VIS 2023

TimeTuner: Diagnosing Time Representations for Time-Series Forecasting with Counterfactual Explanations


Jianing Hao


Qing Shi


Yilin Ye


Wei Zeng

香港科技大学(广州)
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY (GUANGZHOU)

香港科技大學
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Agenda

- Introduction
- Related Work
- Visualization
- Case study
- Conclusion

Introduction

- DL models are often criticized for being **overly complex** and **lacking interpretability**.
- The performance of DL models is highly sensitive to the **representation** of time-series data.
- **Representation learning** becomes increasingly popular.

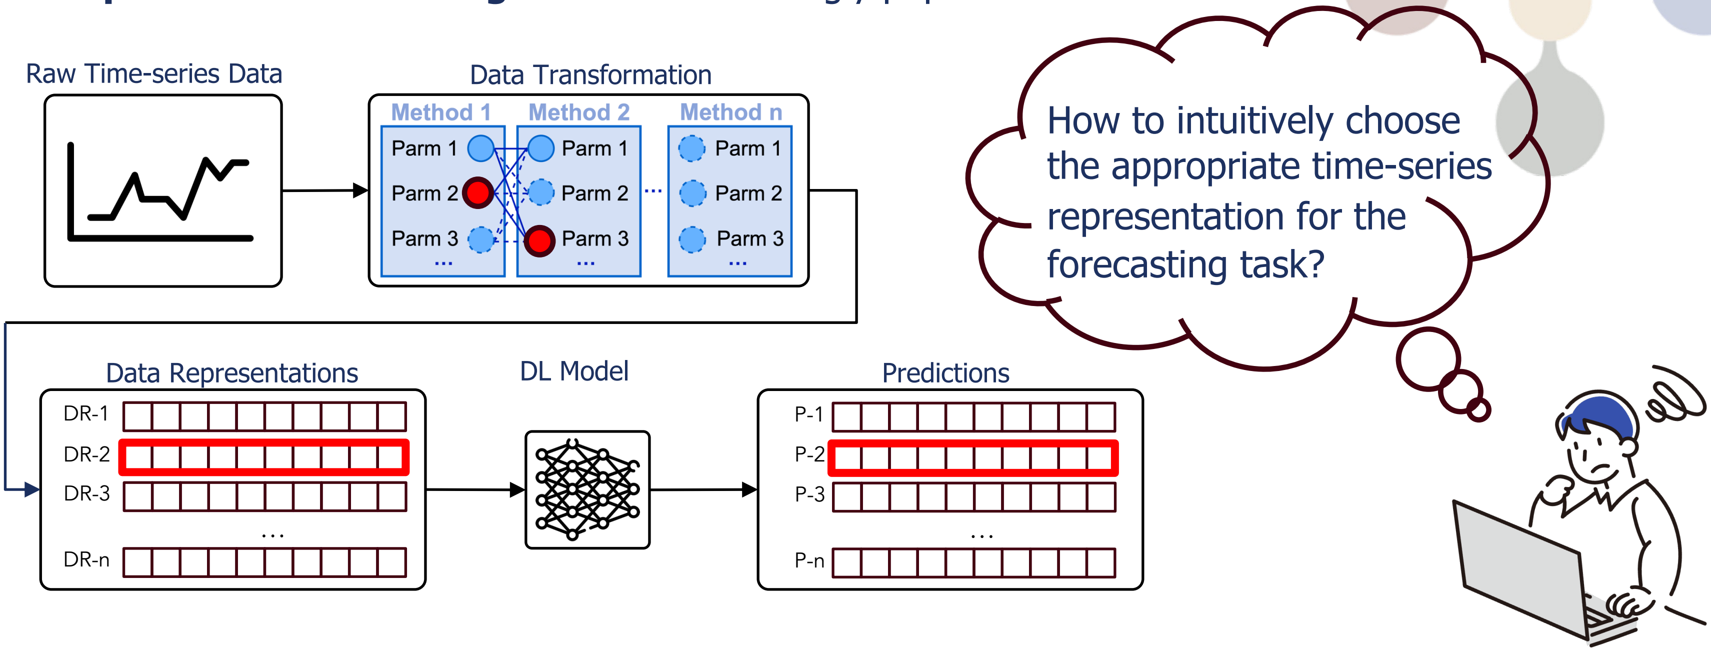

Challenge

Model

- Limited to single numeric metric, such as RMSE.
- Inadequate for uncovering relationships among variables, representations and predictions.
- The lack of interpretability of DL models.

Visualization

- Multivariate variables, diverse transformation methods and large-volume time-series data.
- Potential overlap issues arising from the sliding window mechanism.
- Presenting both explanation and prediction metrics.

Analytical Tasks

T1. Variable-level

- **T1.1** Examine the details of variables
- **T1.2** Identify associations among variables

T2. Representation-level

- **T2.1** Compare different representations
- **T2.2** Analyze the association between representations & variables
- **T2.3** Analyze the association between representations & predictions

T3. Prediction-level

- **T3.1** Overview the overall patterns of predictions
- **T3.2** Detect the prediction outlier

