilistic forecasting.

4.1 Model

A schematic overview of our hierarchical end-to-end architecture can be found in Figure 3.

4.1.1 Multivariate Autoregressive Transformer. The encoder-decoder transformer architecture has been highly successful in advancing the research on multivariate time series forecasting, enabled by its multi-head self-attention mechanism to capture both long- and short-term dependencies, and can be further extended to

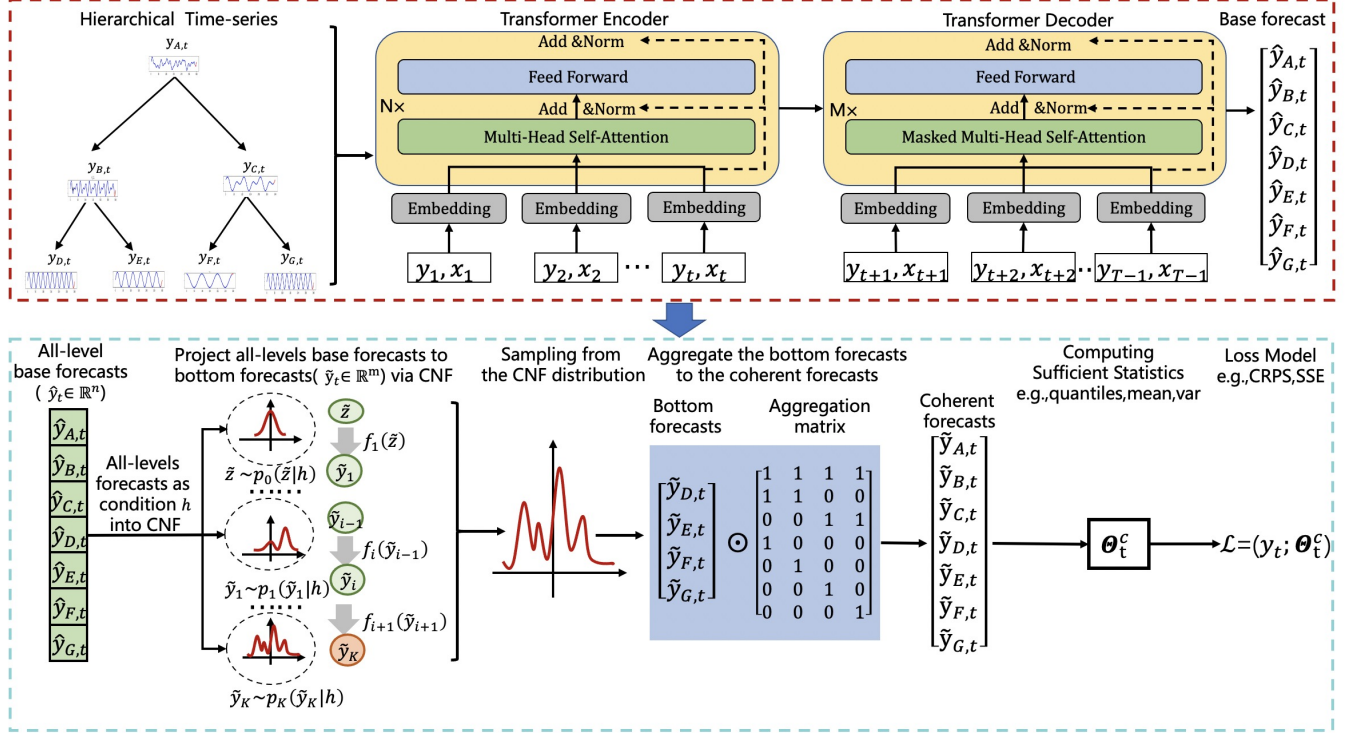


Figure 3: Model Architecture. The red dashed line represents the multivariate autoregressive transformer module (§ 4.1.1) and the blue dashed line highlights the reconciliation method via CNF with bottom-up aggregation matrix (§ 4.1.2). We can compute sufficient statistics from the samples via the empirical distribution to facilitate more complicated optimization objectives.

have autoregressive properties by using *causal masking*. These advantages motivate the use of *multivariate autoregressive transformer* as our building block to better capture patterns (e.g., trends and cycles) inside each individual time series. Note that the global parameters of the transformer are shared across different time series to exploit common patterns over the entire history.

We denote the entities of a hierarchical time series by $y_{t,i} \in \mathbb{R}$ for $i \in \{1, 2, \dots, n\}$, where t is the time index. We consider time series with $t \in [1, T]$, sampled from the complete history of our data, where for training we split the sequence by some context window $[1, t_0]$ and prediction window $[t_0, T]$. We use $x_{t,i}$ to denote time-varying covariate vectors associated with each univariate time series i at time step t .

