Residual Correction in Real-Time Traffic Forecasting

Daejin Kim*, Youngin Cho*, Dongmin Kim,
Cheonbok Park, and Jaegul Choo
(*equal contributions)

Deep Learning for Traffic Forecasting

- Recently, several deep-learning-based models have been introduced for traffic forecasting (*e.g.*, DCRNN, STGCN, Graph-WaveNet, STAWnet…).

- Models take the graph information and the previous speed of each node as the input and predict the future speed of each node.

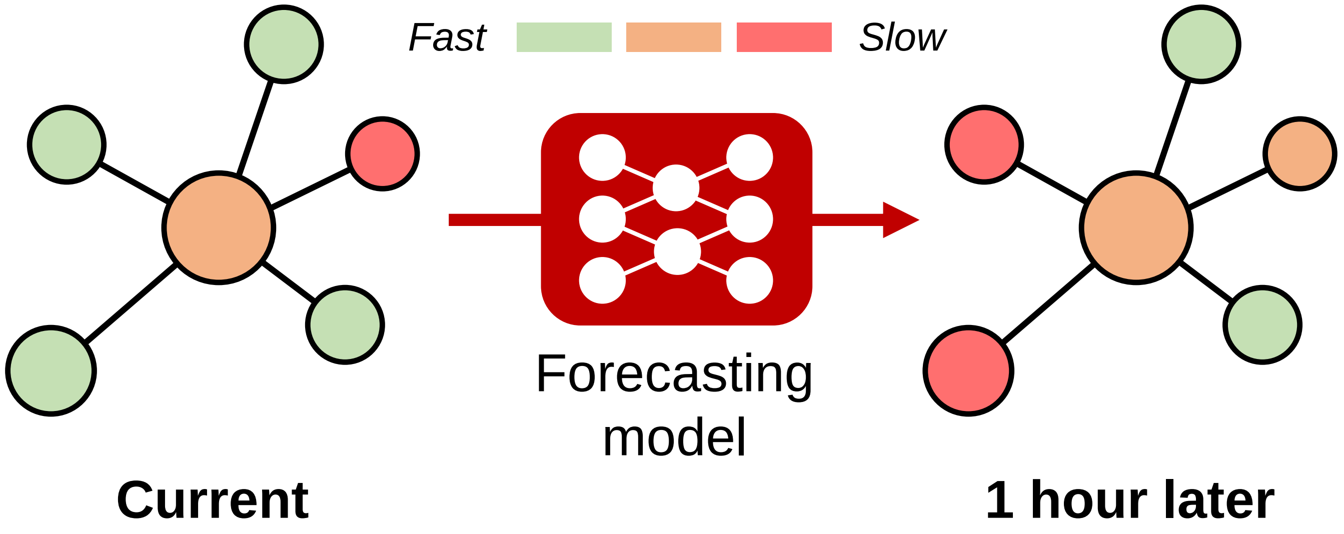

(a) METR-LA (b) PEMS-BAY

Motivation

- Existing deep-learning-based traffic forecasting models do not leverage the ***failures*** of their previous predictions.

"Can't we just tell the model that its *previous prediction failed*?"

Motivation

- It motivates us to design a model that takes the previous residual as the input and corrects the original model.

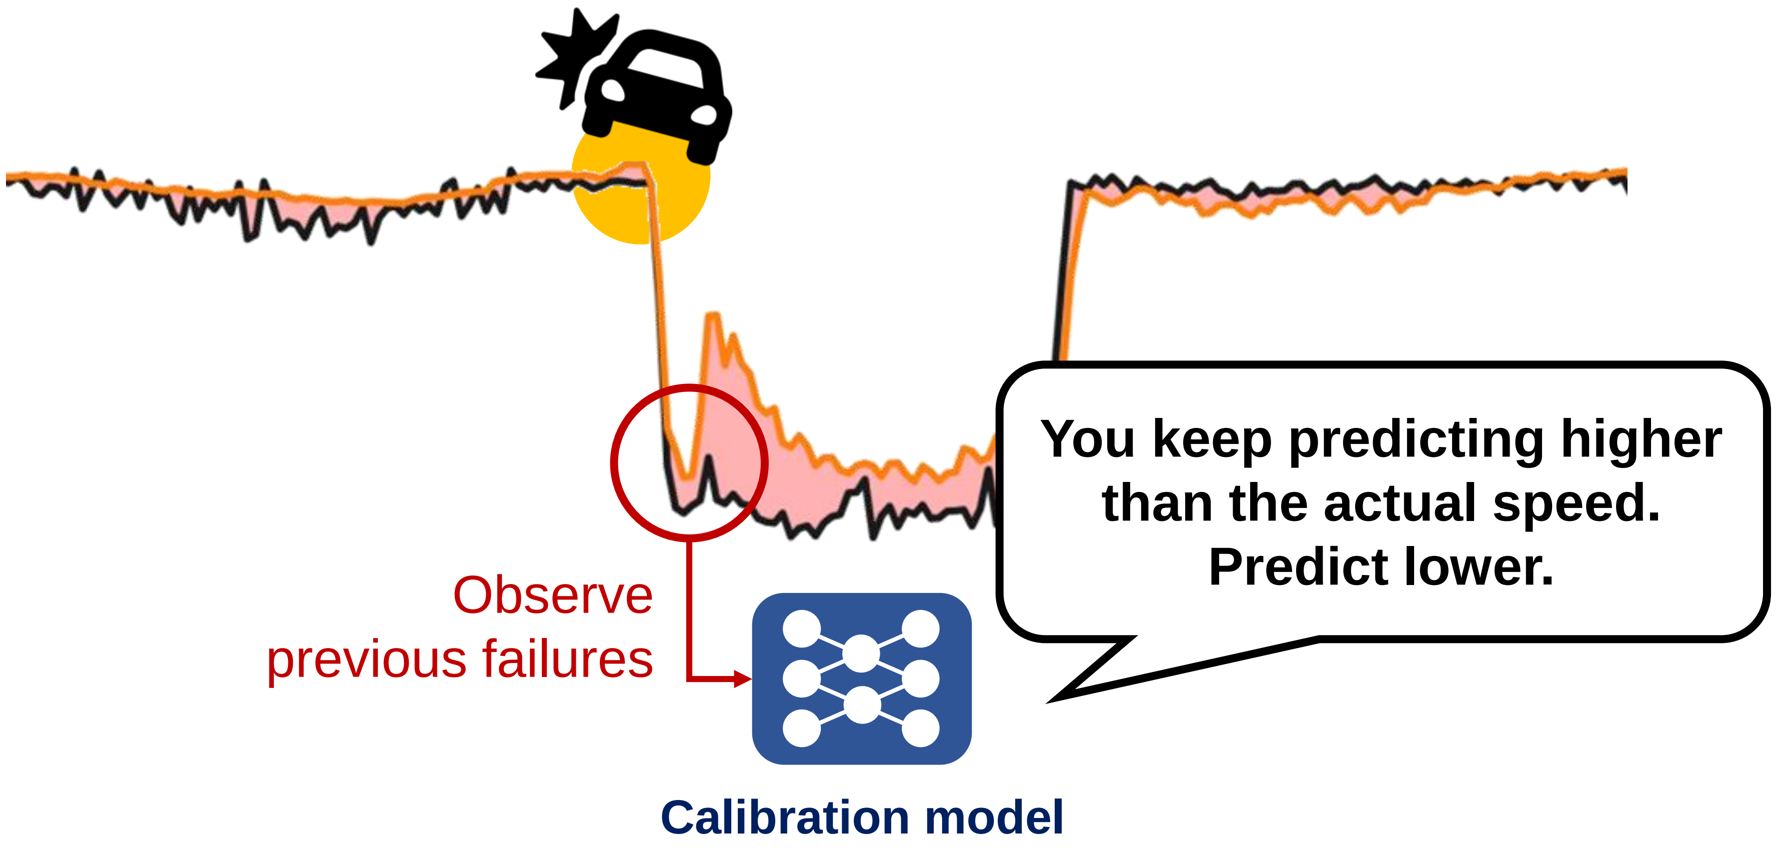

Residual in Traffic Forecasting

- Residuals R^t are defined as: **target values $\boldsymbol{Y^t}$ – estimated values $\boldsymbol{\hat{Y}^t}$**.

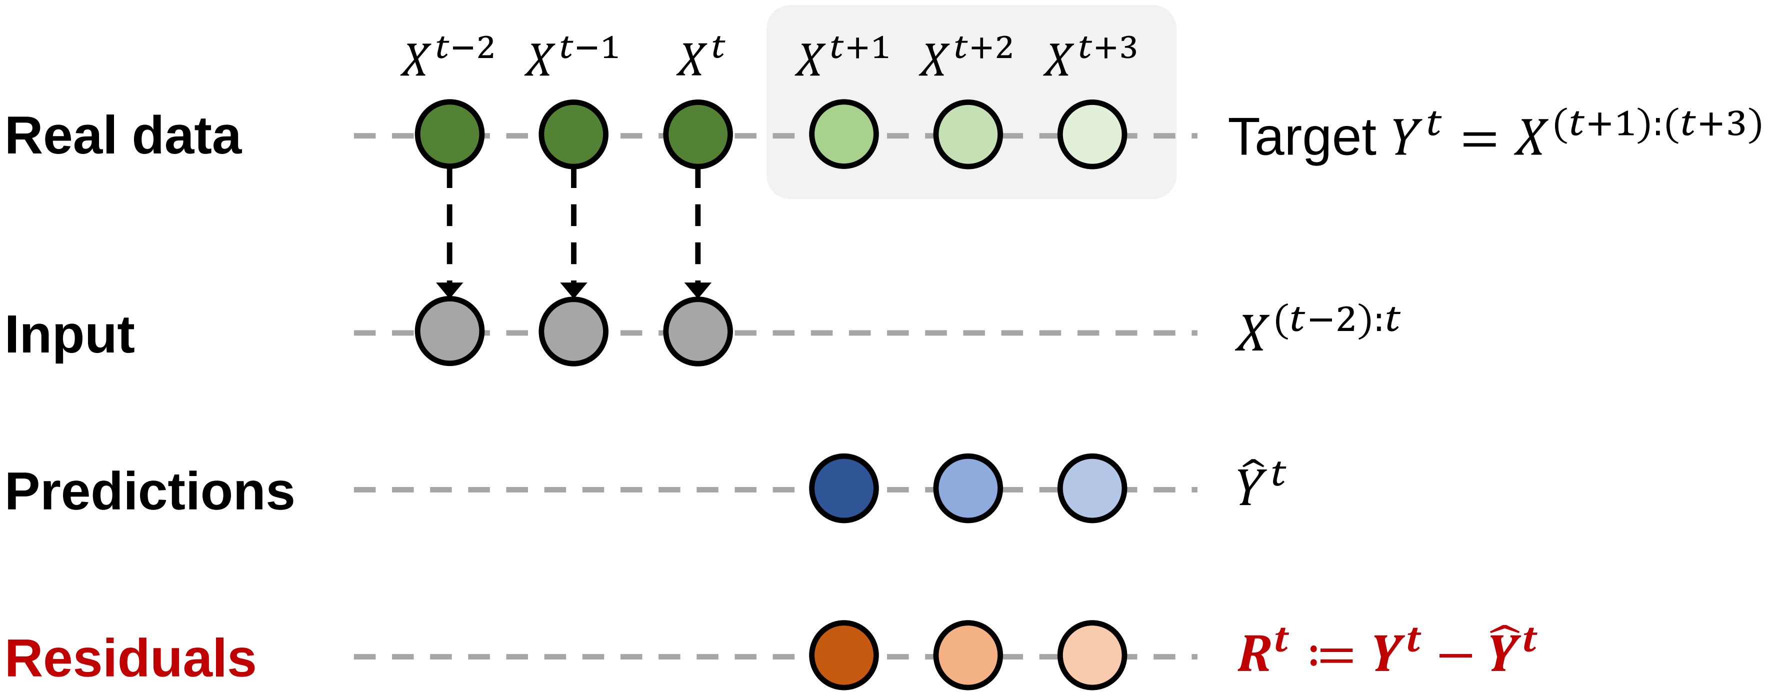

Residual Autocorrelation

- In regression, the residuals (*i.e.*, errors) should be **independent** of each other.
- Otherwise, there is some “***predictable***” information left after forecasting.

