# Methods for Semantic Mining of Spatio-Temporal Data

Steve Gounoue
Data Science and Intelligent Systems Group (DSIS), University of Bonn
Bonn, Germany
Lamarr Institute for Machine Learning and Artificial Intelligence, Bonn, Germany, lamarr-institute.org
steve.gounoue@cs.uni-bonn.de

## ABSTRACT

Network-based spatio-temporal data has various real-world applications, including financial transactions, transportation, traffic management, and social media analysis. In this context, adequate representation learning is crucial to empower various graph-based downstream applications, including fraud detection, transportation demand prediction, traffic flow estimation and social network interactions. Graph representation learning is not trivial and requires encapsulating spatio-temporal knowledge from the data into a numerical vector, potentially involving a loss of relevant information. Current representation learning for spatio-temporal data typically focuses on implicitly encoding spatio-temporal patterns by leveraging complex architectures. However, such architectures can fail to extract relevant spatio-temporal patterns due to data sparsity or pattern-architectural mismatches. To mitigate the loss of relevant information during representation learning, we propose methods for semantic mining of spatio-temporal data. Semantic mining involves discovering and extracting hidden patterns and explicitly including such patterns in the data to support effective representation learning on graphs.

## KEYWORDS

Spatio-temporal data, Semantic mining, Representation learning

**ACM Reference Format:**

Steve Gounoue. 2024. Methods for Semantic Mining of Spatio-Temporal Data. In *Proceedings of ACM Conference (KDD '24)*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

## 1 INTRODUCTION

Understanding spatio-temporal patterns and interactions is critical for various graph-related activities such as trade with supply chain management [20] or traffic with traffic demand forecasting [16]. Also, with infrastructural growth, data networks become more complex, making spatio-temporal data analysis even more challenging [18]. Moreover, the high cost of data collection reduces the data availability and increases sparsity.

However, state-of-the-art deep learning models employed to analyse spatio-temporal data leverage implicit representation learning mechanisms including complex modules such as graph neural networks (GNNs) [11], recurrent neural networks (RNNs) [12], convolutional neural networks (CNNs) [8], or their combinations [1]. Yet, these modules have several limitations: First, they require many training samples, which leads to efficiency problems in large networks [10]. Second, they do not explicitly capture data patterns and the implicitly learned interactions, leading to suboptimal results on downstream tasks. Finally, they lack interpretability although demand for interpretable machine learning is rising [19].

In this thesis, we tackle the problem of leveraging semantic mining of spatio-temporal data to augment the data and reduce the loss of key information while performing representation learning. Moreover, semantic mining can enhance the interpretability of deep learning and the methods employed are illustrated in Figure 1. These methods include 1) Correlation extraction, involving the identification of spatio-temporal correlations and incorporation of context information regarding individual graph nodes, and 2) Concept learning, comprising concept definitions based on the domain knowledge and classifiers to align the data points and patterns with the corresponding concepts. Finally, the goal is to empower various downstream applications that can benefit from the results by producing better-performing, more explainable models.

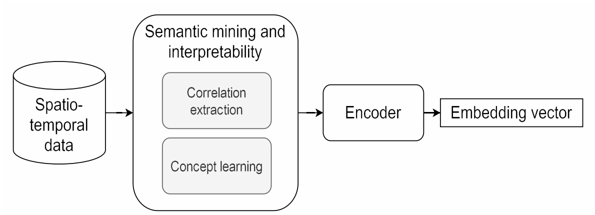


**Figure 1: Mining methods**

## 2 DISSERTATION PLAN

There is a need to capture spatio-temporal interactions efficiently in graph-based data and to achieve better interpretability in deep learning, as presented in Section 1. Thus in this thesis, we propose the *semantic mining* approach constructed around two main semantic mining methods for spatio-temporal data: *correlation extraction* and *concept learning*. In this Section, we discuss the motivation and describe the associated problem statements for each method.

### 2.1 Correlation Extraction

The main motivation is to facilitate more in-depth graph-based data understanding by deep learning methods, providing them with explicit data patterns. Pattern recognition is the principal challenge considering the conjunction of spatial and temporal relationships.