In the encoder-decoder transformer architecture, the encoder embeds $y_{1:t_0-1}$ and the decoder outputs the base forecast of all levels as condition for the density estimations over $y_{t_0:T}$ via a masked attention module:

$$\begin{aligned} h_{t_0} &= \text{TransformerEncoder}(\text{concat}(y_{1:t_0-1}, x_{1:t_0-1}); \theta) \\ \hat{y}_t &= \text{TransformerDecoder}(\text{concat}(y_{t-1}, x_{t-1}, h_{t_0}; \phi), \end{aligned} \quad (7)$$

where θ and ϕ are parameters of the transformer’s encoder and decoder, respectively. These parameters are shared globally, achieving information fusion across all levels in hierarchy to generate the base forecasts \hat{y}_t .

During training, care has to be taken to prevent using information from future. Specifically, to ensure the autoregressive property

of the model, we employ a mask that reflects the causal direction of the progressing time, i.e. to mask out future time points.

Note that, in our case, base forecasts \hat{y}_t do not directly correspond to un-reconciled forecasts for the base series, but rather represent predictions of an unobserved latent states, which can be used for subsequent density estimations. The transformer allows the model to access any part of the historic time series regardless of temporal distance and thus is potentially able to generate better condition $h_t \in \mathbb{R}^H$ for the subsequent density estimations.

4.1.2 Reconciliation via Conditional Normalizing Flow

Next we describe how to obtain the *coherent probabilistic forecasts*, given the base forecasts described above from the autoregressive transformer.

To estimate the probability density of data (in order to obtain probabilistic forecast), one straight-forward method is to use parameterized Gaussian distribution, but as mentioned above, the real-world hierarchical data are mostly non-Gaussian/non-linear. Equipped with the powerful density approximator, NF, we are able to tackle this challenge, capturing the nonlinear relationships among all levels in hierarchy.

According to the definition of Equation 2, the reconciliation takes a projection matrix $P \in \mathbb{R}^{m \times n}$, projecting from the base forecasts (of dimension n) into the bottom-forecasts (of dimension m). In our approach, this projection is replaced using *conditional normalizing flow* (CNF), i.e., the conditional joint distribution $p_Y(\tilde{y}_t | h_t)$, where

\mathbf{h}_t is the condition (base forecasts $\hat{\mathbf{y}}_t$ in our case) and the current $\tilde{\mathbf{y}}_t$ is the reconciled bottom-forecasts (of dimension m).

In the Real-NVP architecture, we extend Equation 6 by concatenating condition \mathbf{h}_t to both the inputs of the scaling and translation function approximators of the coupling layers as follows:

$$\begin{cases} \mathbf{y}^{1:d} = \mathbf{z}^{1:d} \\ \mathbf{y}^{d+1:D} = \mathbf{z}^{d+1:D} \odot \exp(s(\mathbf{z}^{1:d}, \mathbf{h}) + t(\mathbf{z}^{1:d}, \mathbf{h})) \end{cases}, \quad (8)$$

where \mathbf{z} is a noise vector sampled from an isotropic Gaussian, functions s (scale) and t (translation) are usually deep neural networks, which as mentioned above, do not need to be invertible.

To obtain an expressive distribution representation, we can stack K layers of conditional flow modules (Real-NVP), generating the conditional distribution of the future sequences of all time series in hierarchy, given the past time $t \in [1, t_0]$ and the covariates in $t \in [1, T]$. Specifically, it can be written as a product of factors (as an autoregressive model):

$$p_Y(\tilde{\mathbf{y}}_{t_0:T} | \mathbf{y}_{1:t_0-1}, \mathbf{x}_{1:T}; \theta, \phi, \psi) = \prod_{t=t_0}^T p_Y(\tilde{\mathbf{y}}_t | \mathbf{h}_t; \theta, \phi, \psi), \quad (9)$$

where θ and ϕ are parameters of the transformer and ψ is the parameter of CNF.

Then with the power of reparameterization trick [13], we can generate directly a set of Monte Carlo samples from the the above conditional joint distribution $p_Y(\tilde{\mathbf{y}}_t | \mathbf{h}_t)$ as the reconciled bottom forecasts, e.g., in Figure 3, $[\tilde{\mathbf{y}}_{D,t}, \tilde{\mathbf{y}}_{E,t}, \tilde{\mathbf{y}}_{F,t}, \tilde{\mathbf{y}}_{G,t}]$. According to Equation 2 ($\tilde{\mathbf{y}}_t = \mathbf{S}\mathbf{P}\hat{\mathbf{y}}_t$), the bottom forecasts are multiplied by the aggregation matrix \mathbf{S} to obtain the coherent probabilistic forecasts $\tilde{\mathbf{y}}_t$ of all levels.