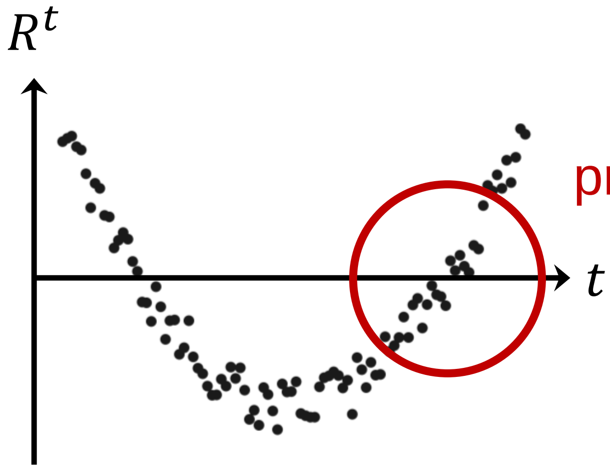

There is still predictable information in residuals.

Correlated Residuals (Bad model)

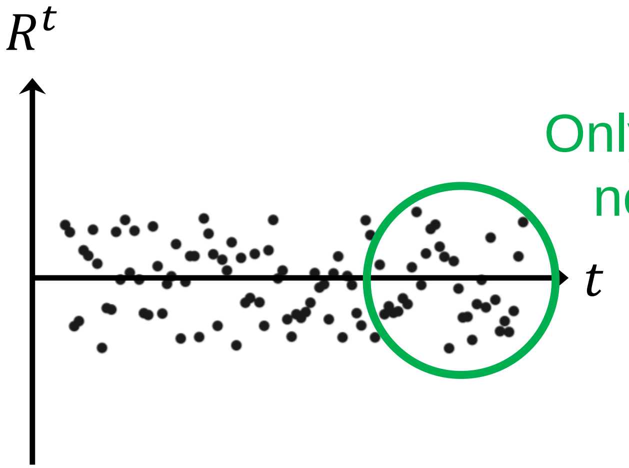

Only unpredictable noise remains.

Uncorrelated Residuals (Good model)

Residual Diagnostics

- Autocorrelation function (ACF) plots are used for analyzing the residuals.
- ACF plots show the correlation between residuals and their lagged versions.

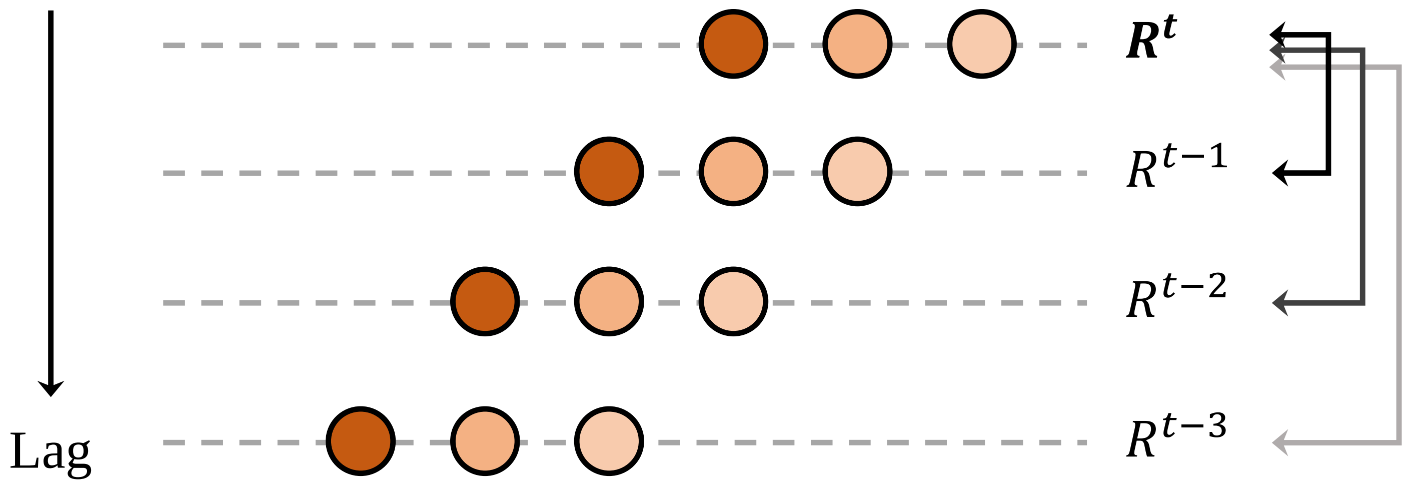

"How much are they correlated?"

$$r_k = \frac{\sum_{t=k+1}^{T}(y_t - \bar{y})(y_{t-k} - \bar{y})}{\sum_{t=1}^{T}(y_t - \bar{y})^2}$$

Residual Diagnostics

- In existing forecasting models, the residuals are ***highly correlated*** with their lagged versions (*i.e.*, previous residuals).

- Deep-learning-based models still leave some predictable information!

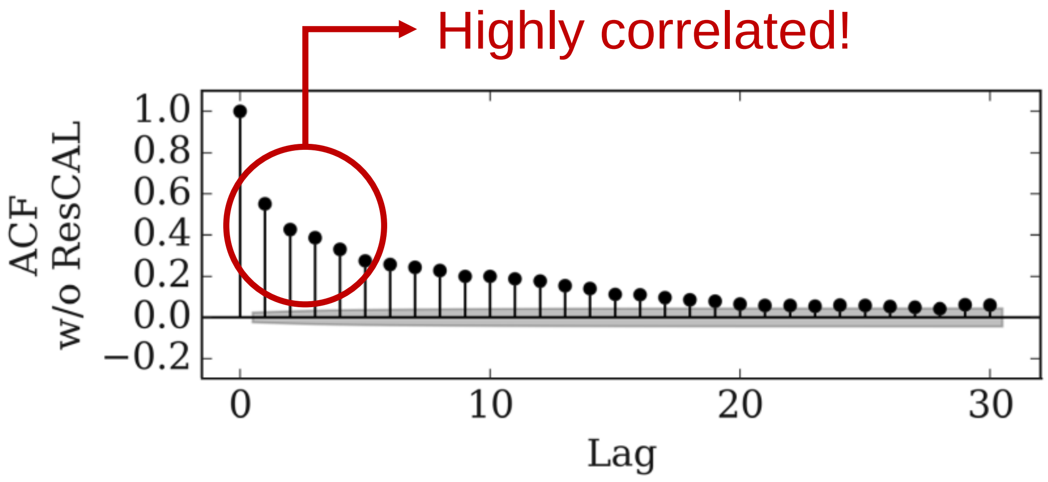

(a) STGCN (Original)

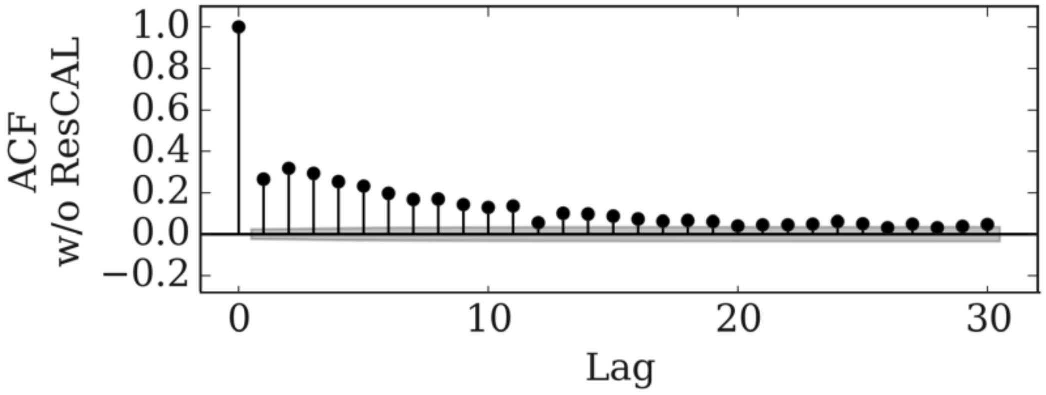

(b) DCRNN (Original)

Residual Diagnostics

- In traffic forecasting, the residuals of the node are not only correlated with its own previous residuals, but also with the previous errors of **neighboring nodes**.

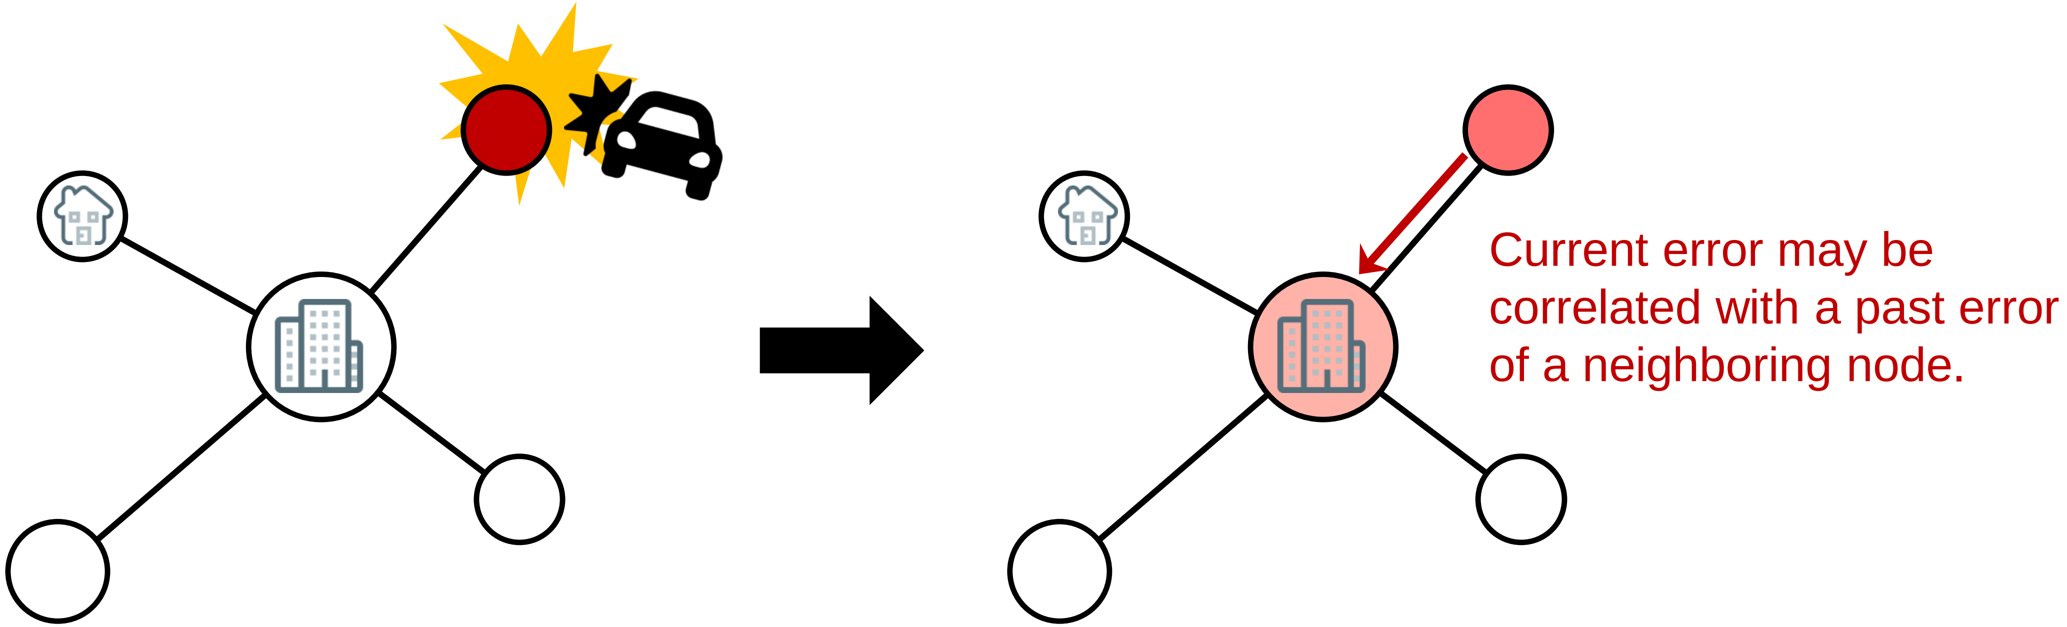

Failure in previous step **Failure in current step**

Residual Diagnostics

- We observed that the current residual of the node is highly correlated with the previous residuals of its neighboring nodes.

