AISTATS 2021

The University of Texas at Austin

SIMULTANEOUSLY RECONCILED QUANTILE FORECASTING OF HIERARCHICALLY RELATED TIME SERIES

XING HAN, JOYDEEP GHOSH
The University of Texas at Austin

SAMBARTA DASGUPTA
Intuit AI

Overview

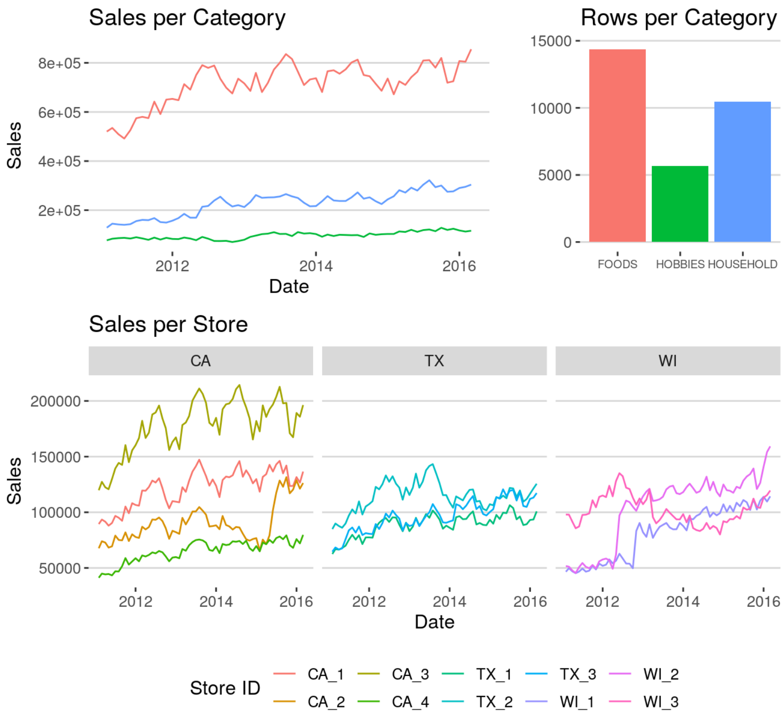

- Hierarchically related time series
  - Grouped by categorical features
    - States/provinces
    - Departments/colleges
    - Race/gender/age
  - Grouped by time period
    - Daily/weekly/monthly/quarterly
- Data properties
  - Skewed noise distribution
  - Irregularly sampled
  - Variable length
  - Sparsity

Challenges

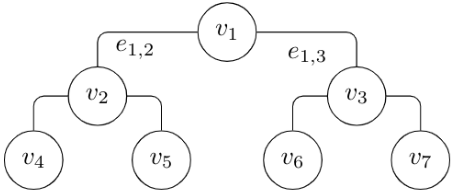

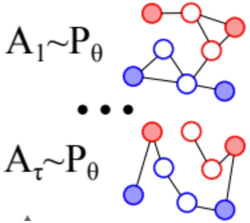

- Accurate and coherent forecasts provide valuable insights
  - Forecast without considering hierarchy, reconcile afterwards **Computationally expensive!**
  - Learn the data relationship in model training **Learned bad graph representations!**

- Further technical challenges for a high-quality forecasting pipeline
  - Precise and consistent uncertainty characterization
  - Interpretable results across multiple levels

Problem Formulation

- Move the reconciliation into model training
  - Use information from adjacent levels
  - Enforce learning pre-specified hierarchical structures

- Predicting quantile confidence intervals
  - Robust to noisy data with outliers
  - User specified uncertainty characterization

- Coherent uncertainty bounds across hierarchy
  - Stable multi-level forecast without quantile crossing