Our approach combines prediction and reconciliation into a unified process using deep parametric models, facilitating the sharing of information in process described above via global parameters. Our reconciliation method not only ensures hierarchical coherence constraints, but also dynamically revises the base forecast via integrating information of all levels, to improve the overall performance. Unlike current probabilistic method for hierarchical forecasting, we apply the projection matrix \mathbf{P} on the conditional joint distribution of CNF, rather than samples of distribution. In other words, our model conduct reconciliation in the forecast distribution rather than just sample points, which makes the estimation exact and more efficient.

Moreover, our framework can accommodate different loss functions besides log-likelihood, and by sampling from the forecast distribution, we can also obtain sufficient statistics via the empirical distribution to facilitate more complicated optimization objectives.

4.2 Training

During training, our loss function is directly computed on the coherent forecast samples. Specifically, given \mathcal{D} , defined as a batch of time series $Y := \{y_1, y_2, \dots, y_T\}$, and the associated covariates

$X := \{x_1, x_2, \dots, x_T\}$, we can maximize the likelihood given by Equation 9 via SGD using Adam, i.e.

$$\begin{aligned} \mathcal{L} &= \frac{1}{|\mathcal{D}|T} \prod_{x_1:T, y_1:T \in \mathcal{D}} \prod_{t=1}^T p_Y(y_t | y_{1:t-1}; x_{1:t}, \theta, \phi, \psi) \\ &= \frac{1}{|\mathcal{D}|T} \prod_{x_1:T, y_1:T \in \mathcal{D}} \prod_{t=1}^T p_Y(y_t | \mathbf{h}_t, \theta, \phi, \psi) \end{aligned}, \quad (10)$$

where the globally shared parameters (θ, ϕ) and ψ are from the transformer module and the CNF module, respectively.

Please note that we can easily obtain θ, ϕ , and ψ for other loss functions such as quantile loss, CRPS (continuous ranked probability score) or any other metrics preferred in the forecasting community, as long as we can compute the sufficient statistics from the Monte Carlo samples $\{\tilde{\mathbf{y}}_t\}$ via the empirical distribution function.

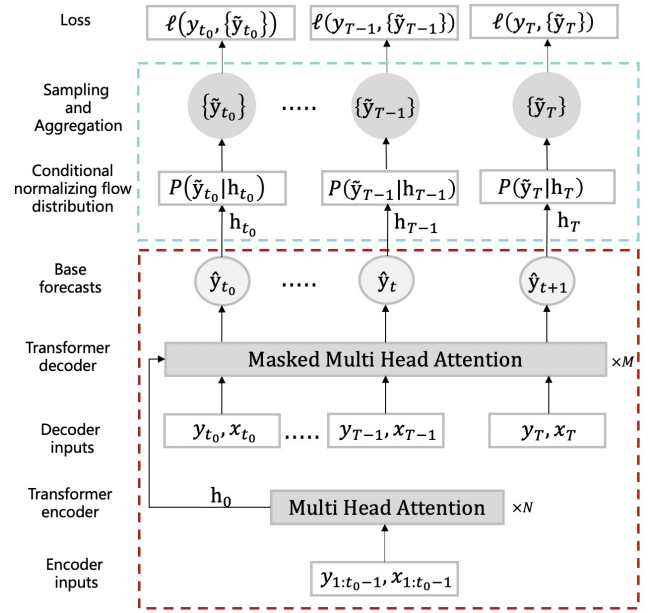


Figure 4: Training Stage. Red dashed line represents multi-variate autoregressive transformer; Reconciliation via CNF with bottom-up aggregation matrix is highlighted

4.3 Inference

During inferencing, we can predict using the autoregressive transformer in a step-by-step fashion over the time horizon. Specifically, we first generate the base forecasts $\hat{\mathbf{y}}_t$ using autoregressive transformer for one time step using the covariate vector $x_{1:t-1}$ and the observed value $y_{1:t-1}$. Then we can incorporate base forecasts $\hat{\mathbf{y}}_t$ from all levels into CNF as additional condition \mathbf{h}_t in the latent space to model conditional joint distribution $p_Y(\tilde{\mathbf{y}}_t | \mathbf{h}_t)$. Lastly, we directly obtain a set of Monte Carlo samples from the distribution $p_Y(\tilde{\mathbf{y}}_t | \mathbf{h}_t)$ as the reconciled bottom forecasts, which are multiplied by the aggregation matrix \mathbf{S} to obtain the coherent probabilistic forecasts $\tilde{\mathbf{y}}_t$. Note that we can repeat the above procedure for the