- The graph structure should also be considered for addressing residual autocorrelation.

Light colors correspond to high correlations.

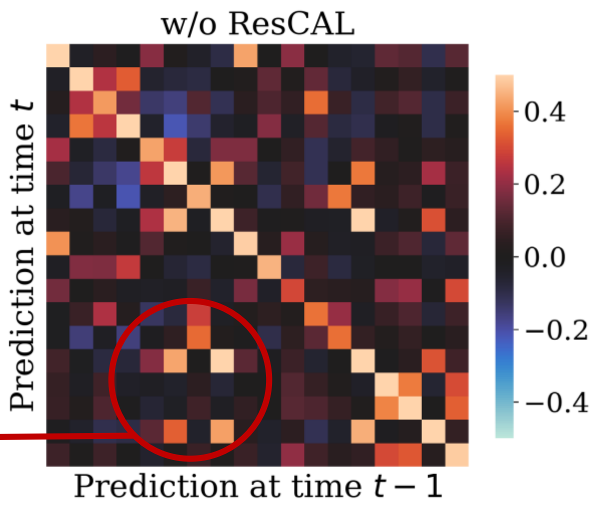

(a) STGCN (Original)

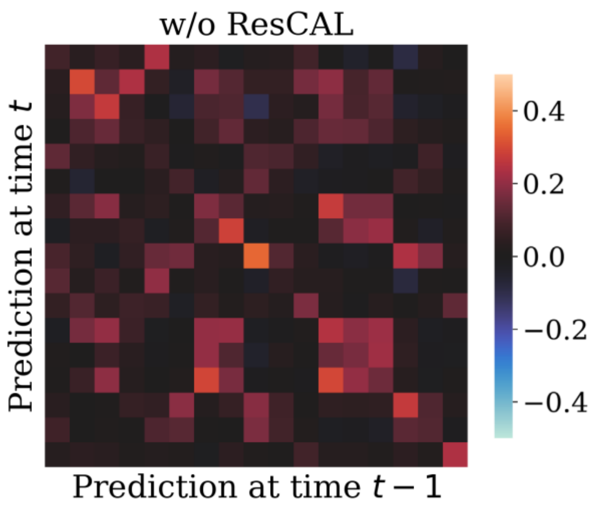

(b) DCRNN (Original)

How to Remove Autocorrelation?

1. Find better forecasting model (New SOTA model!).

- Ideally, if the forecasting model is perfect, we don’t have to consider residuals.
- But… **VERY CHALLENGING!**

How to Remove Autocorrelation?

1. Find better forecasting model (New SOTA model!).

- Ideally, if the forecasting model is perfect, we don’t have to consider residuals.
- But… **VERY CHALLENGING!**

2. Predict residuals and remove correlations DIRECTLY.

- Straightforward and simple.
- Compared to the original task, predicting residual is much easier.
- Compact and model-agnostic solution.

Can We Predict Residuals?

- We already know that residuals are highly correlated. => **Predictable!**
- The residuals do not appear randomly but appear based on previous ones.

In similar events, the model repeats similar failures :(

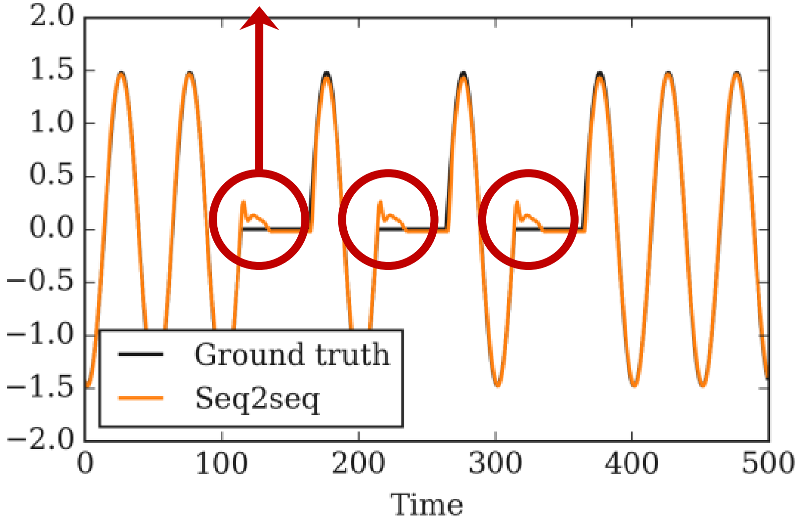

Can We Predict Residuals?

- We already know that residuals are highly correlated. => **Predictable!**
- The residuals do not appear randomly but appear based on previous ones.

In similar events, the model repeats similar failures :(

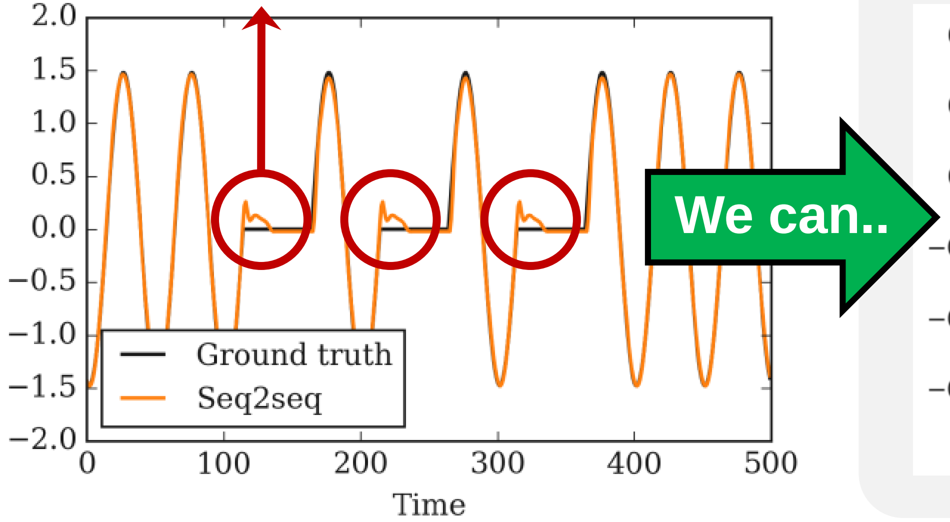

(1) Estimate residuals

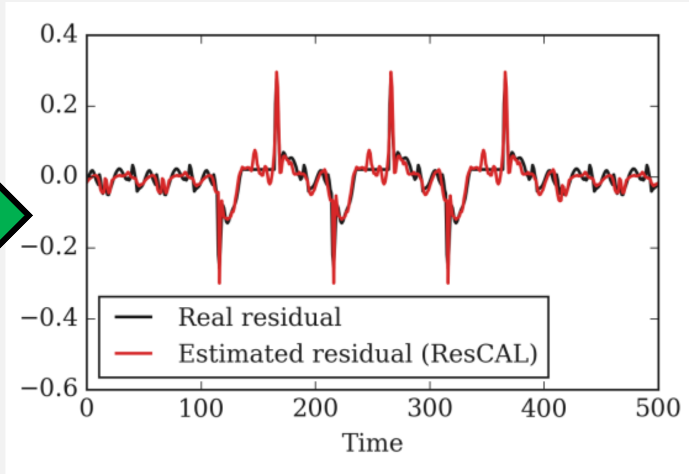

Can We Predict Residuals?

- We already know that residuals are highly correlated. => **Predictable!**
- The residuals do not appear randomly but appear based on previous ones.

In similar events, the model repeats similar failures :(

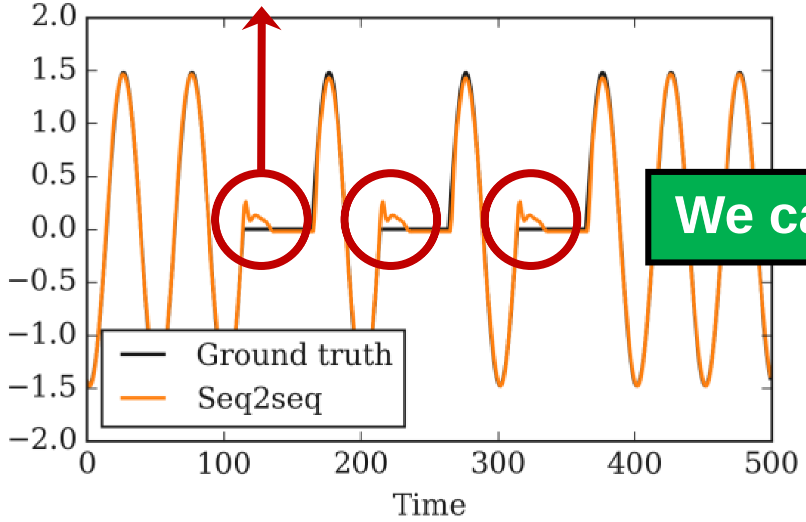

We can..