*KDD '24, August 2024, Barcelona, Spain*

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

To address this challenge, we investigate the research question RQ1: *How to leverage spatio-temporal correlation to enhance speed prediction in traffic networks?*, where the focus is summarising the spatio-temporal relationships in the traffic network. Therefore, we leverage a statistical method, mainly the correlation between graph nodes, considering space and time dimensions. Then, the resulting correlation is integrated into traffic speed prediction models.

## 2.2 Concept Learning

In this semantic mining method, the motivation is to associate prior domain knowledge with deep learning predictive methods. Specifically, we define concepts based on domain knowledge and incorporate classifiers to align the data points and patterns with the corresponding concepts. However, the main challenges here are the high dimensionality of the data and the integration of spatial and temporal features in the concepts.

Thus, we investigate the research question RQ2: *How to learn the patterns that reflect semantic concepts in the data?* The idea is at the earlier stage to define concepts from domain knowledge, and then, learn these concepts by leveraging an appropriate classifier. Considering that this classifier is trained to predict the corresponding concept of each data point, it would allow a flexible inference task in case different models exhibit variable performance depending on the concepts observed in the data. Examples of concepts for the traffic speed prediction task are *rush hour* and *traffic jam (when many cars are delayed in the traffic*).

## 2.3 Semantic Mining and Interpretability

After identifying patterns through semantic mining methods, we are interested in interpreting the results of downstream tasks in light of those patterns. The main challenge is the non-linear relation between the results and the data patterns.

Hence, in this phase, we face the research question RQ3: *What impact do patterns in deep learning models have?* Consequently, We develop techniques leveraging spatio-temporal patterns to understand deep learning models better, without damaging the models' performance. This could be done leveraging a mechanism where each pattern discovered in the data has an associated weight learned as the model is optimised. This weight quantifies the impact of the pattern for the overall result [15]. .

# 3 PRELIMINARY RESULTS

The preliminary results mainly concern the RQ1 in Section 2.1 and RQ2 in Section 2.2. About RQ1 in correlation extraction, we performed an enrichment mechanism to enhance traffic speed prediction and published the first version of the SCANNER approach [6]. In SCANNER, we demonstrate the importance of enriching the data by leveraging temporal correlation between nodes of the traffic network. Moreover, We are working on a second version of the SCANNER approach comprising a neighbourhood-based data enrichment module (NDE). With the NDE module, the enrichment process also considers the spatial distance between nodes in the traffic network and a longer temporal range for the temporal correlations. The initial results of the second version of the SCANNER approach in Table 1 demonstrate that the NDE module benefits different model architectures in traffic speed prediction. Specifically, the results present two modalities for each model: without data enrichment (w/o NDE) and with data enrichment (+ NDE), and show the improvement percentage. The dataset used for the experiments is *metr-la* [9], and the metrics are mean absolute error (MAE) and the root-mean-square error (RMSE).

**Table 1: Traffic speed models with data enrichment (+ NDE) and without data enrichment (w/o NDE).**

| Metrics | Models | w/o NDE | + NDE | Improv. % |
|---|---|---|---|---|
| **MAE** | GRU [3] | 7.46 | 6.96 | 6.69 |
| | STGAT [7] | 4.16 | 3.81 | 8.23 |
| | Wavenet [14] | 3.91 | 3.79 | 3.01 |
| | ST-Norm [4] | 3.53 | 3.35 | 5.29 |
| **RMSE** | GRU [3] | 12.58 | 12.02 | 4.44 |
| | STGAT [7] | 8.13 | 7.45 | 8.27 |
| | Wavenet [14] | 7.88 | 7.69 | 2.37 |
| | ST-Norm [4] | 7.25 | 6.84 | 5.65 |

Concerning RQ2 in concept learning, we propose an ensemble method for traffic speed prediction. The ensemble approach integrates different models, based on the traffic concepts, to benefit from the diversity of models by making them complementary. This is motivated by the fact that the prediction performance of state-of-the-art approaches for traffic speed prediction might vary with the observed traffic concepts. Concretely, we first define three concepts: 1) reoccurring traffic congestion (RTC), 2) traffic speed transition (TST) and 3) other traffic speed (OTS). Then, we design a deep learning classifier to assign the right concept to each traffic data instance. Our preliminary results in Table 2 confirm the importance of considering concepts, given that the best model (**bold**) as well as the worst model (underline) change depending on the concept.

