

Data Driven Decision Making with Time Series and Spatio-temporal Data

Yang, Bin; Liang, Yuxuan; Guo, Chenjuan; Jensen, Christian S.

DOI (link to publication from Publisher):
[10.48550/arXiv.2503.08473](https://doi.org/10.48550/arXiv.2503.08473)

Creative Commons License
CC BY 4.0

Publication date:
2025

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Yang, B., Liang, Y., Guo, C., & Jensen, C. S. (2025). *Data Driven Decision Making with Time Series and Spatio-temporal Data*. arXiv. <https://doi.org/10.48550/arXiv.2503.08473>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Data Driven Decision Making with Time Series and Spatio-temporal Data

Bin Yang

East China Normal University
Shanghai, China
byang@dase.ecnu.edu.cn

Yuxuan Liang

Hong Kong University of
Science and Technology (Guangzhou)
Guangzhou, China
yuxliang@outlook.com

Chenjuan Guo

East China Normal University
Shanghai, China
cjguo@dase.ecnu.edu.cn

Christian S. Jensen

Aalborg University
Aalborg, Denmark
csj@cs.aau.dk

Abstract—Time series data captures properties that change over time. Such data occurs widely, ranging from the scientific and medical domains to the industrial and environmental domains. When the properties in time series exhibit spatial variations, we often call the data spatio-temporal. As part of the continued digitalization of processes throughout society, increasingly large volumes of time series and spatio-temporal data are available. In this tutorial, we focus on data-driven decision making with such data, e.g., enabling greener and more efficient transportation based on traffic time series forecasting.

The tutorial adopts the holistic paradigm of “data-governance-analytics-decision.” We first introduce the data foundation of time series and spatio-temporal data, which is often heterogeneous. Next, we discuss data governance methods that aim to improve data quality. We then cover data analytics, focusing on the “AGREE” principles: Automation, Generalization, Robustness, Explainability, and Efficiency. We finally cover data-driven decision making strategies and briefly discuss promising research directions. We hope that the tutorial will serve as a primary resource for researchers and practitioners who are interested in value creation from time series and spatio-temporal data.

Index Terms—Time series, Spatio-temporal data, Data governance, Data analytics, Data-driven decision making

I. INTRODUCTION

As part of the continued digitalization of processes across society, increasingly large volumes of time series and spatio-temporal data are available. Managing and analyzing heterogeneous, multi-modal time-series and spatio-temporal data enables value creation by supporting data-driven decision-making across domains, such as smart cities, intelligent transportation, smart emergency response, digital energy, AIops, smart agriculture, ocean, and mining [1]–[3].

As an example, path planning is a typical decision-making task in smart cities and intelligent transportation. Consider a scenario in which an autonomous taxi needs to take passengers to an airport. The taxi must select the most “optimal” route from a set of alternatives [4], [5]. Data-driven decision-making facilitates a rational selection process: first, relevant *multi-modal* traffic time series data, e.g., from road-side sensors and in-vehicle Global Navigation Satellite System capabilities, is collected. The quality of the collected raw data may be insufficient, thus requiring *data governance* for enhancement. Subsequently, *data analytics*, such as forecasting, are conducted on the data to predict future traffic conditions. Finally,




Fig. 1: Paradigm underlying Data Driven Decision Making—“Data-Governance-Analytics-Decision”

the optimal route is chosen based on the predicted future traffic using *data-driven decision making*, e.g., favoring the route with the highest probability of an on-time arrival. As another example, consider cloud computing where resource scaling decisions must be made frequently. By gathering and managing historical cloud resource demand data, future demands can be predicted, particularly in the event of unexpected surges, allowing for timely resource auto-scaling to maintain service quality while minimizing energy consumption and carbon emissions [6].

We integrate the aspects of value creation from time series and spatio-temporal data into a data driven decision making paradigm with ‘**Data-Governance-Analytics-Decision**’, as depicted in Figure 1.

Tutorial overview. The tutorial is 1.5 hours long and is presented in a traditional lecture style. We summarize data foundations for multi-modal time series and spatio-temporal data, introduce data governance methods aiming at improving data quality, and cover data analytics methods, focusing on five characteristics: automation, generalization, robustness, explainability, and efficiency, which we summarize as the “AGREE” principles. We then cover various data-driven

decision-making strategies, and finally, we briefly discuss research directions. The tutorial is organized as follows: (1) **Introduction and Data Foundation (5 mins)**. (2) **Data Governance (20 mins)**. (3) **Data Analytics (40 mins)**. (4) **Data-driven Decision Making (15 mins)**. (5) **Conclusion and Future Work (10 mins)**.