(1) Estimate residuals

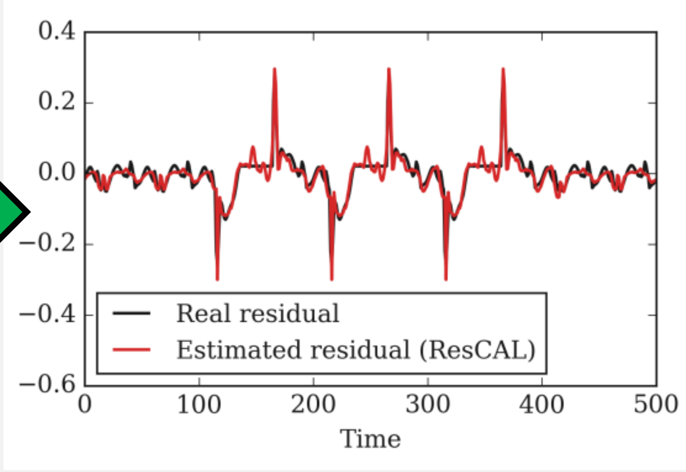

(2) Calibrate predictions

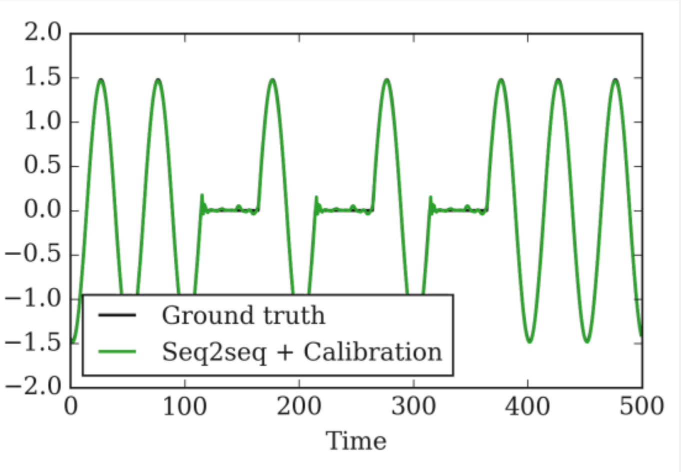

Residual Correction

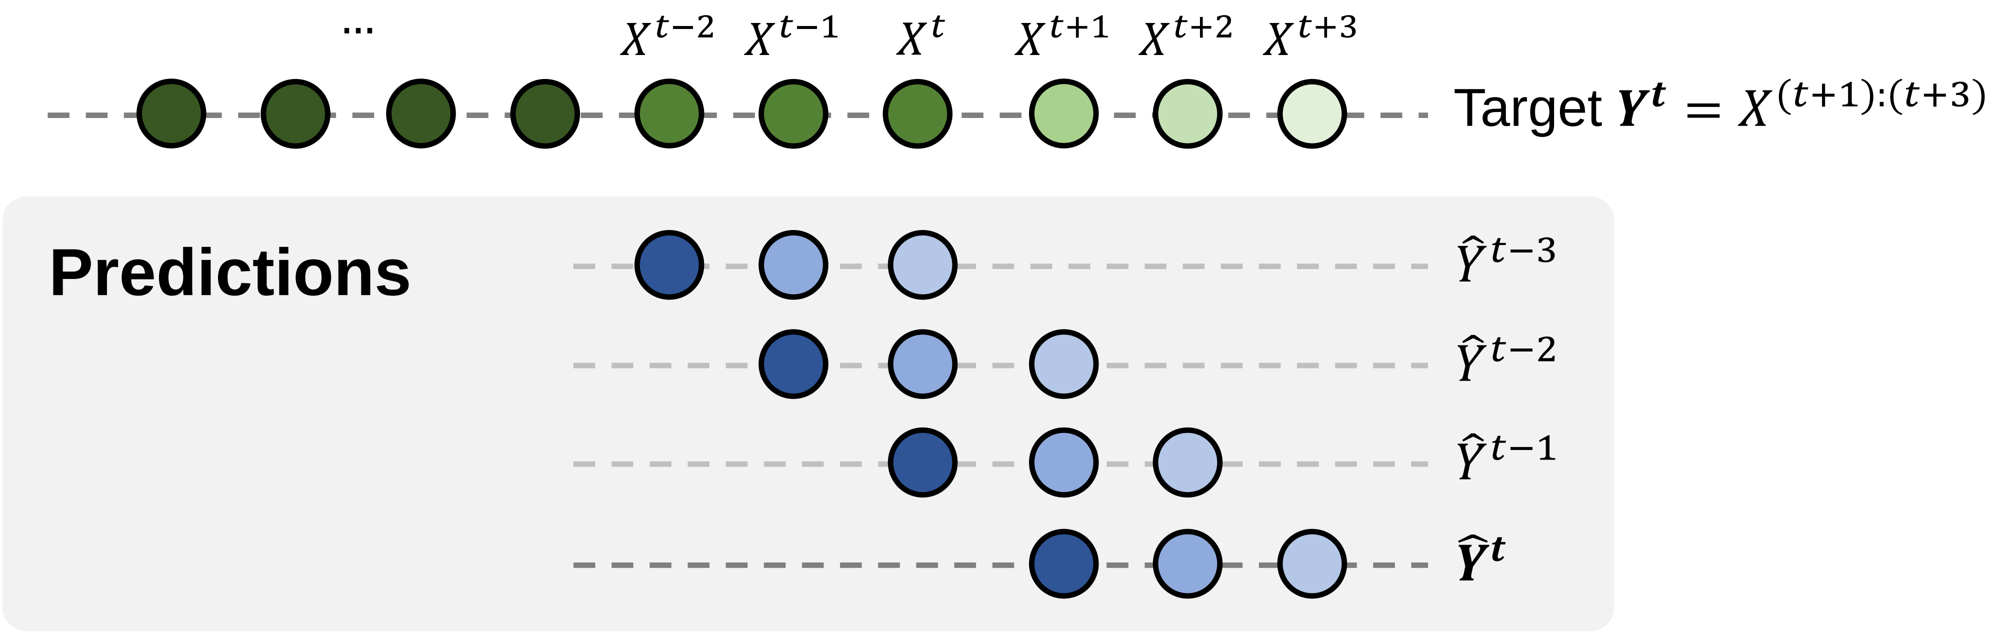

Residual Correction

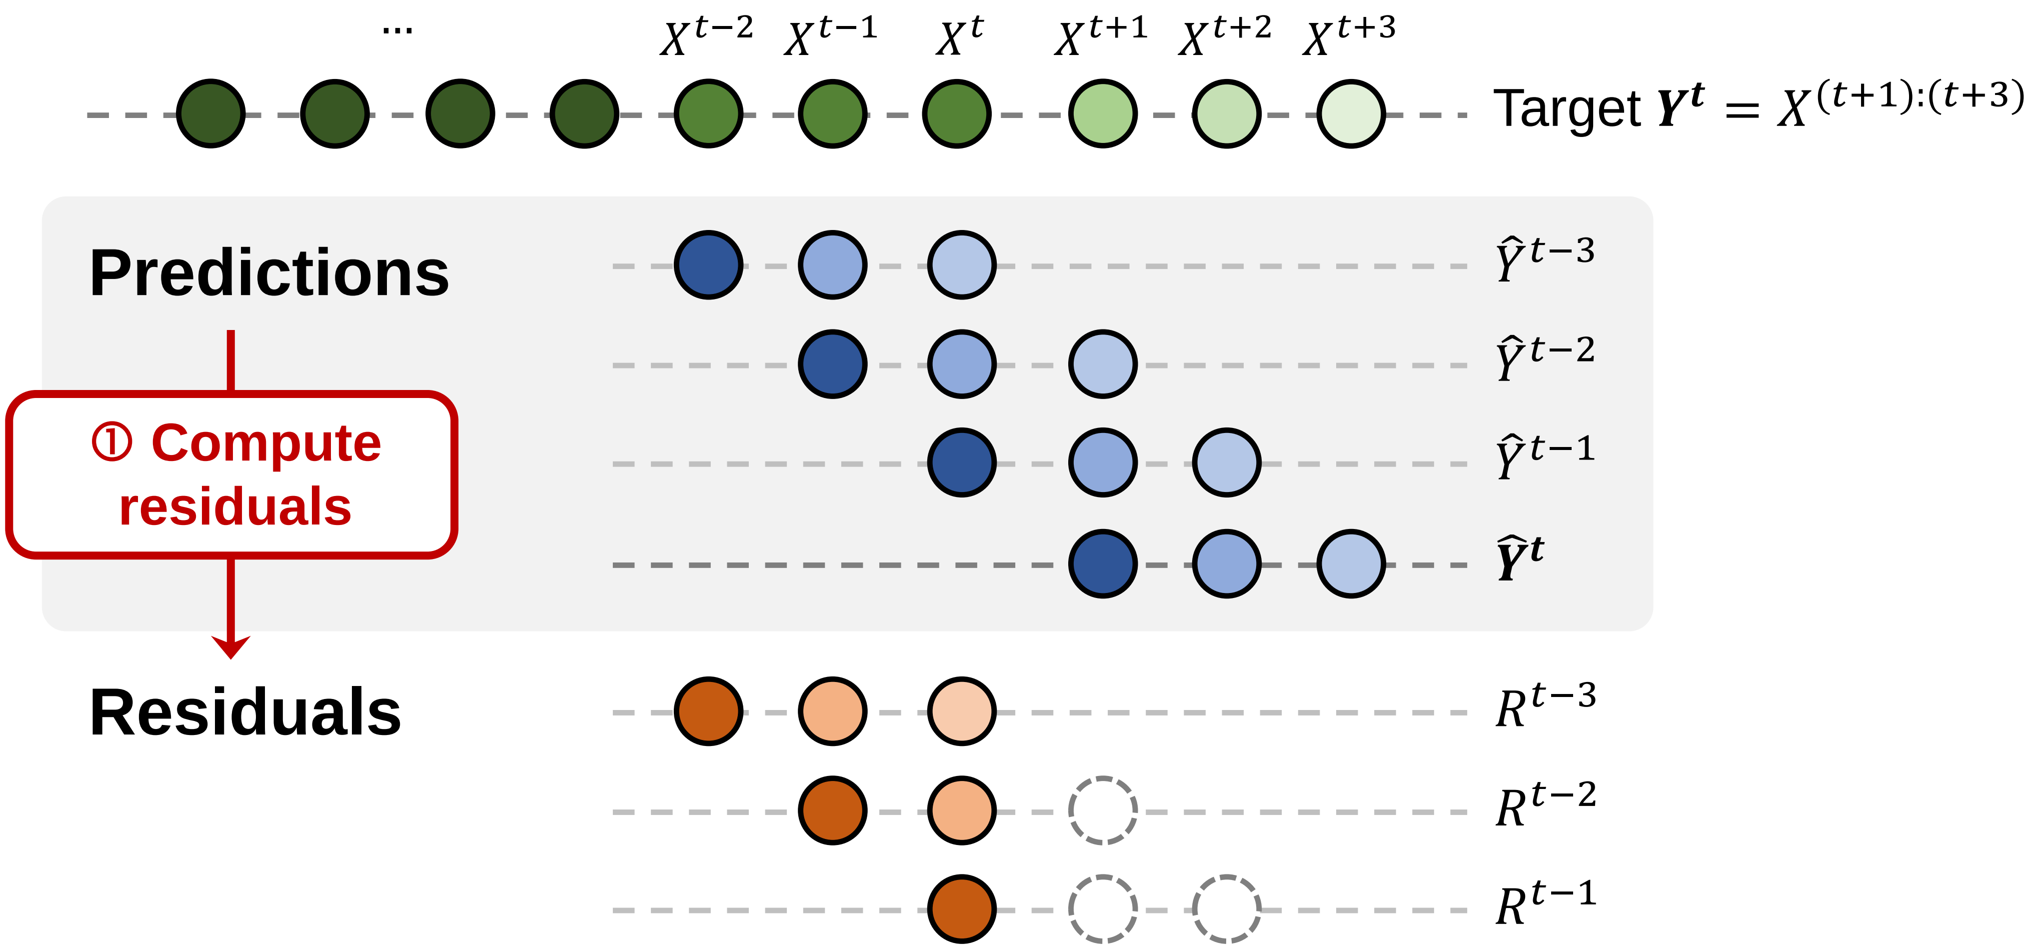

Residual Correction

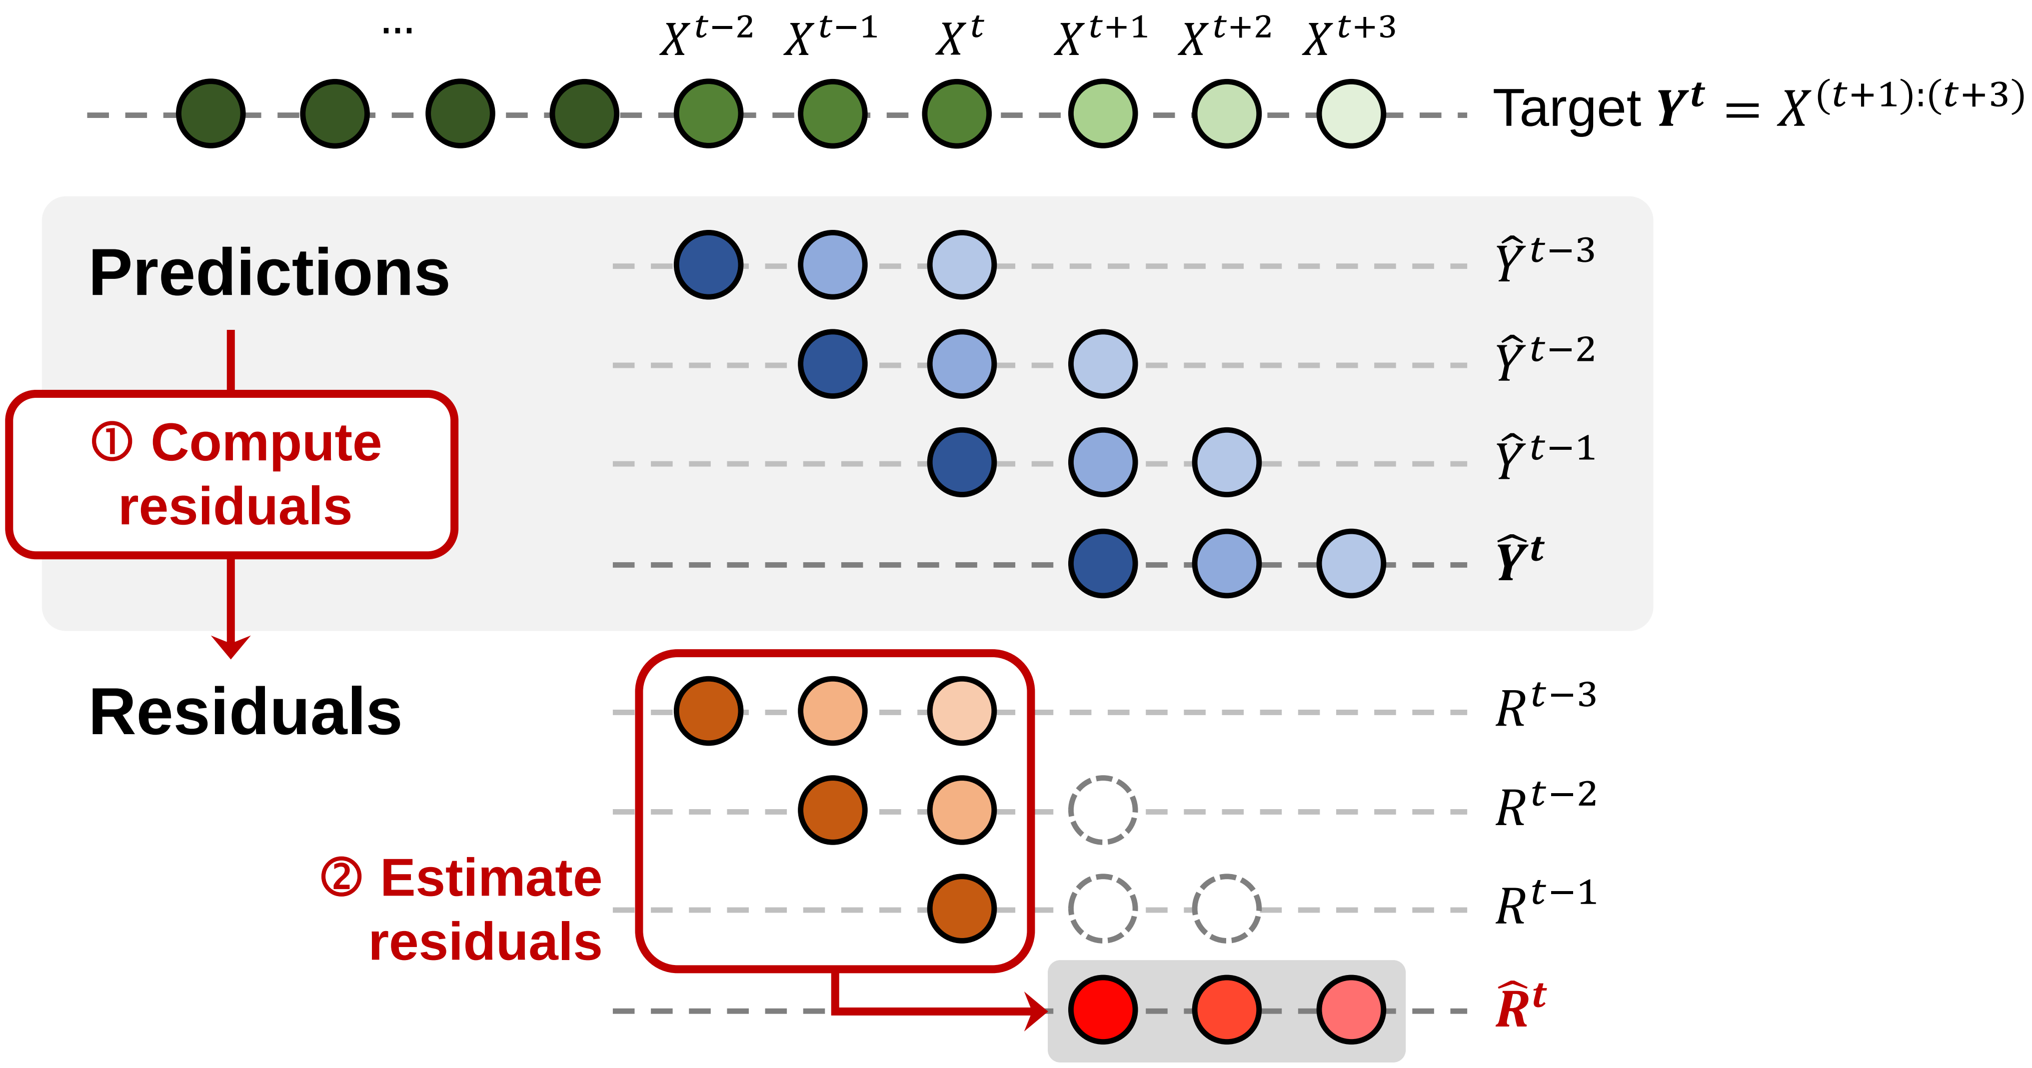