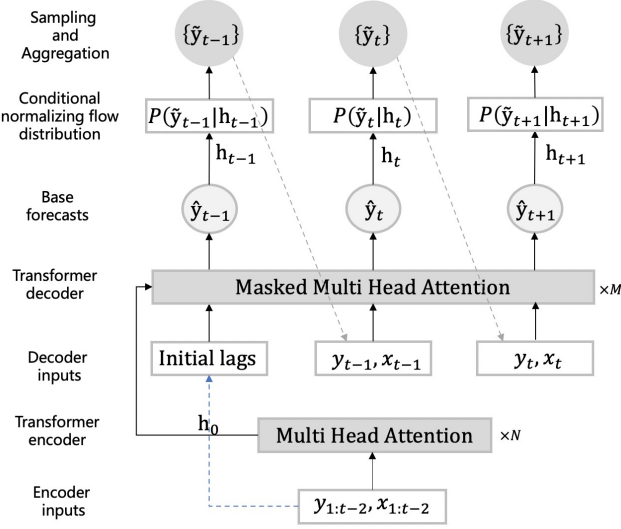


Figure 5: In the multi-step inference, the dashed line indicates that the prediction output from the previous step feeds the inference of the next step via masked attention.

h -period-ahead forecasting to obtain a set of coherent forecasts $\{\tilde{y}_T, \tilde{y}_{T+1}, \dots, \tilde{y}_{T+h}\}$.

5 EXPERIMENTS

In this section, we conduct extensive empirical evaluations on four real-world hierarchical datasets from different industrial domains, including three public datasets and one dataset collected from the application servers of Alipay, which is in the internal review process for releasing publicly.

5.1 Datasets and Analysis

Datasets. The three public datasets include Tourism [3], Tourism-L [28], and Traffic [4]. The new dataset is the data traffic collected from Alipay’s application servers, where our method is deployed for data traffic forecasting. Three real-world hierarchical datasets and a dataset of the web traffic from application servers of Alipay are listed below for our experiments:

- Traffic [4] provides the occupancy rate (between 0 and 1) of 963 car lanes of San Francisco bay area freeways. We aggregate sub-hourly data to obtain daily observations for one year and generate a 207-series hierarchy using the same aggregation strategy as in [2], which is divided into 4 levels. Bottom-level contains 200 series, aggregated-levels contain 7 series in total, and the prediction length is 1.
- Tourism [1, 3] includes an 89-series geographical hierarchy with quarterly observations of Australian tourism flows from 1998 to 2006, which is divided into 4 levels. Bottom-level contains 56 series, aggregated-levels contain 33 series, and the prediction length is 8. This dataset is frequently referenced in hierarchical forecasting studies [9, 26].
- Tourism-L [28] is a larger, more detailed version of Tourism, which contains 555 total series in a grouped structure and 228 observations; This dataset has two hierarchies, i.e., based on geography and based on purpose-of-travel, respectively,

sharing a common root, which is divided into 4 or 5 levels. Bottom-level contains 76 or 304 series, aggregated-level contains 175 series, and the prediction length is 12.

- Server-traffic is the data traffic collected from Alipay’s application servers over the past 90 days, where our method is deployed for data traffic forecasting. As shown in Figure 8, this dataset has the three levels: Bottom-level contains 35 series, aggregated-levels contain 4 series and the prediction length is 8.

Data Analysis. In Figure 6, we plot the relationship between different levels of the Australian tourism dataset using *seaborn*. Country, states, and zones are the upper-level time series, while regions are the bottom-level time series. One *item* (a dot in the Scatterplot) represents the correlation between two time series.

Figure 6 demonstrates not only non-Gaussian distribution of data, but also the nonlinear relationship between data at different levels. Please note that the time series shows nonlinear relationship among the same level and the non-linearity becomes more obvious when the levels of the data are further apart.




Figure 6: Plots of pairwise relationship in the Australian tourism-L dataset using *seaborn*. The diagonal plots are each univariate time series’ distributions (histograms). Plots outside of the diagonal are Scatterplot analysing correlation between the two time series. The color of the item represents the time series of different categories in the same level.

5.2 SOTA Methods and Evaluation Metrics

State-of-the-art Methods. We conduct the performance comparison against state-of-the-art reconciliation algorithms including MinT [28], ERM [2], and Hier-E2E [22], along with other classical baselines, including bottom-up (NaiveBU) approach that generates univariate point forecasts for the bottom-level time series independently followed by the aggregation according to the coherent

constraints to obtain point forecasts for the aggregated series. More details can be found in Appendix B.

Evaluation Metrics. Considering that our method generates probabilistic forecasts instead of point forecasts, it is necessary to evaluate the corresponding probabilistic accuracy with proper metrics, but commonly-used *Mean Absolute Error* (MAE) or *Mean Absolute Percent Error* (MAPE) cannot be directly used for this purpose. On the other hand, *Continuous Probability Ranked Score* (CRPS) generalizes the MAE for probabilistic measurement, making it one of the most widely used accuracy metrics for probabilistic forecasts.

