

444 **A Appendix**

445 To support reproducibility of the results in this paper, we have submitted our code and datasets as  
 446 the supplementary information. Here, we will present the datasets statistics, evaluation metrics,  
 447 implementation details, and more results.

448 **A.1 Datasets Statistics**

449 The dataset used in our experiments (namely PeMSD4 dataset and PeMSD8 dataset) contain the  
 450 traffic flow data measured by road traffic sensors. As introduced in Section 3.1, we formulate the  
 451 traffic forecasting problem on a graph where each node corresponds to a traffic sensor. Our ASTGCN  
 452 can infer spatial proximity from data by DAGG module automatically. Thus it does not require  
 453 pre-defining the adjacent matrix. For graph-based baselines, we reuse the pre-defined graph given in  
 454 [11] to capture spatial correlations. The connectivity between different nodes is determined by the  
 455 actual road network. If two monitors are on the same road, then they are considered connected. The  
 statistics about the two datasets are shown in Table 4.

Table 4: Summary statistics of the PeMSD4 and PeMSD8 dataset

| Dataset | Time Span                | #Nodes | #Edges | #Samples | Data Range | Median |
|---------|--------------------------|--------|--------|----------|------------|--------|
| PeMSD4  | 1/Jan/2018 - 28/Feb/2018 | 307    | 340    | 16992    | 0 ~919     | 180    |
| PeMSD8  | 1/Jul/2016 - 31/Aug/2016 | 170    | 277    | 17856    | 0 ~1147    | 215    |

456

457 **A.2 Evaluation Metrics**

458 We use three evaluation metrics to measure the performance of predictive models. Let  $\mathbf{X}_{:,i} \in R^{N \times 1}$   
 459 be the ground truth traffic of all nodes at time step  $i$ ,  $\mathbf{X}'_{:,i} \in R^{N \times 1}$  be the predicted values, and  $\Omega$  be  
 460 indices of observed samples. The metrics are defined as follows.

Mean Absolute Error (MAE)

$$MAE = \frac{1}{|\Omega|} \sum_{i \in \Omega} |\mathbf{X}_{:,i} - \mathbf{X}'_{:,i}|$$

Root Mean Square Error (RMSE)

$$RMSE = \sqrt{\frac{1}{|\Omega|} \sum_{i \in \Omega} (\mathbf{X}_{:,i} - \mathbf{X}'_{:,i})^2}$$

Mean Absolute Percentage Error (MAPE)

$$MAPE = \frac{1}{|\Omega|} \sum_{i \in \Omega} \left| \frac{\mathbf{X}_{:,i} - \mathbf{X}'_{:,i}}{\mathbf{X}_{:,i}} \right|$$

461 **A.3 Implementation Details**

462 The details of the baselines are as follows:

- 463 • HA: the historical average model operates on each traffic series separately, and it averages  
 464 all the historical traffic at the same time slot to predict current traffic. Historical Average  
 465 does not depend on recent data and thus the performance is invariant for 12 forecasting  
 466 horizons.
- 467 • VAR: we implement the VAR model based on *statsmodel* python package and search the  
 468 number of lags among {1, 3, 6, 9, 12}. The number of lags is set to 12 for both PeMSD4  
 469 and PeMSD8 datasets.