Residual Correction

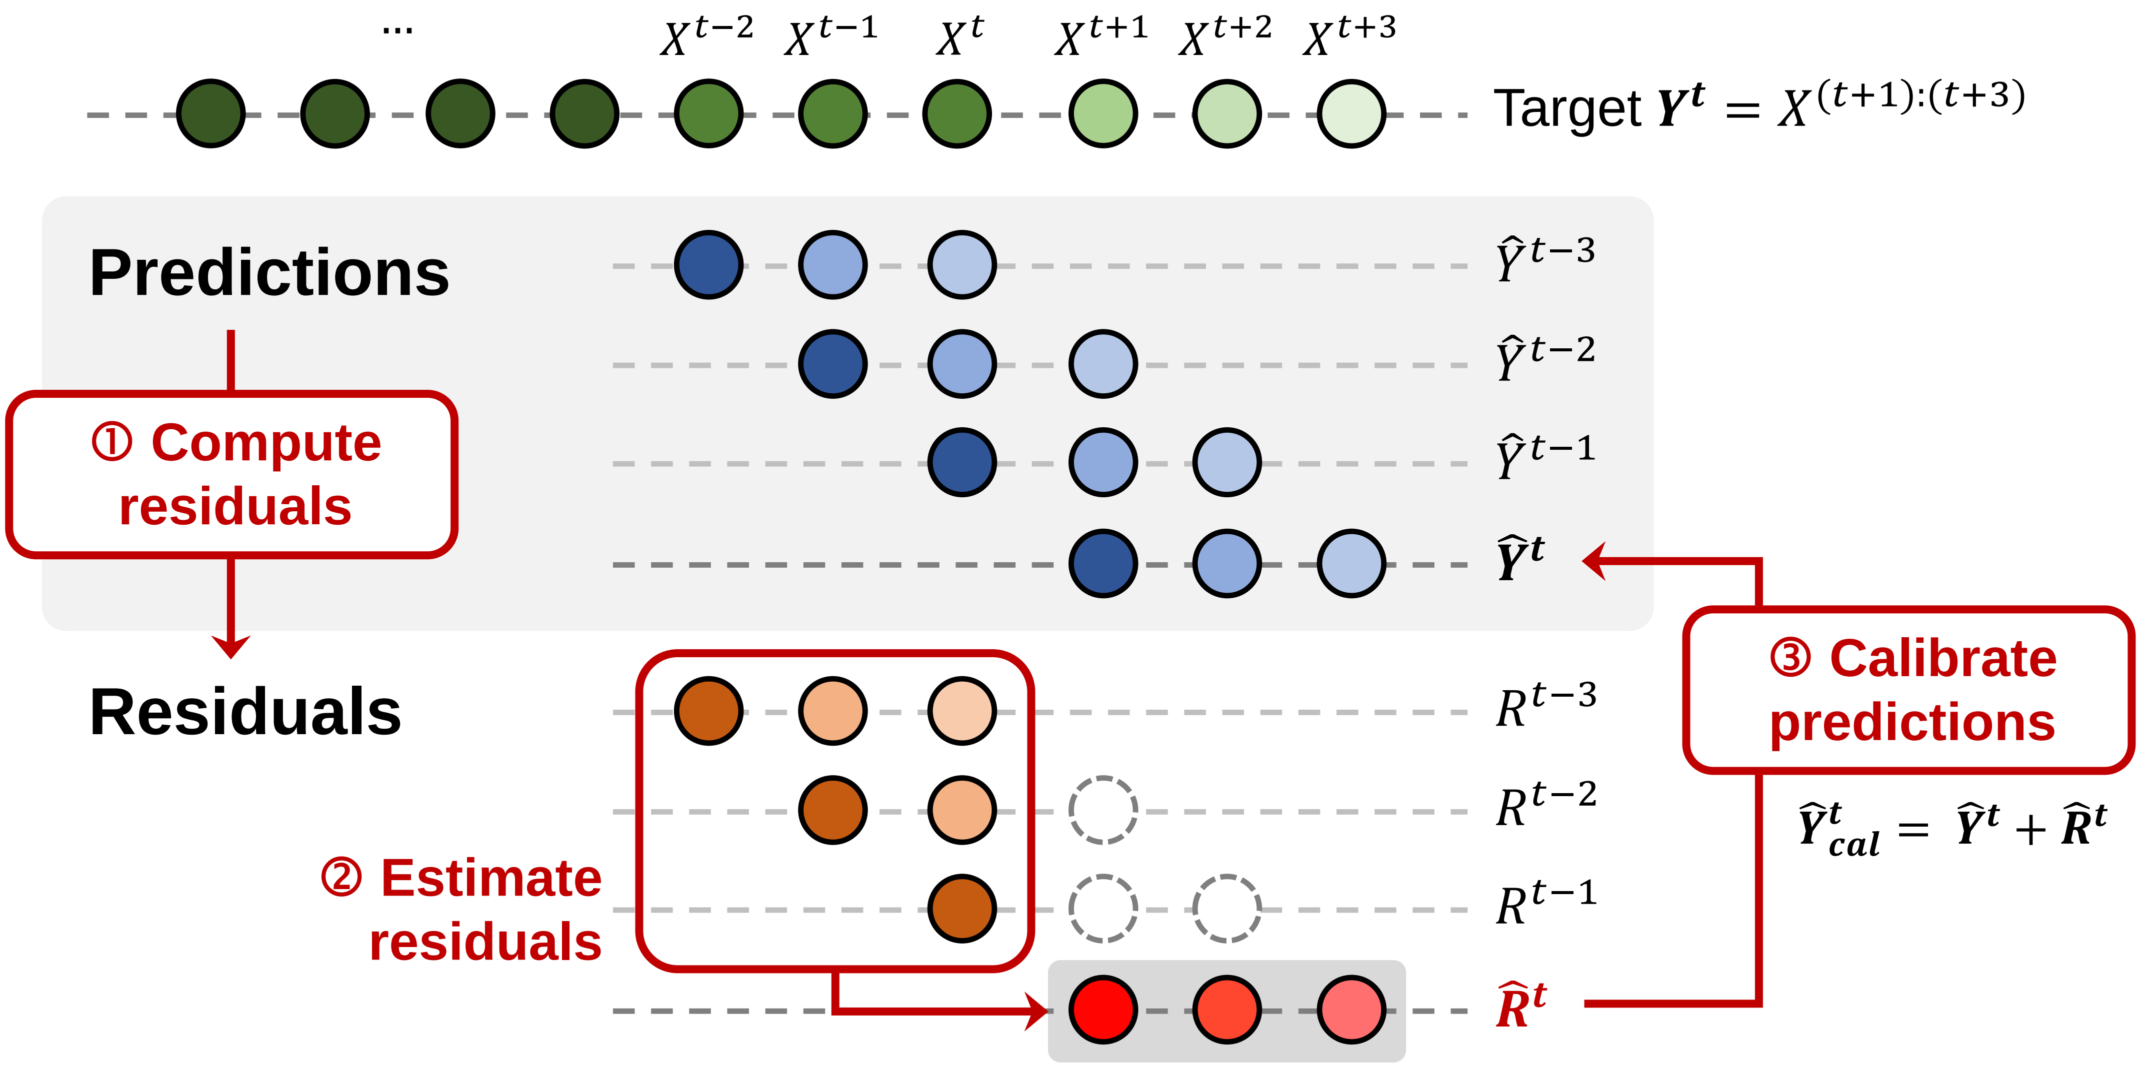

Residual Correction: Model

- Estimating residual is similar to the original traffic forecasting.
- Our **ResCAL**, consisting of spatial-temporal layers, is trained to estimate the future residuals with the previous residuals.