In the stage of evaluation, we will therefore use the *Continuous Ranked Probability Score* (CRPS), which measures the compatibility of a cumulative distribution function $\hat{F}_{t,i}^{-1}$ for time series i against the ground-truth observation $y_{t,i}$, to estimate the accuracy of our forecast distributions. CRPS can be defined as

$$CRPS(\hat{F}_t, y_t) := \sum_i \int_0^1 QS_q(\hat{F}_{t,i}^{-1}(q), y_{t,i}) dq, \quad (11)$$

where QS_q is the quantile score for the q -th quantiles:

$$QS_q = 2(\mathbb{1}\{y_{t,i} \leq \hat{F}_{t,i}^{-1}(q)\} - q)(\hat{F}_{t,i}^{-1}(q) - y_{t,i})$$

We use the discrete version in our experiment, with the integral in Equation 11 replaced by the weighted sum over the quantile set, and we use the quantiles ranging from 0.05 to 0.95 in steps of 0.05.

5.3 Details of SOTA Methods

We conduct the performance comparison against state-of-the-art reconciliation algorithms below:

- **ARIMA-NaiveBU** and **ETS-NaiveBU**: *NaiveBU* generates univariate point forecasts for the bottom-level time series independently and then sums them according to the hierarchical constraint to obtain point forecasts for the aggregate series. Specifically, we use ARIMA and ETS with auto-tuning for base forecast in *NaiveBU* with the R package *hts* [7].
- **ARIMA-MinT-shr**, **ARIMA-MinT-ol**, **ETS-MinT-shr**, **ETS-MinT-ols**: For *MinT*, we use the covariance matrix with shrinkage operator (**MinT-shr**) and the diagonal covariance matrix corresponding to ordinary least squares weights (**MinT-ols**) [28]. In addition, we also use ARIMA and ETS for auto-tuning of base forecasts in *Mint* with the R package *hts* [7].
- **ARIMA-ERM** and **ETS-ERM**: *ERM* is from **Regularized Regression for Hierarchical Forecasting** [2] in related work discussed in the previous section and we also use ARIMA and ETS for auto-tuning of base forecasts in *ERM* with the R package *hts* [7].
- **N-BEATS-NaiveBU**: *N-BEATS* [18] is an interpretable time Series forecasting model, which is a deep neural architecture based on backward and forward residual links and a very deep stack of fully-connected layers. Also We use the above NaiveBU method to ensure coherency. The source code is available at <https://github.com/philipperemy/n-beats>.
- **Transformer-NaiveBU**: *Transformer* [27] is a recently popular method modeling sequences based on self-attention mechanism. We use the native transformer implementation

from pytorch. We use the above NaiveBU method to ensure coherency.

- **Informer-NaiveBU**: *Informer* [30] is a probspare self-attention mechanism based model to enhance the prediction capacity in the Long sequence time-series forecasting problem, which validates the Transformer-like model's potential value to capture individual long-range dependency between long sequence time-series outputs and inputs. Also We use the above NaiveBU method to ensure coherency. The source code is available at <https://github.com/zhouhaoyi/Informer2020>.
- **DeepAR-NaiveBU**: *DeepAR* [24] is a univariate probabilistic forecasting model based rnn, which is widely applied in real-world industrial application. Also We use the above NaiveBU method to ensure coherency. The code is from the Gluonts implementation on <https://github.com/awsmlabs/gluon-ts>.
- **DeepVAR-lowrank-Copula**: *DeepVAR* is a multivariate, nonlinear generalization of classical autoregressive model based rnn [23]. We use the vanilla model with no reconciliation, parameterizing a low-rank plus diagonal covariance via Copula process. We set the rank of low-rank covariance to be 5. The code is from the Gluonts implementation on <https://github.com/awsmlabs/gluon-ts>.
- **Hier-E2E**: In *Hier-E2E* [22], we use the DeepVAR[23] with the diagonal covariance matrix to obtain the base forecast, and use pre-calculated reconciliation matrix M to project into the coherent subspace. The code is from the authors' implementation on <https://github.com/rshyamsundar/gluonts-hierarchical-ICML-2021>.

5.4 Experiment Environment

All experiments run on the Linux server(Ubuntu 16.04) with the Intel(R) Xeon(R) Silver 4214 2.20GHz CPU, 16GB memory, and the single Nvidia V-100 GPU.

5.5 Experiment Results

For evaluation, we generate 200 samples from our probabilistic model to create an empirical predictive distribution. We run our method 5 times and report the mean and standard deviation of CRPS scores. Table 1 shows the performances of different approaches on the four hierarchical datasets, and we can observe that our proposed approach (Hier-Transformer-CNF) achieves all best results, with significant improvements of accuracy on most datasets.

