Green Light Go: Learning to Recommend Signal Plans under Incidents by Hierarchical Models

Weiran Yao Zhen (Sean) Qian

Mobility Data Analytics Center (MAC), Carnegie Mellon University

Abstract

The main question we address in this paper is how to incorporate domain knowledge from the developed incident timing plans, with historical signal engagement records, to learn to recommend contingency signal patterns for incident-induced congestion. The effectiveness of traffic incident management is often limited by the late response time and excessive workload of traffic operators. This paper proposes a novel decision-making framework that learns to recommend incident plans in advance with the outputs from traffic prediction. Specifically, considering the scarcity of engagement records for incident plans, we propose to decompose the end-to-end recommendation task into two hierarchies – traffic prediction and plan association, and learn their connections through metric learning, which reinforces partial-order preferences observed from engagement records.

Background

Incident Signal Timing Plans

- Planned or unplanned incidents (e.g. weather, accidents) on the network can cause *catastrophic traffic gridlocks*;
- Signals engaged during incident-induced congestion, emergency, severe winter conditions and holidays;
- *Manually entered* into system by traffic operators after validating with real time technology (e.g., cameras, calls, loop detectors)

Incident Management Challenges

- *Late response time:* (1) Overhead from verification of incidents and determination of signal plans; (2) Lack of advance awareness of road conditions;
- *Excessive workload:* Traffic operators need to gather and analyze incident information from multiple directives (cameras and travel information platforms)

Research Question: Can we improve efficiency by automating the data analysis process and learns to recommend signal plans even before report of incidents?

Data Sources

1. INRIX Traffic Speed
2. PennDOT RCRS Incident Report
3. Crowdsourced Waze alerts
4. Weather Underground
5. Incident Signal Plan Engagement Records (very few)

Method

We decompose the recommendation task into two subtask models in hierarchy – traffic predictor and signal plan associator, which combines **Machine Learning + Domain Knowledge**

1. Leverage the imbalance between abundant X (road conditions) and few Y (engagement records)
2. Sequential Deep Learning for Traffic Prediction for each road segment 30 min ahead
3. Encode Domain Knowledge models (VISUM Traffic Simulation) to recommender key matrix
4. Zero-Shot (Metric) Learning for signal plan recommendation. We learn the connections through pairwise learning, which reinforces partial-order preferences observed from historical engagement records.