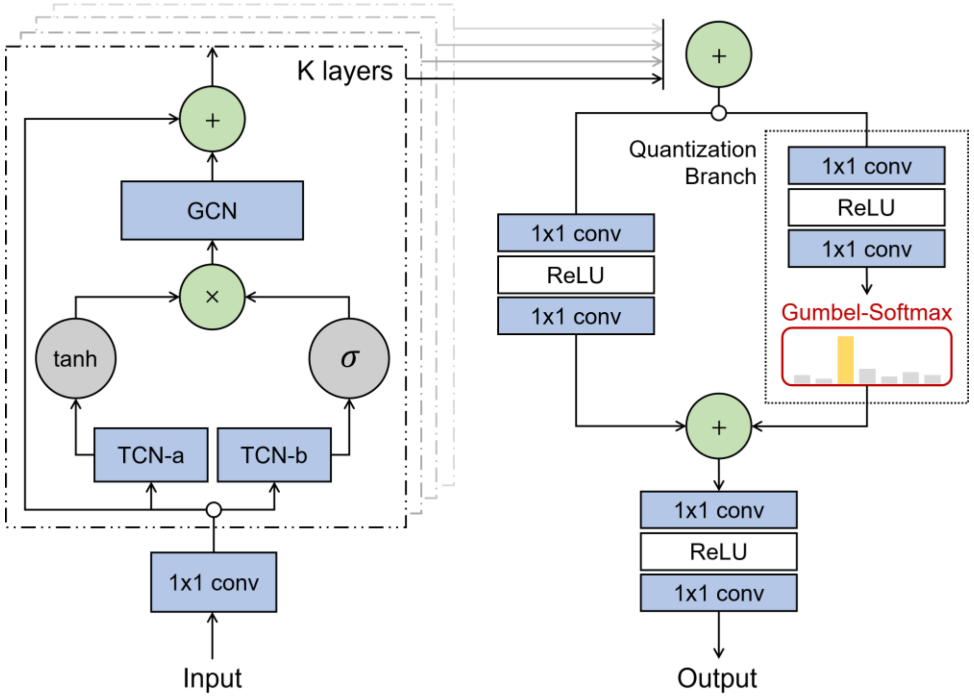

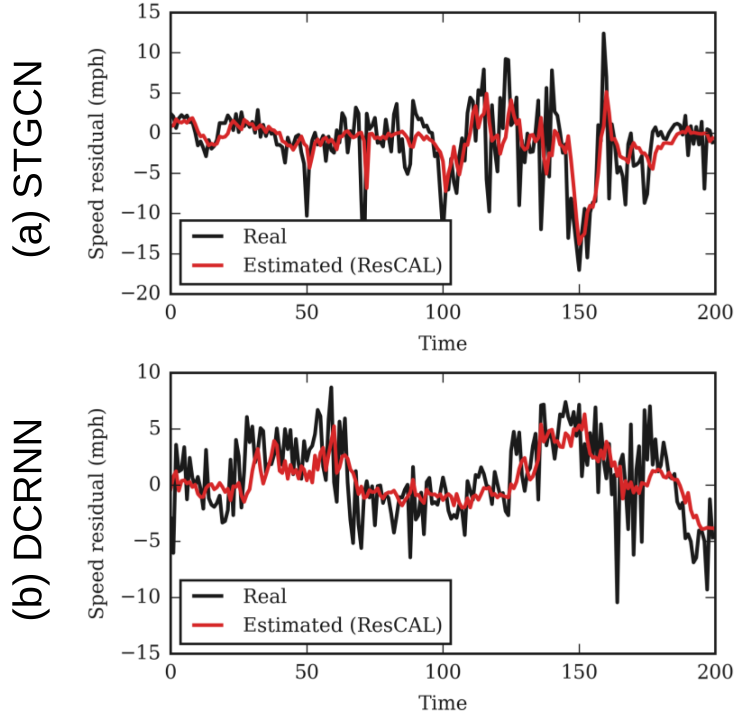

Real residuals

Estimated residuals

Residual Correction: Example

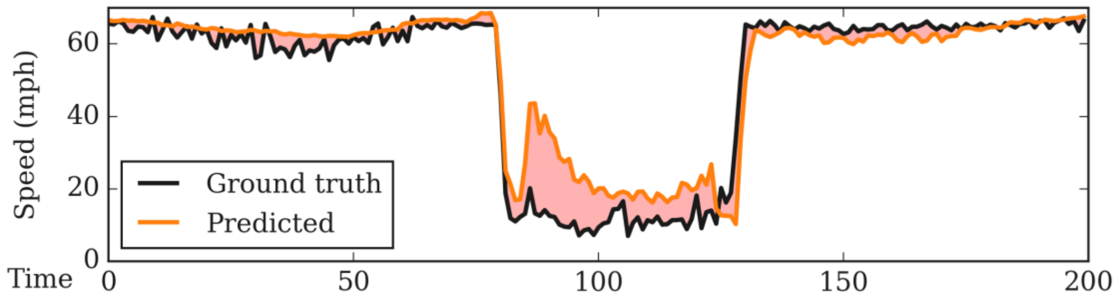

1. The base model (*e.g.*, STGCN) makes the initial predictions (may be wrong).

Residual Correction: Example

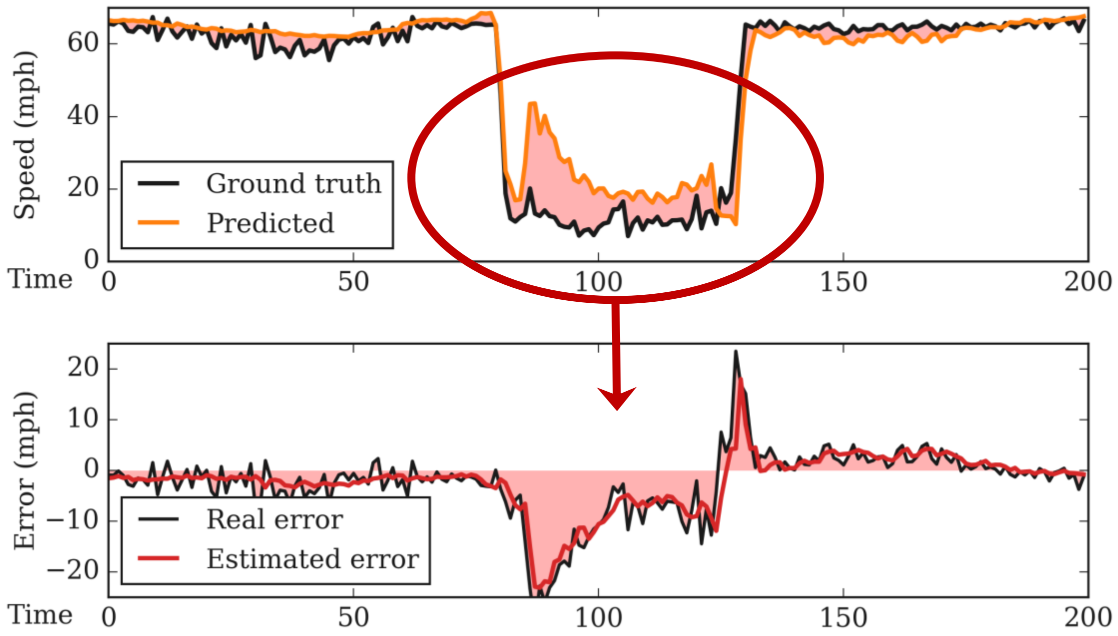

1. The base model (*e.g.*, STGCN) makes the initial predictions (may be wrong).

2. Using the previous residuals, ResCAL **estimates the future residuals.**

Residual Correction: Example

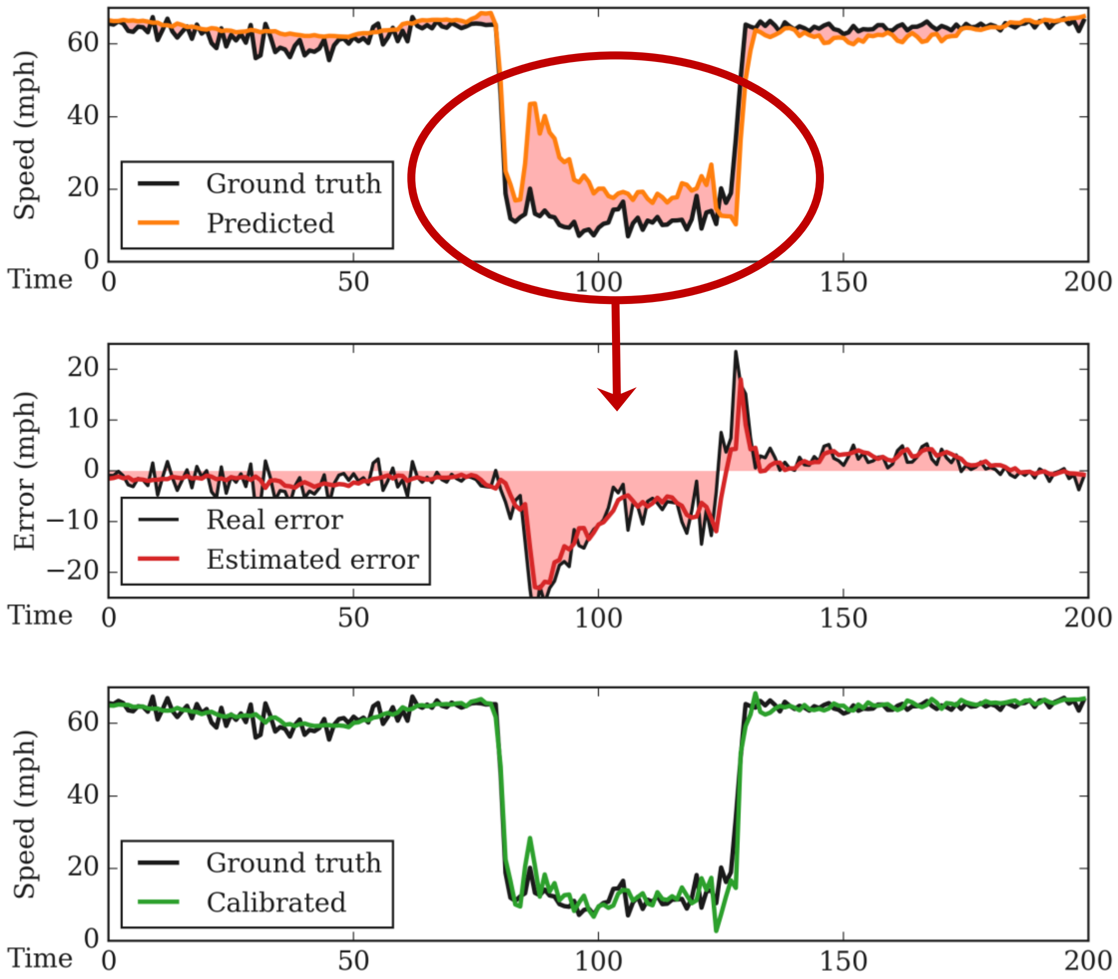

1. The base model (*e.g.*, STGCN) makes the initial predictions (may be wrong).

2. Using the previous residuals, ResCAL **estimates the future residuals.**

3. Now, we can **calibrate** the predictions.

Residual Correction: Results

- To validate that our model successfully ***corrects*** the base model, we focus on the region where the base model ***failed.***