Prior Offerings: This tutorial is newly developed, which has not been presented in other venues.

Target Audience: The tutorial targets researchers specializing in database systems, data mining, machine learning, and data science, particularly those focused on time series and spatio-temporal data analysis. It is also relevant for practitioners engaged in real-world applications such as forecasting, classification, and anomaly detection related to smart cities, AIOps, Internet of Things, etc. The tutorial is data-centric, covering data governance, data analytics, and data-driven decision making, thus well aligning with ICDE’s core topics. The prerequisites include basic knowledge on databases, data mining, and deep learning. Expected benefits include: (1) Gain a deeper understanding of the “data-governance-analytics-decision” paradigm, covering advanced topics and recent progress on data governance, data analytics, and data-driven decision making with time series and spatio-temporal data. (2) Interact with experts in this area and build connections that may lead to collaborations, mentorship, or career opportunities. (3) Understand emerging trends and technologies on time series and spatio-temporal data analytics.

Related Tutorials: We summarize related tutorials at recent top-tier venues. A tutorial at KDD’22 [7] focused on robust time series analysis. A tutorial at VLDB’23 [8] provided a concise selection of methods that can be integrated to effectively address a wide range of time series challenges. A tutorial at KDD’23 [9] focused on utilizing causal theory for temporal data mining. A tutorial at AAAI’24 explore out-of-distribution generalization in time series. A tutorial at KDD’21 [10] provided an overview of physics-guided spatio-temporal data mining. The above tutorials focus on specific technical aspects in machine learning and data mining on time series and spatio-temporal data. In contrast, the present tutorial is data-centric and is based on the new holistic paradigm of “data-governance-analytics-decision,” encompassing data foundations, data governance methods, data analytics, and strategies for data-driven decision making.

II. TUTORIAL OUTLINE

A. Data Foundations

Time series data captures dynamic properties over time, such as temperatures throughout the year. Spatio-temporal data combines time and location to characterize evolving events across temporal and spatial dimensions. Such data is often collected by different sensors, thus being multi-modal.

Definition 1 (Time Series). A time series captures properties over time and is given as $\mathbf{X} = \langle \mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_M \rangle$, where $\mathbf{s}_i \in \mathbb{R}^C$, $1 \leq i \leq M$, is a C -dimensional vector, representing C properties observed at the i -th timestamp.

Definition 2 (Correlated Time Series). A correlated time series is a set of N interconnected time series, given as $\mathbf{T} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N\}$. These time series often exhibit correlations due to the spatial arrangement of the corresponding sensors, and the correlations are often modeled using graphs.

Definition 3 (Trajectories). Trajectories are sequences of (location, time) pairs that capture the positions of moving objects (e.g., humans, animals, and vehicles) over time.

Definition 4 (Image Sequences). An image sequence is given as $\mathbf{V} = \langle \mathbf{I}_1, \mathbf{I}_2, \dots, \mathbf{I}_T \rangle$, where each $\mathbf{I}_i \in \mathbb{R}^{N \times M \times C}$, $1 \leq i \leq T$, is an image that captures a state at the i -th timestamp. Each image comprises $N \times M$ pixels, representing spatial regions. Each pixel includes C observed properties.

B. Data Governance

As raw sensor data suffers from data quality issues, we cover effective data governance approaches that aim to enhance data quality.

Missing Value Imputation: Sensor data collection is prone to missing values due to sensor malfunctions or network disruptions. Within time series and spatio-temporal data management, missing value imputation is categorized into *spatial completion*, *temporal completion*, and *spatio-temporal completion* of missing data. Spatially, weighted graphs are often used to model sensor relationships, and spatially missing values are transformed into graph edge weight completion tasks. We consider graph-based semi-supervised learning [11] and graph neural network autoencoders [12] for this purpose. Temporally, time series imputation and backcast techniques, such as recurrent neural network applications [13], address missing temporal values. Spatio-temporally, integrating graph neural networks with recurrent neural networks [14] enables completion of missing values in Origin-Destination (OD) Matrix over time. In addition, spatio-temporally aware LLMs are employed to complete global significant wave heights using sparse buoy data [2].