Overview

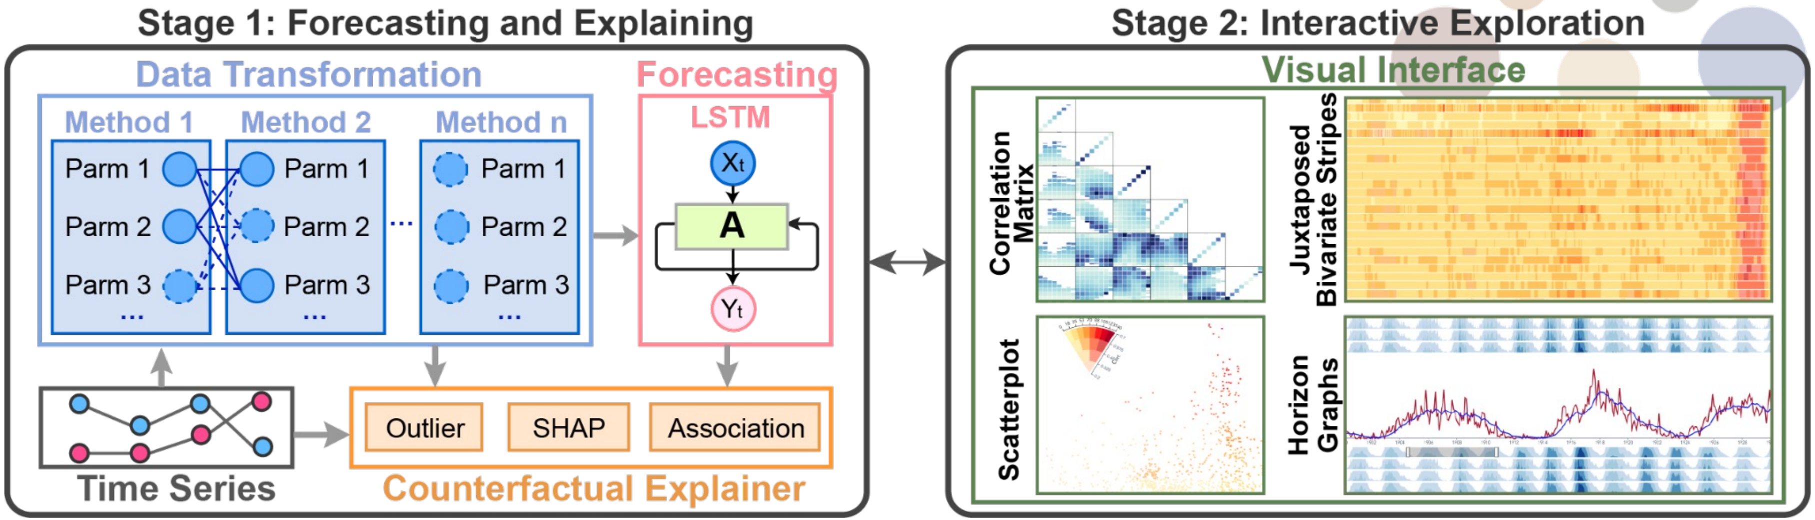

Stage 1: Forecasting and Explaining
Data Transformation
Method 1
Method 2
Method n
Parm 1
Parm 2
Parm 3
Forecasting
LSTM
Xt
A
Yt
Time Series
Outlier
SHAP
Association
Counterfactual Explainer
Stage 2: Interactive Exploration
Visual Interface
Correlation Matrix
Juxtaposed Bivariate Stripes
Scatterplot
Horizon Graphs

Overview

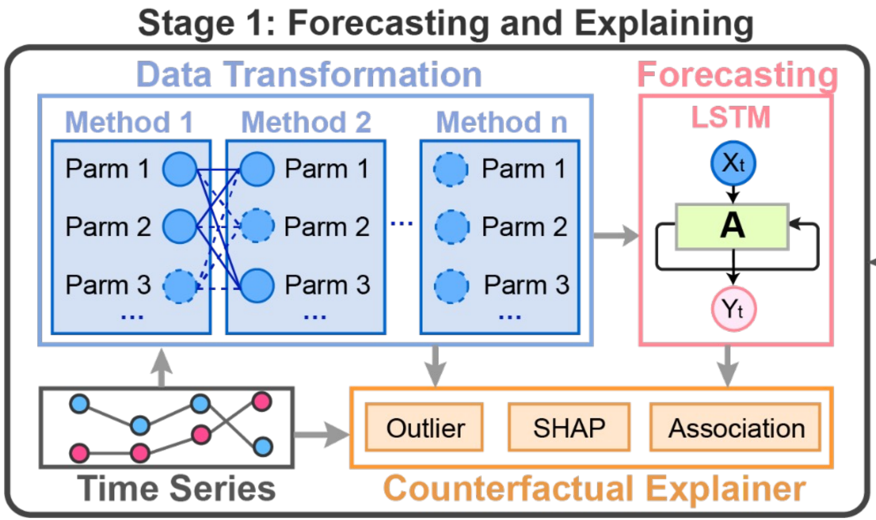

Transformation Methods

- Smoothing

$$MA_m = \frac{1}{m} \sum_{i=n-m+1}^{n} X_i$$

$$WMA_m = \frac{2}{m \times (m+1)} \sum_{i=n-m,+1}^{n} X_i \times (n+1-i)$$

- Sampling

  Sliding windows

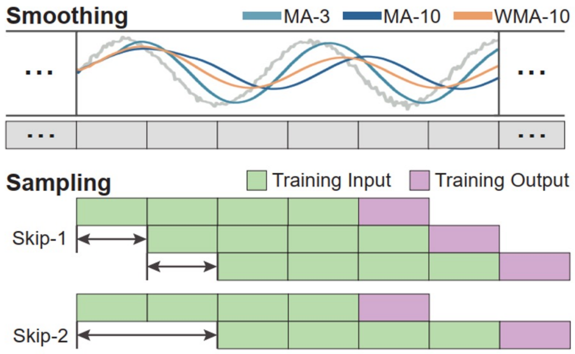

Overview

Coordinated multiple views

- Variable Inspector View
- Representation View
- Prediction Comparator View
- Temporal View

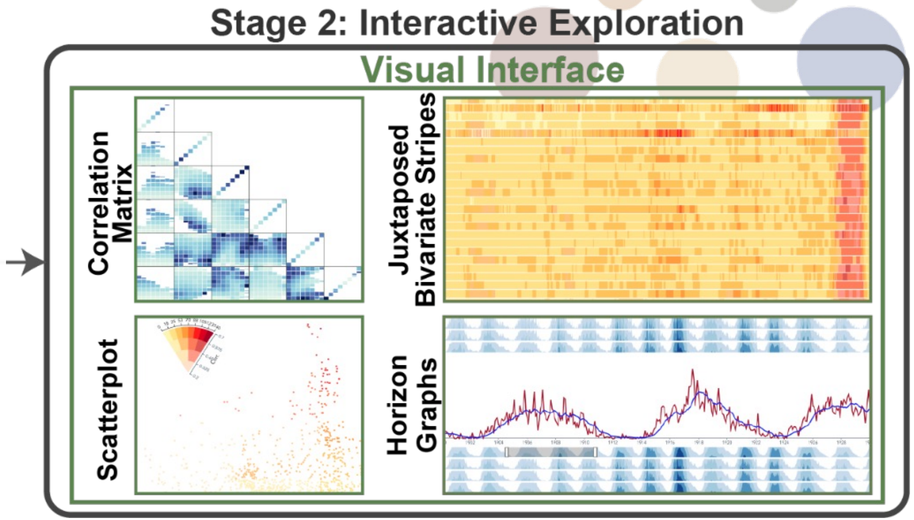

Related Work – *Counterfactual Explanation*

Counterfactual explanations reveal what changes in input alter the outcome of a DL model

- What-If tool [Wexler et al., 2020]
- Partition methods and scales for deep traffic prediction [Zeng et al., 2020]
- DECE [Cheng et al., 2020]

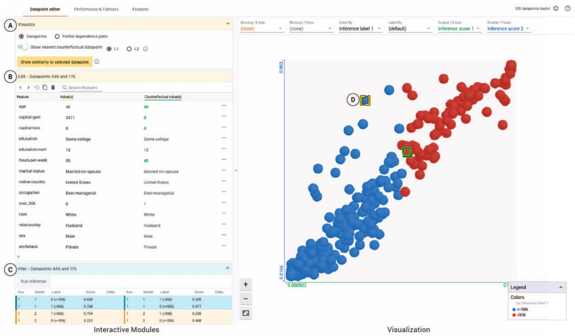

[Wexler et al., 2020]

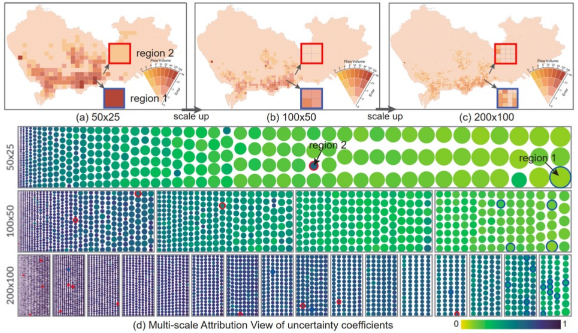

[Zeng et al., 2020]

[Cheng et al., 2020]

Related Work – *Time-series Visualization*

Common visualization methods with time-series data:

- Vertical position or area to encode values: horizon graph [Saito et al,. 2005], stacked graph [Byron and Wattenberg, 2008].
- Pixel-wise elements: heatmaps [Lammarsch et al,. 2009], sparkboxes [Oppermann et al,. 2021], heat stripes [Burch et al,. 2009].

