# Towards Predictive Cities: Modeling Spatio-Temporal Data in the AI Era

Slides for this Talk

**Xia Yutong**
PhD Candidate
Institute of Data Science
National University of Singapore

# Outline

## Background

- What makes cities predictable?
- Spatio-Temporal (ST) Data & Properties

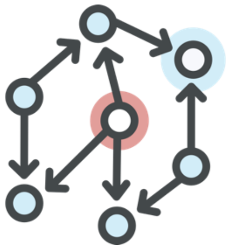

## Spatio-Temporal Graph Forecasting

- What is Spatio-Temporal Graph (STG)?
- What is STG forecasting?
- How we do it?
  - Application-Driven (Air Quality, Traffic, Parking)
  - Theory-Driven (Causality, Uncertainty)

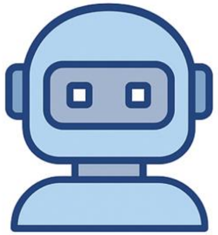

## Beyond Prediction: What’s Next?

- LLMs-powered Agents & Causal Urban Insight

# Outline



**Background**

- What makes cities predictable?
- Spatio-Temporal (ST) Data & Properties

# Cities Are Alive

*Cities are not static structures. They are dynamic organisms — pulsing with people, data, and change.*

# Cities Are Alive — And Becoming Predictive

Nowadays…

## Cities are **Data**-Rich

Recent advances in sensing technologies (e.g., IoT devices, mobile apps, satellite imagery, and urban sensors) have enabled the continuous collection of rich **spatio-temporal data**.

**+**

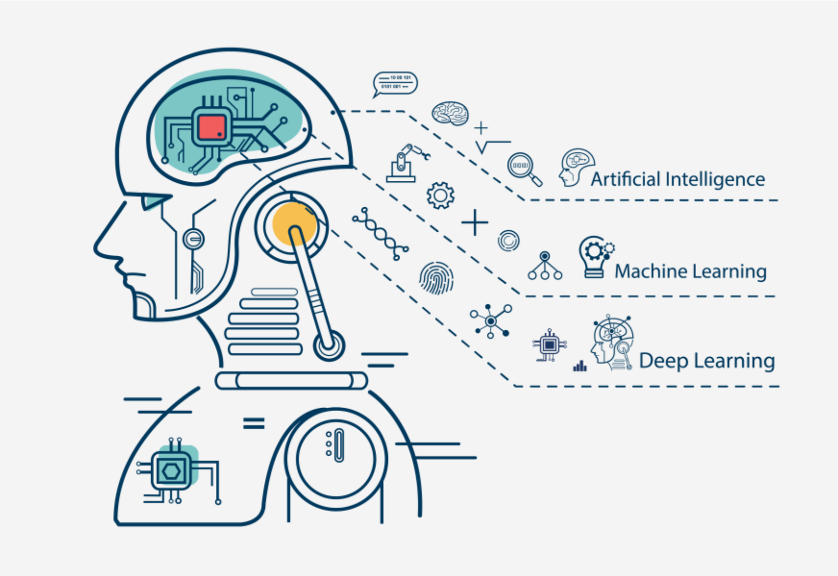


## **AI** is a Powerful Tool

Advances in **AI (ML/DL/LLM)** provide tools to analyze complex patterns, forecast urban dynamics, and support data-driven decision-making.

# Foundation Models

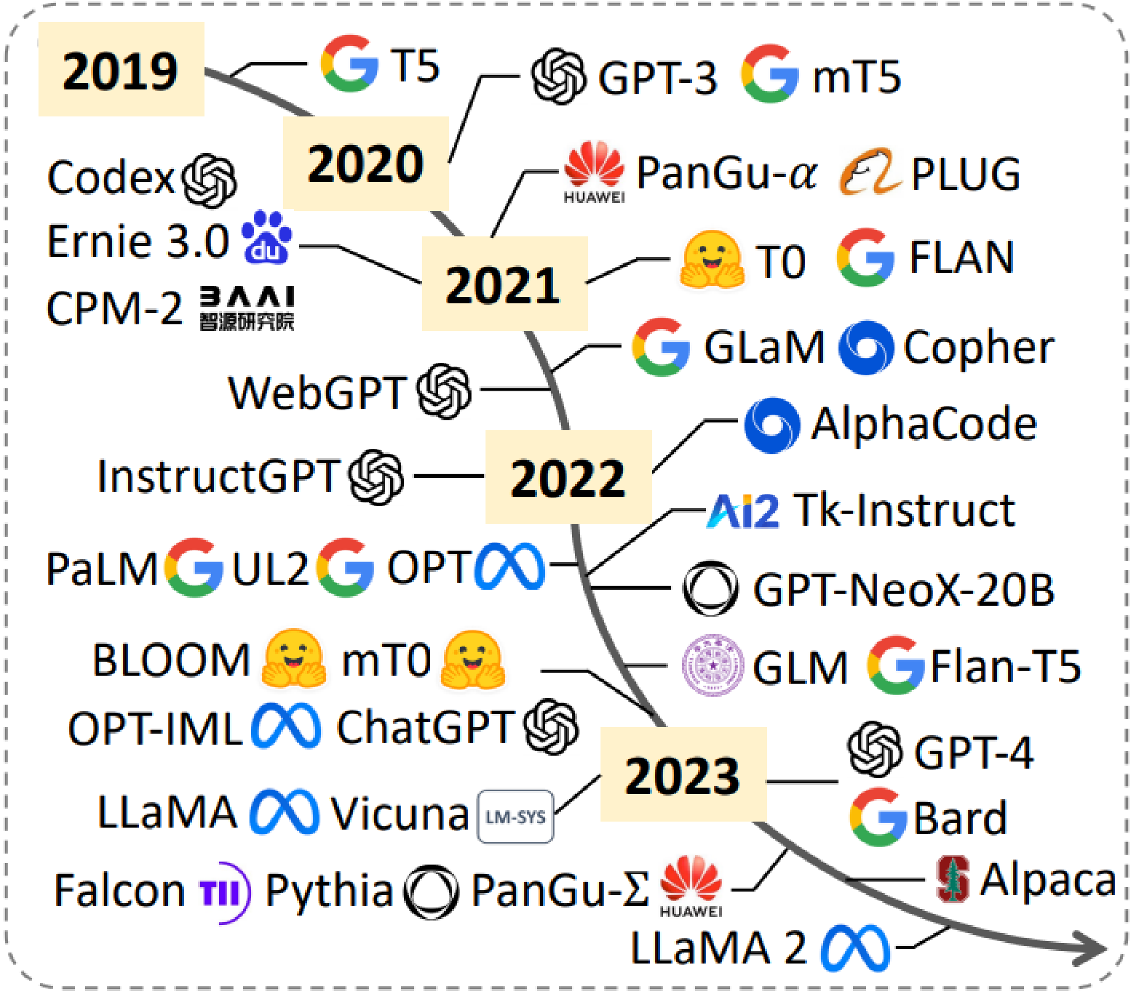

# Foundation Models: From NLP to Multimodal

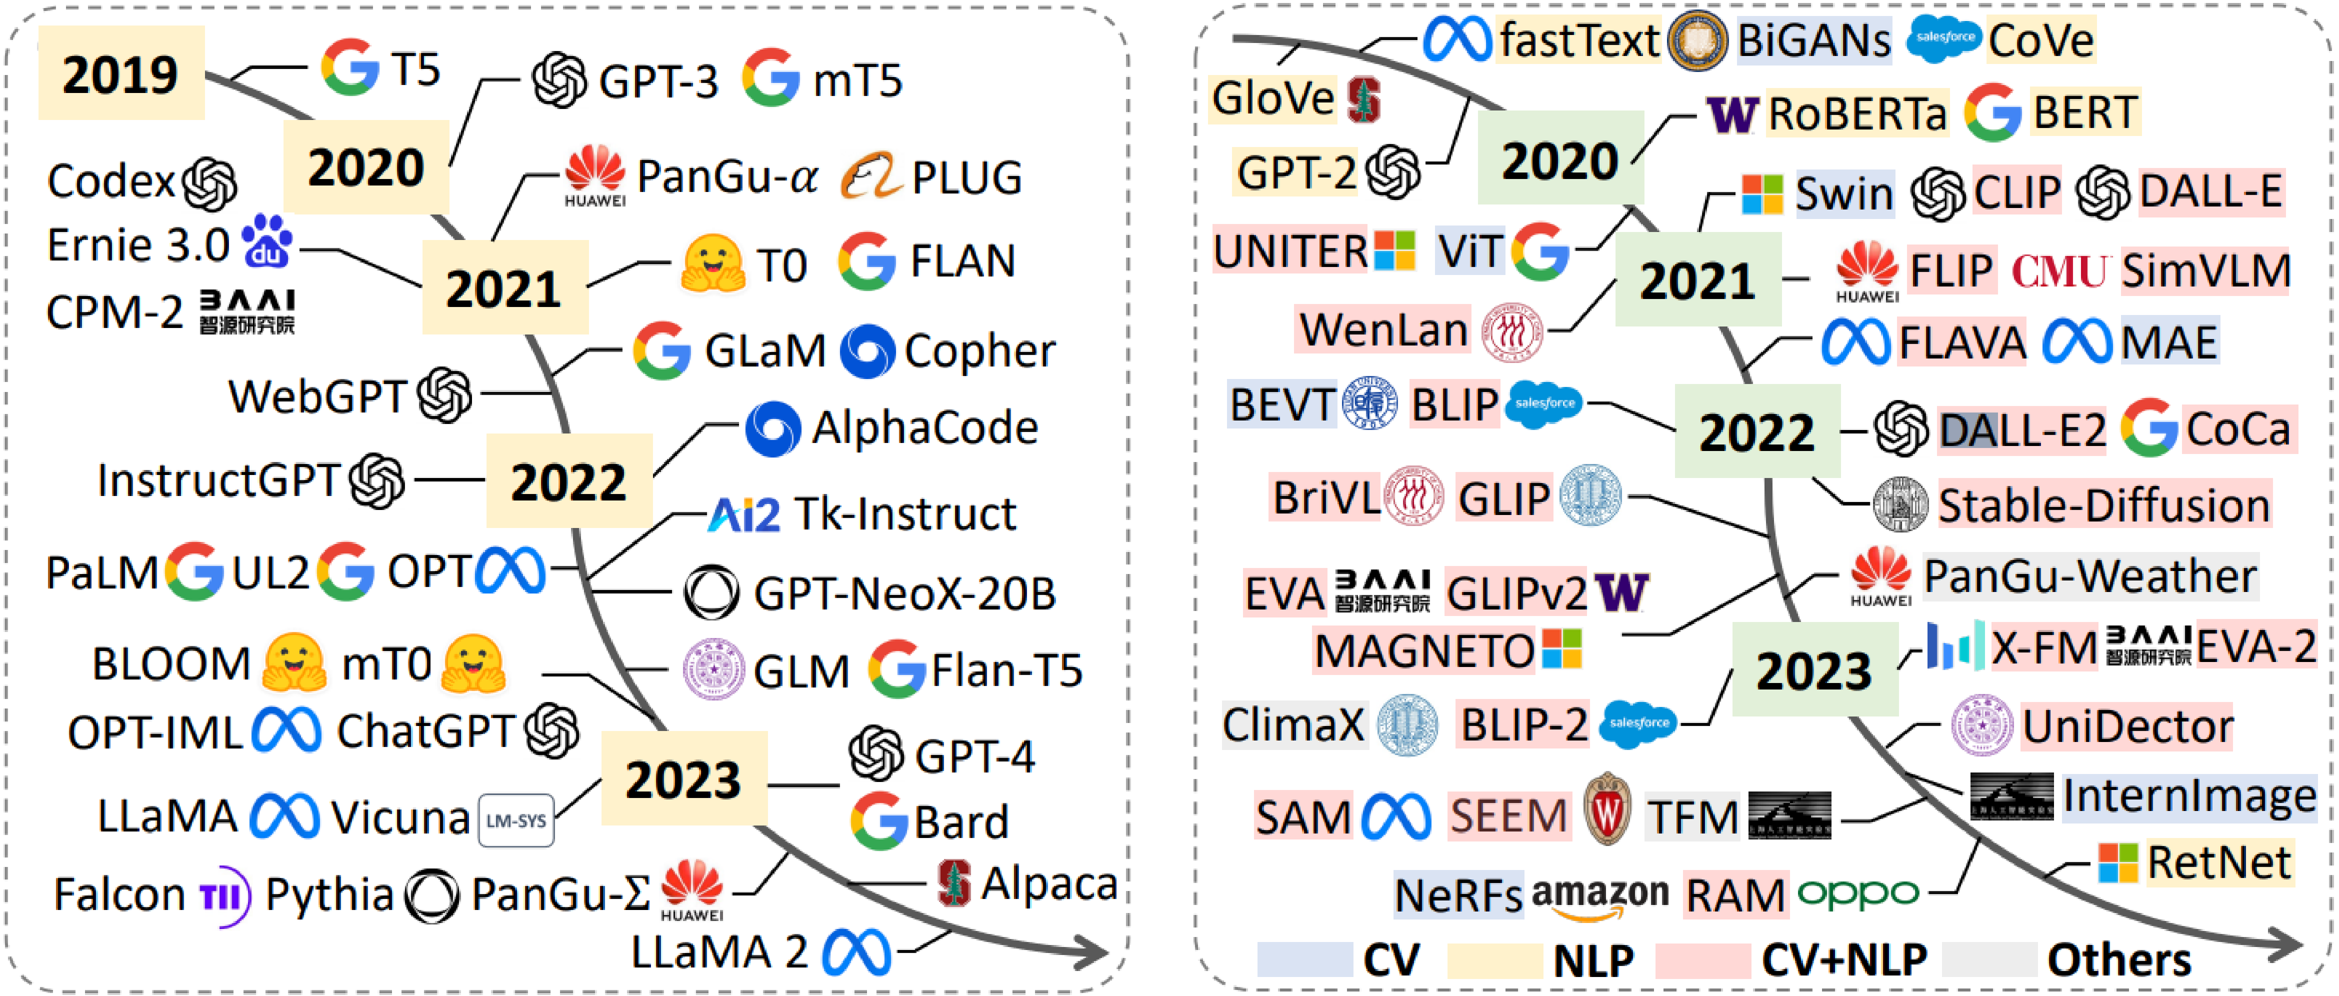


M. Jin et al., Large Models for Time Series and Spatio-Temporal Data: A Survey and Outlook. arXiv 2023.

# Foundation Models: From General to STFMs

**FMs and LLMs** are now capable of supporting the entire urban Spatio-Temporal Data Science lifecycle.

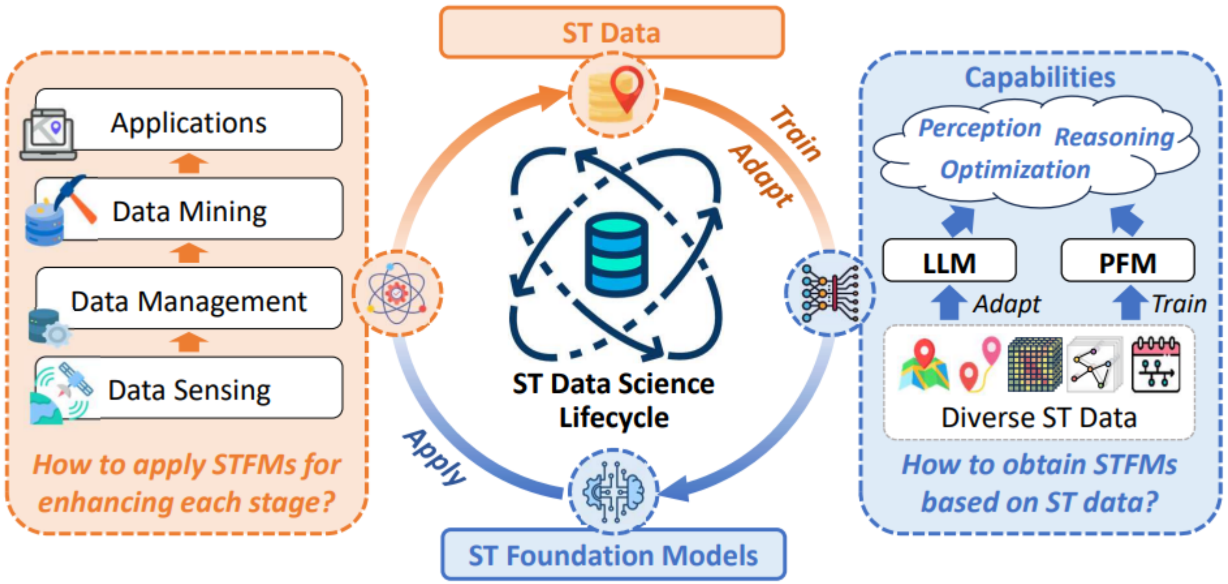


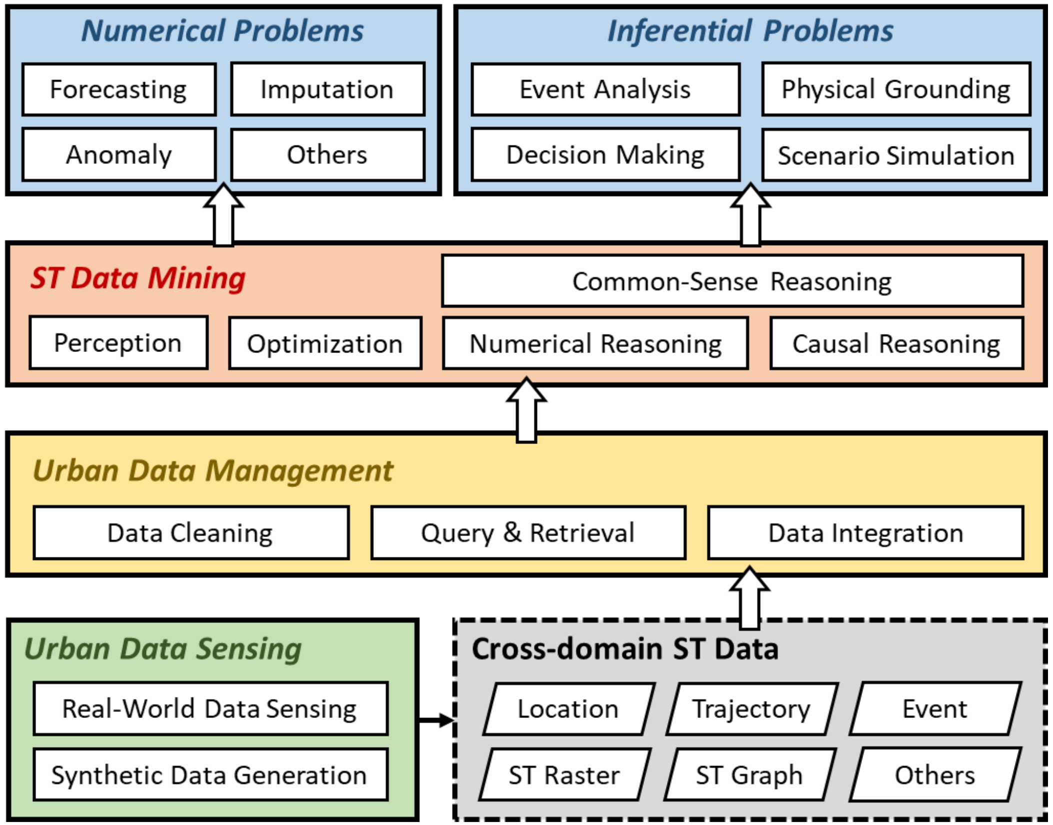


Survey Paper



Y. Liang, H. Wen, Y. Xia et al., Foundation Models for Spatio-Temporal Data Science: A Tutorial and Survey. arXiv 2025.

# Spatio-Temporal (ST) Data

- **Spatio-Temporal (ST) Data** is data that changes both over ***space*** and ***time***. Tells us not just what is happening, but also ***where*** and ***when*** it's happening.
- Type of Urban ST Data