Uncertainty Quantification: This process aims to capture uncertainty related to time and space using probability distributions. Histograms or Gaussian mixture models are often used to represent distributions due to their capacity to approximate distributions without assumptions on the type of distribution. We cover a *dynamic, uncertainty graph model* in the context of intelligent transportation. Traffic flow is inherently uncertain and dynamic, influenced by factors such as accidents and traffic volume. This can be modeled using (I, D) pairs, where travel speed follows distribution D over time interval I . We discuss two paradigms: the *edge-centric paradigm* [15] and *path-centric paradigm* [4]. The former assigns distributions to edges, treating the edge distributions as independent, while the latter captures the distribution correlations along paths. Uncertainty quantification enhances data reliability, providing a robust foundation for further analytics and decision-making, e.g., decision making under uncertainty [6].

Multi-modal Fusion: Time series and spatio-temporal data exhibit multi-source, heterogeneous, multi-modal characteris-

tics, with significant correlations across time and space. Data fusion across modalities improves data quality by consolidating data, laying a robust foundation for further analysis and decision-making. Recent multi-modal fusion approaches can be summarized into four streams: *feature-based*, *alignment-based*, *contrast-based*, and *generation-based* methods [16]. A prime instance of alignment-based fusion is Map Matching [17], which aligns vehicle trajectory data with corresponding map data. This process can eliminate trajectory noise as well as compensate for gaps in trajectories, thereby enhancing data quality and analytical and decision-making processes. Another example is feature-based fusion of multi-modal data, such as historical traffic, weather data, and point of interest, for traffic forecasting [18], [19]. Advances in representation learning and foundational models enable the capture of multi-modal spatio-temporal data as vectors or embeddings [20]–[22]. For example, MM-Path enhances path representations by aligning embeddings of the road network and satellite image of paths [23]. By constructing temporal and spatial indices, we can improve the efficiency of analytics on diverse multi-modal spatio-temporal datasets.

C. Data Analytics

To facilitate data-driven decisions, enabling diverse analytics, e.g., forecasting, anomaly detection, and classification, on time-series and spatio-temporal data are essential. We do not cover specific analytics techniques; instead, we outline five desirable characteristics, which we call the **AGREE** principles—Automation, Generalization, Robustness, Explainability, and Efficiency.

Automation: Traditional time series and spatio-temporal analysis methods are often crafted by human experts, which is resource-intensive. Automated methods offer a practical alternative: a search space is created for a specific task, upon which algorithms can identify and automatically assemble the most fitting models for the specific task and dataset. For instance, automated correlated time series forecasting automates the design of network architectures and hyperparameter selection [24]–[28]. Such automation not only enhances accuracy across diverse datasets and forecasting settings but also facilitates the discovery of optimal models that adhere to additional constraints, e.g., model sizes.

Generalization: Methods using deep learning are capable of superior performance in time-series and spatio-temporal analytics, but they depend on high-quality representations. Representation learning, which requires extensive labeled data, poses challenges in data acquisition and does not generalize across tasks [29]. We consider general representation learning frameworks, pre-trained on abundant unlabeled data, that can be fine-tuned with limited labeled data for different spatio-temporal tasks. These include unsupervised contrastive learning [30], weakly-supervised contrastive learning [31], masked autoencoders with relational inference [32], and cross-modal contrastive learning [23]. For time series tasks, we propose general models for forecasting [33] and anomaly detection [34] to eliminate the need for target-domain-specific training data.

These frameworks demonstrate generality and reduce the dependence on large amounts of labeled data. With the emergence of Artificial General Intelligence, we also consider how to leverage Large Language Models (LLMs) to enhance the zero-shot adaptability for time series forecasting [20], [35] and spatio-temporal data mining [21], [22].

Robustness: Robustness in time-series and spatio-temporal analysis encompasses handling *anomalous data*, *data imbalance*, *distribution shifts*, and *multi-scale adaptation*. Traditional unsupervised anomaly detection algorithms assume implicitly that training occurs on fully-clean data, which is rarely available in practice. We thus cover robust anomaly detection [36]–[39] that operates effectively on noisy training data. Considering data imbalance, we cover an adversarial domain adaptation method [40] that models unbalanced data patterns despite data size discrepancies. Addressing distribution shifts, e.g., due to changes in the physical world such as new roads, we cover continuous learning [41], [42] tailored to spatio-temporal data dynamics. In addition, we cover means to model heterogeneous temporal patterns caused by temporal distribution shifts, e.g., in streaming data or non-stationary time series [43]. Lastly, we cover multi-scale modeling using pathways [44] and ensemble learning strategies [45], [46] that enhance robustness by adaptively selecting and combining multiple scales for analytics.