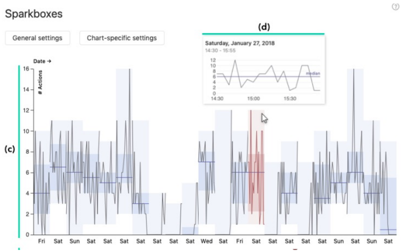

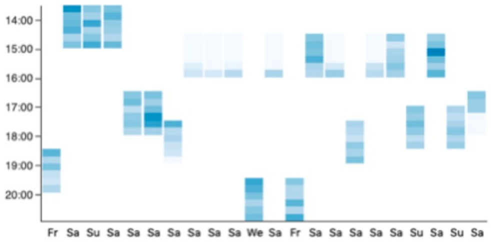

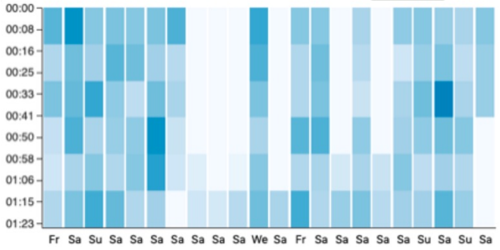

Existing visualizations, however, has limitations in:

- The presenting of two-dimension metrics.
- The extensive range of transformation methods.
- Overlap issues arising from sliding window mechanism.

Visualization

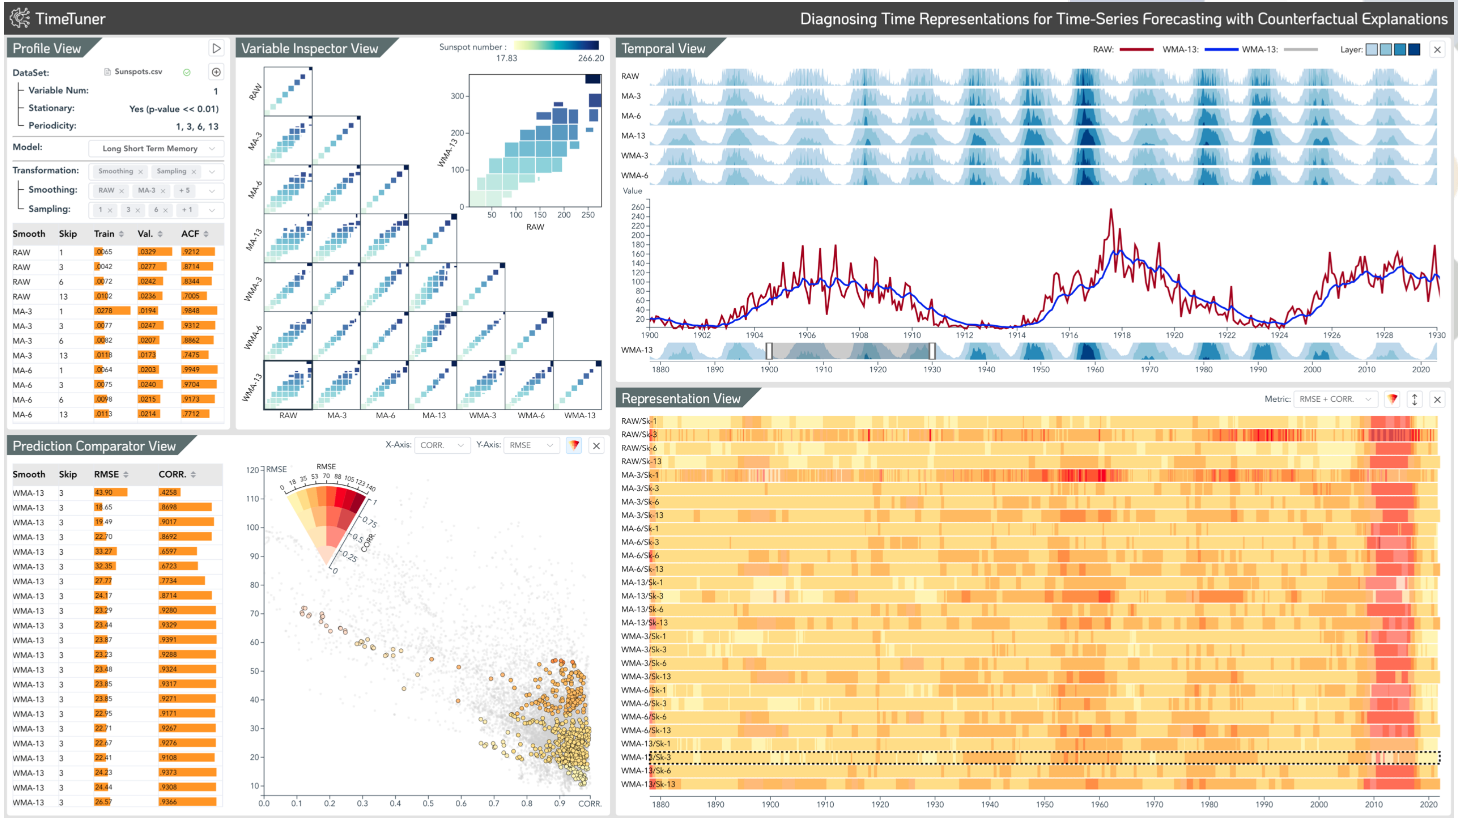

Visualization

Visualization

[Variable, Representation]

Visualization

[*Variable, Representation*]