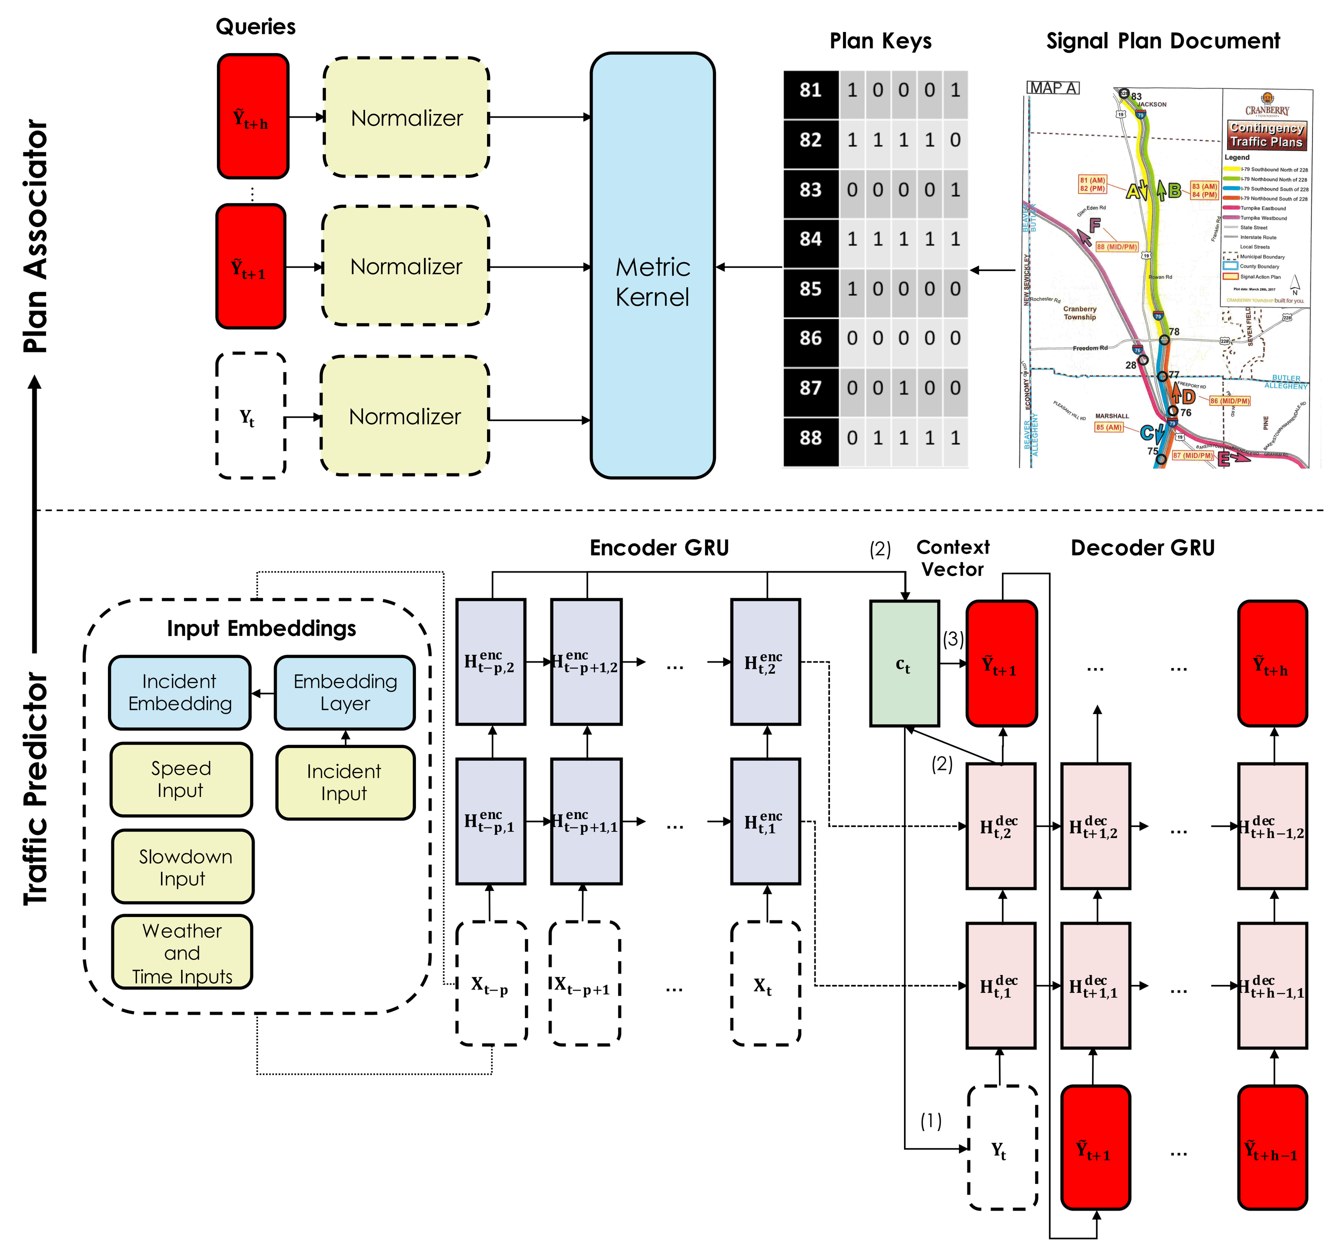

Figure: Incident plan recommendation architecture: traffic predictor is an encoder-decoder recurrent neural network with attention mechanism. We translate the plan triggering conditions into a matrix of plan attributes (keys). The plan associator then generates traffic queries from current and 30-min predicted future speed series and their closeness with plan keys are evaluated with self-defined metrics.

Metric Learning for Plan Associator: RankLR [1]

$$\min_{\mathbf{w}} \frac{1}{P}\sum_{i=1}^{P}\left[y^p \log(\mathbf{w}^T\mathbf{x}_{ij}^p)) + (1-y^p)\log(1-\mathbf{w}^T\mathbf{x}_{ij}^p)\right] + C\|\mathbf{w}\|_1 \quad (1)$$

Results

Traffic Predictor

Table: MAPE prediction error of our encoder-decoder-attention model against other baselines for different prediction horizons on the test samples.