Explainability: Effective decision-making requires both accuracy and explainability in data analytics. Although deep learning exhibits state-of-the-art accuracy at time series and spatio-temporal analyses, its “black box” nature hampers explainability. Considering this, we focus on two important aspects: how to quantify the explainability of different methods and how to improve the explainability of deep learning methods. We present a posthoc explainability metric [37] for the assessment of autoencoder-based anomaly detection models. Furthermore, we consider how to leverage neural networks for feature extraction and integrate extracted features with interpretable models to augment overall prediction clarity [47]. Additionally, we consider how to improve the explainability by tracking temporal associations among time series [48], [49] and employing causal models to predict future correlations [50]. We integrate physical equations into data-driven models to improve the explainability of air quality prediction [51]. Lastly, we cover how LLMs offer remarkable explainability in the realm of spatio-temporal data mining [20], [21], enabling the injection of domain knowledge into language forms and facilitating the understanding of patterns within time series and spatio-temporal data.

Efficiency: Due to increasing privacy concerns, analytics are increasingly performed on personal edge devices rather than having to upload personal data to potentially untrustworthy servers. Contemporary analysis methods prioritize accuracy, resulting in the reliance on computationally intensive model architectures that are unsuitable for resource-constrained edge devices. Such methods also do not align with green computing principles. This calls for data analytics with high resource efficiency. To cover this aspect, we present LightPath [32],

a knowledge distillation-based framework for efficient spatio-temporal data representation, and LightTS [52], an automated, lightweight temporal analysis framework capable of adapting quantization levels to align with memory limitations of edge devices. QCore [53] enables continuous calibrations for quantized models, addressing both resource efficiency and robustness on distribution shifts. EnhancingDFQ [54] improves the accuracy of data-free model quantization by increasing sample-data diversity. TimeDC studies a time series dataset condensation method that is able to compress large time series into a smaller counterpart while maintaining key properties [55]. This facilitates data analytics in resource constrained environments.

Finally, as a wide variety of data analytics approaches have been proposed, it is essential to be able to compare such approaches empirically in a comprehensive and fair manner, thus calling for benchmarking. We will cover recent benchmarking on time series and spatio-temporal data analytics, focusing on forecasting and anomaly detection tasks implemented through end-to-end models [6] and foundation models [56].

D. Data-Driven Decision Making

Decision making involves selecting among options. In our context, we consider temporal and spatial aspects while making decisions. Different scenarios require distinct strategies. For scheduling, it is about optimizing resource use to meet demands efficiently and sustainably. Eco-driving focuses on reducing emissions through informed driving practices. Predictive maintenance aims to preempt equipment failure to ensure uninterrupted operation. Instead of delving into the specifics of particular decision making scenarios, we provide an overview of data-driven strategies that are applicable to different decision making tasks.

Decision Making under Uncertainty: Decision making based on time-series and spatio-temporal data is often subject to uncertainty. This uncertainty can be quantified through data governance. Moreover, spatio-temporal analysis methods, such as predictive models, inherently capture uncertainty, typically using confidence intervals and probability distributions. At the same time, different stakeholders may have different risk preferences, including being risk-loving, risk-averse, and risk-neutral. By employing utility functions, we can encode different risk preferences, and then use expected utility to identify the most favorable options. We cover a novel pruning approach grounded in stochastic dominance [57]–[59], enabling rapid identification of optimal choices across utility functions that encode different risk profiles.

Multi-objective Decision Making: Decision making often involves multiple criteria. For instance, cloud resource scheduling balances operational costs with quality of service (QoS), while fleet path planning considers travel distance and time, fuel consumption, and emissions. Multi-objective decision-making can be categorized into two classes: the first employs Pareto optimality to identify a set of non-dominated options, where no single option outperforms any other option across all objectives [15]; the second consolidates multiple

objectives into a single unified objective via a preference function, facilitating the selection of a single best option [60].

Personalized Decision Making: This approach tailors decisions to individual preferences, which may include personalized risk profiles or preferences on multi-objective trade-offs. The challenge lies in selecting the most suitable preference for a given context—encompassing temporal, spatial, and other contexts like weather or events [29], [61]. Efficiently accommodating large numbers of preferences is critical.