- 470 • GRU-ED: we implement an encoder-decoder model based on GRU with Pytorch. GRU-ED  
471 contains two layers of GUR for both encoder and decoder; each layer has 128 hidden units.  
472 A fully-connected layer projects the output of the decoder at each time step to a prediction.  
473 We set the batch size to 64, learning rate to 0.001, and the loss function to L1 when training  
474 the model.
- 475 • DSANet: we reuse the code released in the original paper and tune the parameters carefully  
476 for our dataset according to the validation error. We set the CNN filter size to 3, number  
477 of CNN kernels to 64, number of attention blocks to 3, dropout probability to 0.1, and the  
478 learning rate to 0.001.
- 479 • DCRNN: similar to GRU-ED, the DCRNN model also deploys the ecoder-decoder frame-  
480 work for multi-step traffic forecasting. It contains two-layers DCGRU for both encoder and  
481 decoder. We set the number of GRU hidden units to 64, the maximum step of randoms  
482 walks to 3, the initial learning rate to 0.01. We decrease the learning rate tby  $\frac{1}{10}$  every 20  
483 epochs starting from  $10_{th}$  epochs.
- 484 • STGCN: STGCN contains two spatial-temporal convlutional blocks, one temporal convo-  
485 lutional layer and one output layer. Different from the original STGCN, we implement  
486 the output layer to generate prediction for all horizons at one time (instead of one step per  
487 time). Following the practice of STGCN, we set the size of temporal kernel to 2, the order of  
488 Chebyshev polynomials to 1, and the filter number to 64 for both CNN and GCN. Besides,  
489 We set the learning rate to 0.003 for the PeMSD4 dataset and 0.001 for the PeMDS8 dataset.
- 490 • ASTGCN: The original ASTGCN model ensembles three bolocks to process the recent, daily-  
491 periodic, and weekly-periodic segments for capturing multi-scale temporal correlations. We  
492 take its recent component that only uses recent input segments for a fair comparison. For  
493 implementation, we reuse the code and parameters released in the original paper and train  
494 the model with a L1 loss function.
- 495 • STSGCN: We reuse the results reported in the original paper directly for our overall com-  
496 parison as it conducts experiments on the PeMSD4 and PeMSD8 datasets with the same  
497 evaluation metrics.

498 **AGCRN:** Our model stacks two layers AGCRN to capture the node-specific spatial and temporal  
499 dynamics. The output at the last step is used as the representation of the historical traffic series,  
500 which is directly mapped to the predictions for all horizons by linear transformation . For the  
501 hype-parameters, we set the hidden unit to 64 for all the AGCRN cells and the batch size also to 64.  
502 We search the learning rate among {0.0007, 0.001, 0.003, 0.005, 0.009}, the embedding dimension  
503 among {1, 3, 5, 10, 15, 20, 30} for the PeMSD4 dataset and among {1, 2, 3, 5, 8, 10, 15} for the  
504 PeMSD8 dataset. Finally, the learning rate is set to 0.003 for both datasets, and the embedding  
505 dimension is to 10 for the PeMSD4 dataset and 2 for the PeMSD8 dataset. Besides, we choose L1  
506 Loss as the loss function and do not use any non-mentioned optimization tricks such as learning rate  
507 decay, weights decay, or gradient normalization when training our model.

508 For all the deep learning models, we optimize them with the Adam optimizer for 100 epochs and use  
509 an early stop strategy with the patience of 15 by monitoring the loss in the validation set.

#### 510 A.4 Multi-step Prediction on PeMSD8

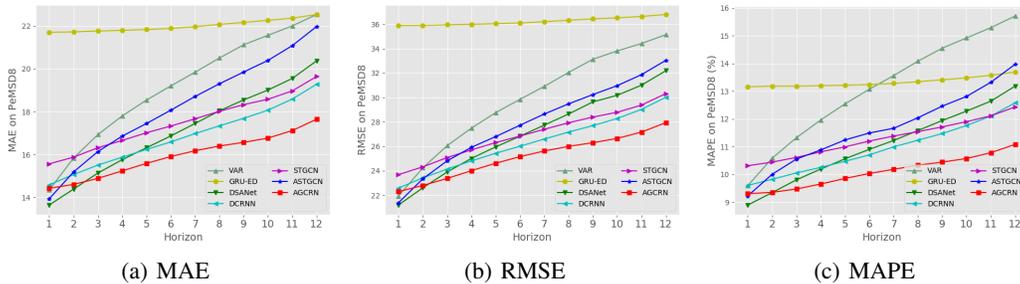
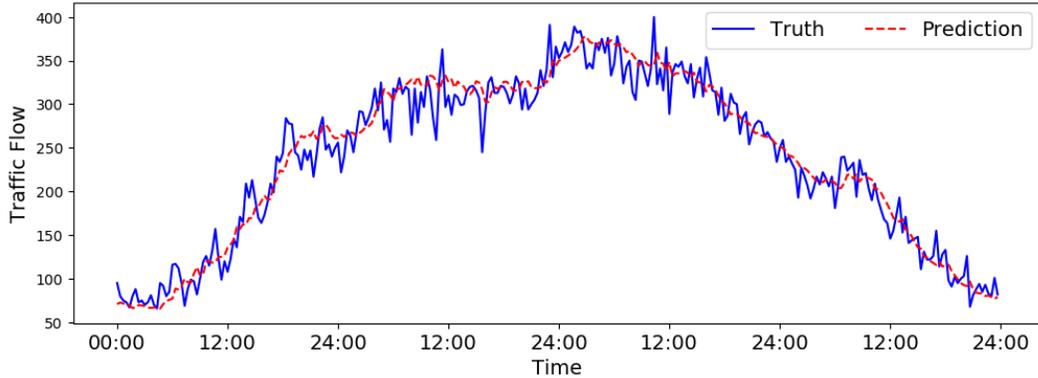