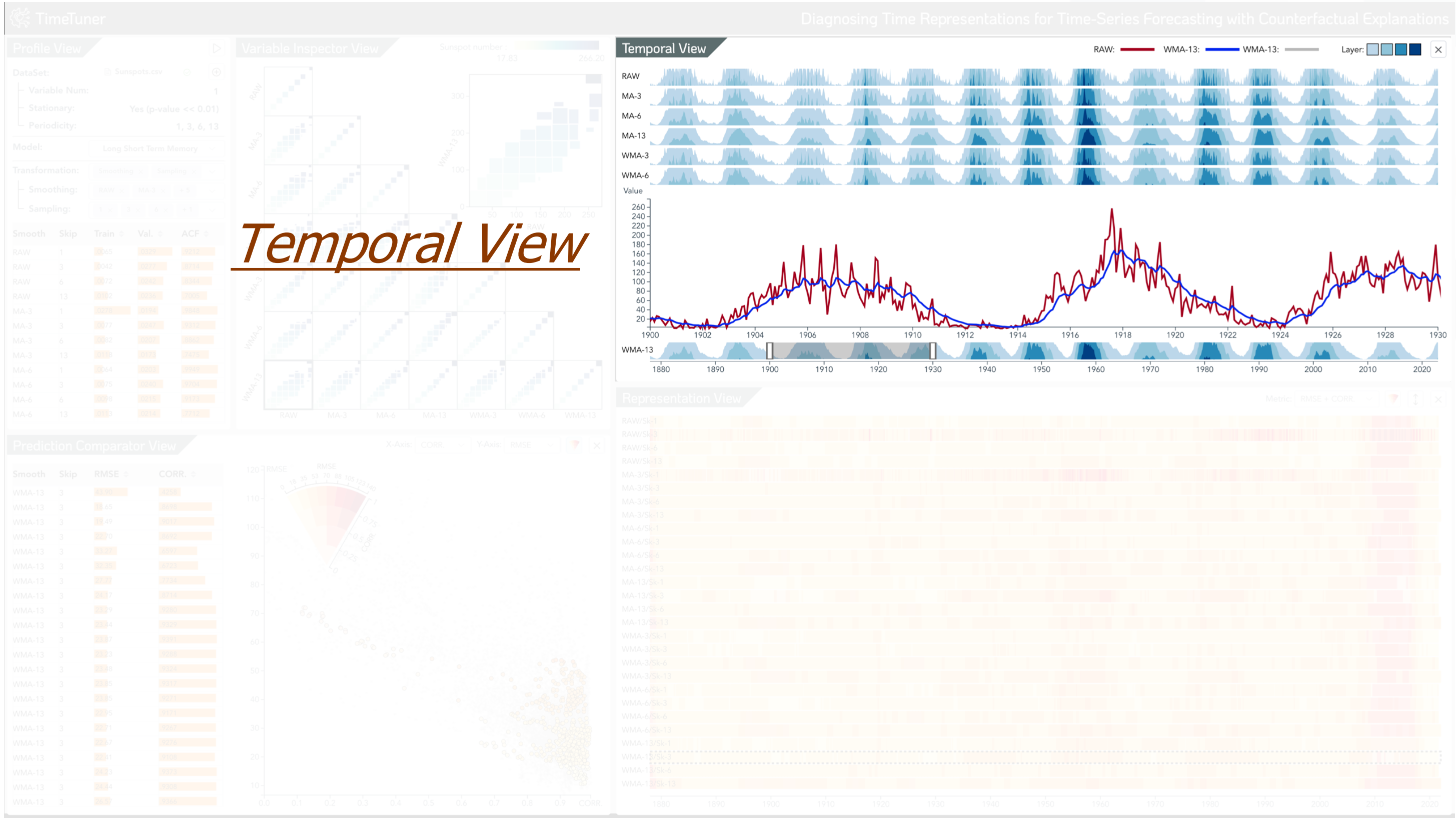

Visualization

[*Representation*]

Representation View

Visualization

[Representation, Prediction]

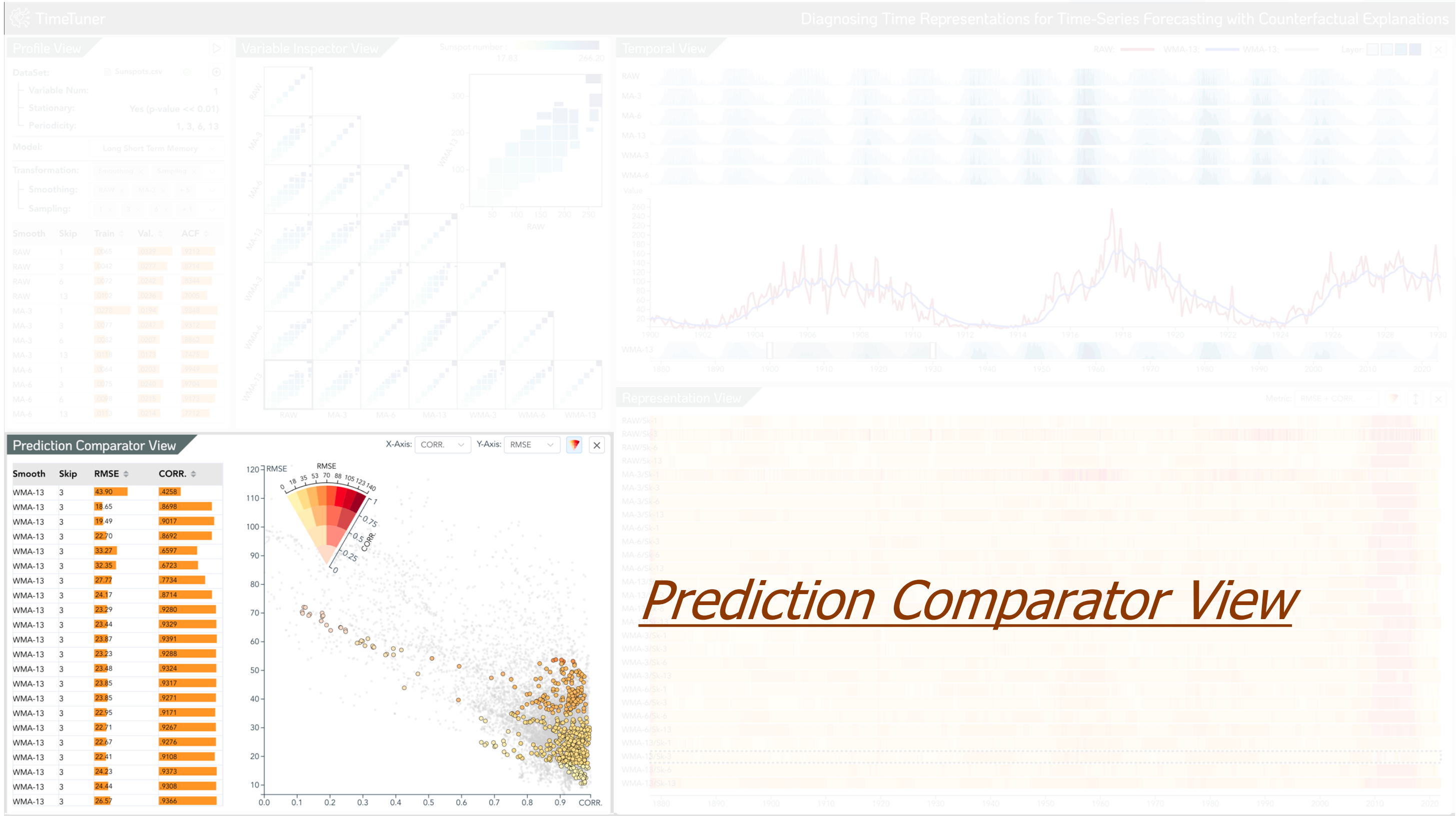

Prediction Comparator View

Variable inspector view

- To support the exploration of relationships between a feature space of continuous or discrete input variables (**T1.2**) and the target variable (**T2,2**)
- A partition-based correlation matrix
- Mosaic plots to overview local relationships of
  - Representations (univariate) - *Left*
  - Pair-wise variables (multivariate) - *Right*