**Spatio-Temporal Static Data (GeoData)**

**Vector**


Points


Lines


Polygons

**Raster**


Satellite Images

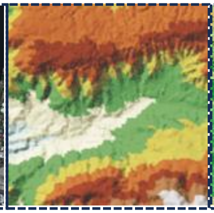
DEMs


NDVI

**Spatial Static, Temporal Dynamic Data (Aggregate-level)**

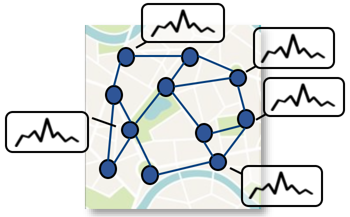
Weather/AQI/Traffic Sensor Data

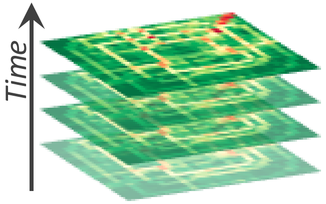

OD Matrix

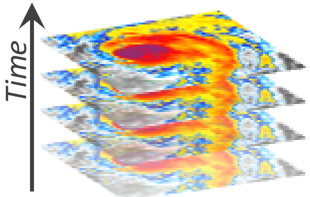

Grided Weather Data

**Spatio-Temporal Dynamic Data (Individual-level)**

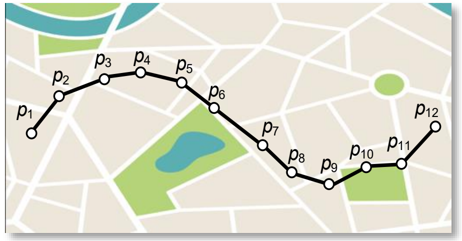

Trajectory Data

# Spatial and Temporal Properties

- **Spatial Properties**

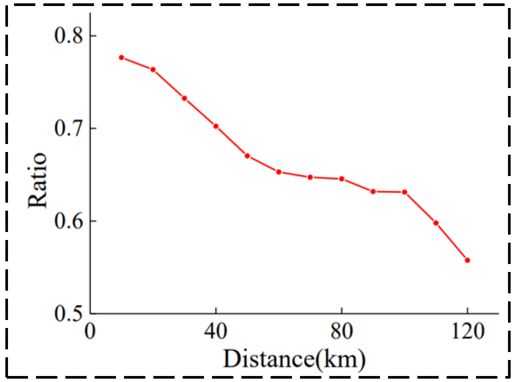


Spatial Closeness

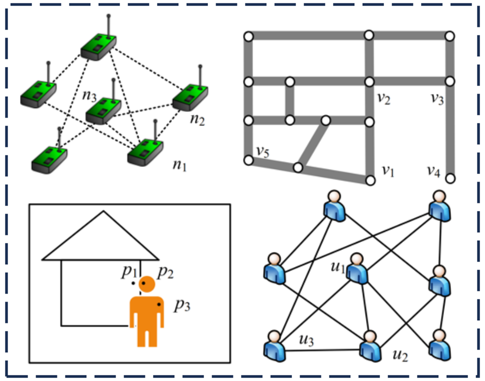


Spatial Distance

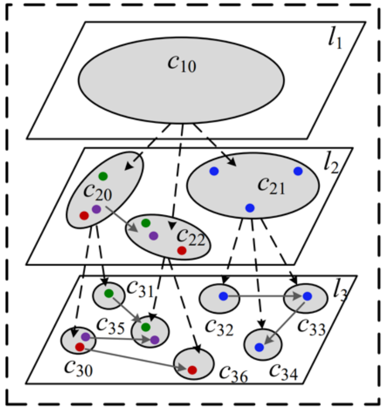


Spatial Hierarchy

- **Temporal Properties**

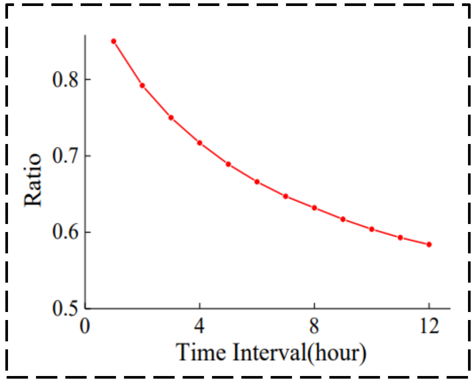


Temporal Closeness

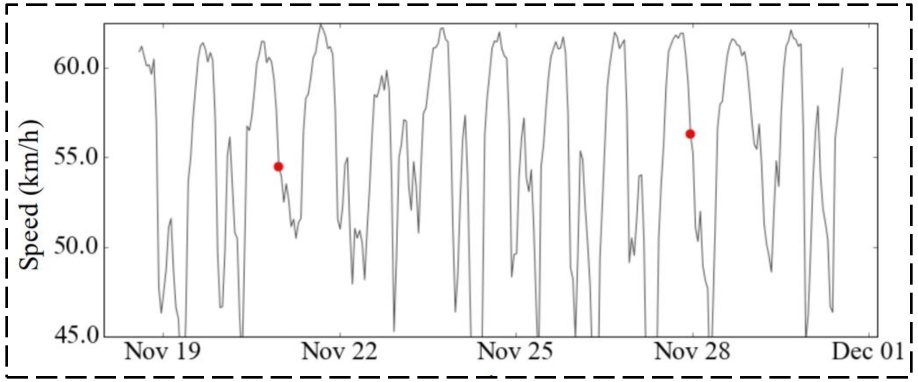


Temporal Periodicity

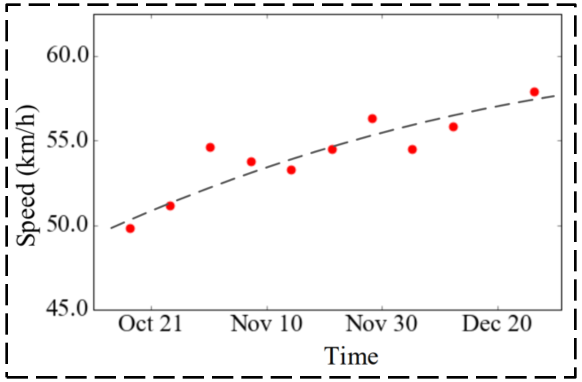


Trend



Y. Zheng. Urban Computing. MIT Press, 2019.

# Spatio-Temporal (ST) Data

- **Spatio-Temporal (ST) Data** is data that changes both over ***space*** and ***time***. Tells us not just what is happening, but also ***where*** and ***when*** it's happening.
- Type of Urban ST Data

**Spatio-Temporal Static Data (GeoData)**

**Vector**

**Raster**

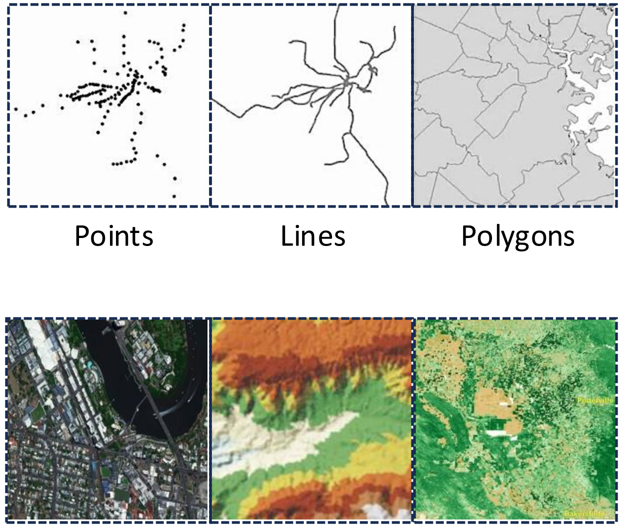

Points | Lines | Polygons

Satellite Images | DEMs | NDVI

**Spatial Static, Temporal Dynamic Data (Aggregate-level)**

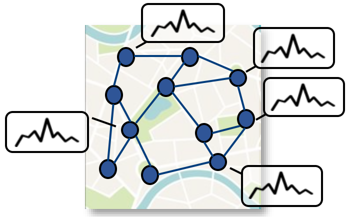

Weather/AQI/Traffic Sensor Data

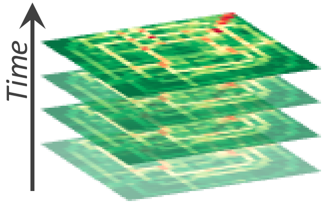


OD Matrix

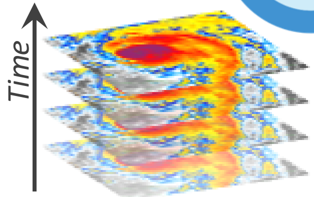


Grided Weather Data

**Spatio-Temporal Dynamic Data (Individual-level)**

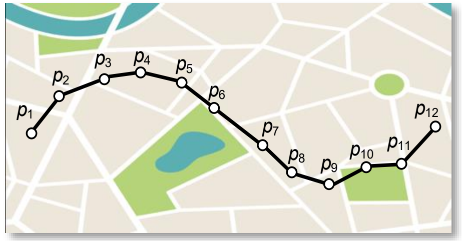


Trajectory Data

# Outline

## Background

- What makes cities predictable?
- Spatio-Temporal (ST) Data & Properties

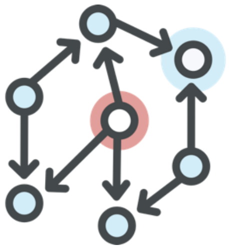

## Spatio-Temporal Graph Forecasting

- What is Spatio-Temporal Graph (STG)?
- What is STG forecasting?
- How we do it?
    - Application-Driven (Air Quality, Traffic, Parking)
    - Theory-Driven (Causality, Uncertainty)

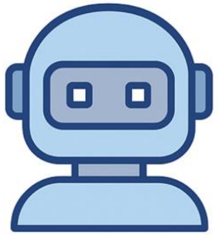

## Beyond Prediction: What’s Next?

- LLMs-powered Agents & Causal Urban Insight

# Outline

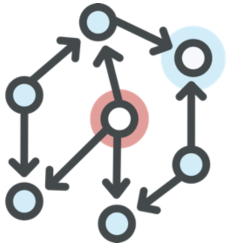

## Spatio-Temporal Graph Forecasting

- What is Spatio-Temporal Graph (STG)?
- What is STG forecasting?
- How we do it?
    - Application-Driven (Air Quality, Traffic, Parking)
    - Theory-Driven (Causality, Uncertainty)

# Spatio-Temporal Graph (STG) Data

- **Spatio-Temporal Graph (STG)** is one type of ST data, which represents the spatial and temporal relationships between nodes or entities.

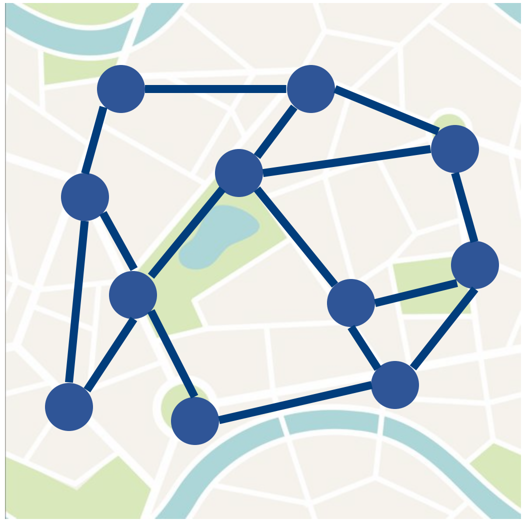

Graph

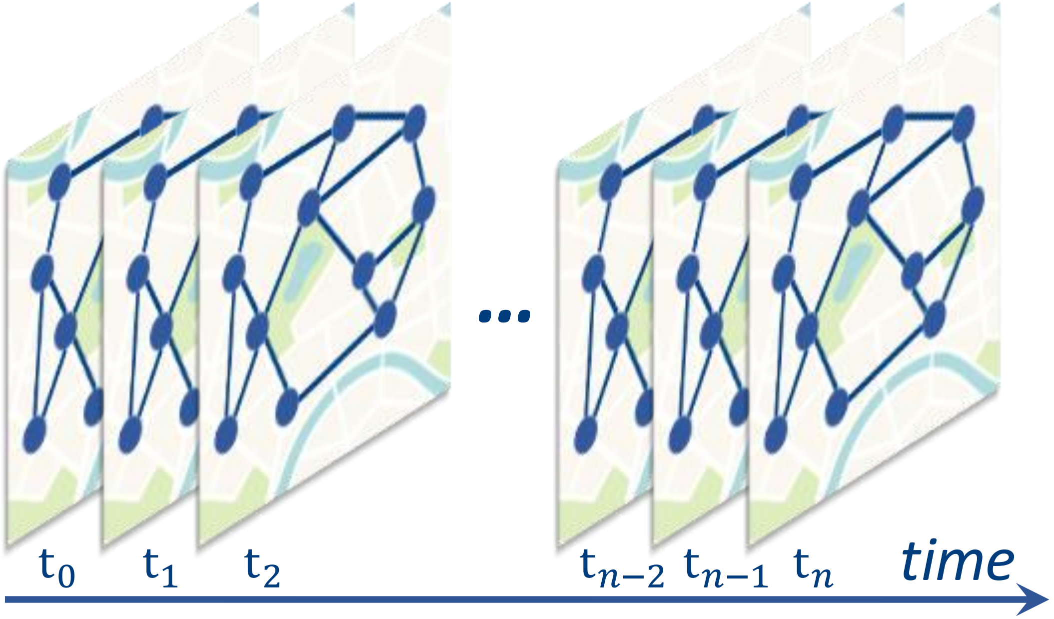


Spatio-Temporal Graph

# Spatio-Temporal Graph (STG) Data

- **Spatio-Temporal Graph (STG)** is one type of ST data, which represents the spatial and temporal relationships between nodes or entities.

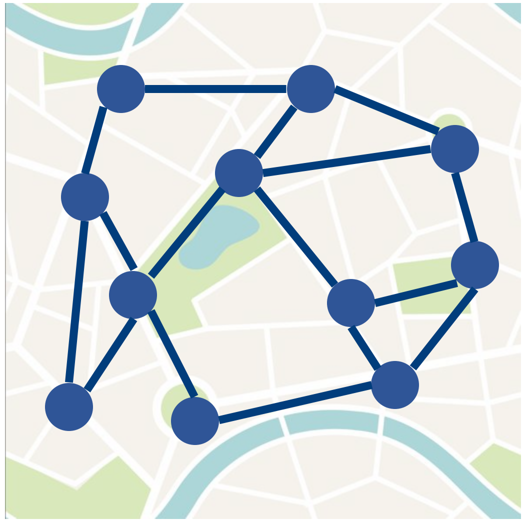

Graph

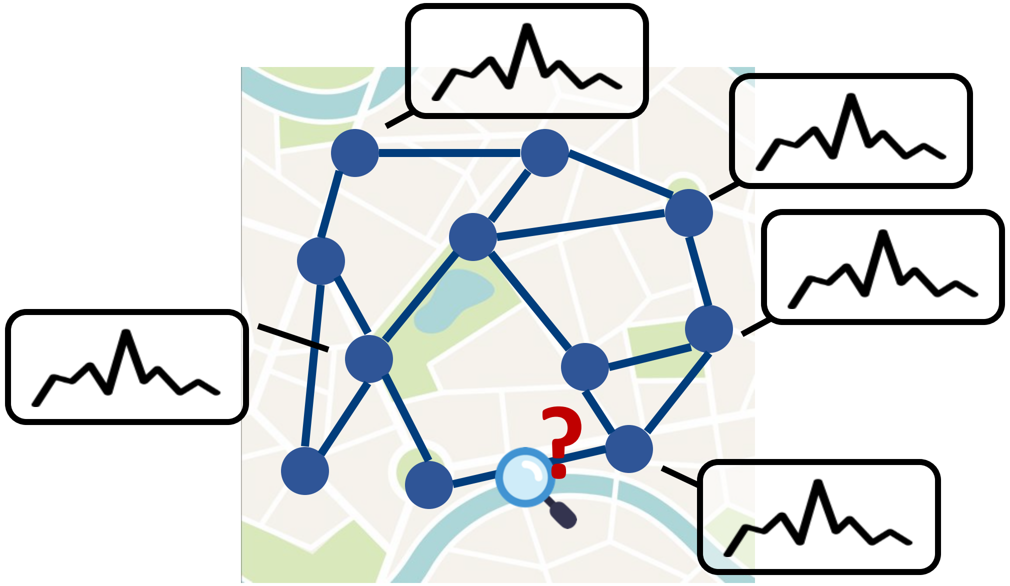


Spatio-Temporal Graph (Series)

# Spatio-Temporal Graph (STG) Data

- **Graph Construction Methods**
  - Topology-based graph

$$a_{ij}^t = \begin{cases} 1, \text{ if } v_i \text{ connects to } v_j \\ 0, \text{ otherwise} \end{cases},$$

  - Distance-based graph

$$a_{ij}^t = \begin{cases} \dfrac{\exp(-\left\|d_{ij}^t\right\|_2)}{\sigma}, \text{if } d_{ij}^t < \epsilon, \\ 0, \quad \text{otherwise} \end{cases}$$

  - Similarity-based graph

$$a_{ij}^t = \begin{cases} \dfrac{\sum_{i=1}^n \left(x_i^{0:t} - \bar{x_i^{0:t}}\right)\left(x_j^{0:t} - \bar{x_j^{0:t}}\right)}{\sqrt{\sum_{i=1}^n \left(x_i^{0:t} - \bar{x_i^{0:t}}\right)^2}\sqrt{\sum_{i=1}^n \left(x_j^{0:t} - \bar{x_j^{0:t}}\right)^2}}, \\ 0, \quad \text{otherwise} \end{cases}$$

  - Interaction-based graph

$$a_{ij}^t = \begin{cases} \dfrac{F_{ij}^t}{\sum_{m \in N(i)} F_{im}^t}, \; if\ F_{ij}^t > 0 \\ 0, \quad \text{otherwise} \end{cases},$$



G. Jin et al. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. TKDE. 2023.

# STG Forecasting Problem

- **STG forecasting** has become crucial in the context of smart cities (e.g. Air quality prediction, traffic flow forecasting…)

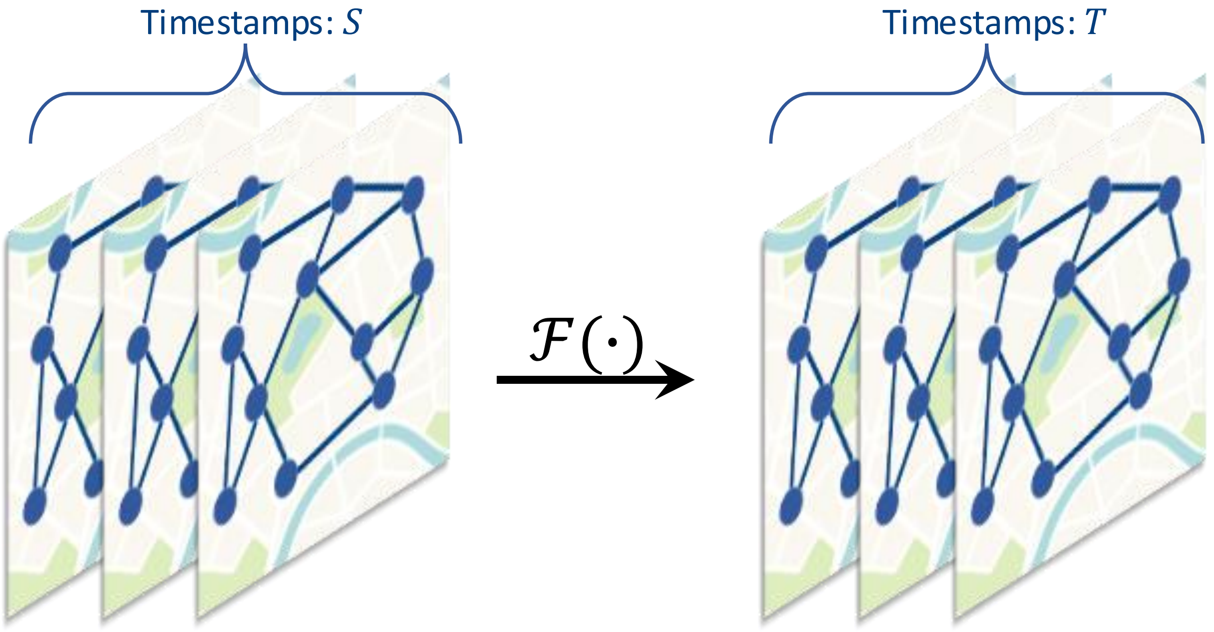


Historical Data $\boldsymbol{X}^{(t-S):t}$
Graph Structure $\boldsymbol{G}$

Future Data $\boldsymbol{Y}^{t:t+T}$

# STG Forecasting Method

*How to learn $\mathcal{F}(\cdot)$?*

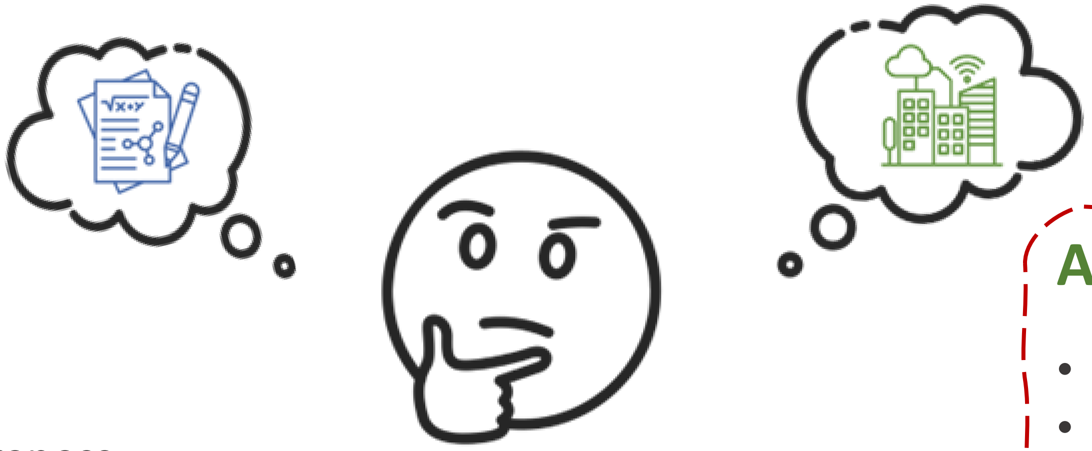

**Theory-Driven**

- Causal Inference
- Uncertainty Awareness

**Application-Driven**

- Air Quality
- Traffic Flow
- Parking Availability

# Application-Driven Method

NUS National University of Singapore

PM.25 **Air Quality Forecasting** | Traffic Flow Forecasting | Parking Availability Forecasting