Learning-based Decision Making: These strategies simulate the decision patterns of experts, enabling non-experts to replicate expert-level decisions. For instance, we mimic expert taxi drivers’ routing strategies by extracting valuable knowledge from their historical trajectories and decision patterns [62]. HiSSD [63] incorporates cooperation temporal learning within a hierarchical architecture, identifying shared and specific skills across tasks, thereby enhancing policy generalization.

E. Research Directions

Most existing studies are highly specialized, focusing on individual cases, specific locations, and time. This specificity leads to models with limited applicability, high costs, and reduced scalability. To broaden the technology’s reach and reduce barriers to its use, the future of temporal and spatio-temporal decision intelligence may include: (1) Pre-training and few-shot learning: Leveraging pre-training on diverse tasks and extensive datasets across various spatio-temporal contexts can yield foundational models with broad applicability. Such models may facilitate few-shot or zero-shot learning when applied to novel tasks and domains. (2) Generative model potential: Advances in generative models may improve performance in temporal and spatio-temporal decision-making. The strong generalization capabilities and precision in data generation of such models could enable better spatio-temporal transferability. (3) Integration with LLMs: Merging temporal and spatio-temporal decision-making with LLMs may enable more intuitive and seamless interactions through natural language.

III. PRESENTERS

Bin Yang is a Chair Professor at East China Normal University, working on data-driven decision intelligence, with a focus on time series and spatio-temporal data. He has published extensively on the tutorial’s topics in database and data mining venues, e.g., ICDE, PVLDB, SIGMOD, and KDD, and machine learning venues, e.g., NERUIPS, ICLR, and IJCAI.

Yuxuan Liang is a tenure-track Assistant Professor at Hong Kong University of Science and Technology (Guangzhou), working on spatio-temporal data mining and urban computing. He has published 70+ papers in prestigious venues (e.g., TPAMI, AI, TKDE, KDD, NeurIPS, and WWW).

Chenjuan Guo is a Professor at East China Normal University, working on AI for time series and spatio-temporal data, in particular, on pretrained foundation models for time series and spatio-temporal data.

Christian S. Jensen is a Professor at Aalborg University, a member of the Academia Europaea, and an ACM and IEEE fellow, working on spatio-temporal data management and analytics. He received ACM SIGMOD Contributions Award and IEEE TCDE Impact Award.