**Table 2: Prediction performance for several traffic-speed concepts using Hanover (Germany) speed data. MAE metric.**

| **Models / Concepts** | **OTS** | **RTC** | **TST** |
|---|---|---|---|
| HA | 4.80 | 13.25 | 7.39 |
| MA | 3.60 | **4.07** | **6.55** |
| GTS [13] | 8.38 | 11.89 | 16.15 |
| G. Wavenet [17] | 2.85 | 4.54 | 14.55 |
| ST-Norm [4] | **2.61** | 4.43 | 13.32 |

# 4 NEXT STEPS

In the next steps, we will proceed with the concept learning method finalising our proposal for the research question RQ2. In this regard, we are currently investigating systematic concept definitions and a dedicated attention mechanism to weigh the contribution of each model. After that, we plan to complete the architecture and propose a solution for the research question RQ3 following state-of-the-art studies [2, 5]. We want to combine global and local patterns by applying attention layers at different levels of the architecture.

***Acknowledgements:*** This work was partially funded by the Federal Ministry for Economic Affairs and Climate Action (BMWK), Germany ("ATTENTION!", 01MJ22012C).

# REFERENCES


[1] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2020. Adaptive Graph Convolutional Recurrent Network for Traffic Forecasting. In *Proceedings of the Annual Conference on Neural Information Processing Systems, NeurIPS, December 6-12, virtual.* https://proceedings.neurips.cc/paper/2020/hash/ce1aad92b939420fc17005e5461e6f48-Abstract.html

[2] Lianliang Chen and Ziyang Shan. 2021. Forecasting PM2.5 and Tracking Spatial Influence Patterns of Traffic Using Interpretable Deep Learning. In *Proceedings of the 4th International Conference on Control and Computer Vision, ICCCV, Macau, SAR, China, August 13 - 15.* 177–182. https://doi.org/10.1145/3484274.3484302

[3] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL.* 1724–1734. https://doi.org/10.3115/V1/D14-1179

[4] Jinliang Deng, Xiusi Chen, Renhe Jiang, Xuan Song, and Ivor W. Tsang. 2021. ST-Norm: Spatial and Temporal Normalization for Multi-variate Time Series Forecasting. In *Proceedings of the 27th Conference on Knowledge Discovery and Data Mining, KDD, Virtual Event, Singapore, August 14-18.* 269–278. https://doi.org/10.1145/3447548.3467330

[5] Malikeh Ehghaghi, Pei-Yuan Zhou, Wendy Yusi Cheng, Sahar Rajabi, Chih-Hao Kuo, and En-Shiun Annie Lee. 2023. Interpretable Disease Prediction from Clinical Text by Leveraging Pattern Disentanglement. In *Proceedings of the International Conference on Biomedical and Health Informatics, BHI, Pittsburgh, PA, USA, October 15-18.* 1–7. https://doi.org/10.1109/BHI58575.2023.10313408

[6] Steve Gounoue, Ran Yu, and Elena Demidova. 2023. SCANNER: A Spatio-temporal Correlation and Neighborhood-based Feature Enrichment for Traffic Prediction. In *Proceedings of the 31st International Conference on Advances in Geographic Information Systems, SIGSPATIAL, Hamburg, Germany, November 13-16.* 103:1–103:4. https://doi.org/10.1145/3589132.3625653

[7] Xiangyuan Kong, Weiwei Xing, Xiang Wei, Peng Bao, Jian Zhang, and Wei Lu. 2020. STGAT: Spatial-Temporal Graph Attention Networks for Traffic Flow Forecasting. *IEEE Access* 8 (2020), 134363–134372. https://doi.org/10.1109/ACCESS.2020.3011186

[8] Neslihan Kose, Okan Köpüklü, Alexander Unnervik, and Gerhard Rigoll. 2019. Real-Time Driver State Monitoring Using a CNN Based Spatio-Temporal Approach. In *Proceedings of the Intelligent Transportation Systems Conference, ITSC, Auckland, New Zealand, October 27-30.* 3236–3242. https://doi.org/10.1109/ITSC.2019.8917460