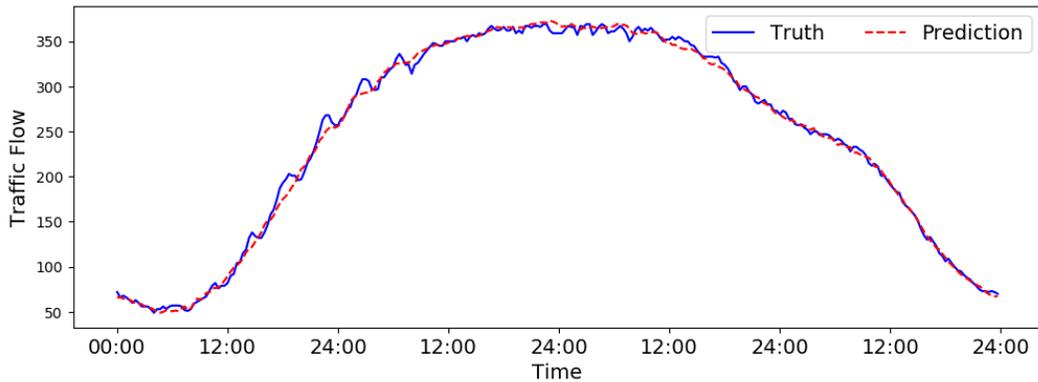
Figure 5: Prediction performance comparison at each horizon on the PeMSD8 dataset.

511 Fig. 5 presents the prediction performance of our AGCRN and baselines at each horizon on the  
512 PeMSD8 dataset. STSGCN is not included because the step-wise results of it are not reported in  
513 [11]. Besides, we omit HA as it's performance is consistent for all 12 horizons. Our AGCRN  
514 model outperforms existing baselines with a significant margin, especially for long-term predictions.  
515 Besides, the performance of AGCRN deteriorates much slower than the other GCN-based models.  
516 The observations are similar on the PeMSD4 dataset.

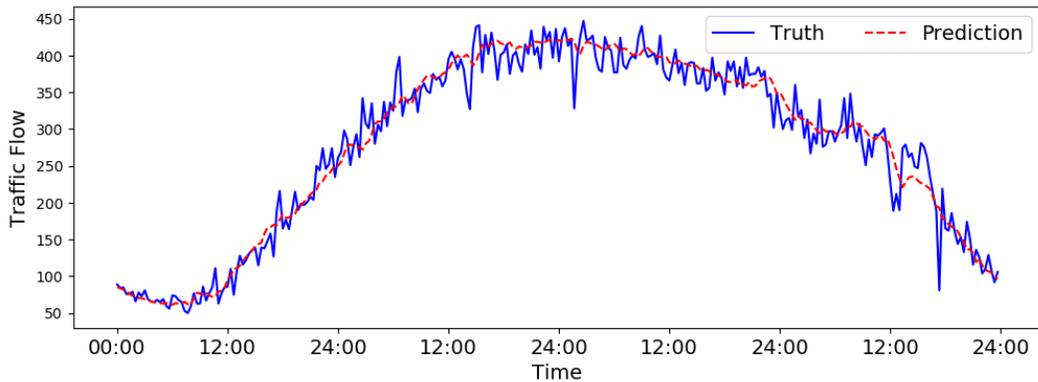
### 517 A.5 Prediction Visualization



(a)