Furthermore the results of experiments proven that effect of traditional classical statistical methods is far inferior to the deep learning model based transformer or rnn. By harnessing the power of deep model, We not only achieve best results but also overcome the shortcomings of statistical models, which rely on any assumption such as unbiased estimates or Gaussian distribution.

5.6 Ablation Study

To further demonstrate the effectiveness of the designs described in Section 4, we conduct ablation studies using the variant of our model and the related models. The results of each aggregation level are shown in Figure 7.

method	Traffic	Tourism	Tourism-L	Server-Traffic
ARIMA-NaiveBU	0.0808	0.1138	0.1741	0.3834
ETS-NaiveBU	0.0665	0.1008	0.1690	0.3899
ARIMA-MinT-shr	0.0770	0.1171	0.1609	0.2713
ARIMA-MinT-ols	0.1116	0.1195	0.1729	0.2588
ETS-MinT-shr	0.0963	0.1013	0.1627	0.3472
ETS-MinT-ols	0.1110	0.1002	0.1668	0.2652
ARIMA-ERM	0.0466	0.5887	0.5635	0.2320
ETS-ERM	0.1027	2.3755	0.5080	0.2501
N-BEATS-NaiveBU	0.0703	0.1062	0.1912	0.1213
Transformer-NaiveBU	0.0621	0.1102	0.1877	0.1121
Informer-NaiveBU	0.0571	0.0911	0.1601	0.1101
DeepAR-NaiveBU	0.0574±0.0026	0.1023±0.0019	0.1816±0.0088	0.1136±0.0073
DeepVAR-lowrank-Copula	0.0583±0.0071	0.0991±0.0083	0.1781±0.0093	0.1125±0.0041
Hier-E2E	0.0376±0.0060	0.0834±0.0052	0.1520±0.0032	0.0530±0.0127
Hier-Transformer-CNF(<i>Ours</i>)	0.0217 ± 0.0055	0.0611 ± 0.0077	0.1135 ± 0.0059	0.0314 ± 0.0067
Multivariate Autoregressive Transformer	0.0377	0.0815	0.1471	0.0417

Table 1: CRPS values(lower is better) averaged over 5 runs. We conduct the ablation study in Multivariate Autoregressive Transformer. State-of-the-art methods except for DeepVAR-lowrank-Copula and Hier-E2E produce consistent results over multiple runs.

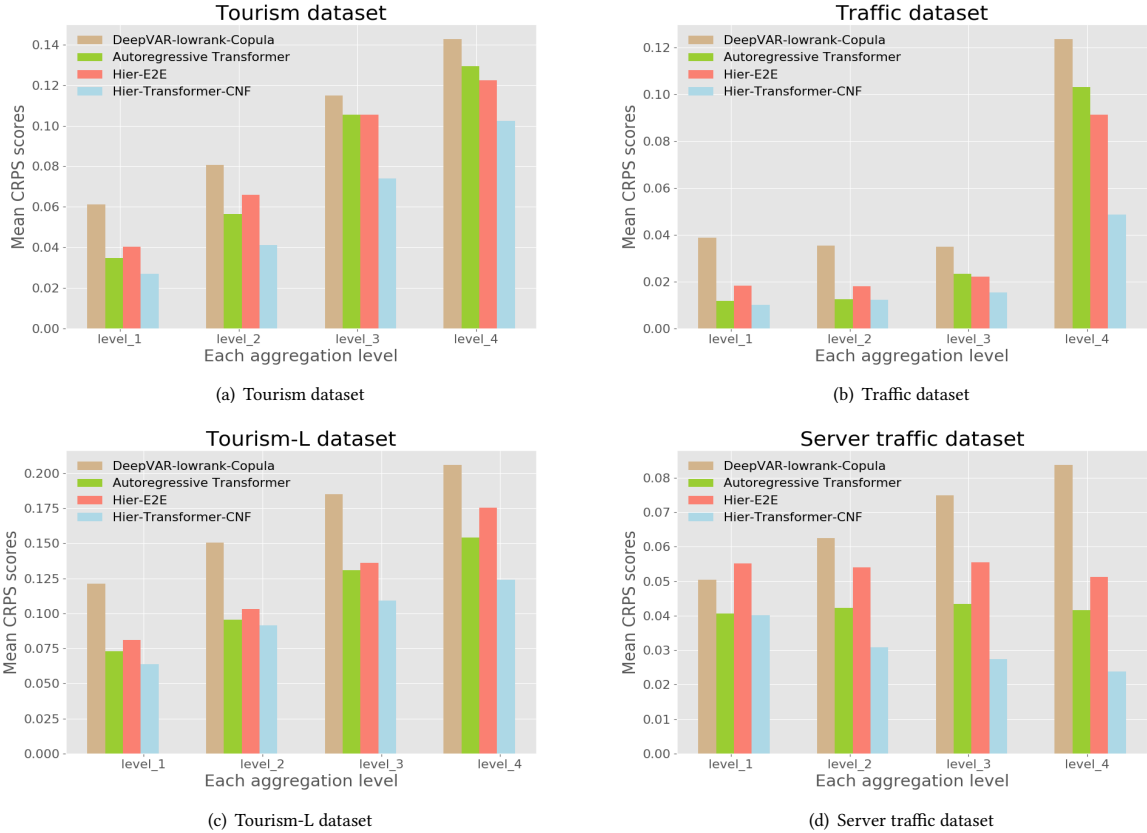


Figure 7: Mean CRPS scores (lower is better) of all aggregation levels in the Tourism, Traffic datasets.