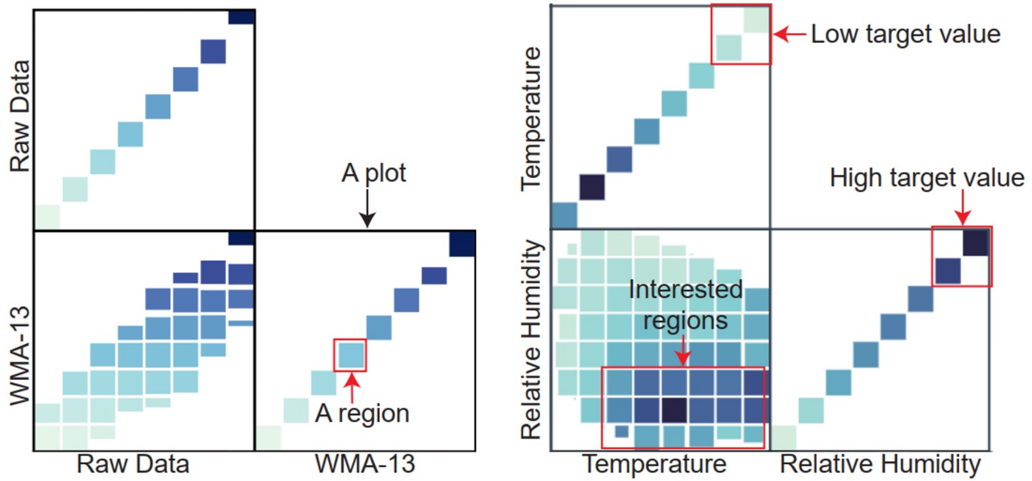

Temporal View

- To visualize the details of different time series representations for univariate data (**T2.1**) and variables for multivariate data (**T1.1**).

- A line chart for the details

- Horizon graphs for the context information

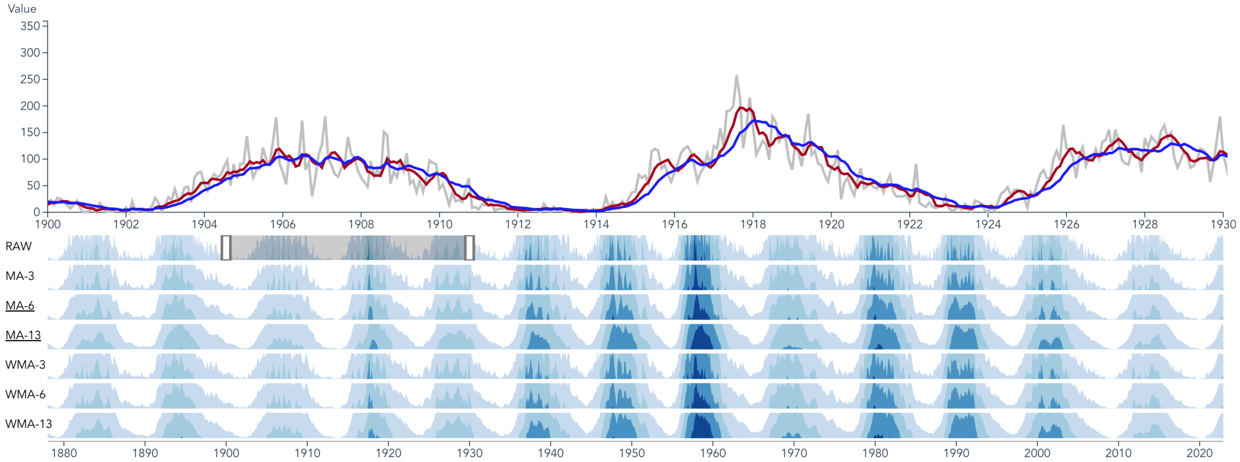

Representation view

- To compare different representations (**T2.1**), and analyze the association between representations & predictions (**T2.3**).

- Horizon juxtaposed bivariate stripes

VSUP

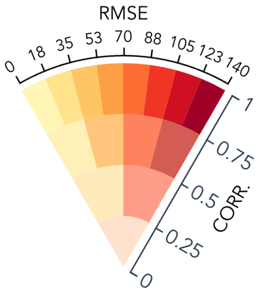

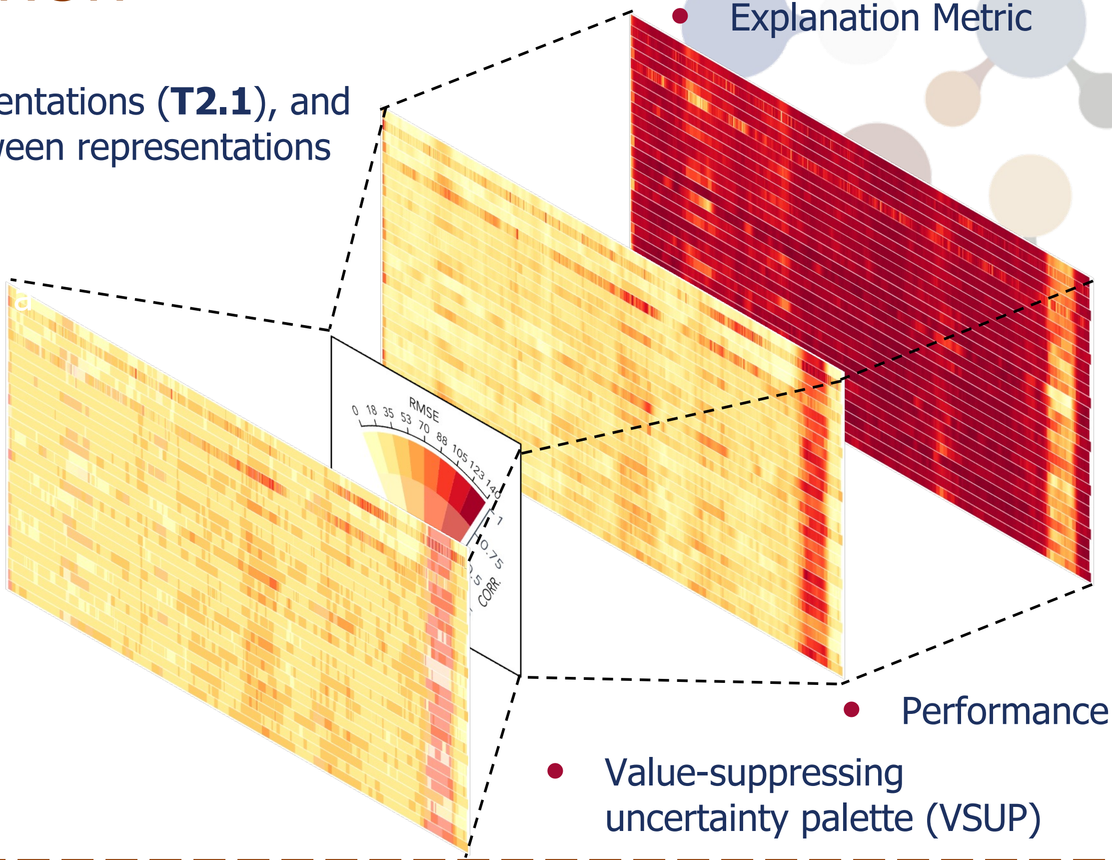

- Explanation Metric
- Performance Metric
- Value-suppressing uncertainty palette (VSUP)

Representation view *- Alternative Design*

juxtaposed bivariate stripes.