## **AirFormer** (AAAI'23)

- **Challenges**: Inefficiency & Uncertainty
- **Solution:**
  - Bottom-up deterministic stage
    - Dartboard Spatial-MSA (DS-MSA)
    - Causal Temporal-MSA (CT-MSA)
  - Top-down stochastic stage
    - Generation model

$$p_\theta\left(\mathcal{Z}_t|\mathbf{X}_{1:t-1}\right)=\prod_{n=1}^{N} p_\theta\left(\left\{\mathbf{z}_{t,n}^1,\ldots,\mathbf{z}_{t,n}^L\right\}|\mathbf{X}_{1:t-1}\right)$$

$$=\prod_{n=1}^{N} p_\theta\left(\mathbf{z}_{t,n}^L|\mathbf{h}_{t-1,n}^L\right)\prod_{l=1}^{L-1} p_\theta\left(\mathbf{z}_{t,n}^l|\mathbf{z}_{t,n}^{l+1},\mathbf{h}_{t-1,n}^l\right),$$

    - Inference model

$$q_\phi\left(\mathcal{Z}_t|\mathbf{X}_{1:t}\right)=\prod_{n=1}^{N} q_\phi\left(\mathbf{z}_{t,n}^L|\mathbf{h}_{t,n}^L\right)\prod_{l=1}^{L-1} q_\phi\left(\mathbf{z}_{t,n}^l|\mathbf{z}_{t,n}^{l+1},\mathbf{h}_{t,n}^l\right)$$

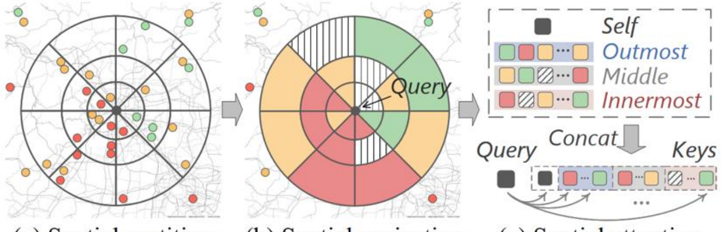


**DS-MSA**

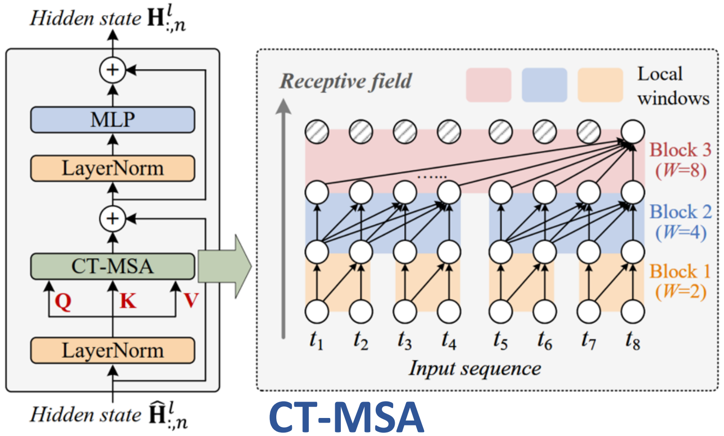


**CT-MSA**

$\mathcal{O}(N^2C)$ MSA

Reduced — $N$: #stations, $M$: #regions

$\mathcal{O}(NMC)$ our DS-MSA

$\mathcal{O}(T^2C)$ MSA

Reduced — $T$: #time stamps, $W$: #windows

$\mathcal{O}(TWC)$ our DS-MSA

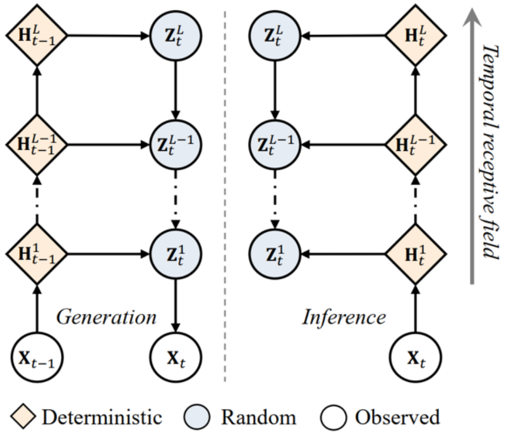




Y. Liang, **Y. Xia** et al., AirFormer: Predicting Nationwide Air Quality in China with Transformers. **AAAI 2023.**

# Application-Driven Method

Air Quality Forecasting | **Traffic Flow Forecasting** | Parking Availability Forecasting

## LargeST (NeurIPS’23)

### Motivation – Limitation in Existing Dataset

***Scale***: Existing datasets (e.g., PeMS03, PeMS04) contain only hundreds of sensors, not reflecting real-world traffic network scales.

***Temporal Coverage***: Typically cover less than 6 months, hindering the study of long-term patterns.

***Metadata***: Often lack comprehensive sensor metadata, affecting data reliability and interpretability

### LargeST - A new large-scale dataset

***Larger Graph Size***: **8,600** sensors across California.

***Higher Temporal Coverage***: **5** years of data (2015–2019) with a **5-minute** sampling rate.

***Richer Node Metadata***: Includes sensor ID, location, highway category, number of lanes, and direction.

***Subsets***: Provides regional subsets for Greater Los Angeles (GLA), Greater Bay Area (GBA), and San Diego (SD).

**Dataset & Code**

| Released Source | Dataset | Nodes | Edges | Degree | Meta | Time Range | Frames | Data Points |
|---|---|---|---|---|---|---|---|---|
| Yu et al. [34] | PeMSD7(M) | 228 | 1,664 | 7.3 | 6 | 05/01/2012 – 06/30/2012 | 12,672 | 2.89M |
| | PeMSD7(L) | 1,026 | 14,534 | 14.2 | 0 | 05/01/2012 – 06/30/2012 | 12,672 | 13.00M |
| Li et al. [19] | METR-LA | 207 | 1,515 | 7.3 | 3 | 03/01/2012 – 06/27/2012 | 34,272 | 7.09M |
| | PEMS-BAY | 325 | 2,369 | 7.3 | 3 | 01/01/2017 – 06/30/2017 | 52,116 | 16.94M |
| Song et al. [30] | PEMS03 | 358 | 546 | 1.5 | 1 | 09/01/2018 – 11/30/2018 | 26,208 | 9.38M |
| | PEMS04 | 307 | 338 | 1.1 | 0 | 01/01/2018 – 02/28/2018 | 16,992 | 5.22M |
| | PEMS07 | 883 | 865 | 1.0 | 0 | 05/01/2017 – 08/06/2017 | 28,224 | 24.92M |
| | PEMS08 | 170 | 276 | 1.6 | 0 | 07/01/2016 – 08/31/2016 | 17,856 | 3.04M |
| LargeST (ours) | CA | 8,600 | 201,363 | 23.4 | 9 | 01/01/2017 – 12/31/2021 | 525,888 | 4.52B |
| | GLA | 3,834 | 98,703 | 25.7 | 9 | 01/01/2017 – 12/31/2021 | 525,888 | 2.02B |
| | GBA | 2,352 | 61,246 | 26.0 | 9 | 01/01/2017 – 12/31/2021 | 525,888 | 1.24B |
| | SD | 716 | 17,319 | 24.2 | 9 | 01/01/2017 – 12/31/2021 | 525,888 | 0.38B |

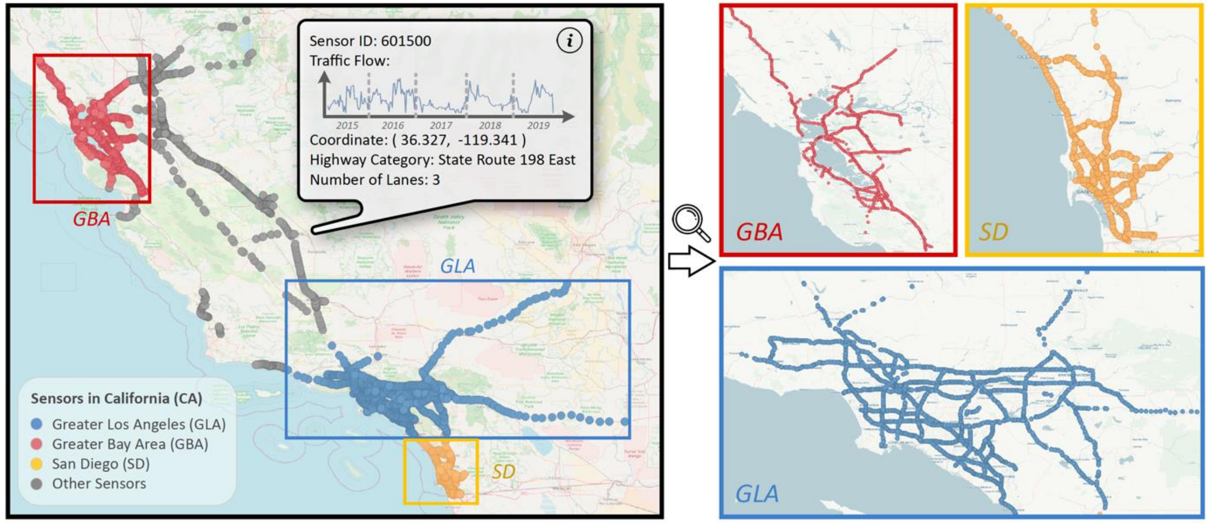


(a) Overview of the LargeST dataset (b) Fine-grained distribution of sensors

X. Liu, **Y. Xia** et al., LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. **NeurIPS 2023.**

# Application-Driven Method

## **DeepPA** + **SINPA** (IJCAI'24)

Predicting Parking Availability in Singapore with Cross-Domain Data

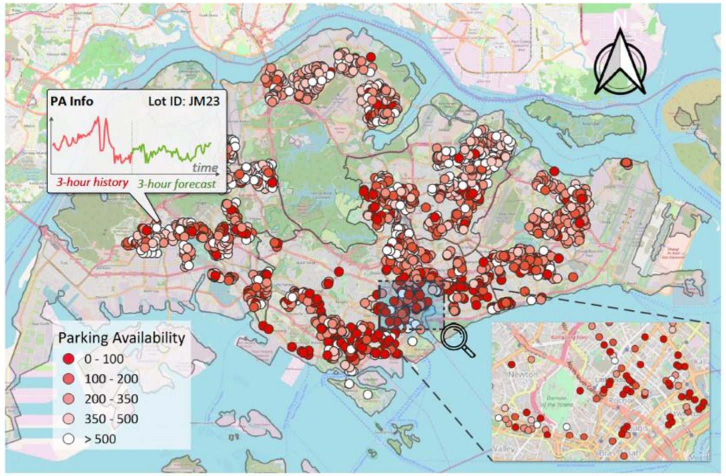


**Parking Availability Data**

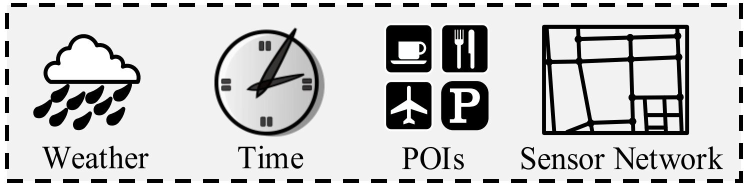


**Cross-domain Data**

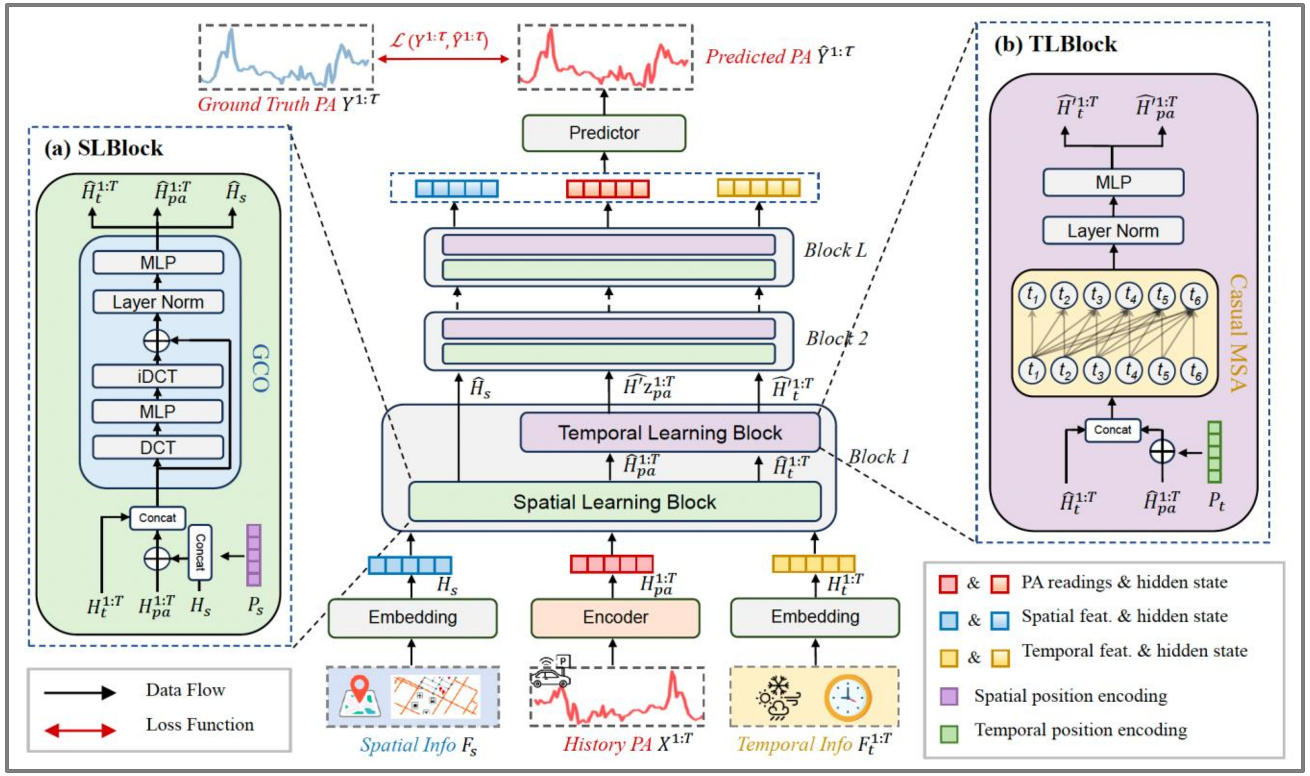


Dataset & Code

***A New Dataset + A Novel Data-driven Method***



H. Zhang*, Y. Xia* et al., Predicting Carpark Availability in Singapore with Cross-Domain Data: A New Dataset and A Data-Driven Approach. **IJCAI 2024.**

# Application-Driven Method

For application, we also deploy our cutting-edge AI solutions on large-scale cloud platforms

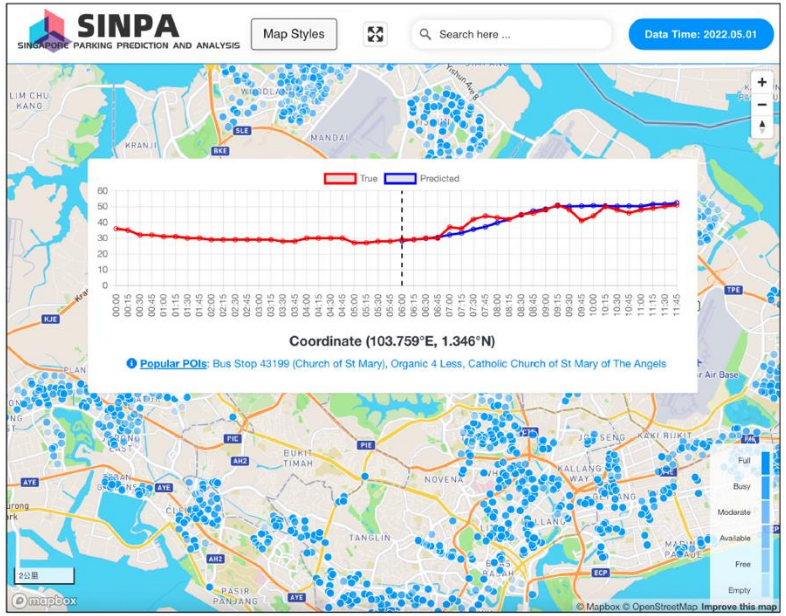


Parking Availability Prediction in Singapore

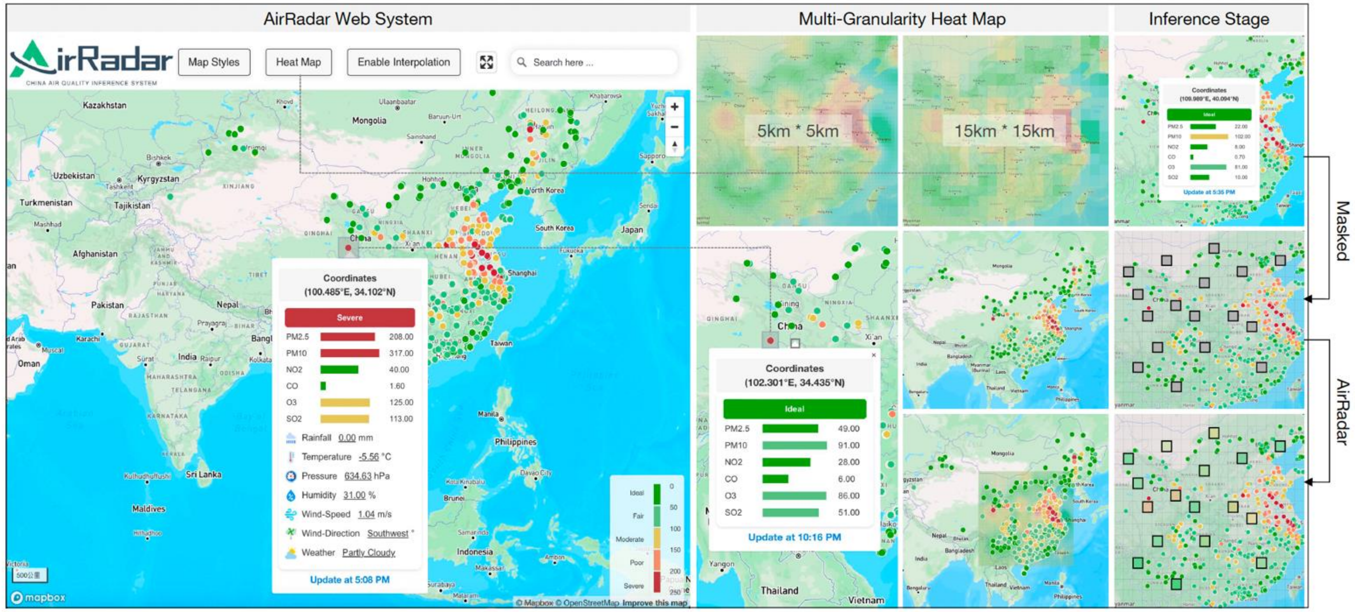


Air Quality Inference in China

H. Zhang*, **Y. Xia***, et al., Predicting Parking Availability in Singapore with Cross-Domain Data: A New Dataset and A Data-Driven Approach. (IJCAI'24)
Q. Wang, **Y. Xia**, et al., AirRadar: Inferring Nationwide Air Quality in China with Deep Neural Networks. (AAAI'25)