DeepVAR-lowrank-copula is an RNN-based model using the lowrank Gaussian coupla [23] and **Autoregressive Transformer**

is based the encoder-decoder transformer architecture [11], enabled by its multi-head self-attention mechanism to capture both long-

and short-term dependencies in time series data. **Hier-E2E** is the end-to-end model combining DeepVAR(Gaussian distribution) and projection matrix [22]

It is obvious that transformer-based models outperform RNN-based models, especially in the upper-level data. The transformer allows the access to any part of the historic data regardless of temporal distance and is thus capable of generating better conditioning for NF head, evidenced by our experiments.

The non-Gaussian data distribution and nonlinear correlations depicted in Figure 7 attest the need for more expressive density estimators, i.e., CNF, which makes our approach significantly outperforming other methods especially in the bottom-level data. Moreover, the end-to-end reconciliation method also beats direct prediction model in the forecasting at all levels.

In summary, significant improvements at all levels prove the effectiveness of our approach.

6 INDUSTRIAL APPLICATIONS

Our method is successfully deployed on application servers of Alipay, the world’s leading company of payment technology, to forecast server traffic. Alipay has a giant cluster of application servers to support its complex financial Internet business, separated into numerous *Internet Data centers* (IDC), which is a facility where an organization or a service provider rely on as computation infrastructure. In order to improve the efficiency of resource utilization, pre-scheduling application services is essential, for which we need to predict the traffic to facilitate resource management decisions. In Alipay’s IDCs, the deployment of application services adopts a hierarchical structure. Specifically, each deployment consists of three levels: the upper layer is called *app deployment unit*, and the bottom deployment unit is called *zone*, which is divided into Gzone, Rzone, and Czone. Each zone is then divided into multiple groups, as shown in Figure 8 below. In practice, the hierarchical prediction for the

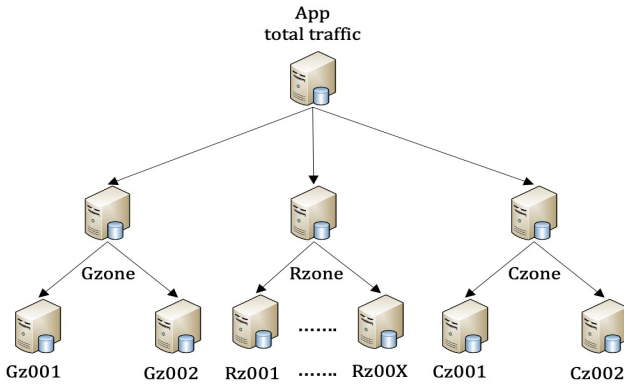


Figure 8: A deployment example of Alipay’s application service

service traffic at each deployment level facilitates the precise management and scheduling [29]. In the experiments above as shown in Table 1, we perform $T + 8h$ predictions on the data collected from real-world servers in Alipay over the past 90 days to verify the effectiveness of our method. The prediction of server traffic can be used not only for efficient resource management/allocation in

IDCs, but also for achieving **energy savings** and **reducing carbon dioxide emissions** through server scheduling management [21] to promote **carbon neutrality**. Our method was deployed before the Double 11 shopping festival and according to the certification of *China Environmental United Certification Center* (CEC), a reduction of **394 tons of carbon dioxide** and a saving of **640,000 kWh of electricity** were achieved during the shopping festival, marking a critical step for Alipay’s goal towards *carbon neutrality* by 2030. Our forecasting algorithm has been running robustly and continuously since then, contributing to the energy efficiency of Alipay’s cloud service.

7 CONCLUSION

In this paper, we proposed a novel end-to-end approach that tackles forecasting and reconciliation simultaneously for hierarchical time series, without requiring any explicit post-processing step, by combining *autoregressive transformer* and *conditioned normalizing flow* to generate the coherent probabilistic forecasts. We conducted extensive empirical evaluations on real-world datasets and demonstrated the competitiveness of our method under various conditions against other state-of-the-art methods. Our ablation study proved the efficacy of every design component we employed. We have also successfully deployed our method in Alipay as continuous and robust forecasting services for server traffic prediction to promote energy efficiency.