REFERENCES

- [1] C. Guo, C. S. Jensen, and B. Yang, “Towards total traffic awareness,” *SIGMOD Record*, vol. 43, no. 3, pp. 18–23, 2014.
- [2] Z. Li, R. Xu, J. Hu, Z. Peng, X. Lu, C. Guo, and B. Yang, “Ocean significant wave height estimation with spatio-temporally aware large language models,” in *CIKM*, 2024, pp. 3892–3896.
- [3] S. A. Pedersen, B. Yang, and C. S. Jensen, “Anytime stochastic routing with hybrid learning,” *PVLDB*, vol. 13, no. 9, pp. 1555–1567, 2020.
- [4] B. Yang, J. Dai, C. Guo, C. S. Jensen, and J. Hu, “PACE: a path-centric paradigm for stochastic path finding,” *VLDB J.*, vol. 27, no. 2, pp. 153–178, 2018.
- [5] C. Guo, R. Xu, B. Yang, Y. Yuan, T. Kieu, Y. Zhao, and C. S. Jensen, “Efficient stochastic routing in path-centric uncertain road networks,” *Proc. VLDB Endow.*, vol. 17, no. 11, pp. 2893–2905, 2024.
- [6] Z. Pan, Y. Wang, Y. Zhang, S. B. Yang, Y. Cheng, P. Chen, C. Guo, Q. Wen, X. Tian, Y. Dou, Z. Zhou, C. Yang, A. Zhou, and B. Yang, “Magicscaler: Uncertainty-aware, predictive autoscaling,” *PVLDB*, vol. 16, no. 12, pp. 3808–3821, 2023.
- [7] Q. Wen, L. Yang, T. Zhou, and L. Sun, “Robust time series analysis and applications: An industrial perspective,” in *KDD (tutpril)*, 2022.
- [8] E. Keogh, “Time series data mining: A unifying view,” *PVLDB*, vol. 16, no. 12, pp. 3861–3863, 2023.
- [9] C. Gong, D. Yao, C. Zhang, W. Li, J. Bi, L. Du, and J. Wang, “Causal discovery from temporal data,” in *SIGKDD*, 2023, pp. 5803–5804.
- [10] R. Yu, P. Perdikaris, and A. Karpats, “Physics-guided ai for large-scale spatiotemporal data,” in *SIGKDD*, 2021, pp. 4088–4089.
- [11] B. Yang, M. Kaul, and C. S. Jensen, “Using incomplete information for complete weight annotation of road networks,” *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 5, pp. 1267–1279, 2014.
- [12] J. Hu, C. Guo, B. Yang, and C. S. Jensen, “Stochastic weight completion for road networks using graph convolutional networks,” in *ICDE*, 2019, pp. 1274–1285.
- [13] C. Costa, A. Charalampous, A. Konstantinidis, D. Zeinalipour-Yazti, and M. F. Mokbel, “Decaying telco big data with data postdiction,” in *MDM*, 2018, pp. 106–115.
- [14] J. Hu, B. Yang, C. Guo, C. S. Jensen, and H. Xiong, “Stochastic origin-destination matrix forecasting using dual-stage graph convolutional, recurrent neural networks,” in *ICDE*, 2020, pp. 1417–1428.
- [15] B. Yang, C. Guo, C. S. Jensen, M. Kaul, and S. Shang, “Stochastic skyline route planning under time-varying uncertainty,” in *ICDE*, 2014, pp. 136–147.
- [16] X. Zou, Y. Yan, X. Hao, Y. Hu, H. Wen, E. Liu, J. Zhang, Y. Li, T. Li, Y. Zheng *et al.*, “Deep learning for cross-domain data fusion in urban computing: Taxonomy, advances, and outlook,” *arXiv:2402.19348*, 2024.
- [17] P. Newson and J. Krumm, “Hidden markov map matching through noise and sparseness,” in *SIGSPATIAL*, 2009, pp. 336–343.
- [18] J. Zhang, Y. Zheng, and D. Qi, “Deep spatio-temporal residual networks for citywide crowd flows prediction,” in *AAAI*, vol. 31, no. 1, 2017.
- [19] Z. Lin, J. Feng, Z. Lu, Y. Li, and D. Jin, “Deepstn+: Context-aware spatial-temporal neural network for crowd flow prediction in metropolis,” in *AAAI*, vol. 33, no. 01, 2019, pp. 1020–1027.
- [20] M. Jin, S. Wang, L. Ma, Z. Chu, J. Y. Zhang, X. Shi, P.-Y. Chen, Y. Liang, Y.-F. Li, S. Pan, and Q. Wen, “Time-LLM: Time series forecasting by reprogramming large language models,” in *ICLR*, 2024.
- [21] Y. Yan, H. Wen, S. Zhong, W. Chen, H. Chen, Q. Wen, R. Zimmermann, and Y. Liang, “When urban region profiling meets large language models,” *arXiv:2310.18340*, 2023.
- [22] X. Hao, W. Chen, Y. Yan, S. Zhong, K. Wang, Q. Wen, and Y. Liang, “Urbanvlp: A multi-granularity vision-language pre-trained foundation model for urban indicator prediction,” *arXiv:2403.16831*, 2024.
- [23] R. Xu, H. Cheng, C. Guo, H. Gao, J. Hu, S. B. Yang, and B. Yang, “Mm-path: Multi-modal, multi-granularity path representation learning,” in *SIGKDD*, 2025.
- [24] X. Wu, D. Zhang, C. Guo, C. He, B. Yang, and C. S. Jensen, “AutoCTS: Automated correlated time series forecasting,” *PVLDB*, vol. 15, no. 4, pp. 971–983, 2022.
- [25] X. Wu, D. Zhang, M. Zhang, C. Guo, B. Yang, and C. S. Jensen, “AutoCTS+: Joint neural architecture and hyperparameter search for correlated time series forecasting,” *Proc. ACM Manag. Data*, vol. 1, no. 1, pp. 97:1–97:26, 2023.