# STG Forecasting Method

*How to learn* $\mathcal{F}(\cdot)$*?*

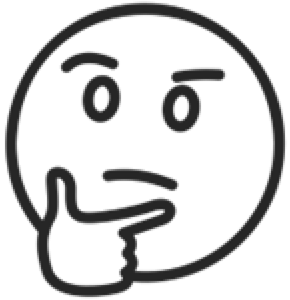

**Theory-Driven**

- Causal Inference
- Uncertainty Awareness

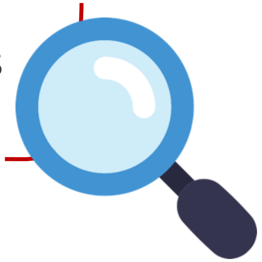

**Application-Driven**

- Air Quality
- Traffic Flow
- Parking Availability

# STG Forecasting Method

*How to learn* $\mathcal{F}(\cdot)$*?*

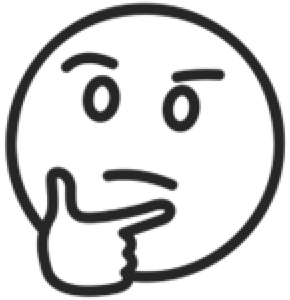

### Theory-Driven

- **Causal Inference**
- Uncertainty Awareness

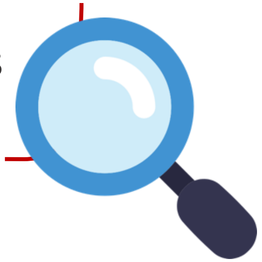

### Application-Driven

- Air Quality
- Traffic Flow
- Parking Availability

# Theory-Driven - *Causal Inference*



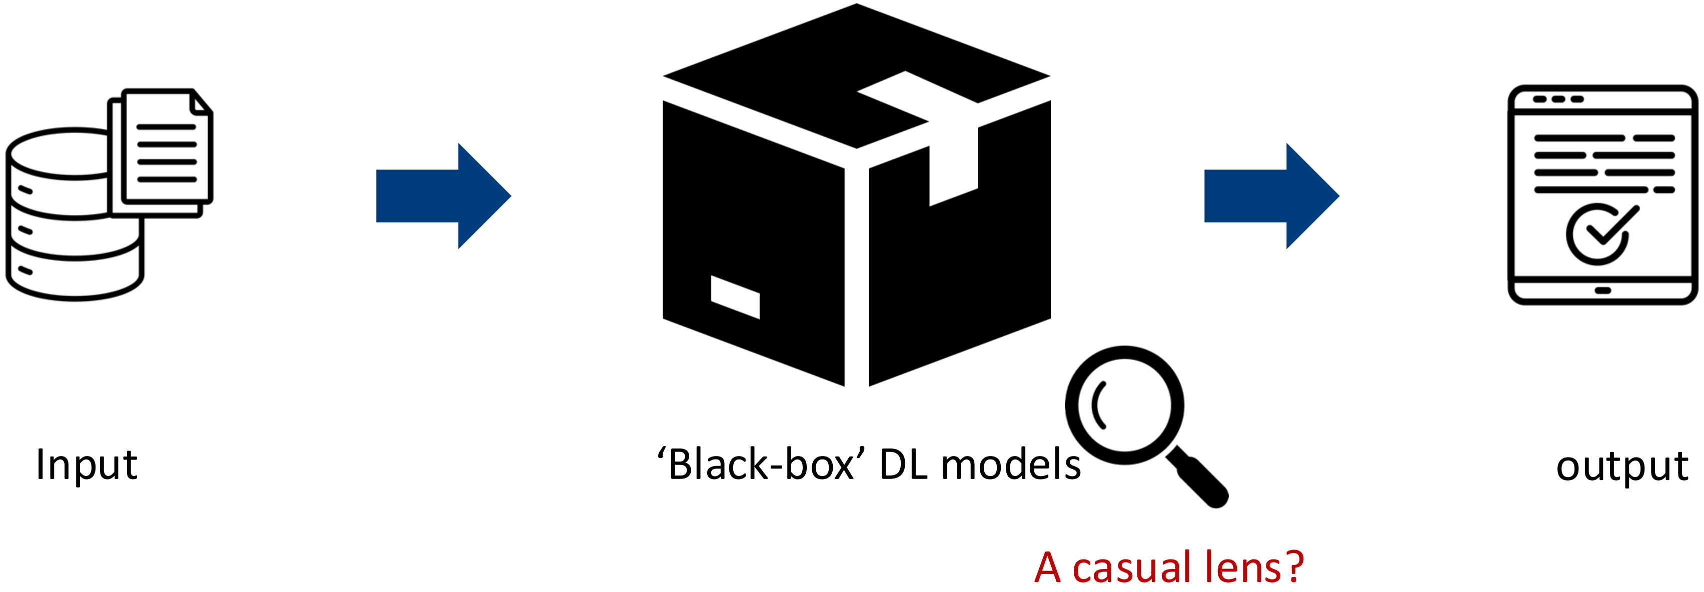

# Theory-Driven - *Causal Inference*



- **Causal Inference** is the process of determining if a relationship between two things is a cause-and-effect relationship.
- *Correlation doesn't mean causality.*

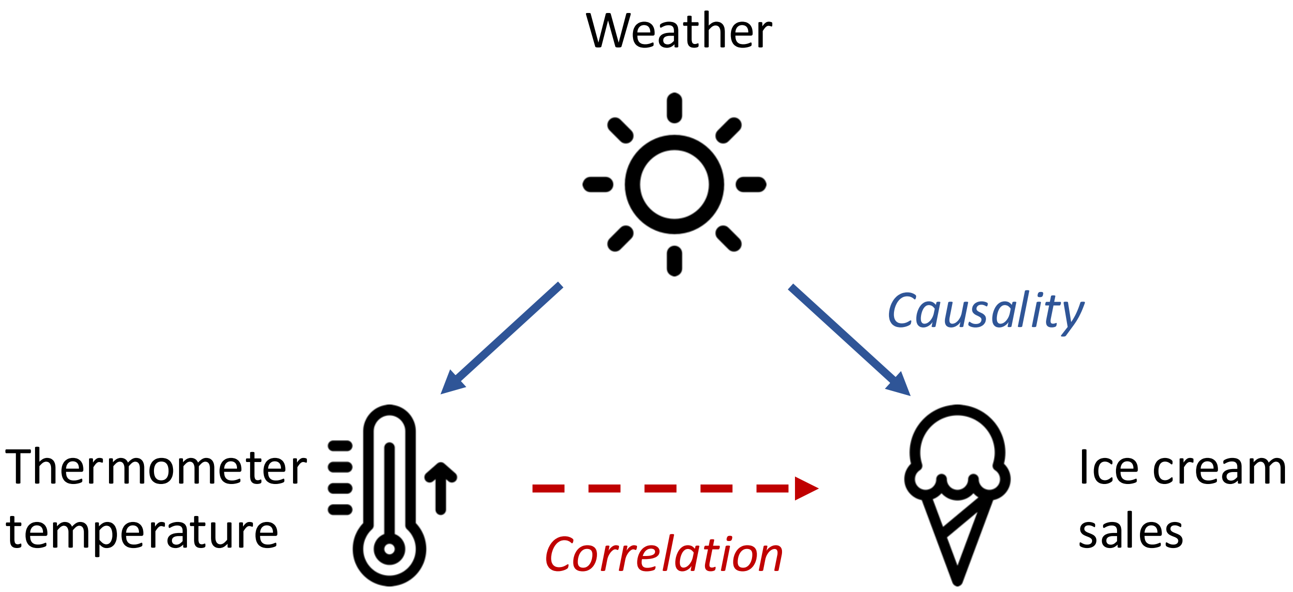

# Theory-Driven - *Causal Inference*



**Integrating deep learning with causal inference, we craft models that are not only accurate but truly understand the real world.**

- ***Why?*** - Advantage of a causal lens
  - Improved Interpretability
  - Real-world insights for better model design
  - Enhanced generalization
- ***How?*** – A big picture

## Structure Causal Model


Traditional model

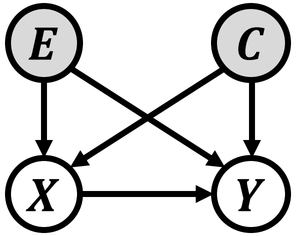


Underlying causal system



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*

## Structure Causal Model

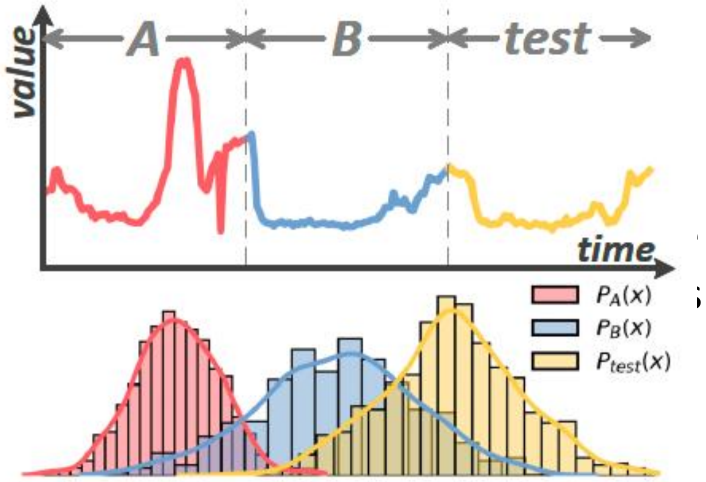


$P_A(X) \neq P_B(X) \neq P_{test}\ (X)$

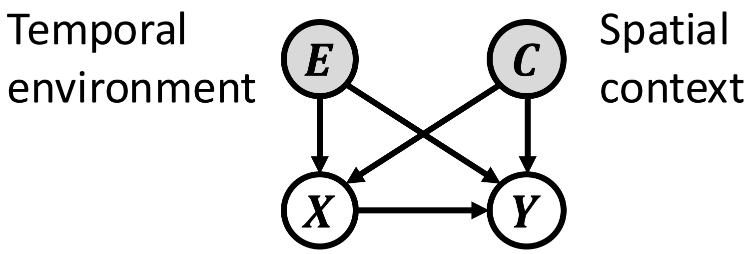


Underlying causal system

- $X \leftarrow E \rightarrow Y$: The ***temporal Out-of-Distribution (OoD)*** can arise due to changes in external variables over time (e.g., weather can affect traffic flow observations).
- $X \leftarrow C \rightarrow Y$**:** $X$ and $Y$ are intrinsically affected by the surrounding spatial context, comprising both ***spurious*** and ***genuine*** causal components.
- $X \rightarrow Y$**:** Our primary prediction goal.



Y. Xia, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



## Structure Causal Model

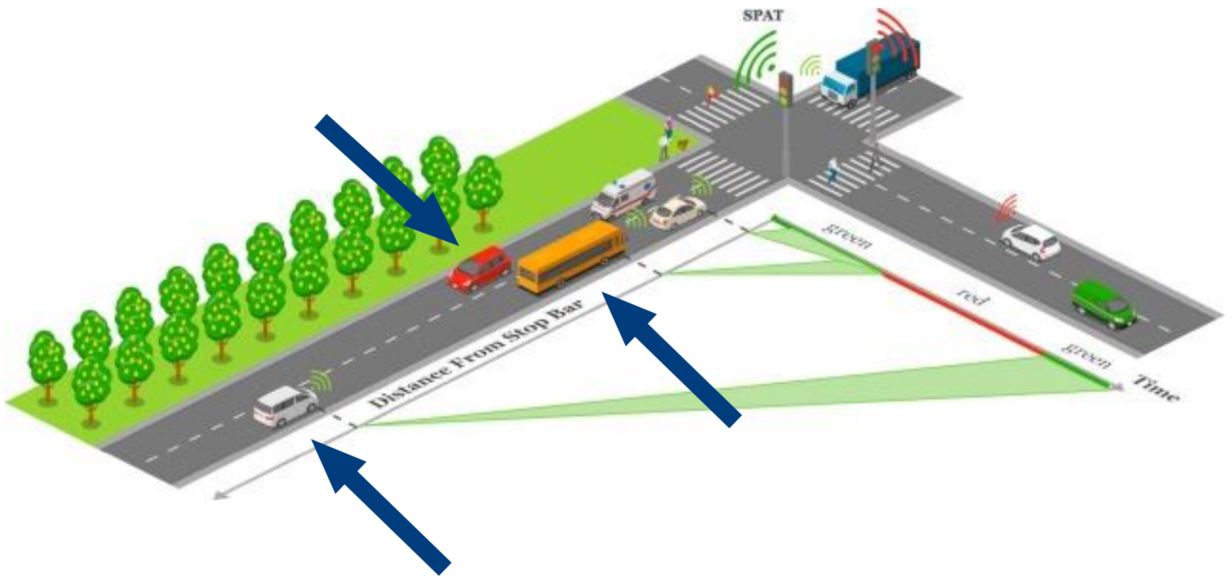


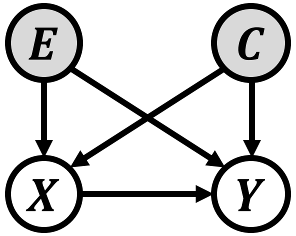


Underlying causal system

- $X \leftarrow E \rightarrow Y$: The ***temporal Out-of-Distribution (OoD)*** can arise due to changes in external variables over time (e.g., weather can affect traffic flow observations).
- $X \leftarrow C \rightarrow Y$: $X$ and $Y$ are intrinsically affected by the surrounding spatial context, comprising both ***spurious*** and ***genuine*** causal components.
- $X \rightarrow Y$: Our primary prediction goal.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

## Structure Causal Model

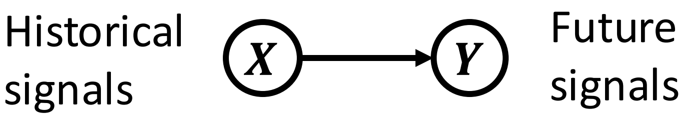


Traditional model

***Confounding factors***

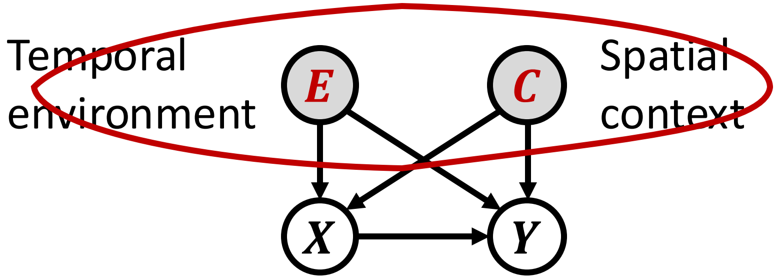


Underlying causal system

***Backdoor paths***

- $X \leftarrow E \rightarrow Y$: The ***temporal Out-of-Distribution (OoD)*** can arise due to changes in external variables over time (e.g., weather can affect traffic flow observations).
- $X \leftarrow C \rightarrow Y$: $X$ and $Y$ are intrinsically affected by the surrounding spatial context, comprising both ***spurious*** and ***genuine*** causal components.
- $X \rightarrow Y$: Our primary goal.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