These designs are **less scalable than** juxtaposed bivariate stripes.

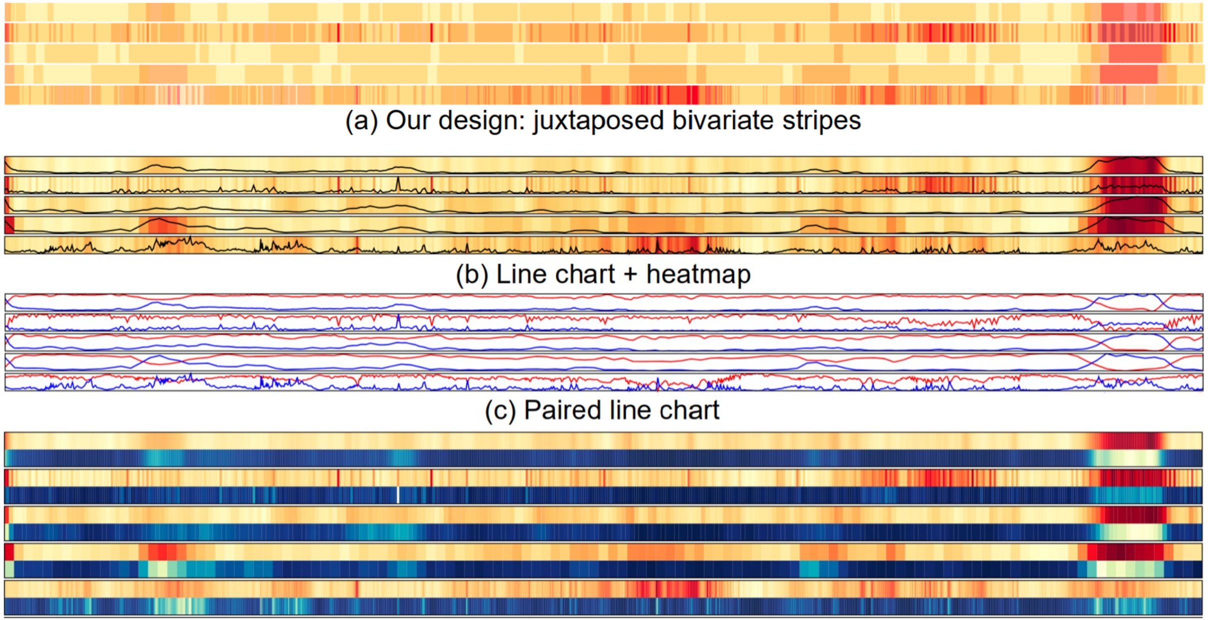

(a) Our design: juxtaposed bivariate stripes

(b) Line chart + heatmap

(c) Paired line chart

(d) Paired heatmaps

Prediction comparator view

- To reveal the relationship between explanation metrics and performance metrics (**T2.3**) and subsequently analyze prediction outliers (**T3.1**)

- X-Axis: Explanation Metric

$$CORR. = \frac{n\times\sum XY - \sum X \times \sum Y}{\sqrt{(n\times\sum X^2 - (\sum X)^2)\times(n\times\sum Y^2 - (}}$$

$$SHAP\ value: y_i = y_{base} + f(X_i, 1) + f(X_i, 2) + \cdots$$

- Y-Axis: Performance Metric

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{n}}$$

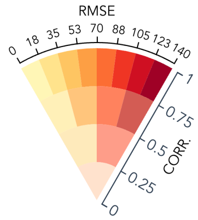

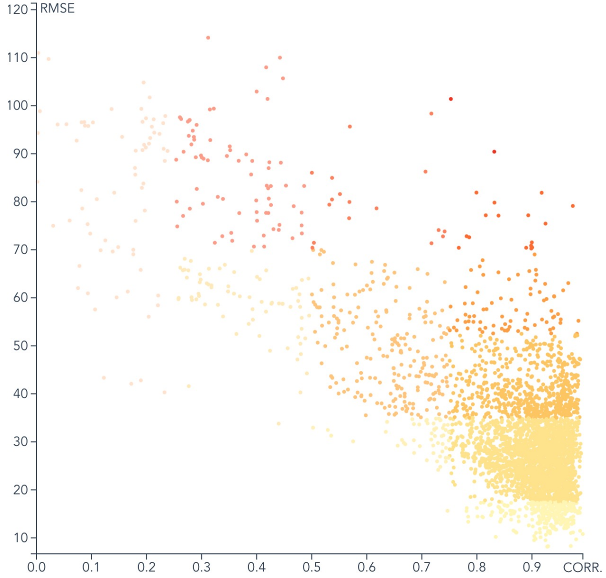

Case 1: univariate sunspots forecasting

Dataset: Sunspots

- Time: Jan. 1, 1818 – Nov. 30, 2022
- Periodicity (month): 1, 3, 6, 13

Relationship & Insight

- The highest prediction accuracy:
  WMA-13/Sk-3
- The largest training error:
  MA-3/Sk-1
- Outlier Period:
  2009 is the period of minimum solar activity and the model does not capture the weakening of sunspot activity.