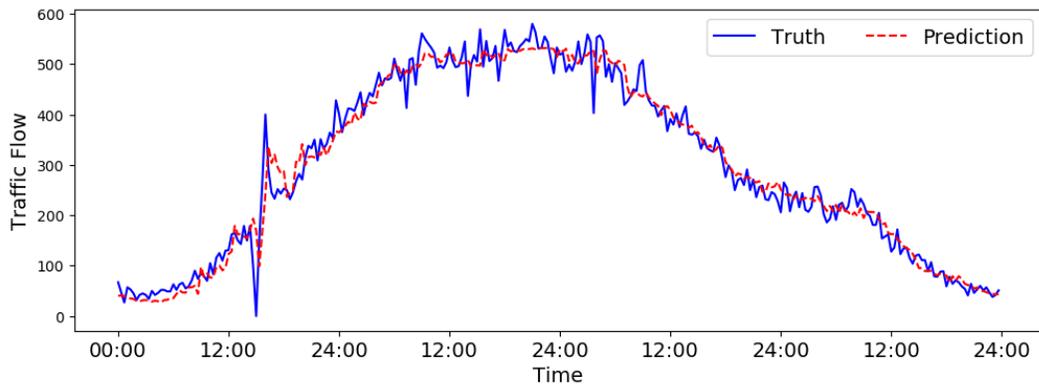


(b)

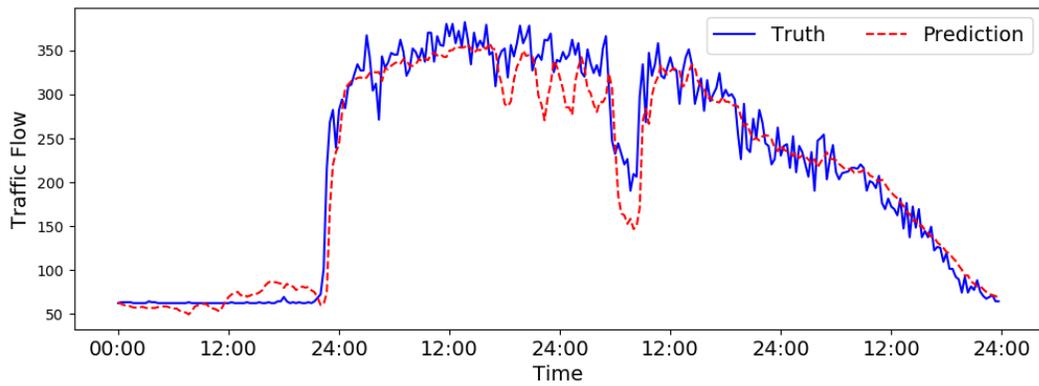


(c)

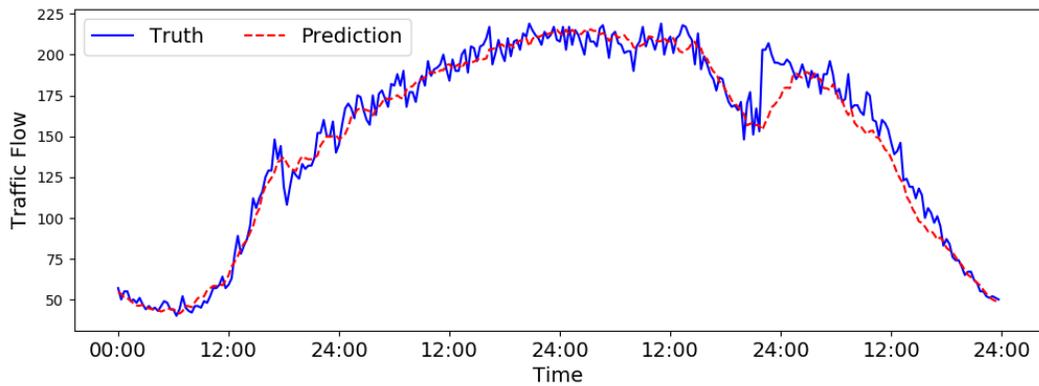
Figure 6: Traffic forecasting visualization.



(a)



(b)



(c)

Figure 7: Traffic forecasting visualization.