**Assumption**: $\boldsymbol{E}$ and $\boldsymbol{C}$ are independent

**Aim**: $P(Y|do(X))$

**Step 1**. Back-door adjustment for $\boldsymbol{E}$

$$\sum_{e} P(Y|X, E = e)P(E = e)$$

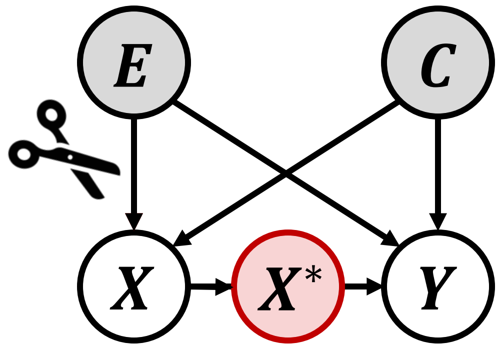


**Step 2**. Front-door adjustment for $\boldsymbol{C}$

$$\sum_{x^*} \sum_{x'} P(X^* = x^*|X)\, P(Y|X^* = x^*, X = x')P(X = x')$$



Y. Xia, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*

- To achieve **Back-door adjustment** …

$$\sum_{e} P(Y|X, E = e)P(E = e)$$

1) Separate the entity and the environments
2) Discretizing the environment

- To achieve **Front-door adjustment** …
  - Obtain a surrogate entity

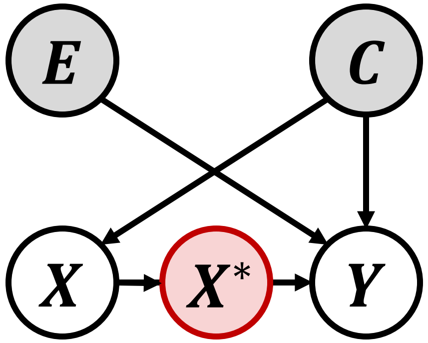




**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

## Causal Spatio-Temporal neural network (CaST)

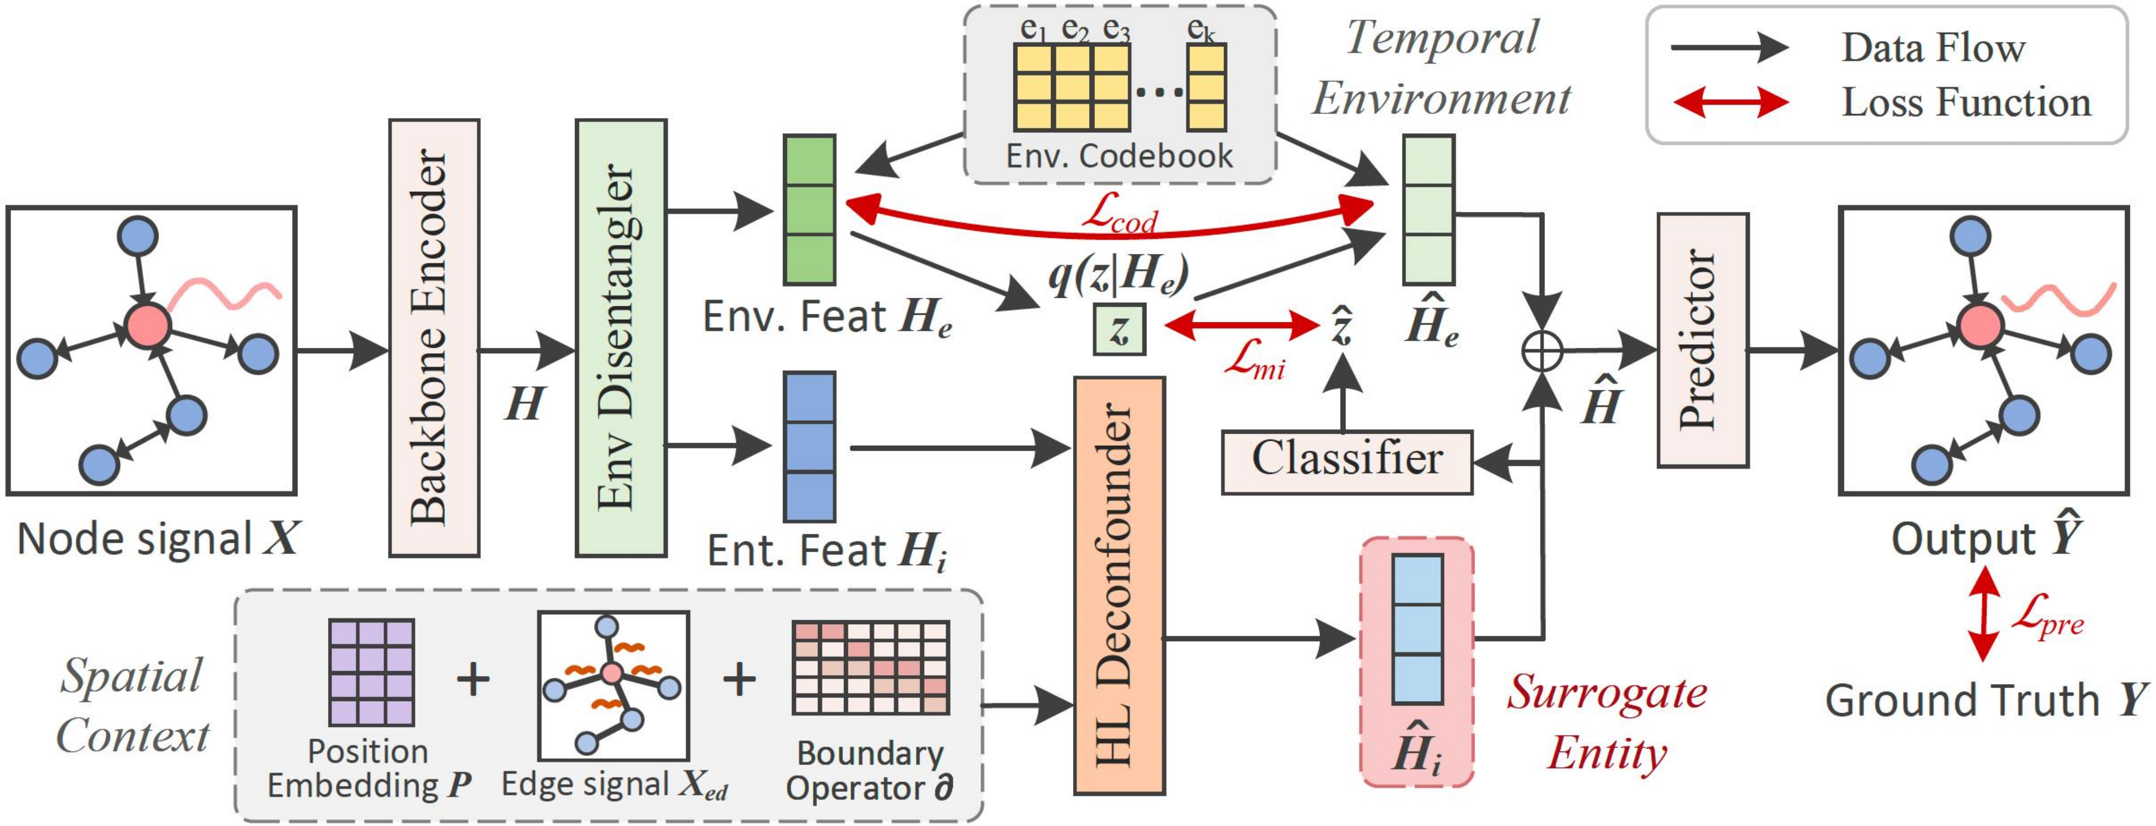


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*

## Causal Spatio-Temporal neural network (CaST)

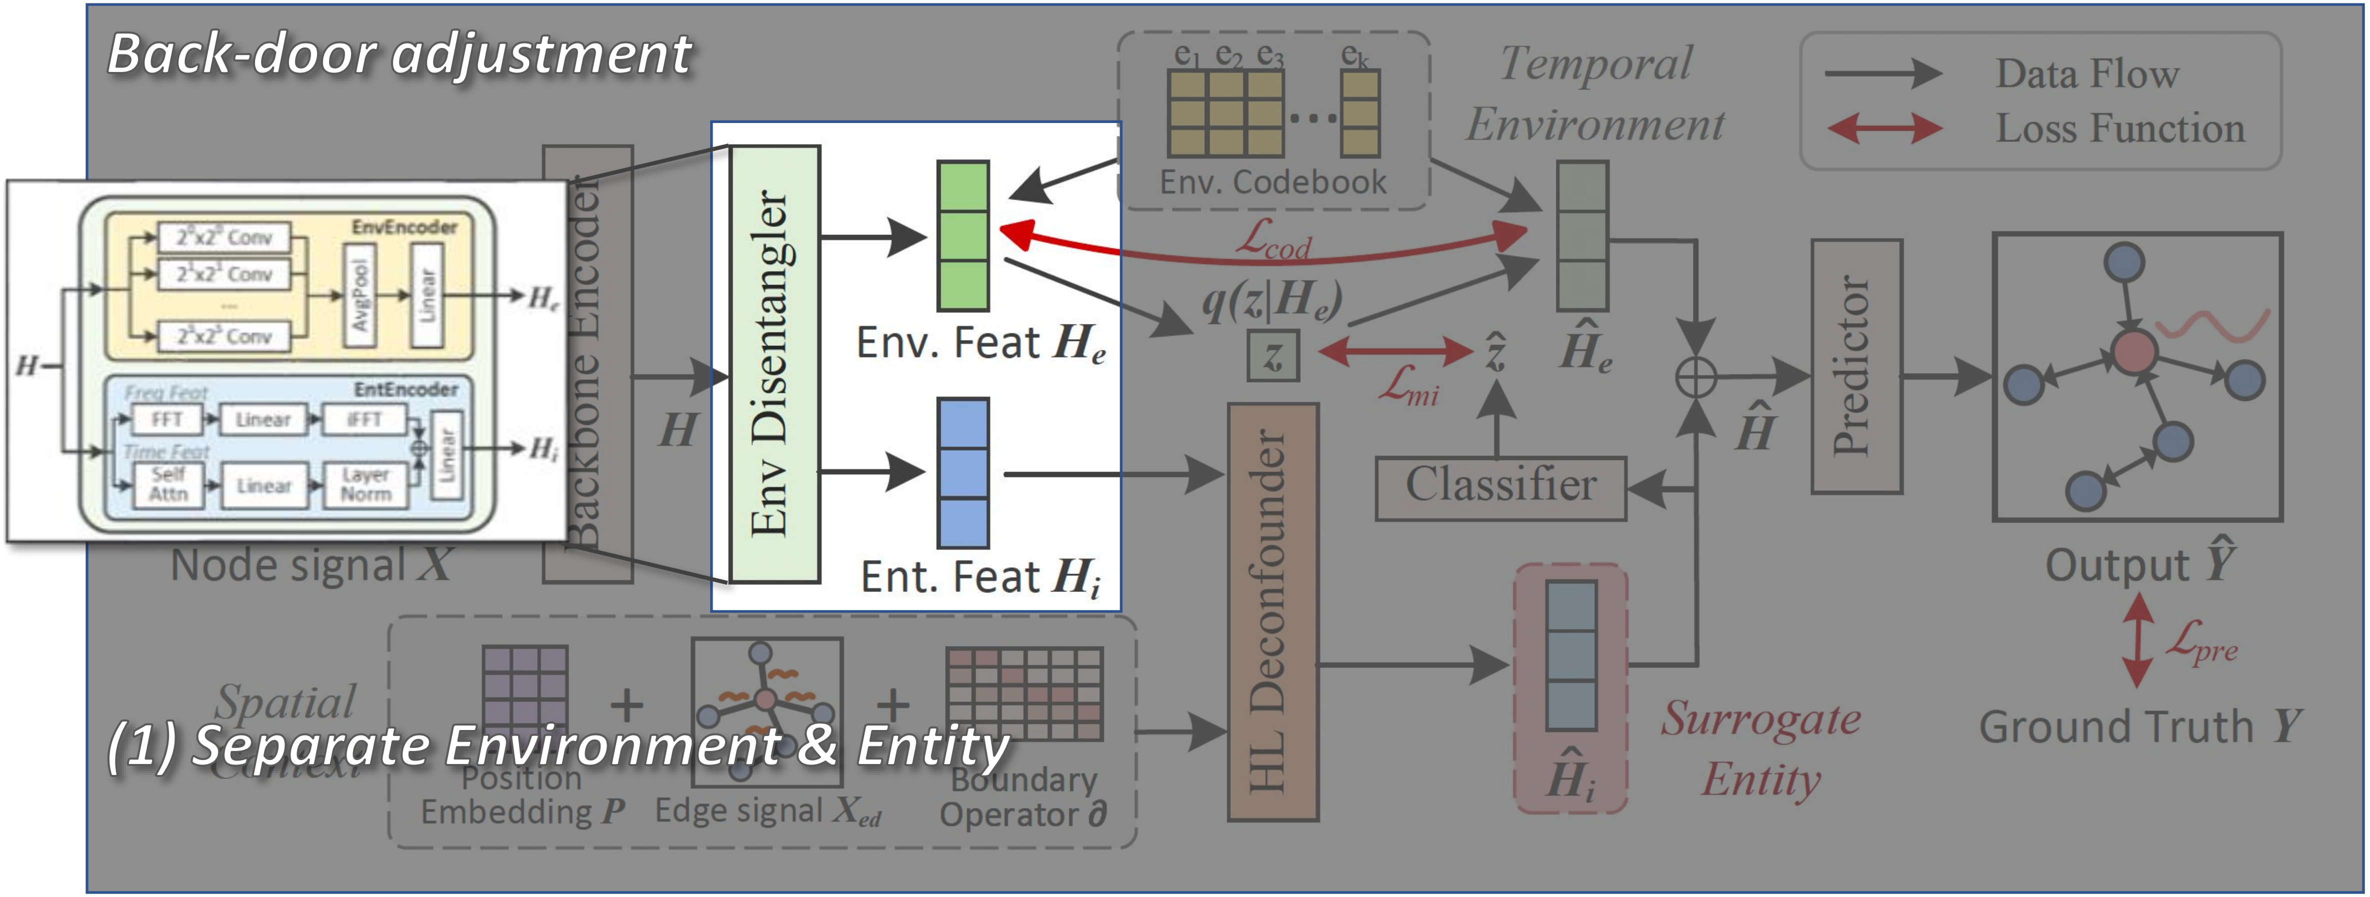


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



## Causal Spatio-Temporal neural network (CaST)

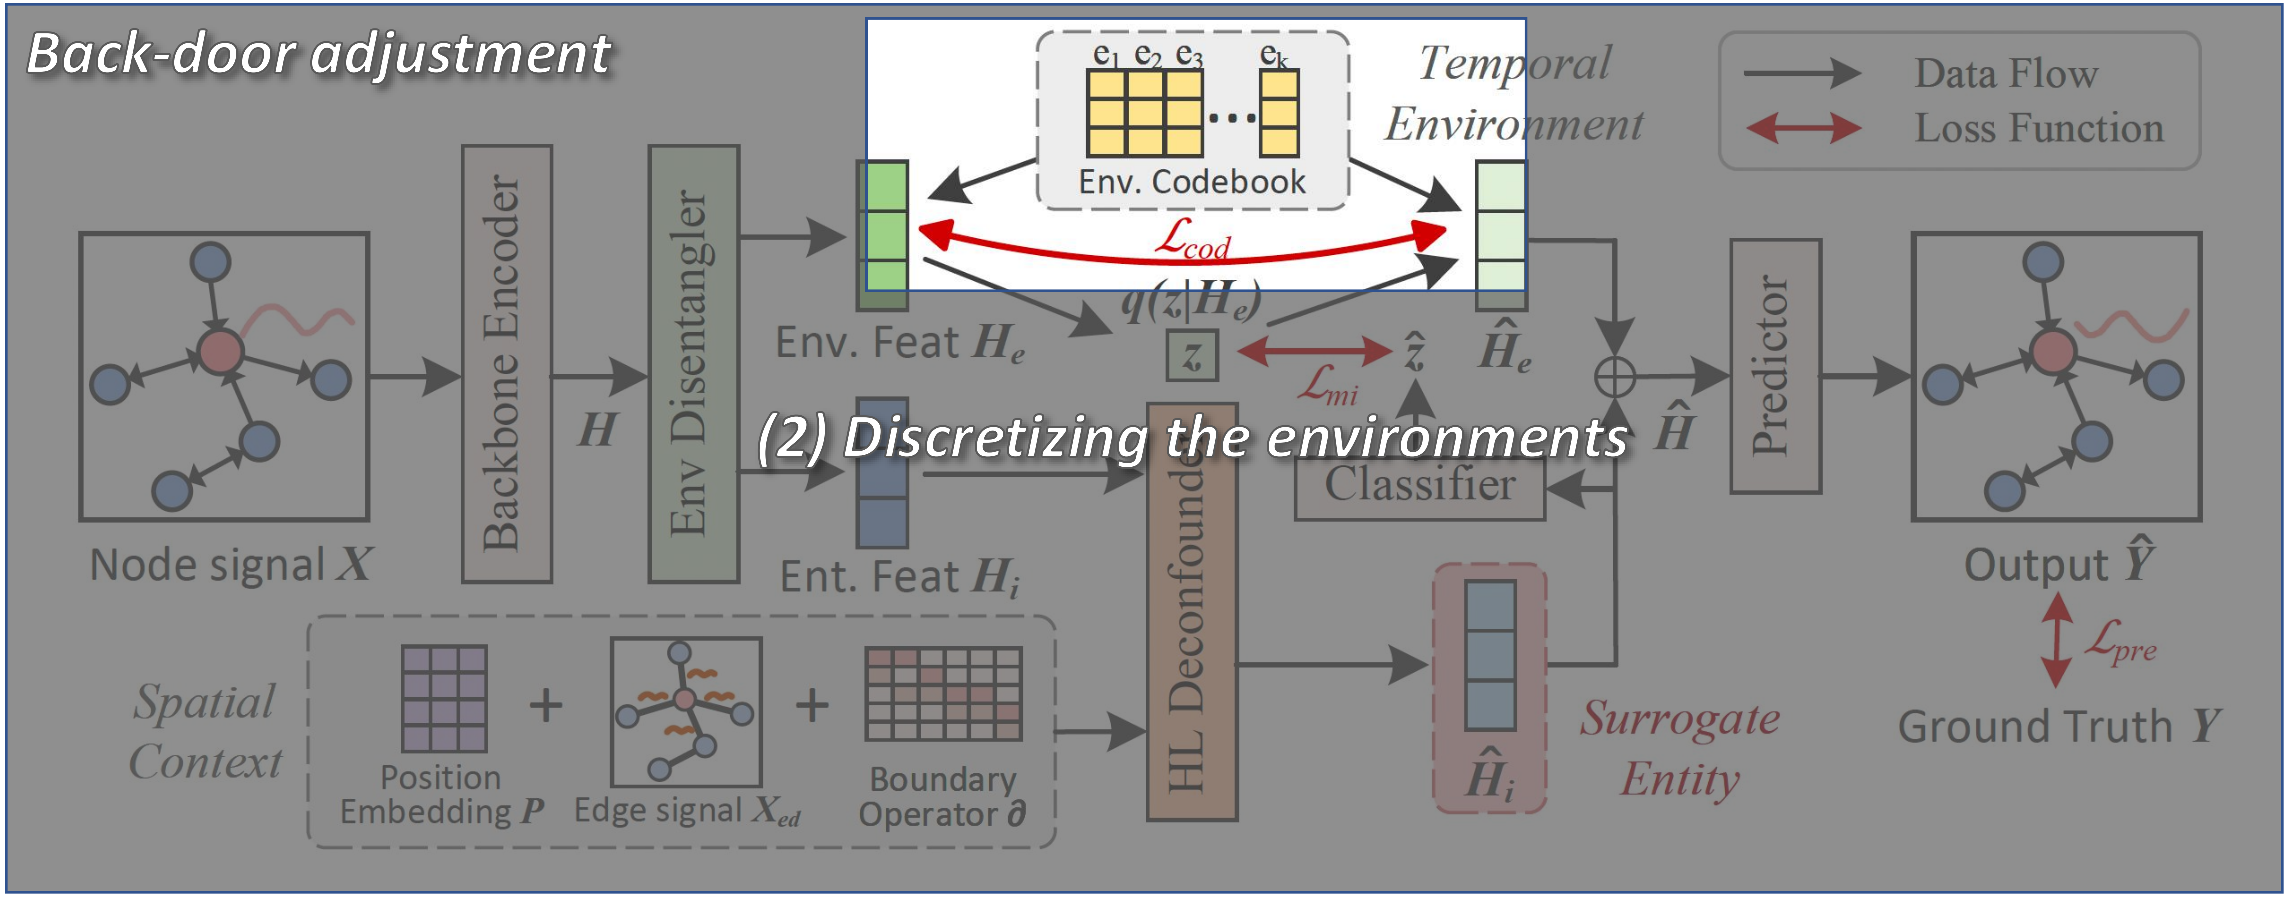


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

## Causal Spatio-Temporal neural network (CaST)

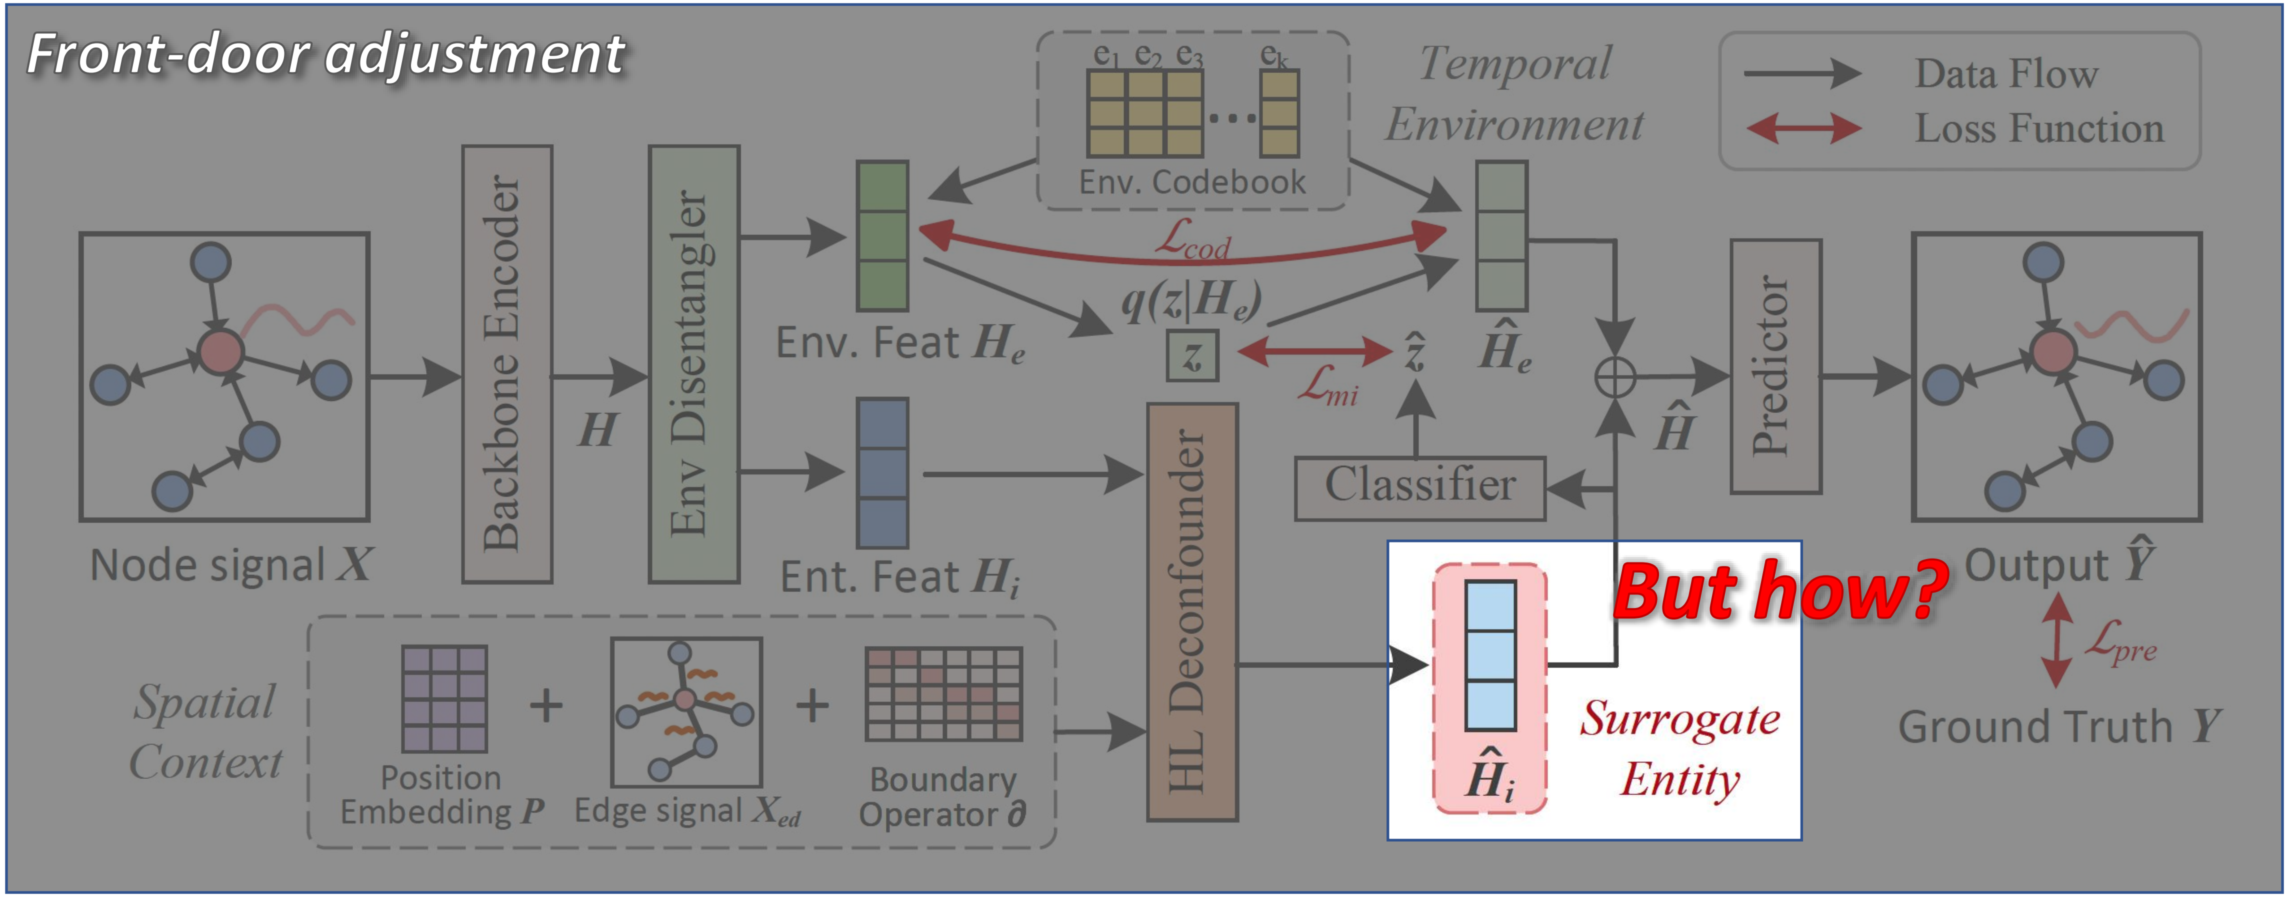


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



- To achieve **Back-door adjustment** …

$$\sum_e P(Y|X, E = e)P(E = e)$$

1) Separate the entity and the environments
2) Discretizing the environment

- To achieve **Front-door adjustment** …
  - Obtain a surrogate entity – ***But how?***
    - **Data/task-specific challenge** – Causation's ripple effect
    - Solution – Graph convolution networks?