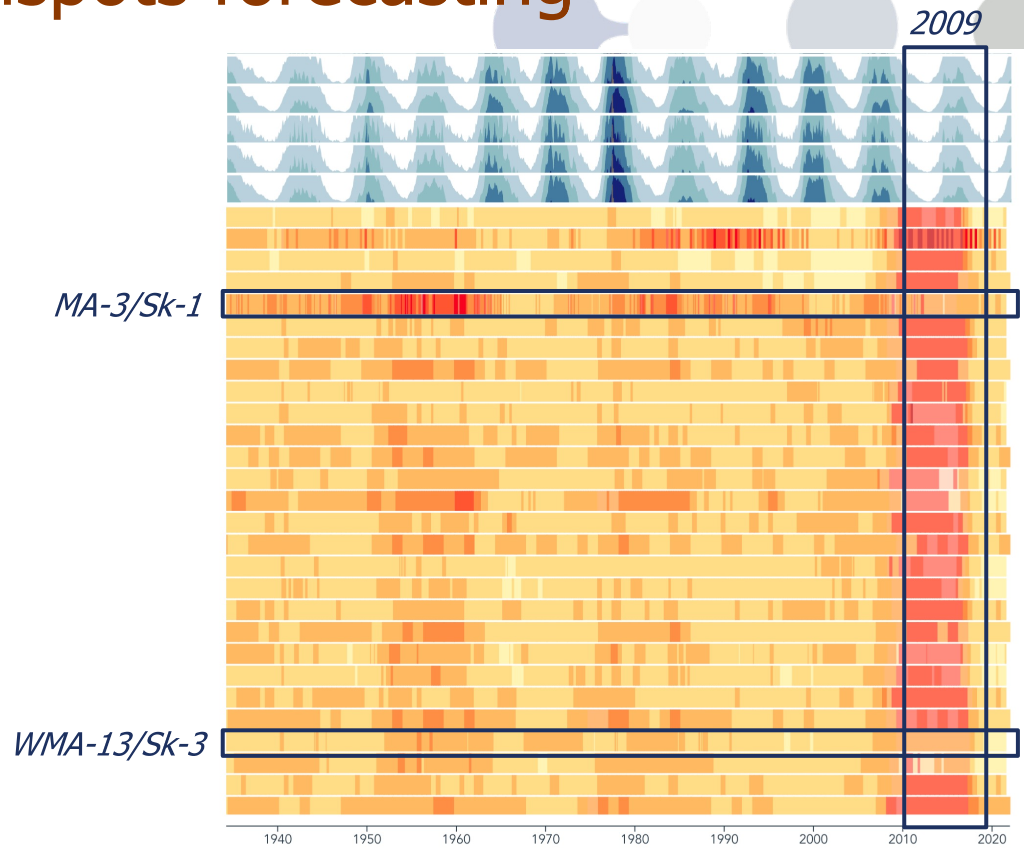

Case 2: multivariate PM2.5 forecasting

Dataset: Air pollution
- Time: Jan. 1 00:00, 2013 – Mar. 31 23:00, 2014
- Periodicity (hour): 1, 6, 12, 24
- Variable: Pm2.5, Temp, Ph, Psfc, Wnd_dir, Wnd_spd

Relationship & Insight

- A negative correlation

  Pm2.5 & Temperature

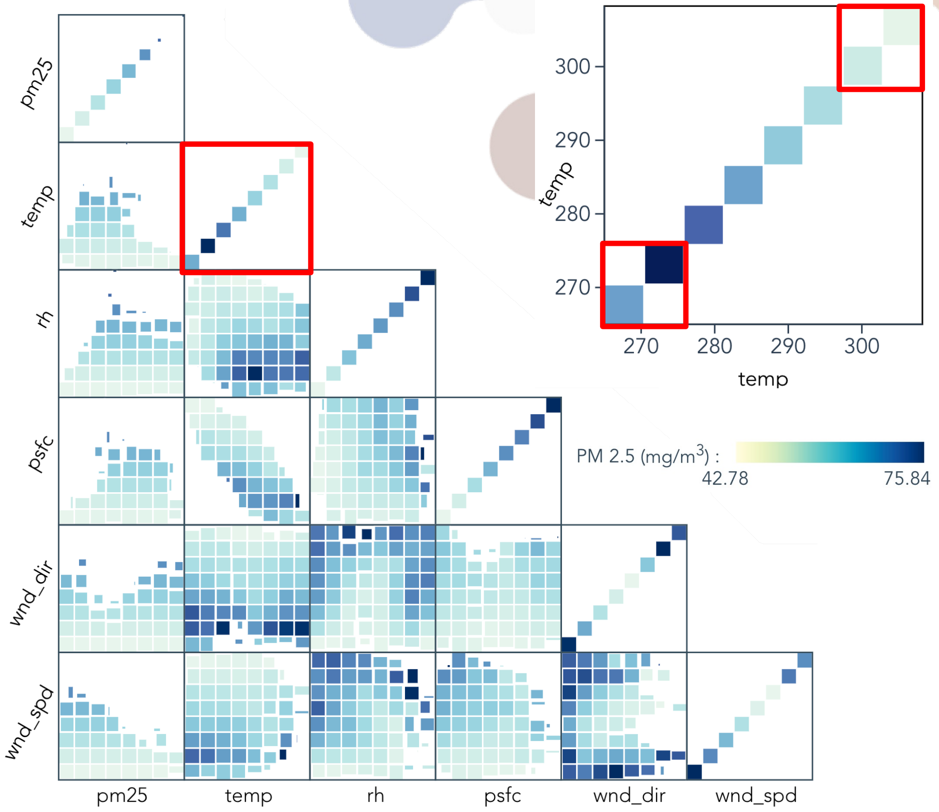

Case 2: multivariate PM2.5 forecasting

Dataset: Air pollution

- Time: Jan. 1 00:00, 2013 – Mar. 31 23:00, 2014
- Periodicity (hour): 1, 6, 12, 24
- Variable: Pm2.5, Temp, Ph, Psfc, Wnd_dir, Wnd_spd

Relationship & Insight

- A negative correlation

  Pm2.5 & Temperature

- A positive correlation

  Pm2.5 & Relative Humidity

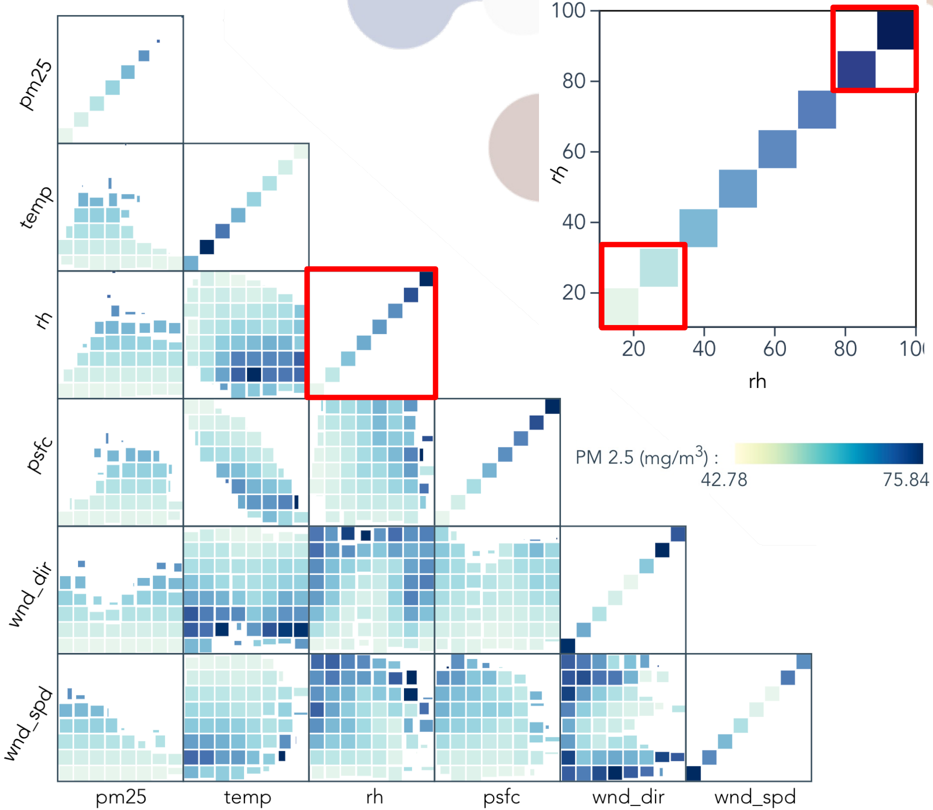

Case 2: multivariate PM2.5 forecasting

Dataset: Air pollution

- Time: Jan. 1 00:00, 2013 – Mar. 31 23:00, 2014
- Periodicity (hour): 1, 6, 12, 24
- Variable: Pm2.5, Temp, Ph, Psfc, Wnd_dir, Wnd_spd