[9] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *Proceedings of the 6th International Conference on Learning Representations, ICLR, Vancouver, BC, Canada, April 30 - May 3.* https://openreview.net/forum?id=SJiHXGWAZ

[10] Bin Lu, Xiaoying Gan, Weinan Zhang, Huaxiu Yao, Luoyi Fu, and Xinbing Wang. 2022. Spatio-Temporal Graph Few-Shot Learning with Cross-City Knowledge Transfer. In *Proceedings of the 28th Conference on Knowledge Discovery and Data Mining, SIGKDD, Washington, DC, USA, August 14 - 18.* 1162–1172. https://doi.org/10.1145/3534678.3539281

[11] Maedeh Nasri, Zhizhou Fang, Mitra Baratchi, Gwenn Englebienne, Shenghui Wang, Alexander Koutamanis, and Carolien Rieffe. 2023. A GNN-Based Architecture for Group Detection from Spatio-Temporal Trajectory Data. In *Proceedings of 21st International Symposium on Intelligent Data Analysis, IDA, Louvain-la-Neuve, Belgium, April 12-14 (Lecture Notes in Computer Science, Vol. 13876).* 327–339. https://doi.org/10.1007/978-3-031-30047-9_26

[12] Soumyasundar Pal, Liheng Ma, Yingxue Zhang, and Mark Coates. 2021. RNN with Particle Flow for Probabilistic Spatio-temporal Forecasting. In *Proceedings of the 38th International Conference on Machine Learning, ICML, 18-24 July, Virtual Event (Proceedings of Machine Learning Research, Vol. 139).* 8336–8348. http://proceedings.mlr.press/v139/pal21b.html

[13] Chao Shang, Jie Chen, and Jinbo Bi. 2021. Discrete Graph Structure Learning for Forecasting Multiple Time Series. In *Proceedings of the 9th International Conference on Learning Representations, ICLR, Virtual Event, Austria, May 3-7.* https://openreview.net/forum?id=WEHSlH5mOk

[14] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. 2016. WaveNet: A Generative Model for Raw Audio. In *The 9th ISCA Speech Synthesis Workshop, Sunnyvale, CA, USA, 13-15 September 2016.* 125. http://www.isca-speech.org/archive/SSW_2016/abstracts/ssw9_DS-4_van_den_Oord.html

[15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Proceedings of the Annual Conference on Neural Information Processing Systems, NeurIPS, December 4-9, Long Beach, CA, USA.* 5998–6008. https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

[16] Yiling Wu, Xinfeng Zhang, and Yaowei Wang. 2023. Adaptive Graph Neural Diffusion for Traffic Demand Forecasting. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM, Birmingham, United Kingdom, October 21-25.* 4350–4354. https://doi.org/10.1145/3583780.3615153

[17] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI, Macao, China, August 10-16.* 1907–1913. https://doi.org/10.24963/IJCAI.2019/264

[18] Mengran Xu, Demin Li, Guanglin Zhang, Mengqi Cao, and Shuya Liao. 2018. A Path Planning Approach with Maximum Traffic Flow and Minimum Breakdown Probability in Complex Road Network. In *Proceedings of the 7th International Conference, CSoNet, Shanghai, China, December 18-20 (Lecture Notes in Computer Science, Vol. 11280).* 92–102. https://doi.org/10.1007/978-3-030-04648-4_8

[19] Han Yuan, Jin Wee Lee, Mingxuan Liu, Siqi Li, Chenglin Niu, Jun Wen, and Feng Xie. 2023. Interpretable Machine Learning-Based Risk Scoring with Individual and Ensemble Model Selection for Clinical Decision Making. In *The First Tiny Papers Track at ICLR 2023, Tiny Papers @ ICLR 2023, Kigali, Rwanda, May 5, 2023.* https://openreview.net/pdf?id=RNlfw6KXJey

[20] Lingmian Zhang. 2023. A Data-driven Optimization Method for Supply Chain Management Accounting Demand Forecasting. In *Proceedings of the 5th International Conference on Big Data Engineering, BDE, Zhuhai, China, November 17-20.* 35–41. https://doi.org/10.1145/3640872.3640878