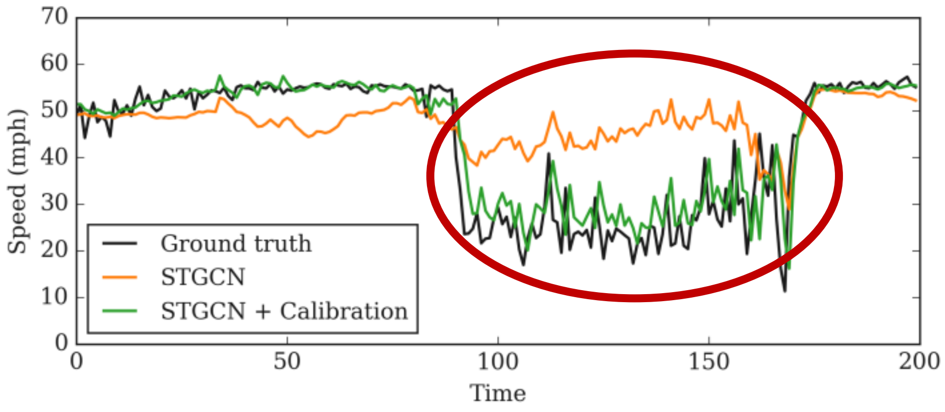

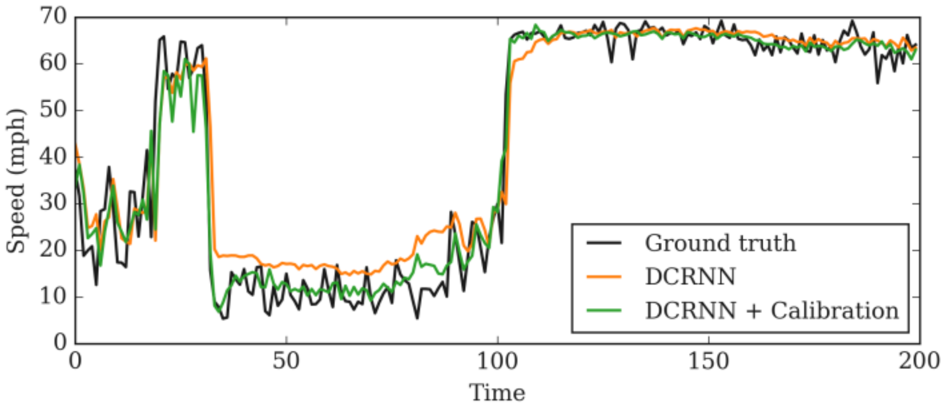

Data	Models	15min			30min			60min		
		MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
METR-LA	DCRNN [17]	14.53	16.31	48.63%	16.98	18.93	60.79%	20.03	22.00	75.27%
	+ Calibration	**13.46**	**15.44**	**44.22%**	**16.26**	**18.38**	**57.37%**	**19.45**	**21.54**	**73.13%**
	STGCN [27]	15.29	16.93	53.19%	17.94	19.78	65.91%	21.17	23.00	82.76%
	+ Calibration	**12.77**	**15.07**	**43.55%**	**15.79**	**18.18**	**57.50%**	**19.10**	**21.44**	**73.93%**
	Graph WaveNet [26]	13.39	15.06	41.49%	16.15	17.99	54.92%	19.08	20.93	69.24%
	+ Calibration	**13.28**	**14.99**	**41.45%**	**16.03**	**17.93**	**54.43%**	**18.94**	**20.85**	**68.34%**
	STAWnet [21]	13.56	15.30	43.75%	16.15	18.10	56.00%	19.00	21.00	68.44%
	+ Calibration	**13.17**	**14.98**	**42.24%**	**15.87**	**17.89**	**54.88%**	**18.80**	**20.84**	**68.21%**
PEMS-BAY	DCRNN [17]	4.41	6.07	10.51%	5.95	8.45	15.36%	7.50	10.53	20.39%
	+ Calibration	**4.25**	**5.93**	**10.07%**	**5.65**	**8.15**	**14.37%**	**6.94**	**10.00**	**18.60%**
	STGCN [27]	6.42	7.95	15.88%	7.50	9.44	19.39%	8.74	11.03	23.50%
	+ Calibration	**3.88**	**5.91**	**9.55%**	**5.41**	**8.00**	**14.36%**	**6.84**	**9.78**	**18.94%**
	Graph WaveNet [26]	4.36	5.89	10.45%	5.72	7.97	14.46%	6.89	9.50	18.07%
	+ Calibration	**4.28**	**5.87**	**10.07%**	**5.63**	**7.93**	**14.21%**	**6.81**	**9.46**	**18.02%**
	STAWnet [21]	4.38	5.96	10.34%	5.71	7.91	14.51%	6.75	9.29	18.07%
	+ Calibration	**4.27**	**5.87**	**10.07%**	**5.62**	**7.88**	**14.23%**	**6.70**	**9.27**	**17.77%**

Revisit: Residual Diagnostics

- Originally, the residuals of the existing models were ***highly correlated*** with their lagged versions (*i.e.*, previous residuals).

- ***After calibration,*** the residuals are almost ***independent*** of the previous ones!

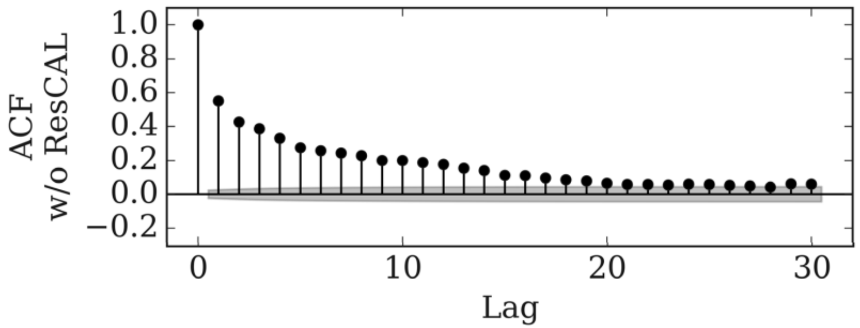

(a) STGCN (After calibration)

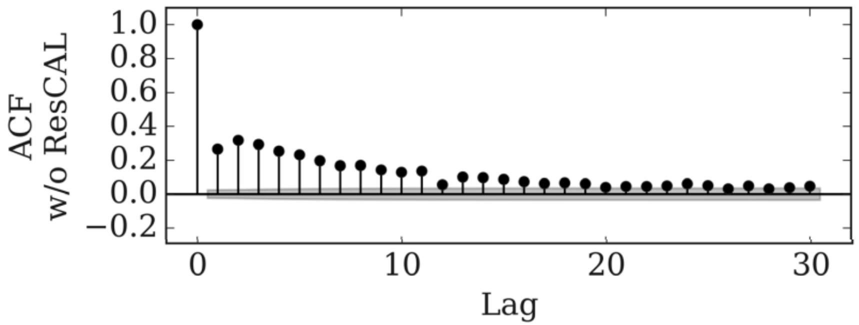

(b) DCRNN (After calibration)

Revisit: Residual Diagnostics

- Originally, the residuals of the existing models were ***highly correlated*** with their lagged versions (*i.e.*, previous residuals).

- ***After calibration,*** the residuals are almost ***independent*** of the previous ones!

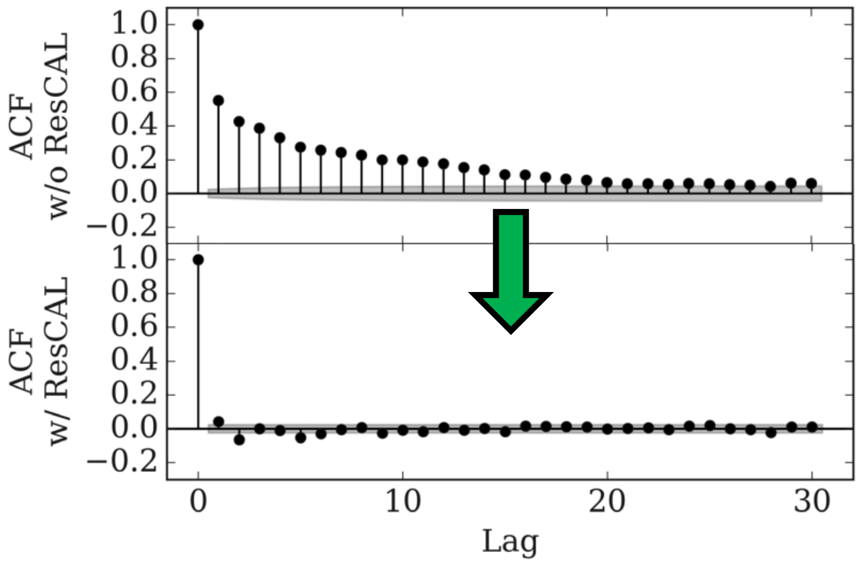

(a) STGCN (After calibration)

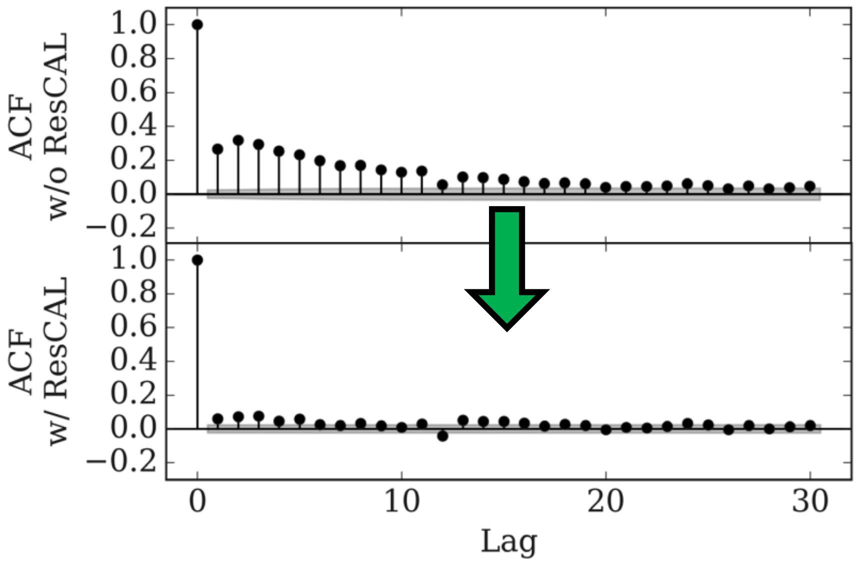

(b) DCRNN (After calibration)

Revisit: Residual Diagnostics

- Originally, the current residual of the node is highly correlated with the previous residuals of its neighboring nodes.

- ***After calibration,*** the correlation with the previous residuals of other nodes almost disappears.