**<u>Y. Xia</u>**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



- To achieve **Back-door adjustment** …

$$\sum_{e} P(Y|X, E = e)P(E = e)$$

1) Separate the entity and the environments
2) Discretizing the environment

- To achieve **Front-door adjustment** …
  - Obtain a surrogate entity – ***But how?***
    - **Data/task-specific challenge** – Causation's ripple effect
    - Solution – ~~Graph~~ Edge convolution networks



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

## Causal Spatio-Temporal neural network (CaST)

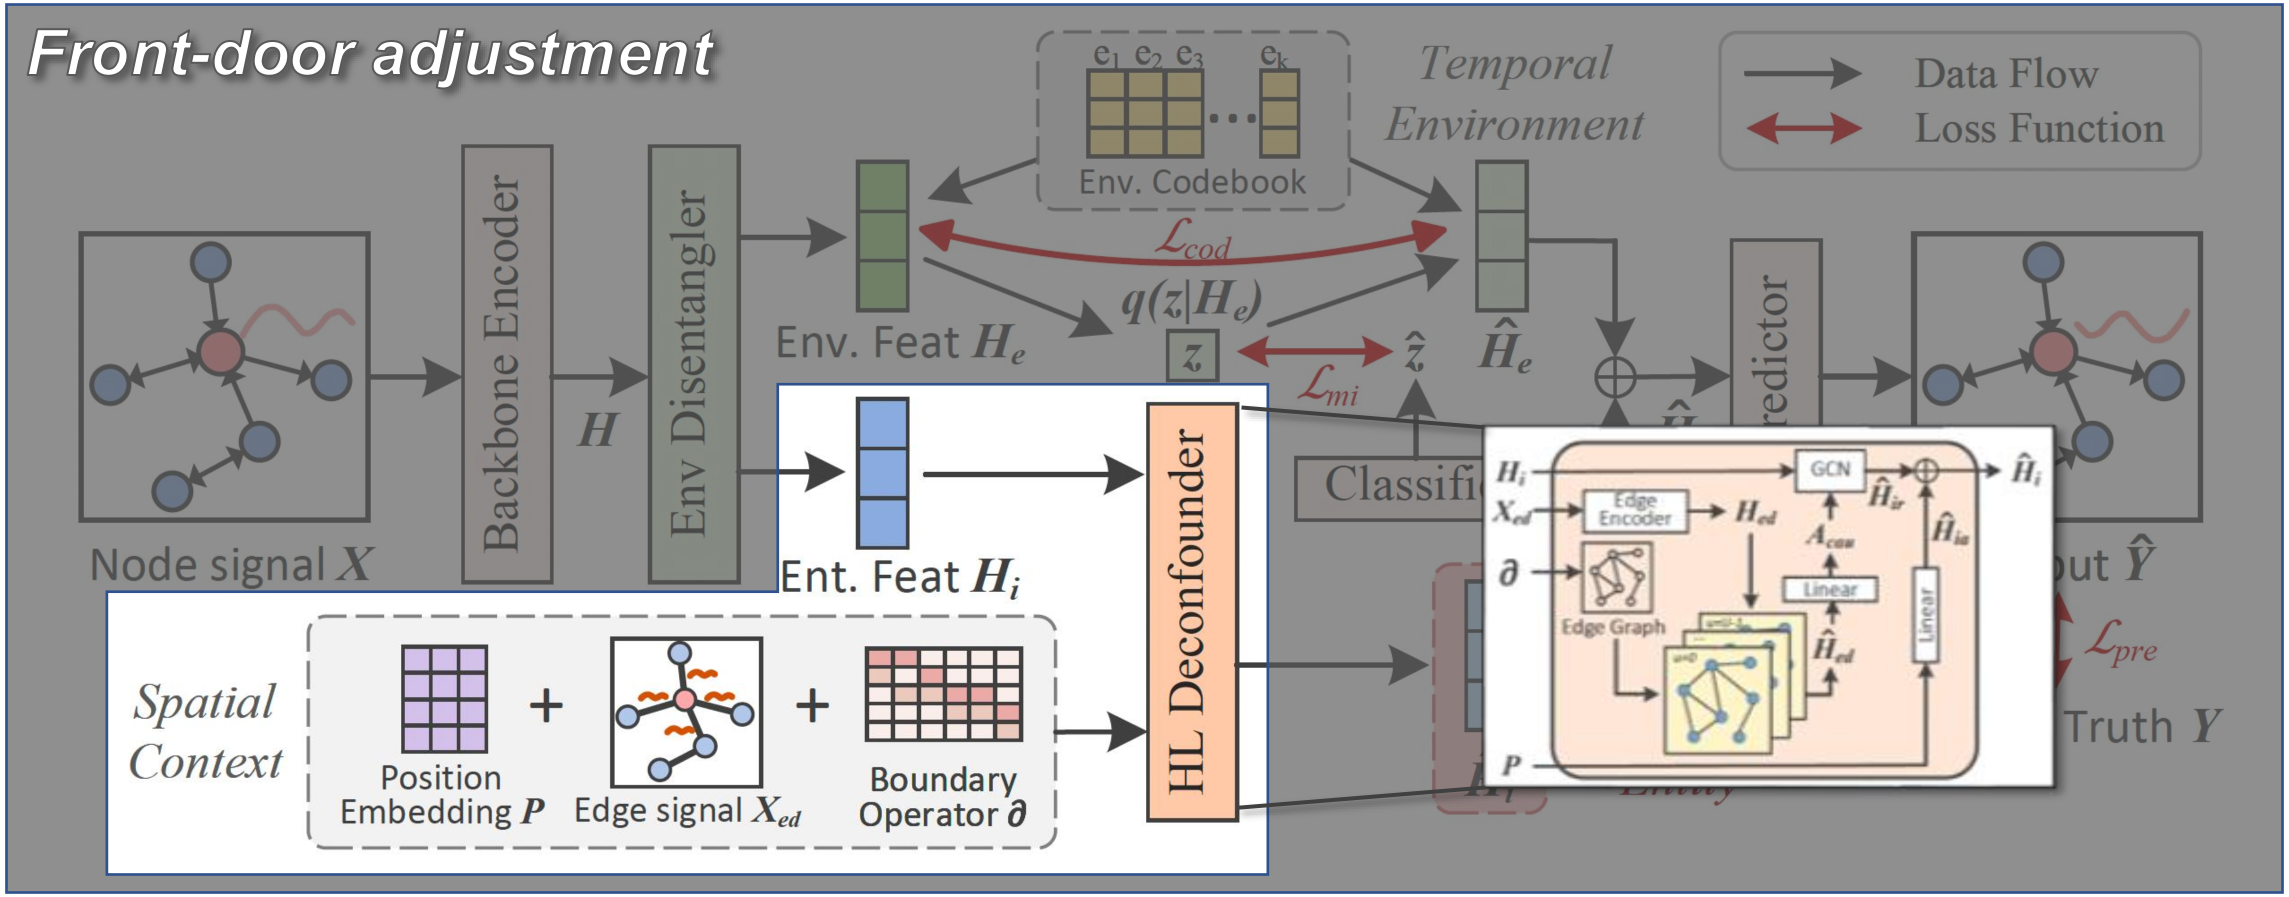


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

## Causal Spatio-Temporal neural network (CaST)

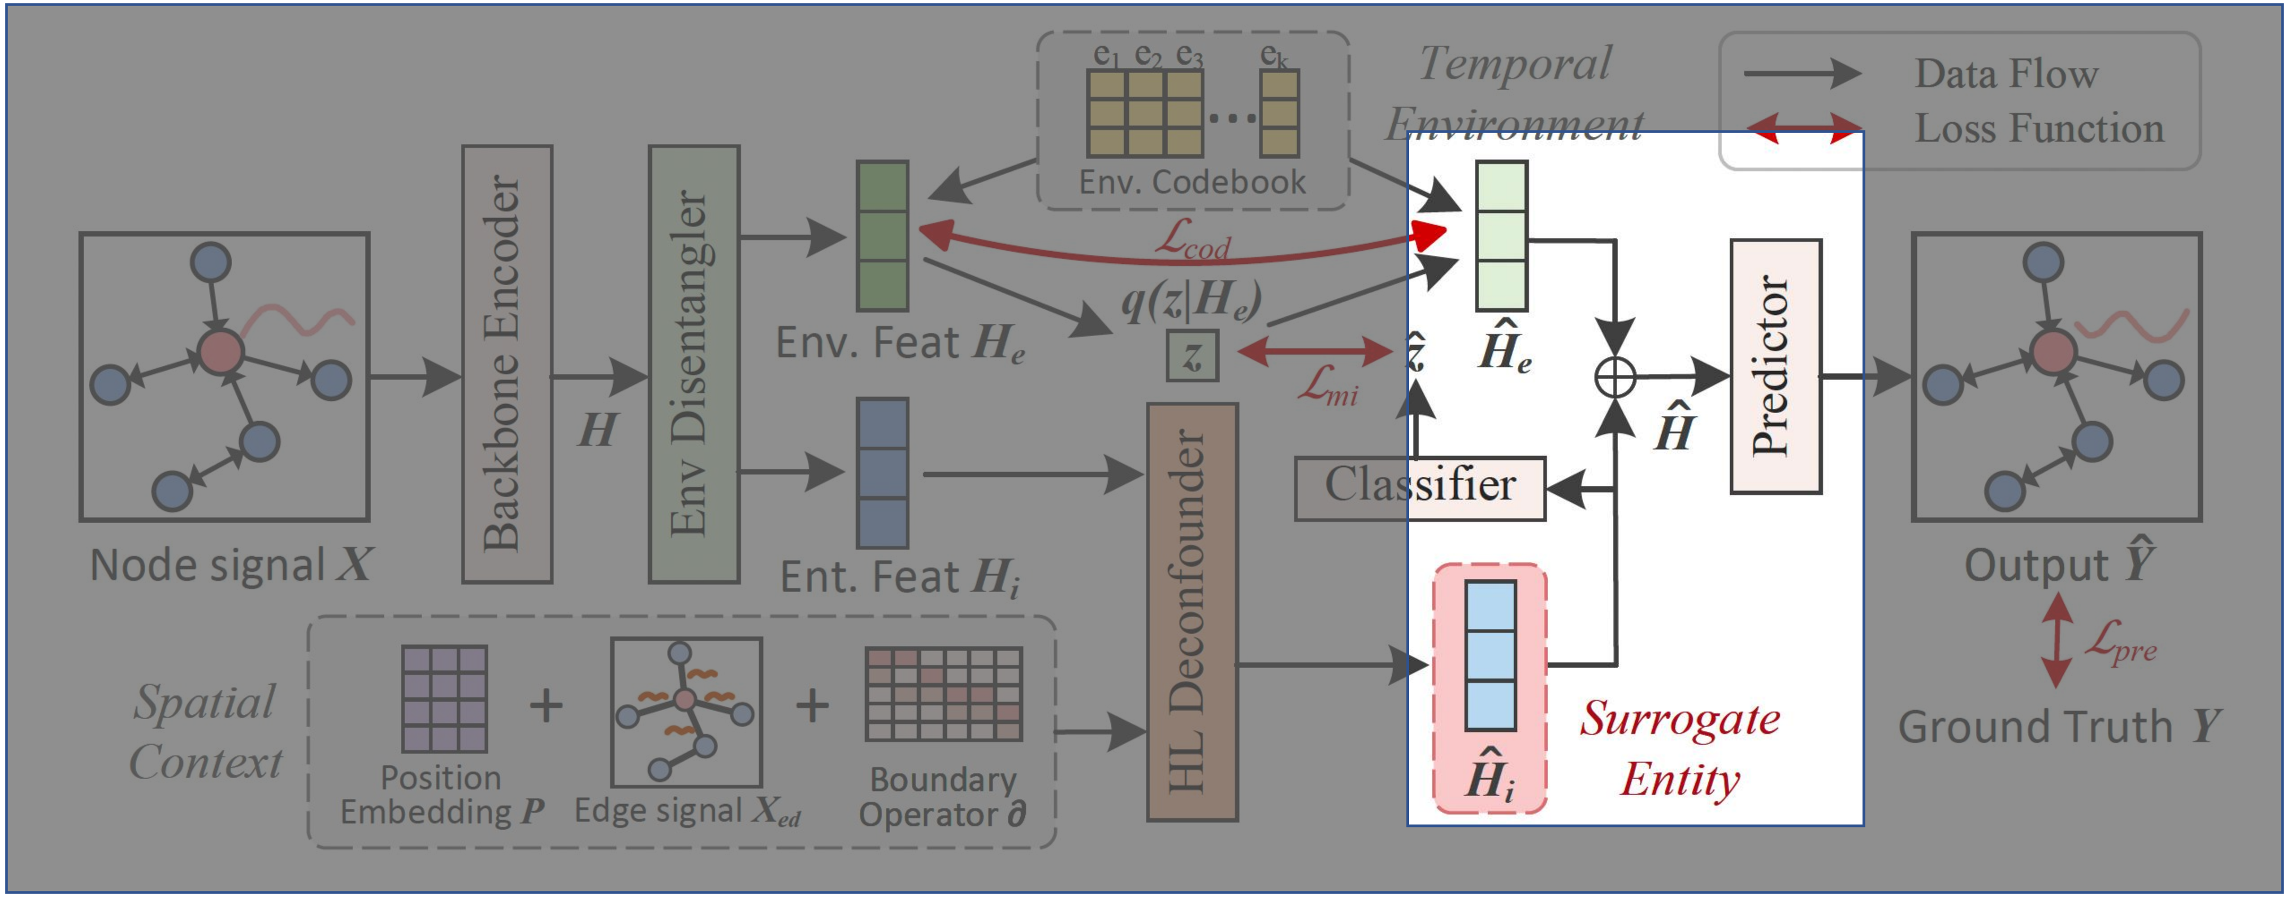


Figure 3: The pipeline of CaST. Env: Environment. Ent: Entity. Feat: Feature.



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



- Datasets: PEMS08, AIR-BJ, AIR-GZ
- Experiment settings: predict over next 24 steps given past 24 steps
- Evaluation metrics: MAE, RMSE

Table 1: 5-run error comparison. The bold/underlined font means the best/the second-best result.

| Model | PEMS08 (24→24) | | AIR-BJ (24→24) | | AIR-GZ (24→24) | |
|---|---|---|---|---|---|---|
| | MAE | RMSE | MAE | RMSE | MAE | RMSE |
| HA(2017) | 58.83 | 81.96 | 32.12 | 43.95 | 19.56 | 25.77 |
| VAR(1991) | 37.04 | 53.08 | 29.79 | 42.04 | 14.97 | 20.61 |
| DCRNN(2017) | 22.10 ± 0.45 | 33.96 ± 0.59 | 23.72 ± 0.36 | 35.84 ± 0.56 | 12.99 ± 0.26 | 18.27 ± 0.41 |
| STGCN(2018) | 18.60 ± 0.08 | 28.44 ± 0.15 | 23.71 ± 0.21 | 36.30 ± 0.58 | 12.69 ± 0.04 | 17.66 ± 0.09 |
| ASTGCN(2019) | 20.36 ± 0.48 | 30.87 ± 0.55 | 23.78 ± 0.22 | 35.91 ± 0.11 | 12.91 ± 0.15 | 18.02 ± 0.27 |
| MTGNN(2020) | 18.13 ± 0.10 | 28.85 ± 0.12 | 24.35 ± 0.74 | 38.97 ± 1.81 | 12.43 ± 0.11 | 17.99 ± 0.18 |
| AGCRN(2020) | 17.06 ± 0.14 | 26.80 ± 0.15 | 23.43 ± 0.29 | 35.66 ± 0.57 | 12.74 ± 0.01 | 17.49 ± 0.01 |
| GMSDR(2022) | 18.34 ± 0.68 | 28.36 ± 1.01 | 25.92 ± 0.52 | 39.60 ± 0.44 | 13.47 ± 0.31 | 19.04 ± 0.46 |
| STGNCDE(2022) | 17.55 ± 0.30 | 27.28 ± 0.36 | 24.35 ± 0.31 | 35.91 ± 0.48 | 13.70 ± 0.10 | 19.15 ± 0.07 |
| CaST (ours) | **16.44** ± 0.10 | **26.61** ± 0.15 | **22.90** ± 0.09 | **34.84** ± 0.11 | **12.36** ± 0.01 | **17.25** ± 0.05 |

***Traffic Flow*** ***Air Quality***

Y. Xia, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# Theory-Driven - *Causal Inference*



## Dynamic Spatial Causation

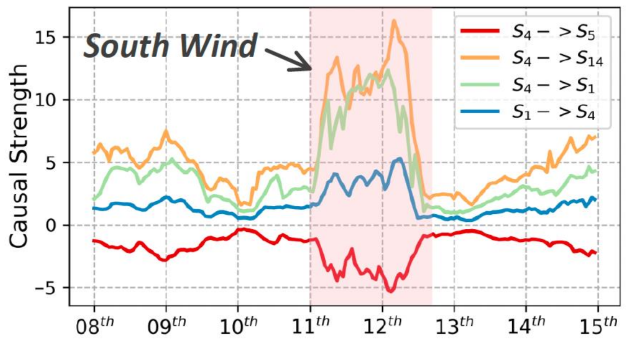


(b) Dynamic Causal Relations

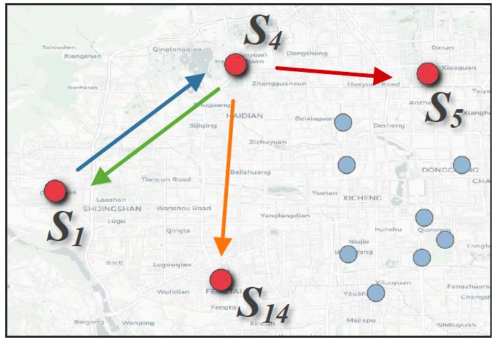


(c) Geographical Distribution

## Temporal Environments

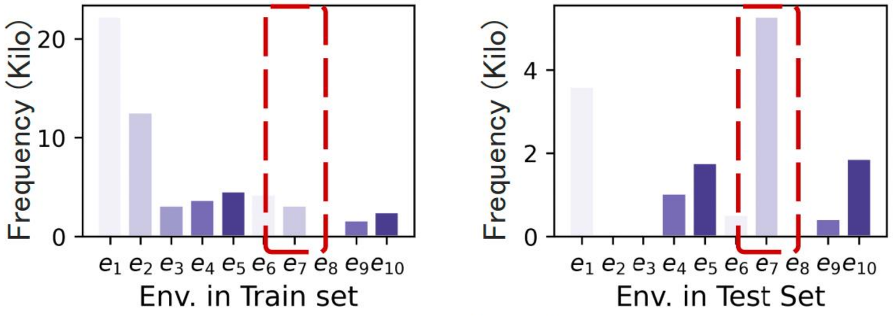


(b) Frequency of Environments

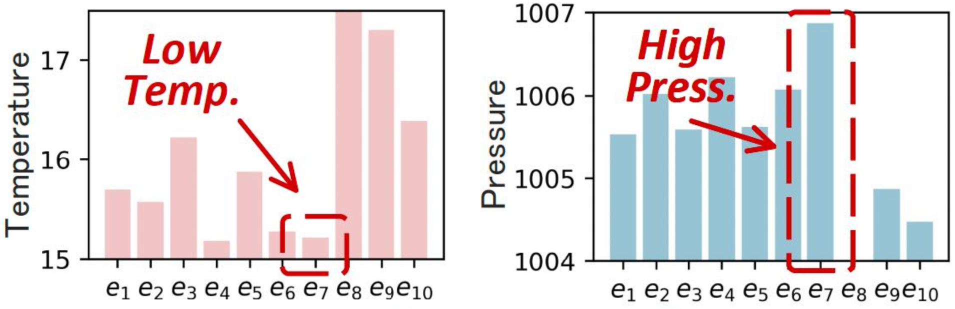


(c) External Factors of Environments



**Y. Xia**, Y. Liang et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.

# STG Forecasting Method

*How to learn* $\mathcal{F}(\cdot)$*?*

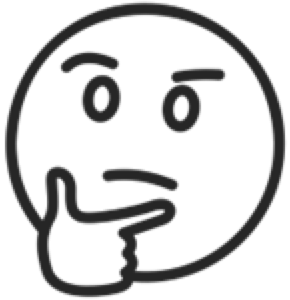

## Theory-Driven

- Causal Inference
- **Uncertainty Awareness**

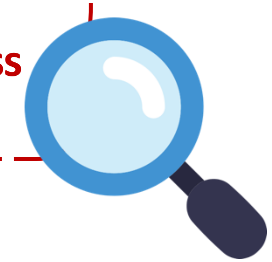

## Application-Driven

- Air Quality
- Traffic Flow
- Parking Availability

# Theory-Driven - *Uncertainty*

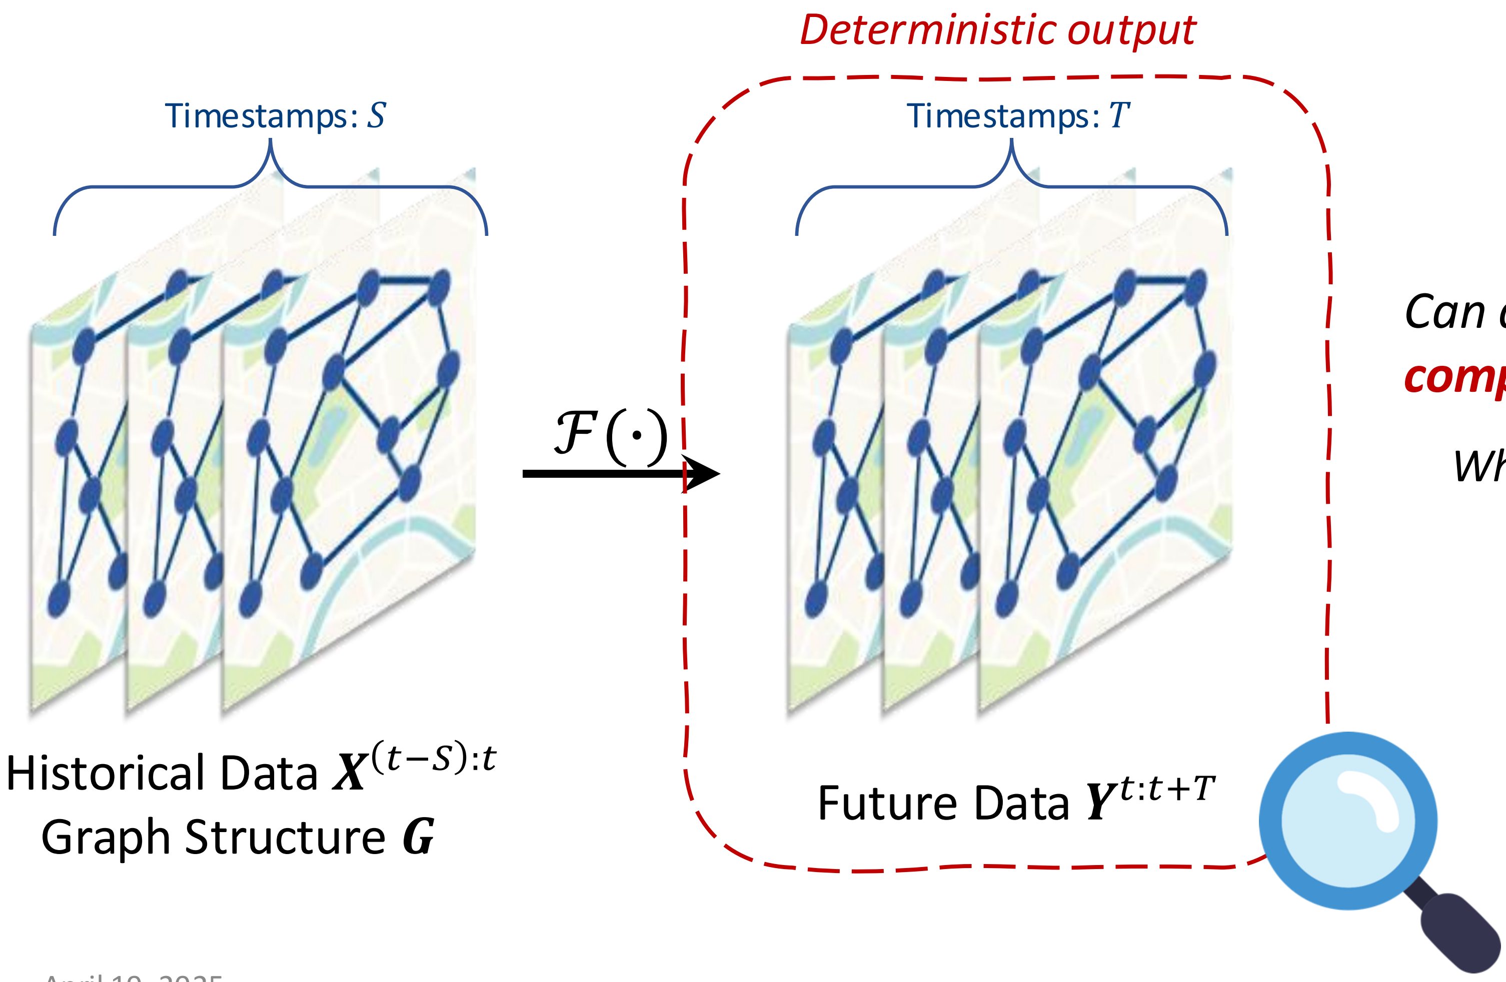


Urban dynamics are inherently **uncertain**.