Relationship & Insight

- A negative correlation

  Pm2.5 & Temperature

- A positive correlation

  Pm2.5 & Relative Humidity

- Special Case

  Exhibiting elevated *Pm2.5* levels with *high-speed south winds* may stem from significant pollutant sources in the southern areas

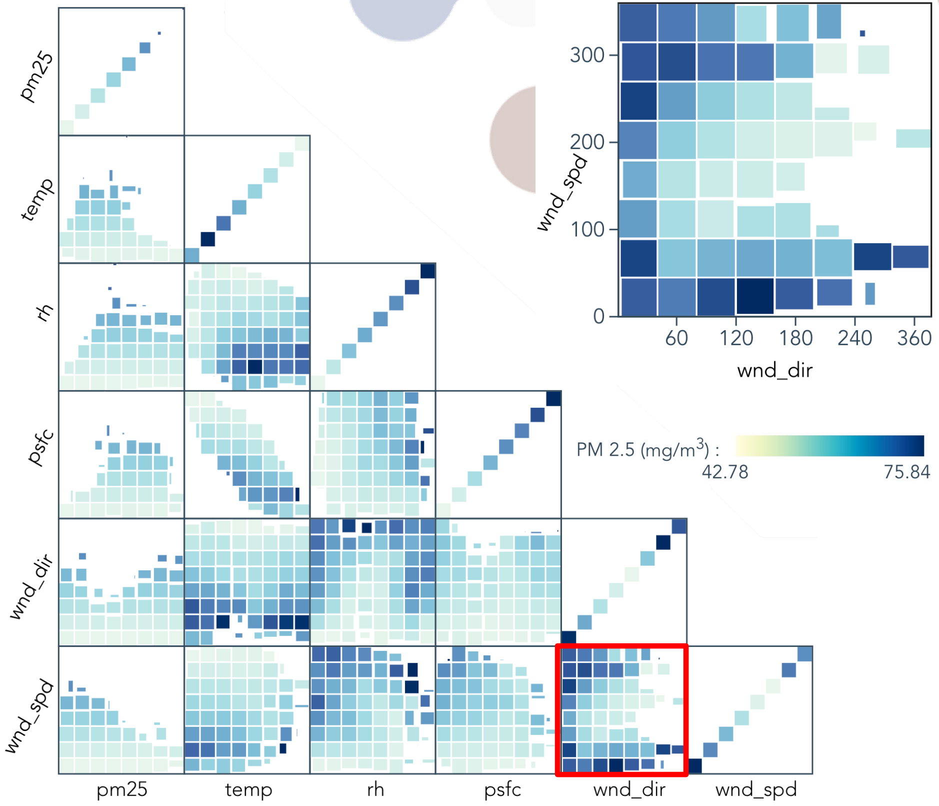

Case 2: multivariate PM2.5 forecasting

Jan. 30 2014
–
Feb. 2 2014

Dataset: Air pollution

- Time: Jan. 1 00:00, 2013 – Mar. 31 23:00, 2014
- Periodicity (hour): 1, 6, 12, 24
- Variable: Pm2.5, Temp, Ph, Psfc, Wnd_dir, Wnd_spd

Relationship & Insight

- Outliers Case

  Jan. 30 2014 – Feb. 2 2014

- Reason

  Fireworks of Chinese New Year

- Explanation

  Outliers cased by *specific events* are difficult to be captured and learned by models.

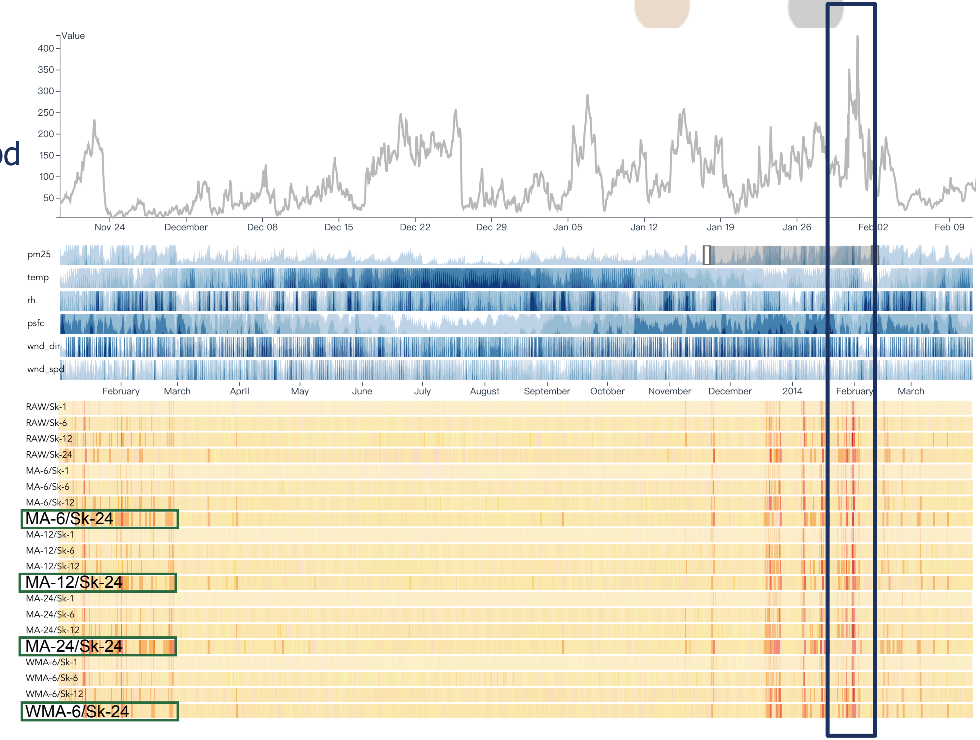

Conclusion

Feedback

- The combination of counterfactual explanations with visual analytics.
- The exploration of both univariate and multivariate data.

Analysis

- Providing useful hints for selecting appropriate transformation methods.
- Explaining why models perform poorly.

Domain

- Highlighting the significant potential of counterfactual explanations combined with visual analytics in advancing the field of data representation learning.

VIS 2023

Thank you!

Project Homepage

Jianing Hao

Email: jhao768@connect.hkust-gz.edu.cn

Please feel free to contact me
if you have any questions.

Jianing Hao Homepage

HKUST CIVAL