Model	5 min	10 min	15 min	20 min	25 min	30 min
our model	**1.54%**	**1.88%**	**1.81%**	**1.47%**	**1.43%**	**1.91%**
Historical-average	11.73%	11.73%	11.73%	11.73%	11.73%	11.73%
Latest-observation	9.05%	15.80%	18.92%	20.35%	21.22%	21.94%
LASSO	8.05%	13.43%	15.63%	16.02%	17.46%	17.52%
GRU-no-attention	8.13%	11.05%	11.48%	12.30%	13.09%	12.05%

Plan Associator

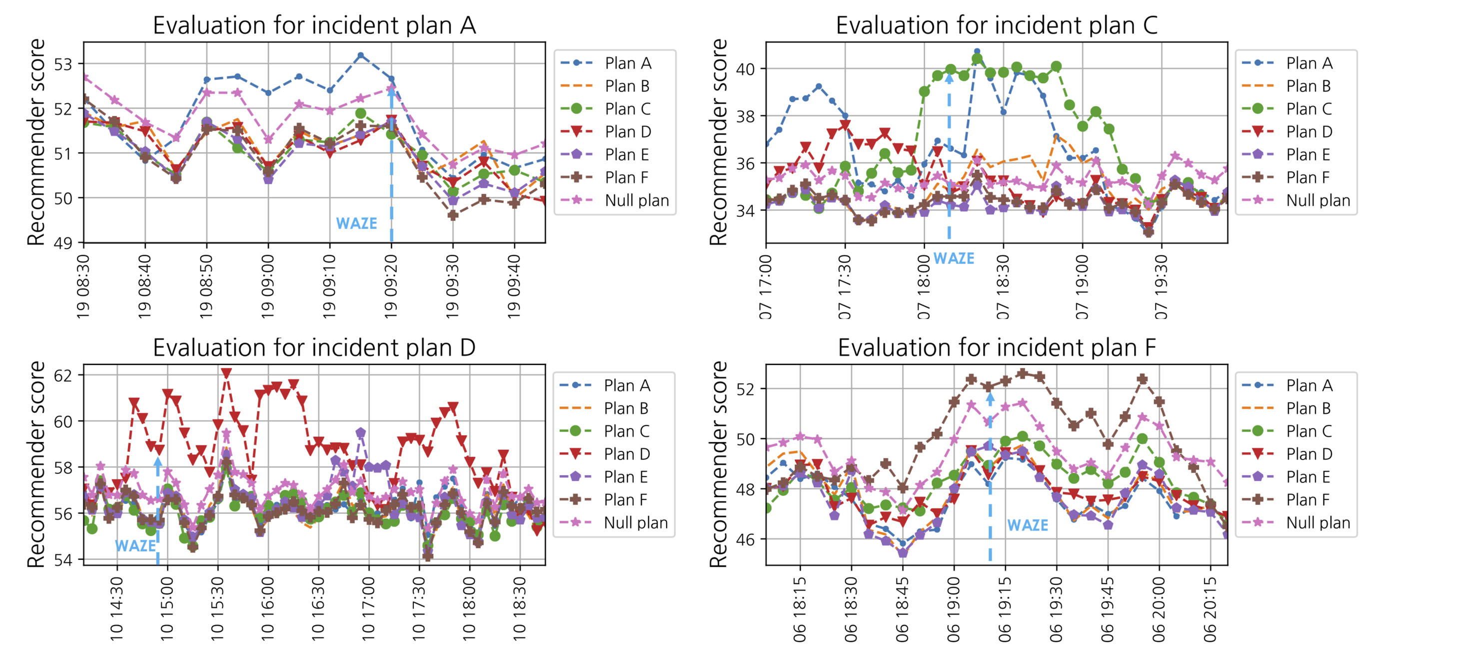

Figure: Visualization of plan recommender zero-shot model performances: our model can trigger recommendations 10 minutes ahead for plan C, 35 minutes ahead for plan F, 30 minutes for plan A and 15 minutes for plan D.

Discussion

Our proposed framework is expected to give traffic operators **significant time** to access the condition and react appropriately. In addition, our recommender has been shown to effectively recommend **unseen plans** in training. This generalization property makes our method an appropriate **initializer for cold-start recommendation** of new incident plans without engagement records, which may be created recently for expansion of signalized intersections.

References

[1] Thorsten Joachims.
Optimizing search engines using clickthrough data.
In Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, pages 133–142. ACM, 2002.

Acknowledgements

Funding provided by