*Can deterministic models handle real-world* ***complexity****?*

*What about …*

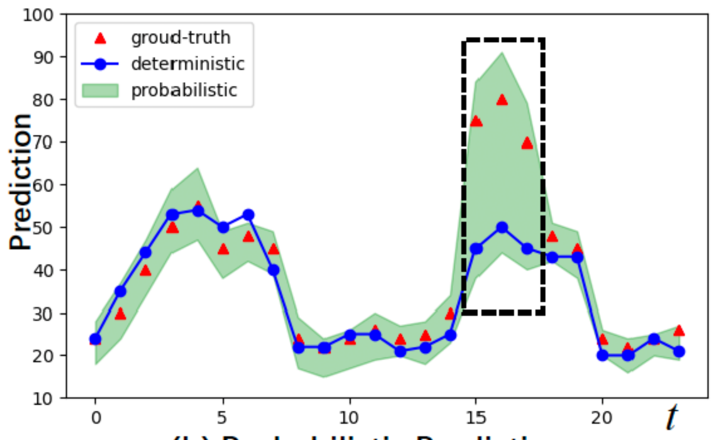


(b) Probabilistic Prediction

# Theory-Driven - *Uncertainty*



## Denoising diffusion probabilistic models (DDPM) *- A powerful generative model*

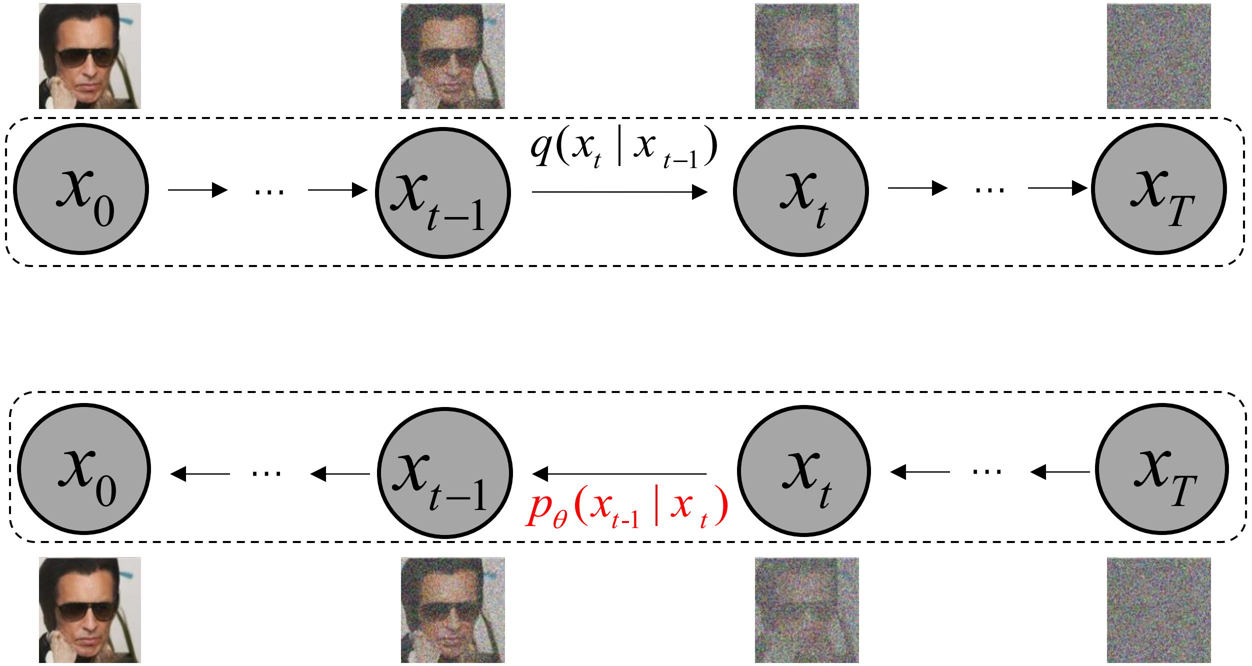


Forward Process
- Markov chain
- Fixed process
- add noise

Reverse Process
- Markow chain
- remove noise

J. Ho et al., Denoising diffusion probabilistic models. NeurIPS 2020.

# Theory-Driven - *Uncertainty*



1. How to generalize ***DDPM*** to stochastic STG forecasting?

2. How to capture the ***ST-correlation*** in $p_\theta$?

3. How to make it ***efficient*** in the reverse process?

# Theory-Driven - *Uncertainty*



*1. Generalize DDPM to STG forecasting -* ***DiffSTG***

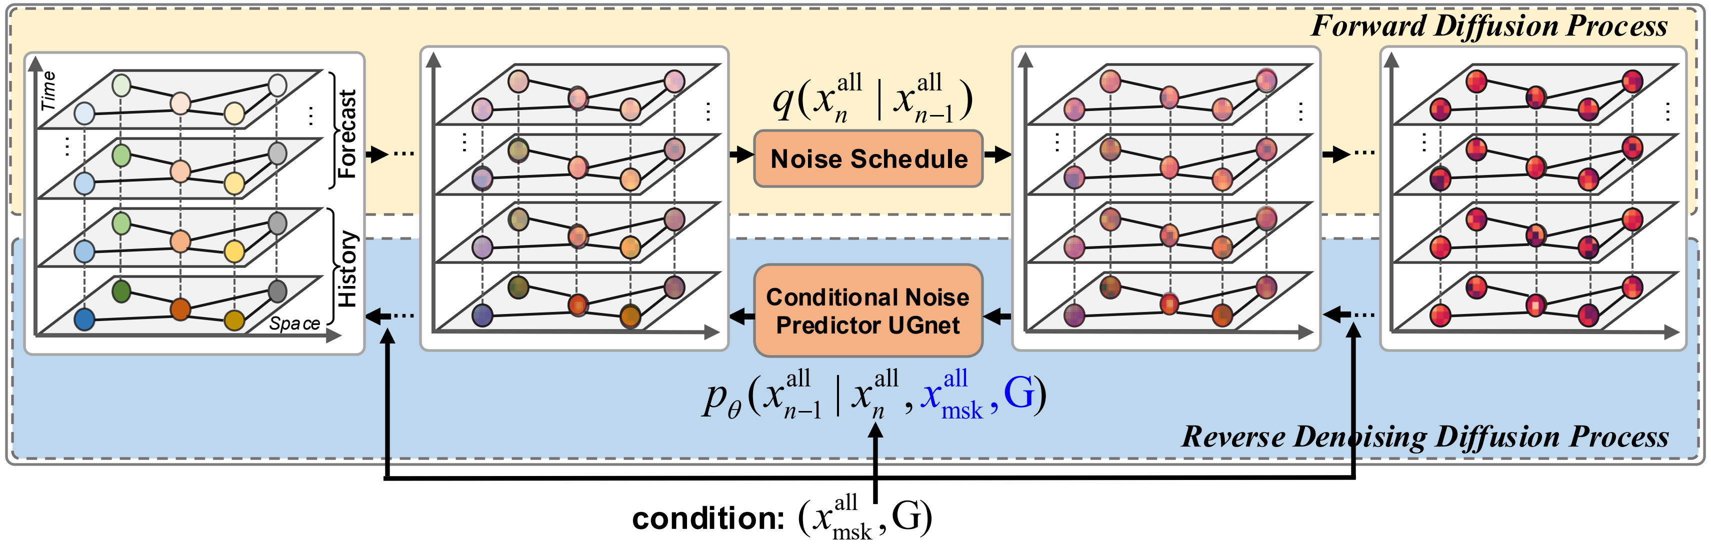


*2. Capture the ST-correlation in $p_\theta$ -* ***UGnet***

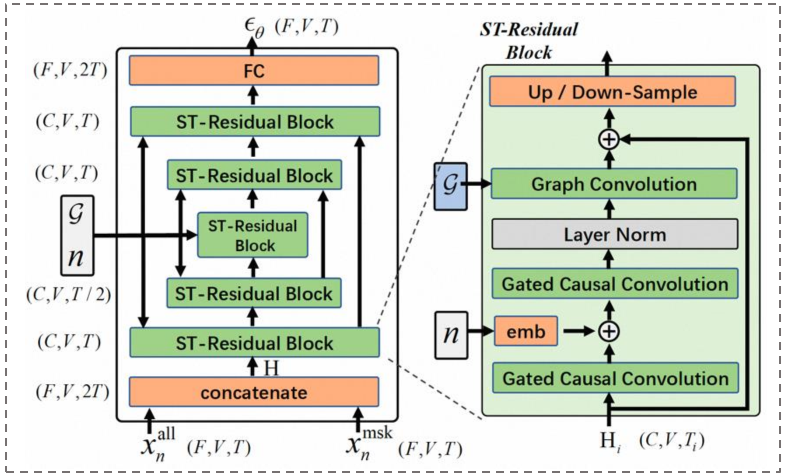


*3. Efficient in the reverse process –* ***One diffusion loop for all predictions + skip diffusion steps***

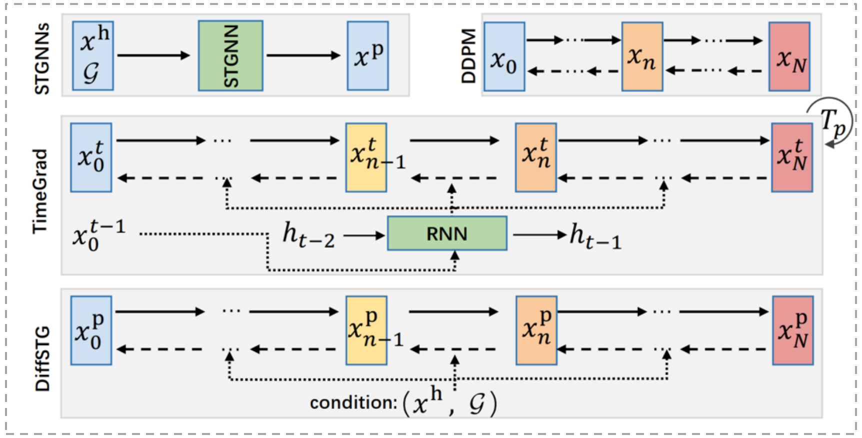




H. Wen, Lin Y, **Y. Xia** et al., DiffSTG: Probabilistic Spatio-Temporal Graph Forecasting with Denoising Diffusion Models. SIGSPATIAL'23.

# Theory-Driven - *Uncertainty*



## Datasets

| Dataset | Nodes | $F$ | Data Type | Time interval | #Samples |
|---|---|---|---|---|---|
| PEMS08 | 170 | 1 | Traffic flow | 5 minutes | 17,856 |
| AIR-BJ | 34 | 1 | $PM_{2.5}$ | 1 hour | 8,760 |
| AIR-GZ | 41 | 1 | $PM_{2.5}$ | 1 hour | 8,760 |

## Metrics

$$\mathrm{CRPS}(F, x) = \int_{\mathbb{R}} (F(z) - \mathbb{I}\{x \le z\})^2 \, \mathrm{d}z,$$

$$\mathrm{MAE}(Y, \hat{Y}) = \frac{1}{|Y|} \sum_{i=1}^{|Y|} \left| Y_i - \hat{Y}_i \right|,$$

## Results

| Method | AIR-BJ | | | AIR-GZ | | | PEMS08 | | |
|---|---|---|---|---|---|---|---|---|---|
| | MAE | RMSE | CRPS | MAE | RMSE | CRPS | MAE | RMSE | CRPS |
| Latent ODE [30] | 20.61 | 32.27 | 0.47 | 12.92 | 18.76 | 0.30 | 26.05 | 39.50 | 0.11 |
| DeepAR [31] | 20.15 | 32.09 | 0.37 | 11.77 | 17.45 | <u>0.23</u> | 21.56 | 33.37 | <u>0.07</u> |
| CSDI [37] | 26.52 | 40.33 | 0.50 | 13.75 | 19.40 | 0.28 | 32.11 | 47.40 | 0.11 |
| TimeGrad [26] | <u>18.64</u> | <u>31.86</u> | <u>0.36</u> | 12.36 | 18.15 | 0.25 | 24.46 | 38.06 | 0.09 |
| MC Dropout [44] | 20.80 | 40.54 | 0.45 | <u>11.12</u> | <u>17.07</u> | 0.25 | <u>19.01</u> | <u>29.35</u> | 0.07 |
| DiffSTG (ours) | **17.88** | **29.60** | **0.34** | **10.95** | **16.66** | **0.22** | **18.60** | **28.20** | **0.06** |
| Error reduction | -4.1% | -7.1% | -5.6% | -1.5% | -2.4% | -4.3% | -2.2% | -3.9% | -14.3% |

H. Wen, Lin Y, **<u>Y. Xia</u>** et al., DiffSTG: Probabilistic Spatio-Temporal Graph Forecasting with Denoising Diffusion Models. SIGSPATIAL'23.

# Outline

## Background

- What makes cities predictable?
- Spatio-Temporal (ST) Data & Properties

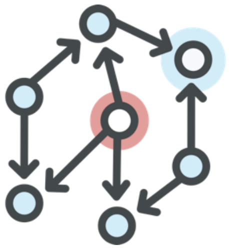

## Spatio-Temporal Graph Forecasting

- What is Spatio-Temporal Graph (STG)?
- What is STG forecasting?
- How we do it?
  - Application-Driven (Air Quality, Traffic, Parking)
  - Theory-Driven (Causality, Uncertainty)

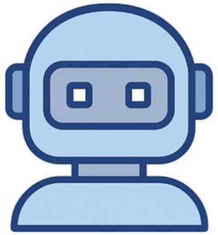

## Beyond Prediction: What's Next?

- LLMs-powered Agents & Causal Urban Insight

# Outline

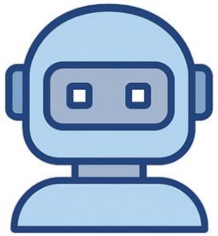

## Beyond Prediction: What's Next?

- LLMs-powered Agents & Causal Urban Insight

# The Road Ahead

*What next?*

*We don’t just want to predict cities — we want to understand them.*

**Beyond Prediction:**

Toward a More Intelligent
and Accessible Urban Causal Analysis

# Urban Causal Research


**Causal or coincidence?**

Well, I'll just order GrabFood instead.


*Do **food court closures** cause more people to **order delivery**?*


***Causal or coincidence?***


*Do **food court closures** cause more people to **order delivery**?*

**Confounding factors**: time of year, ongoing promotions, income levels, weather

This is where **causal inference** tools become necessary.

$$\underbrace{\boxed{y_i}}_{\text{Outcome}} = \overbrace{\boxed{T_i}}^{\text{Treatment}}\underbrace{\boxed{\beta_i}}_{\text{Causal Effect}} + \underbrace{\boxed{X_i}}_{\text{Control Variable}}\delta_i + \overbrace{\boxed{U_i}}^{\text{Unobersved Confounders}} + \underbrace{\boxed{e_i}}_{\text{Stochastic Error}}$$

# Urban Causal Research

Urban Phenomenon

+

Causal Inference


*Causal or coincidence?*


*Do **food court closures** cause more people to **order delivery**?*

**Confounding factors**: time of year, ongoing promotions, income levels, weather

This is where **causal inference** tools become necessary.

$$\underbrace{\boxed{y_i}}_{\text{Outcome}} = \overbrace{\boxed{T_i}}^{\text{Treatment}}\underbrace{\boxed{\beta_i}}_{\text{Causal Effect}} + \underbrace{\boxed{X_i}}_{\text{Control Variable}}\delta_i + \overbrace{\boxed{U_i}}^{\text{Unobersved Confounders}} + \boxed{e_i},\ \text{Stochastic Error}$$

# Current Landscape of Urban Causal Research

## Systematic Review

- Journal: *Cities*
- # total urban-related papers: **2,428** articles
- # total causal inference papers: **219** articles
- Timespan: **2012–2021**