- [26] Z. Pan, S. Ke, X. Yang, Y. Liang, Y. Yu, J. Zhang, and Y. Zheng, “Autostg: Neural architecture search for predictions of spatio-temporal graph,” in *The Web Conference*, 2021, pp. 1846–1855.
- [27] X. Wu, X. Wu, B. Yang, L. Zhou, C. Guo, X. Qiu, J. Hu, Z. Sheng, and C. S. Jensen, “AutoCTS++: zero-shot joint neural architecture and hyperparameter search for correlated time series forecasting,” *VLDB J.*, vol. 33, no. 5, pp. 1743–1770, 2024.
- [28] X. Wu, X. Wu, D. Zhang, M. Zhang, C. Guo, B. Yang, and C. S. Jensen, “Fully automated correlated time series forecasting in minutes,” *Proc. VLDB Endow.*, vol. 18, 2024.
- [29] S. B. Yang, C. Guo, and B. Yang, “Context-aware path ranking in road networks,” *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 7, pp. 3153–3168, 2022.
- [30] S. B. Yang, C. Guo, J. Hu, J. Tang, and B. Yang, “Unsupervised path representation learning with curriculum negative sampling,” in *IJCAI*, 2021, pp. 3286–3292.
- [31] S. B. Yang, C. Guo, J. Hu, B. Yang, J. Tang, and C. S. Jensen, “Weakly-supervised temporal path representation learning with contrastive curriculum learning,” in *ICDE*, 2022, pp. 2873–2885.
- [32] S. B. Yang, J. Hu, C. Guo, B. Yang, and C. S. Jensen, “Lightpath: Lightweight and scalable path representation learning,” in *SIGKDD*, ACM, 2023, pp. 2999–3010.
- [33] Y. Wang, Y. Qiu, P. Chen, K. Zhao, Y. Shu, Z. Rao, L. Pan, B. Yang, and C. Guo, “Rose: Register assisted general time series forecasting with decomposed frequency learning,” *arXiv preprint arXiv:2405.17478*, 2024.
- [34] Q. Shentu, B. Li, K. Zhao, Y. Shu, Z. Rao, L. Pan, B. Yang, and C. Guo, “Towards a general time series anomaly detector with adaptive bottlenecks and dual adversarial decoders,” in *ICLR*, 2025.
- [35] X. Liu, J. Hu, Y. Li, S. Diao, Y. Liang, B. Hooi, and R. Zimmermann, “Unitime: A language-empowered unified model for cross-domain time series forecasting,” *arXiv:2310.09751*, 2023.
- [36] T. Kieu, B. Yang, C. Guo, R. Cirstea, Y. Zhao, Y. Song, and C. S. Jensen, “Anomaly detection in time series with robust variational quasi-recurrent autoencoders,” in *ICDE*, 2022, pp. 1342–1354.
- [37] T. Kieu, B. Yang, C. Guo, C. S. Jensen, Y. Zhao, F. Huang, and K. Zheng, “Robust and explainable autoencoders for unsupervised time series outlier detection,” in *ICDE*, 2022, pp. 3038–3050.
- [38] X. Wu, X. Qiu, Z. Li, Y. Wang, J. Hu, C. Guo, H. Xiong, and B. Yang, “Catch: Channel-aware multivariate time series anomaly detection via frequency patching,” in *ICLR*, 2025.
- [39] Z. Zhuang, Y. Zhang, K. Zhao, C. Guo, B. Yang, Q. Wen, and L. Fan, “Noise matters: Cross contrastive learning for flink anomaly detection,” *VLDB Endowment*, vol. 18, no. 4, pp. 1159–1168, 2024.
- [40] Y. Cheng, P. Chen, C. Guo, K. Zhao, Q. Wen, B. Yang, and C. S. Jensen, “Weakly guided adaptation for robust time series forecasting,” *PVLDB*, 2024.
- [41] H. Miao, Y. Zhao, C. Guo, B. Yang, Z. Kai, F. Huang, J. Xie, and C. S. Jensen, “A unified replay-based continuous learning framework for spatio-temporal prediction on streaming data,” *ICDE*, 2024.
- [42] D. Kieu, T. Kieu, P. Han, B. Yang, C. S. Jensen, and B. Le, “Team: Topological evolution-aware framework for traffic forecasting,” *Proc. VLDB Endow.*, vol. 18, 2024.
- [43] X. Qiu, X. Wu, Y. Lin, C. Guo, J. Hu, and B. Yang, “Duet: Dual clustering enhanced multivariate time series forecasting,” in *SIGKDD*, 2025.
- [44] P. Chen, Y. Zhang, Y. Cheng, Y. Shu, Y. Wang, Q. Wen, B. Yang, and C. Guo, “Pathformer: Multi-scale transformers with adaptive pathways for time series forecasting,” *ICLR*, 2024.
- [45] T. Kieu, B. Yang, C. Guo, and C. S. Jensen, “Outlier detection for time series with recurrent autoencoder ensembles,” in *IJCAI*, 2019, pp. 2725–2732.
- [46] D. Campos, T. Kieu, C. Guo, F. Huang, K. Zheng, B. Yang, and C. S. Jensen, “Unsupervised time series outlier detection with diversity-driven convolutional ensembles,” *PVLDB*, vol. 15, no. 3, pp. 611–623, 2022.
- [47] Y. Cheng, C. Guo, K. Chen, K. Zhao, B. Yang, J. Xie, C. S. Jensen, F. Huang, and K. Zheng, “A pattern discovery approach to multivariate time series forecasting,” *arXiv:2212.10306*, 2022.