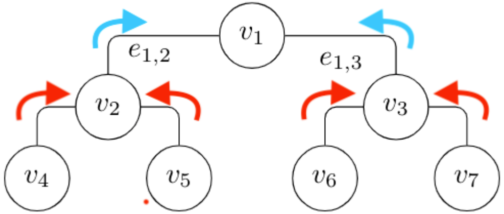

Regularize higher-level forecasts using lower-level predictions

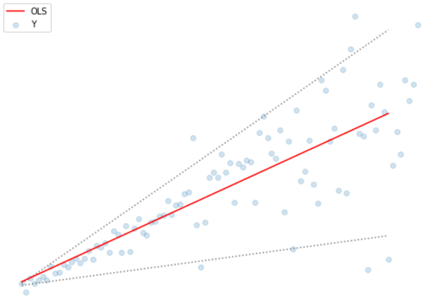

Quantile regression with 5% and 95% quantiles

SHARQ

- Simultaneously reconcile multiple time series and their uncertainty bounds in a **bottom-up** procedure
- Learn pre-specified graph structure in model training and generate **multiple quantiles** as forecasting intervals
- Optimal solution of balancing accuracy and coherency
- **Variance reduction** at higher aggregation levels in model training

Algorithm 1 SHARQ

Input: Training data $\mathcal{I}_1 = \{X_i, Y_i\}_{i=1}^{T_1}$, testing data $\mathcal{I}_2 = \{X_i, Y_i\}_{i=1}^{T_2}$.
Process:
train each leaf node independently without consistency regularization
for each vertex v in upper level l **do**
  train vertex v at level l
end for
Reconciled Median Forecast **MF** $\leftarrow$ **Models**($\mathcal{I}_2$)
for each vertex v in upper level l **do**
  train vertex v at level l and **MF**
end for
Output: Reconciled forecasts at user-specified quantiles.

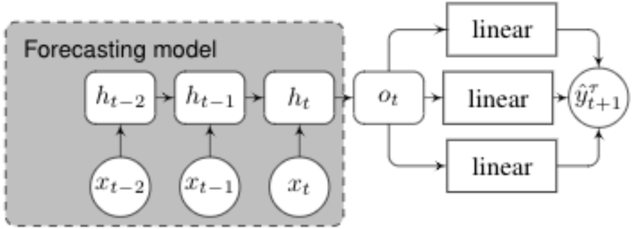

Results & Future Works

- Quantile forecast is robust to multiple types of noise (log normal, gamma)
- SHARQ is **model agnostic**
- Performed better than other baseline methods, particularly at **higher aggregation levels**, and **less parsimonious** forecasting models
- **Reduced inference time** from post-processing reconciliation methods

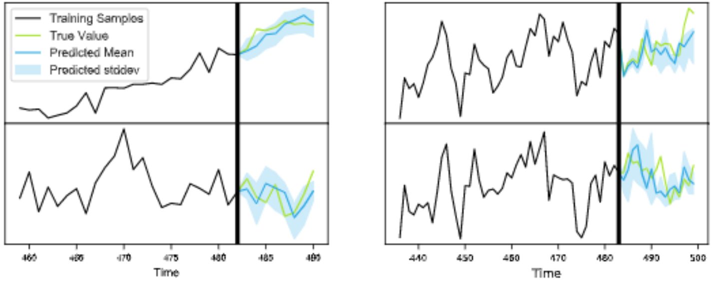

(a) Labour (SHARQ) (b) FTSE (SHARQ)

Time (s)	FTSE		Labour		M5		Wikipedia	
	training	inference	training	inference	training	inference	training	inference
Base	115.96	0.01	68.35	0.00	181.58	0.00	205.47	0.01
BU	65.83	0.03	57.06	0.00	105.45	0.00	142.53	0.01
MinT-sam	106.55	1,784.77	72.24	430.42	172.11	1,461.81	208.26	1,106.70
MinT-shr	104.35	1,148.49	60.83	317.02	175.83	1,039.53	198.16	788.31
MinT-ols	103.23	1,129.45	64.14	310.13	163.24	977.88	196.88	702.02
ERM	547.66	0.05	497.88	0.01	551.60	0.01	1,299.30	0.04
SHARQ	121.84	0.01	99.96	0.00	201.40	0.00	241.97	0.01

Full Paper & Poster Session

Session 1: April 13 at 14:00-16:00 PDT

Contact: aaronhan223@utexas.edu