REFERENCES

- [1] George Athanasopoulos, Roman A Ahmed, and Rob J Hyndman. 2009. Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting* 25, 1 (2009), 146–166.
- [2] Souhaib Ben Taieb and Bonsoo Koo. 2019. Regularized regression for hierarchical forecasting without unbiasedness conditions. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1337–1347.
- [3] Robyn Bushnell, Gary M Prosser, Herbert William Faulkner, and Jafar Jafari. 2001. Tourism research in Australia. *Journal of travel research* 39, 3 (2001), 323–326.
- [4] Marco Cuturi. 2011. Fast Global Alignment Kernels. In *ICML*. 929–936.
- [5] Byron J Dangerfield and John S Morris. 1992. Top-down or bottom-up: Aggregate versus disaggregate extrapolations. *International journal of forecasting* 8, 2 (1992), 233–241.
- [6] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803* (2016).
- [7] Lee A. Wang E. Hyndman, R. and S. Wickramasuriya. 2020. hts R package version 6.0.1. <https://CRAN.R-project.org/package=hts>.
- [8] Rob J Hyndman, Roman A Ahmed, George Athanasopoulos, and Han Lin Shang. 2011. Optimal combination forecasts for hierarchical time series. *Computational statistics & data analysis* 55, 9 (2011), 2579–2589.
- [9] Rob J Hyndman and George Athanasopoulos. 2018. *Forecasting: principles and practice*. OTexts.
- [10] Jooyoung Jeon, Anastasios Panagiotelis, and Fotios Petropoulos. 2018. Reconciliation of probabilistic forecasts with an application to wind power. *arXiv preprint arXiv:1808.02635* (2018).
- [11] 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*. PMLR, 5156–5165.
- [12] Diederik P Kingma and Prafulla Dhariwal. 2018. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039* (2018).
- [13] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [14] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. 2020. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [15] Iryna Korshunova, Yarin Gal, Joni Dambre, and Arthur Gretton. 2018. Conditional bruno: A deep recurrent process for exchangeable labelled data. In *Bayesian Deep Learning NeurIPS Workshop*. 72371B.

- [16] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyu Zhou, Wenhui 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).
- [17] Zitao Liu, Yan Yan, and Milos Hauskrecht. 2018. A flexible forecasting framework for hierarchical time series with seasonal patterns: A case study of web traffic. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 889–892.
- [18] Boris N Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. 2019. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437* (2019).
- [19] Anastasios Panagiotelis, Puwasala Gamakumara, George Athanasopoulos, Rob J Hyndman, et al. 2020. Probabilistic forecast reconciliation: Properties, evaluation and score optimisation. *Monash econometrics and business statistics working paper series* 26 (2020), 20.
- [20] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. 2019. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762* (2019).
- [21] Haoran Qiu, Subho S Banerjee, Saurabh Jha, Zbigniew T Kalbarczyk, and Ravihankar K Iyer. 2020. {FIRM}: An Intelligent Fine-grained Resource Management Framework for {SLO-Oriented} Microservices. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. 805–825.
- [22] Syama Sundar Rangapuram, Lucien D Werner, Konstantinos Benidis, Pedro Mercado, Jan Gasthaus, and Tim Januschowski. 2021. End-to-End Learning of Coherent Probabilistic Forecasts for Hierarchical Time Series. In *International Conference on Machine Learning*. PMLR, 8832–8843.
- [23] David Salinas, Michael Bohlke-Schneider, Laurent Callot, Roberto Medico, and Jan Gasthaus. 2019. High-dimensional multivariate forecasting with low-rank gaussian copula processes. *arXiv preprint arXiv:1910.03002* (2019).
- [24] 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.
- [25] Souhaib Ben Taieb, James W Taylor, and Rob J Hyndman. 2017. Coherent probabilistic forecasts for hierarchical time series. In *International Conference on Machine Learning*. PMLR, 3348–3357.
- [26] Souhaib Ben Taieb, James W Taylor, and Rob J Hyndman. 2021. Hierarchical probabilistic forecasting of electricity demand with smart meter data. *J. Amer. Statist. Assoc.* 116, 533 (2021), 27–43.
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [28] Shanika L Wickramasuriya, George Athanasopoulos, and Rob J Hyndman. 2019. Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *J. Amer. Statist. Assoc.* 114, 526 (2019), 804–819.
- [29] Siqiao Xue, Chao Qu, Xiaoming Shi, Cong Liao, Shiyi Zhu, Xiaoyu Tan, Lintao Ma, Shiyu Wang, Shijun Wang, Yun Hu, et al. 2022. A Meta Reinforcement Learning Approach for Predictive Autoscaling in the Cloud. *arXiv preprint arXiv:2205.15795* (2022).
- [30] 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 *Proceedings of AAAI*.