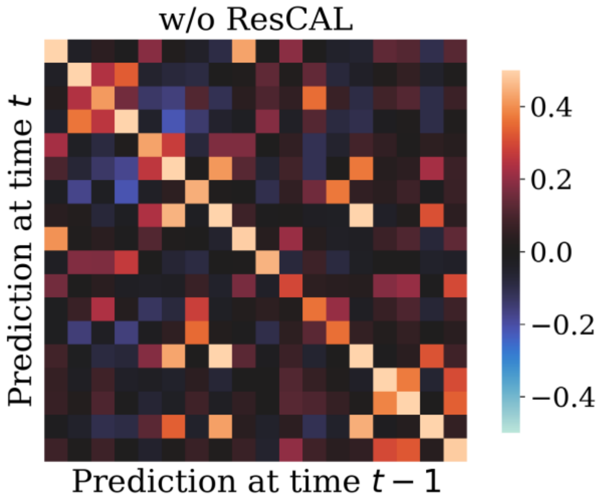

(a) STGCN (After calibration)

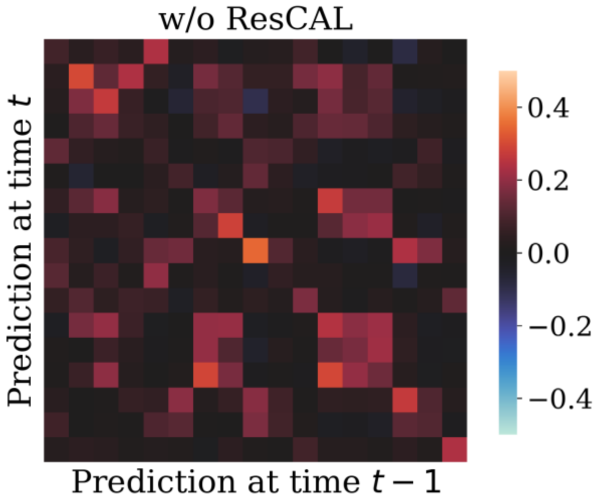

(b) DCRNN (After calibration)

Revisit: Residual Diagnostics

- Originally, the current residual of the node is highly correlated with the previous residuals of its neighboring nodes.

- ***After calibration,*** the correlation with the previous residuals of other nodes almost disappears.

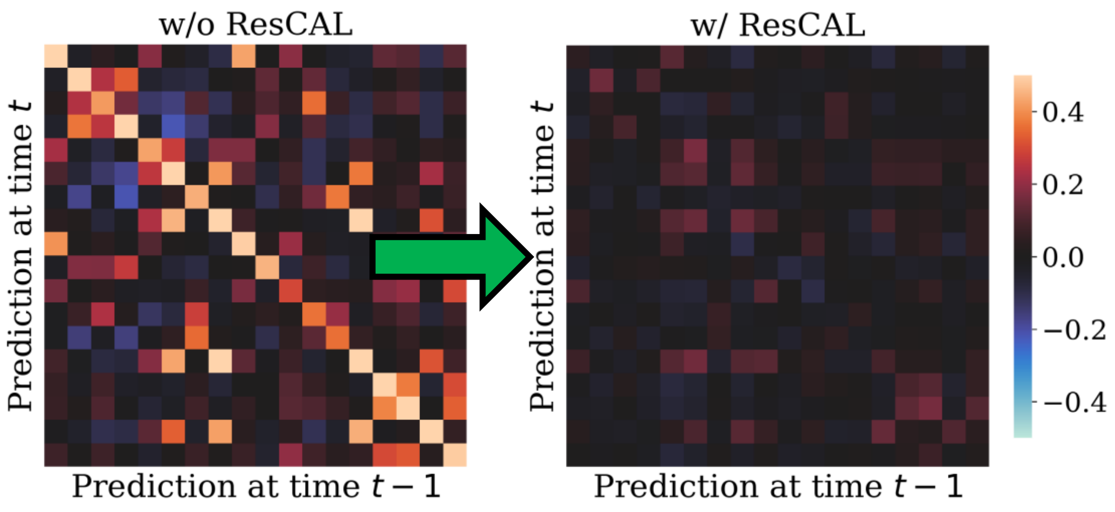

(a) STGCN (After calibration)

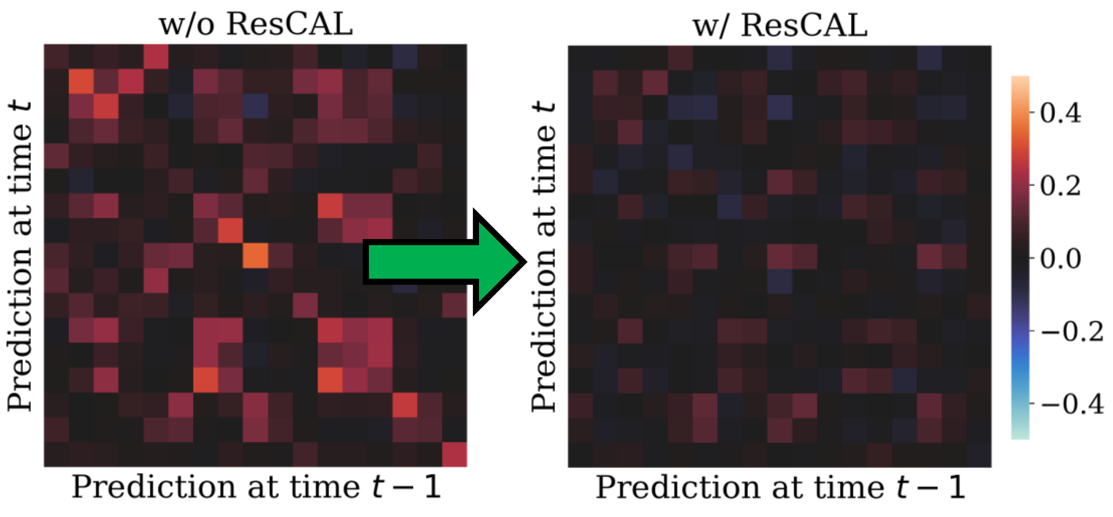

(b) DCRNN (After calibration)

Conclusion

- We observed that the residual autocorrelation could not be properly addressed by existing deep-learning-based models.

- In this work, we proposed a novel method for calibrating the predictions of traffic forecasting models using their previous residuals.

- In the experiments, we confirmed that our method leaves only unpredictable noise and removes the residual autocorrelation.

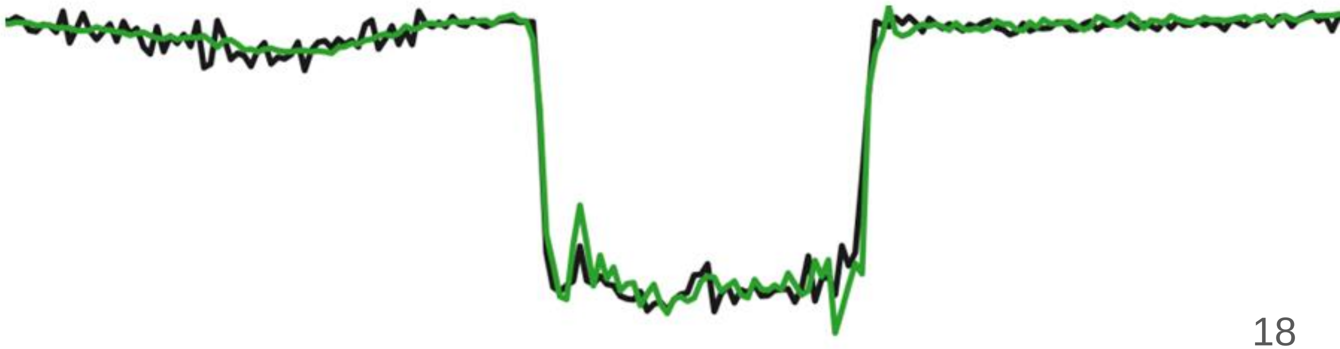

Thank you

Please refer to our paper for more details!

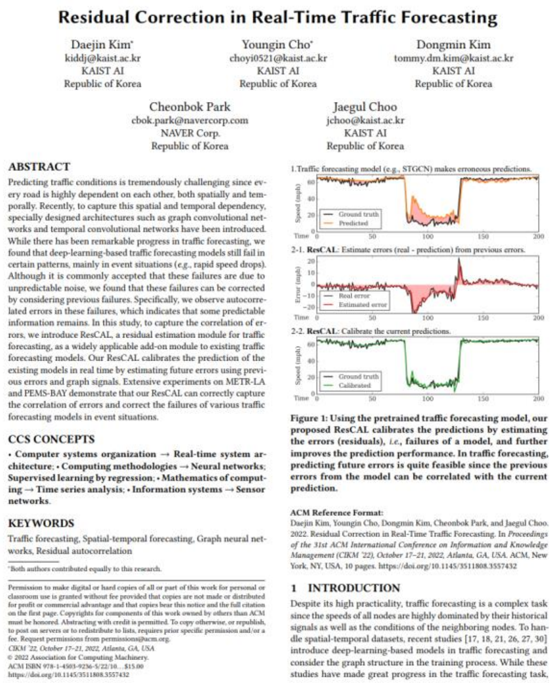

Residual Correction in Real-Time Traffic Forecasting

Daejin Kim*
kiddj@kaist.ac.kr
KAIST AI
Republic of Korea

Youngin Cho*
choyi0521@kaist.ac.kr
KAIST AI
Republic of Korea

Dongmin Kim
tommy.dm.kim@kaist.ac.kr
KAIST AI
Republic of Korea

Cheonbok Park
cbok.park@navercorp.com
NAVER Corp.
Republic of Korea

Jaegul Choo
jchoo@kaist.ac.kr
KAIST AI
Republic of Korea

ABSTRACT

Predicting traffic conditions is tremendously challenging since every road is highly dependent on each other, both spatially and temporally. Recently, to capture this spatial and temporal dependency, specially designed architectures such as graph convolutional networks and temporal convolutional networks have been introduced. While there has been remarkable progress in traffic forecasting, we found that deep-learning-based traffic forecasting models still fail in certain patterns, mainly in event situations (*e.g.*, rapid speed drops). Although it is commonly accepted that these failures are due to unpredictable noise, we found that these failures can be corrected by considering previous failures. Specifically, we observe autocorrelated errors in these failures, which indicates that some predictable information remains. In this study, to capture the correlation of errors, we introduce ResCAL, a residual estimation module for traffic forecasting, as a widely applicable add-on module to existing traffic forecasting models. Our ResCAL calibrates the prediction of the existing models in real time by estimating future errors using previous errors and graph signals. Extensive experiments on METR-LA and PEMS-BAY demonstrate that our ResCAL can correctly capture the correlation of errors and correct the failures of various traffic forecasting models in event situations.

CCS CONCEPTS

• Computer systems organization → Real-time system architecture; • Computing methodologies → Neural networks; Supervised learning by regression; • Mathematics of computing → Time series analysis; • Information systems → Sensor networks.

KEYWORDS

Traffic forecasting, Spatial-temporal forecasting, Graph neural networks, Residual autocorrelation

*Both authors contributed equally to this research.

CIKM '22, October 17–21, 2022, Atlanta, GA, USA

ACM ISBN 978-1-4503-9236-5/22/10...$15.00
https://doi.org/10.1145/3511808.3557432

Figure 1: Using the pretrained traffic forecasting model, our proposed ResCAL calibrates the predictions by estimating the errors (residuals), *i.e.*, failures of a model, and further improves the prediction performance. In traffic forecasting, predicting future errors is quite feasible since the previous errors from the model can be correlated with the current prediction.

ACM Reference Format:
Daejin Kim, Youngin Cho, Dongmin Kim, Cheonbok Park, and Jaegul Choo. 2022. Residual Correction in Real-Time Traffic Forecasting. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22), October 17–21, 2022, Atlanta, GA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3511808.3557432

1 INTRODUCTION

Despite its high practicality, traffic forecasting is a complex task since the speeds of all nodes are highly dominated by their historical signals as well as the conditions of the neighboring nodes. To handle spatial-temporal datasets, recent studies [17, 18, 21, 26, 27, 30] introduce deep-learning-based models in traffic forecasting and consider the graph structure in the training process. While these studies have made great progress in the traffic forecasting task,

Residual Correction in
Real-Time Traffic Forecasting