- [48] R.-G. Cirstea, T. Kieu, C. Guo, B. Yang, and S. J. Pan, “EnhanceNet: Plugin neural networks for enhancing correlated time series forecasting,” in *ICDE*, 2021, pp. 1739–1750.
- [49] R. Cirstea, B. Yang, C. Guo, T. Kieu, and S. Pan, “Towards spatio-temporal aware traffic time series forecasting,” in *ICDE*, 2022, pp. 2900–2913.
- [50] K. Zhao, C. Guo, P. Han, M. Zhang, Y. Cheng, and B. Yang, “Multiple time series forecasting with dynamic graph modeling,” *PVLDB*, 2024.
- [51] J. Tian, Y. Liang, R. Xu, P. Chen, C. Guo, A. Zhou, L. Pan, Z. Rao, and B. Yang, “Air quality prediction with physics-guided dual neural odes in open systems,” in *ICLR*, 2025.
- [52] D. Campos, M. Zhang, B. Yang, T. Kieu, C. Guo, and C. S. Jensen, “LightTS: Lightweight time series classification with adaptive ensemble distillation,” *Proc. ACM Manag. Data*, vol. 1, no. 2, pp. 171:1–171:27, 2023.
- [53] D. Campos, B. Yang, T. Kieu, M. Zhang, C. Guo, and C. S. Jensen, “Qcore: Data-efficient, on-device continual calibration for quantized models,” *Proc. VLDB Endow.*, vol. 17, no. 11, pp. 2708–2721, 2024.
- [54] K. Zhao, Z. Zhuang, M. Zhang, C. Guo, Y. Shu, and B. Yang, “Enhancing diversity for data-free quantization,” in *CVPR*. IEEE, 2025.
- [55] H. Miao, Z. Liu, Y. Zhao, C. Guo, B. Yang, K. Zheng, and C. S. Jensen, “Less is more: Efficient time series dataset condensation via two-fold modal matching,” *Proc. VLDB Endow.*, vol. 18, 2024.
- [56] Z. Li, X. Qiu, P. Chen, Y. Wang, H. Cheng, Y. Shu, J. Hu, C. Guo, A. Zhou, Q. Wen *et al.*, “Foundts: Comprehensive and unified benchmarking of foundation models for time series forecasting,” *arXiv preprint arXiv:2410.11802*, 2024.
- [57] J. Hu, B. Yang, C. Guo, and C. S. Jensen, “Risk-aware path selection with time-varying, uncertain travel costs: a time series approach,” *VLDB J.*, vol. 27, no. 2, pp. 179–200, 2018.
- [58] S. A. Pedersen, B. Yang, and C. S. Jensen, “Fast stochastic routing under time-varying uncertainty,” *VLDB J.*, vol. 29, no. 4, pp. 819–839, 2020.
- [59] S. A. Pedersen, B. Yang, C. S. Jensen, and J. Møller, “Stochastic routing with arrival windows,” *ACM Trans. Spatial Algorithms Syst.*, vol. 9, no. 4, pp. 30:1–30:48, 2023.
- [60] B. Yang, C. Guo, Y. Ma, and C. S. Jensen, “Toward personalized, context-aware routing,” *VLDB J.*, vol. 24, no. 2, pp. 297–318, 2015.
- [61] C. Guo, B. Yang, J. Hu, C. S. Jensen, and L. Chen, “Context-aware, preference-based vehicle routing,” *VLDB J.*, vol. 29, no. 5, pp. 1149–1170, 2020.
- [62] C. Guo, B. Yang, J. Hu, and C. S. Jensen, “Learning to route with sparse trajectory sets,” in *ICDE*, 2018, pp. 1073–1084.
- [63] S. Liu, Y. Shu, C. Guo, and B. Yang, “Learning generalizable skills from offline multi-task data for multi-agent cooperation,” in *ICLR*, 2025.