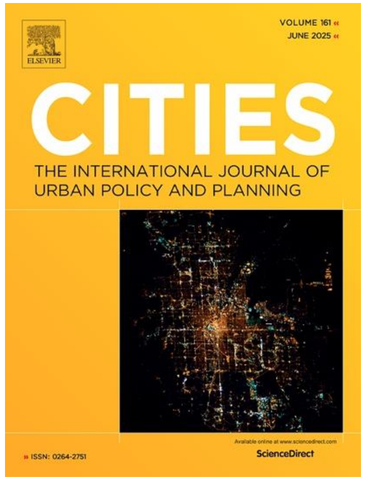


## Key Findings:

- **Trend**: growing adoption

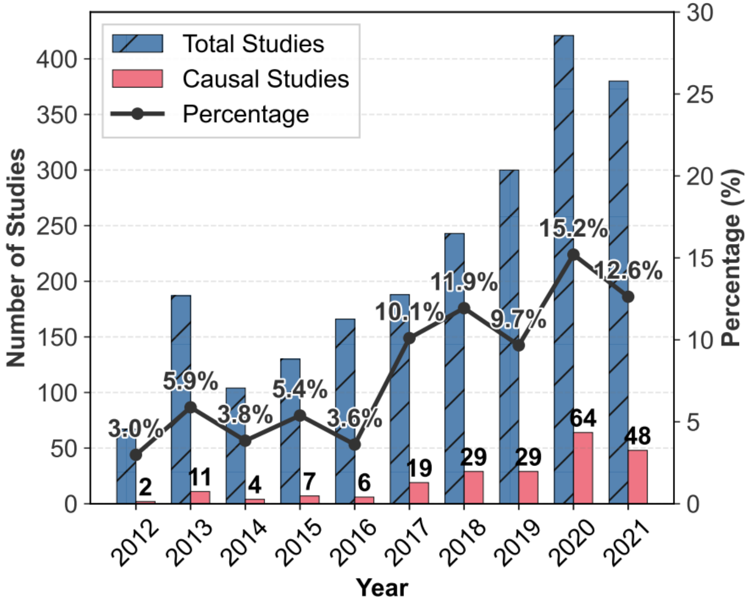




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# Curre

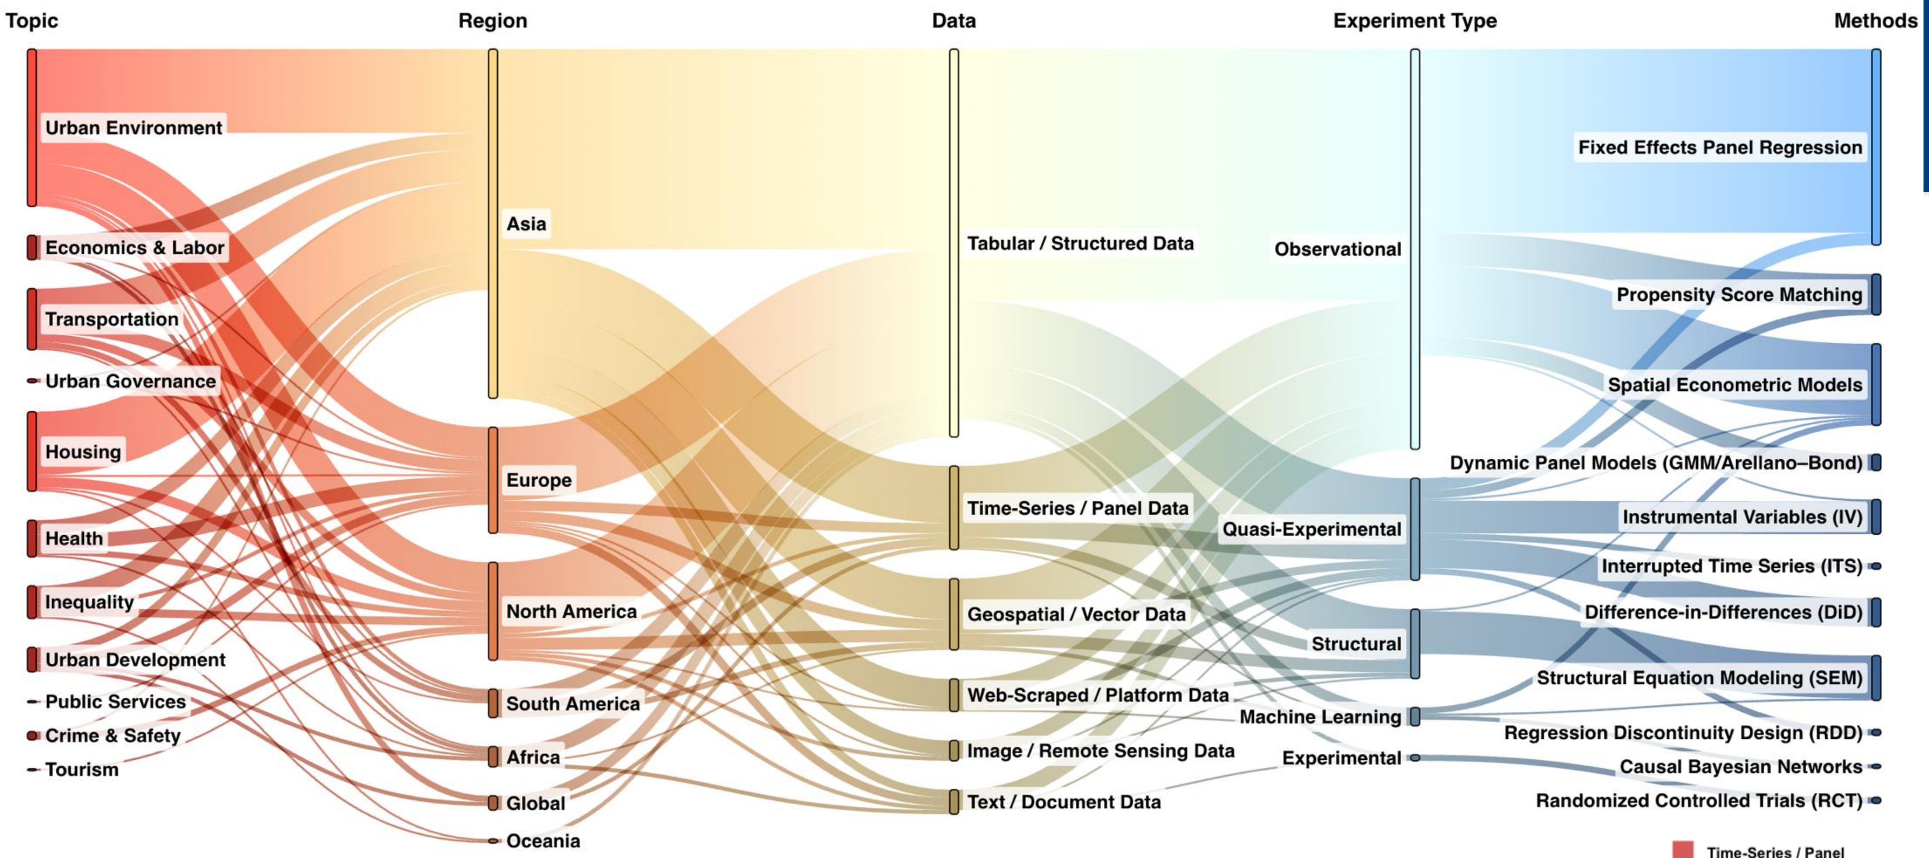


## Syste

•

•

•

•

## Key F

•

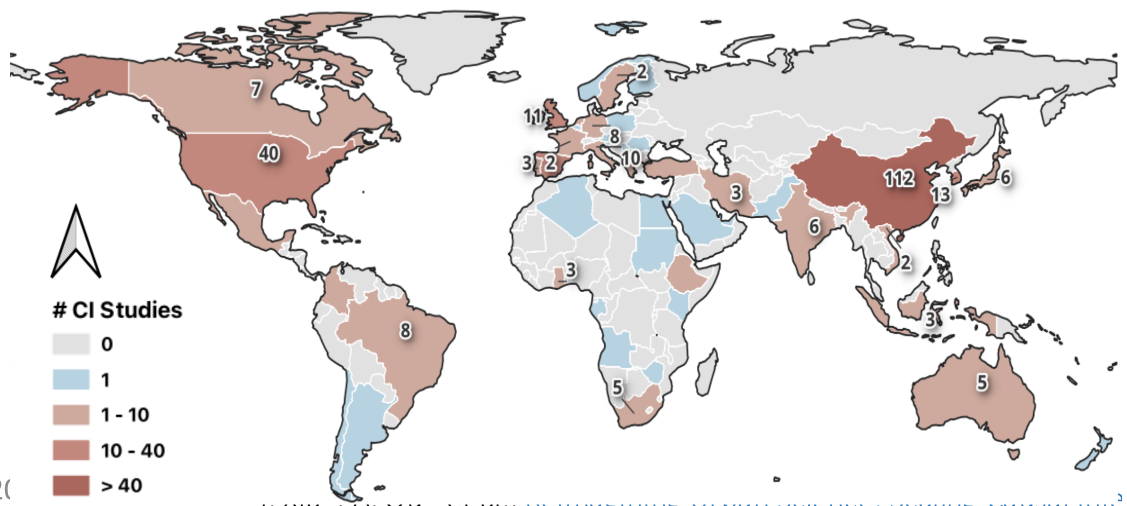


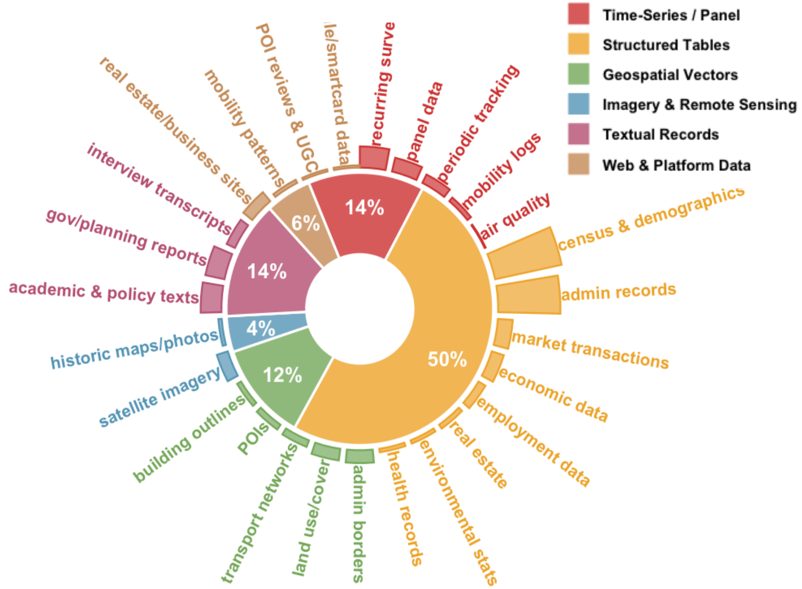

# Current Landscape of Urban Causal Research

## Systematic Review

- Journal: *Cities*
- # total urban-related papers: **2,428** articles
- # total causal inference papers: **219** articles
- Timespan: **2012–2021**

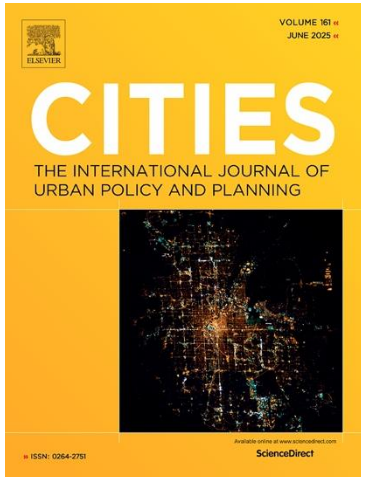


## Key Findings:

- **Trend**: growing adoption
- **Region**: geographical imbalances
- **Data**: a heavy reliance on structured data
- **Method**: limited methodological diversity
- **Code**: poor reproducibility



Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# Manual Urban Causal Research

## Manual Workflow Limitzations

**Phase 1**
Generating Hypotheses

- Challenges in timely and novel question formulation
- Subjectivity in hypothesis formation

**Phase 2**
Assembling Urban Data

- Multimodal, multisource, and multiresolution data complexity
- Repetitive manual data integration and cleaning

**Phase 3**
Executing CI Experiments

- Difficulty selecting the most suitable causal inferece methods
- Ensuring methodological rigor

**Phase 4**
Drawing Policy Insights

- Translating findings into generalizable, policy-relevant insights

# Manual Urban Causal Research

## Manual Workflow Limitzations

**Phase 1**
Generating Hypotheses

- Challenges in timely and novel question formulation
- Subjectivity in hypothesis formation

**Phase 2**
Assembling Urban Data

- Multimodal, multisource, and multiresolution data complexity
- Repetitive manual data integration and cleaning

**Phase 3**
Executing CI Experiments

- Difficulty selecting the most suitable causal inferece methods
- Ensuring methodological rigor

**Phase 4**
Drawing Policy Insights

- Translating findings into generalizable, policy-relevant insights

## *Multiple Sources*

**Geographical Data**
POI/AOI; Satellite Image; Street-view Image

**Traffic Data**
Trajectory; Road Network; Traffic Flow; Logistic

**Social Media Data**
Geo-textual Data; Geo-tagged Photo; Geo-tagged Video

**Demographic Data**
Population; Crime Data; Land Use Data

**Environment Data**
Meteorological Data; Greenery Data; Air Quality Data

## *Multiple Modalities*

**Spatio-temporal Modality**
POI; GPS Point; Time series; Trajectory...

**Visual Modality**

Satellite Image

Street-view Image

**Textual Modality**
Social Media Text:
Users' comment : *Arrived at an impressive park surrounded with...*
Geo-information Text:
*City A is located at the subtropics with ...*

**Other Modalities**
*Audio; Video; Hyper-spectrum..*



Zou, et al. (2025) Deep Learning for Cross-Domain Data Fusion in Urban Computing: Taxonomy, Advances, and Outlook. *Information Fusion*, *113*, 102606.

# Manual Urban Causal Research

NUS National University of Singapore

## Manual Workflow Limitzations

**Phase 1**
Generating Hypotheses

- Challenges in timely and novel question formulation
- Subjectivity in hypothesis formation

**Phase 2**
Assembling Urban Data

- Multimodal, multisource, and multiresolution data complexity
- Repetitive manual data integration and cleaning

**Phase 3**
Executing CI Experiments

- Difficulty selecting the most suitable causal inferece methods
- Ensuring methodological rigor

**Phase 4**
Drawing Policy Insights

- Translating findings into generalizable, policy-relevant insights

***Can LLMs Help?***

# LLMs & Agents

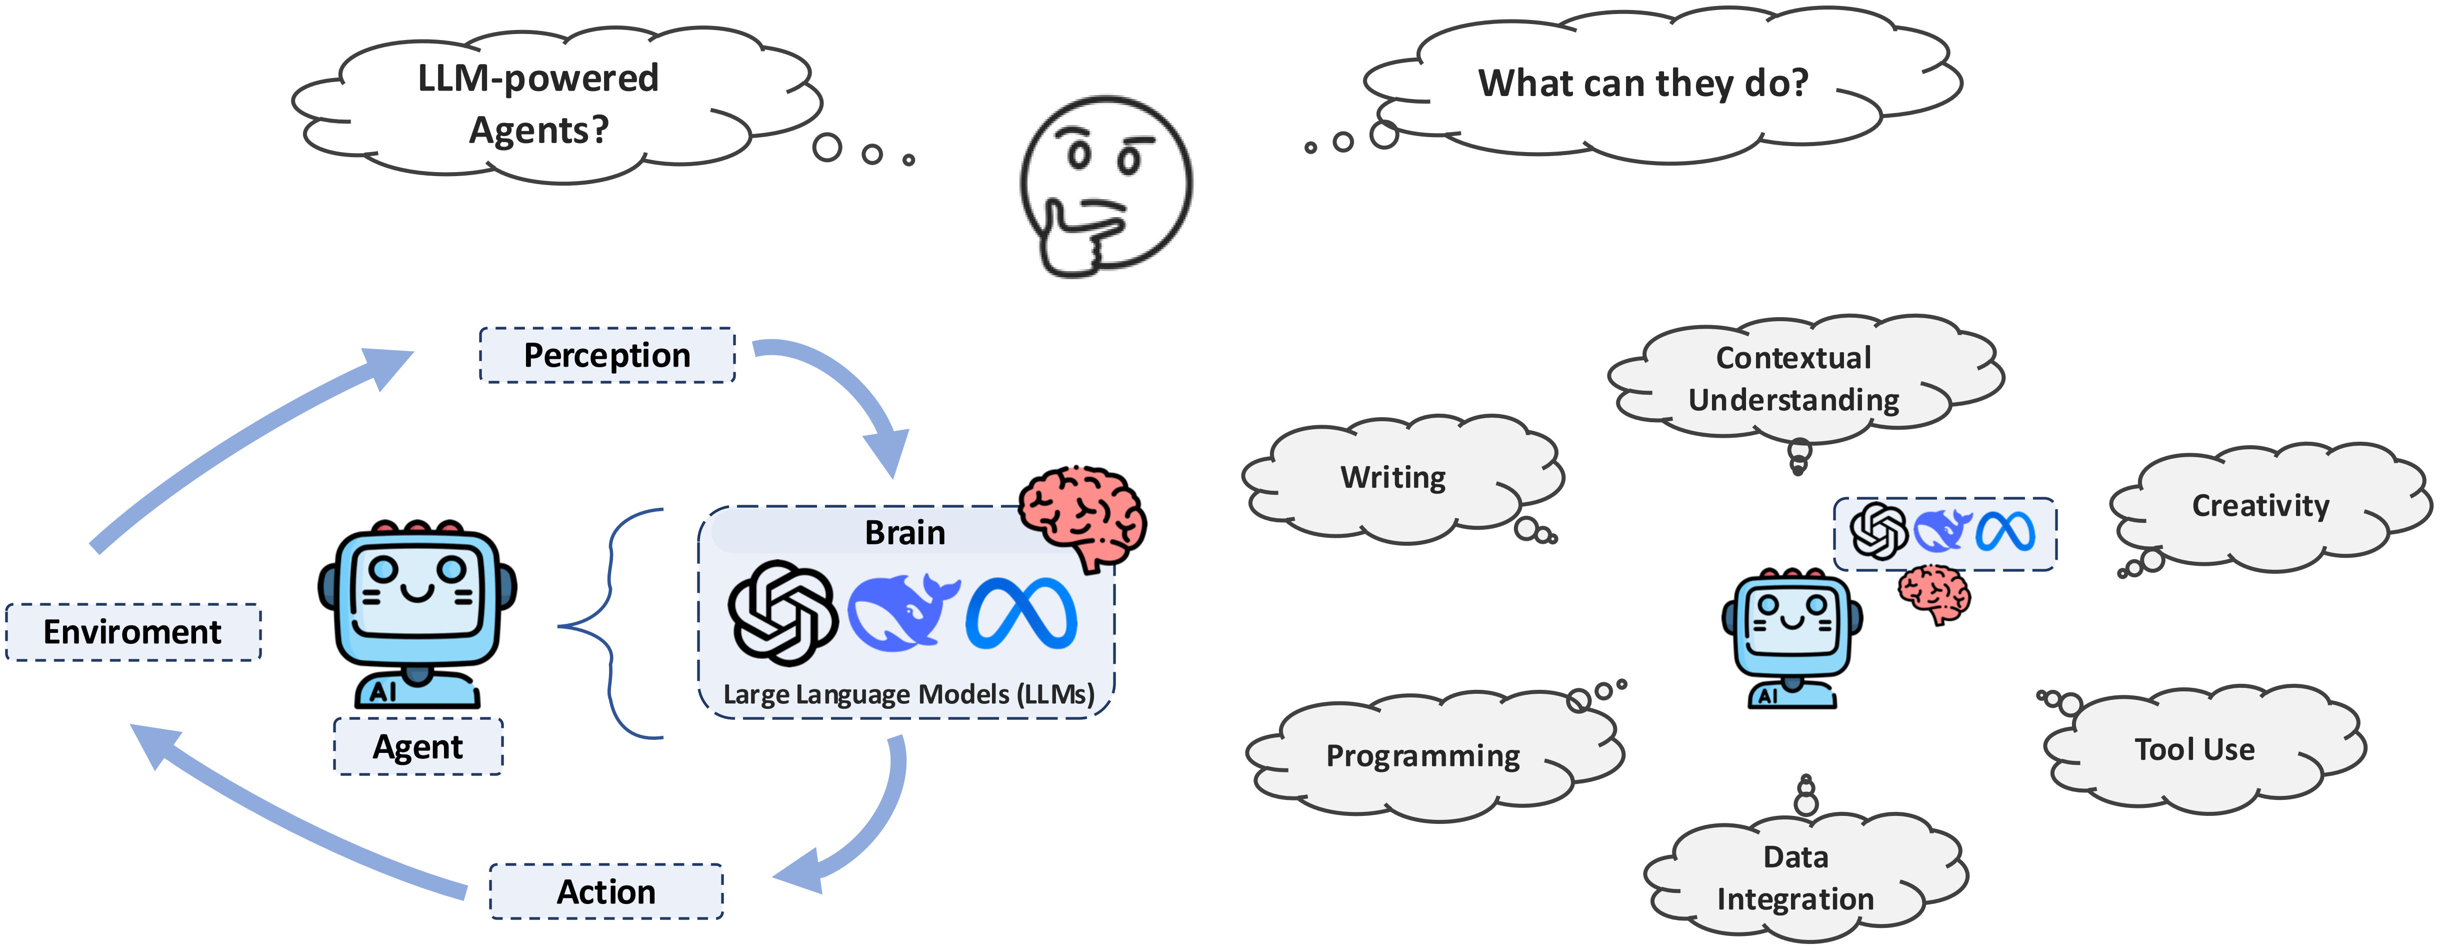

LLM-powered Agents?
What can they do?
Perception
Enviroment
Agent
Brain
Large Language Models (LLMs)
Action
Contextual Understanding
Writing
Creativity
Programming
Tool Use
Data Integration

# A Multi-Agents System - Smallville Town

What happens when putting **multiple agents** together in an environment?


J. Park et al., Generative Agents: Interactive Simulacra of Human Behavior. UIST 2023.

What if we created **a team of agents** to simulate an **urban research lab**?

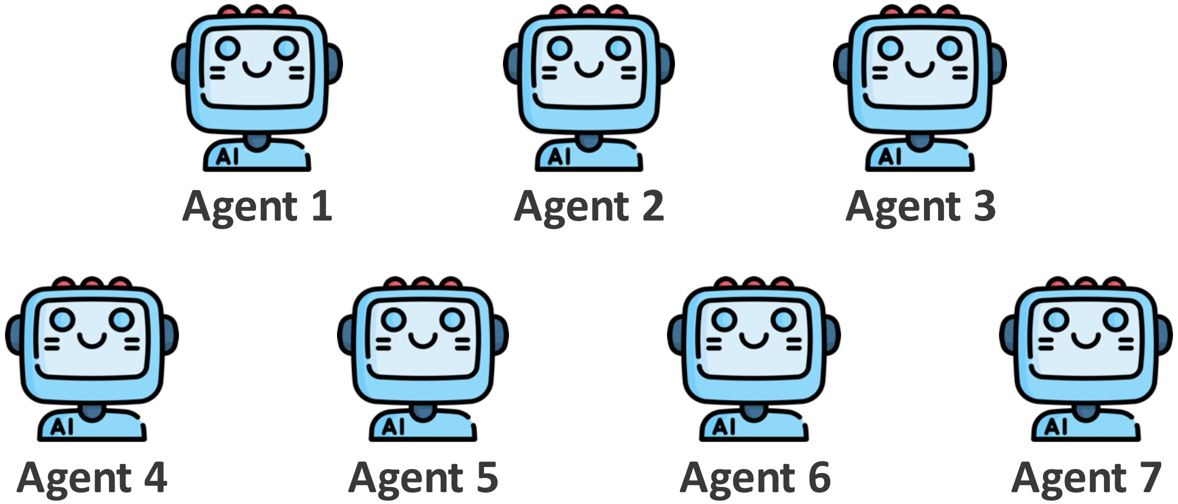

What if we created **a team of agents** to simulate an **urban research lab**?

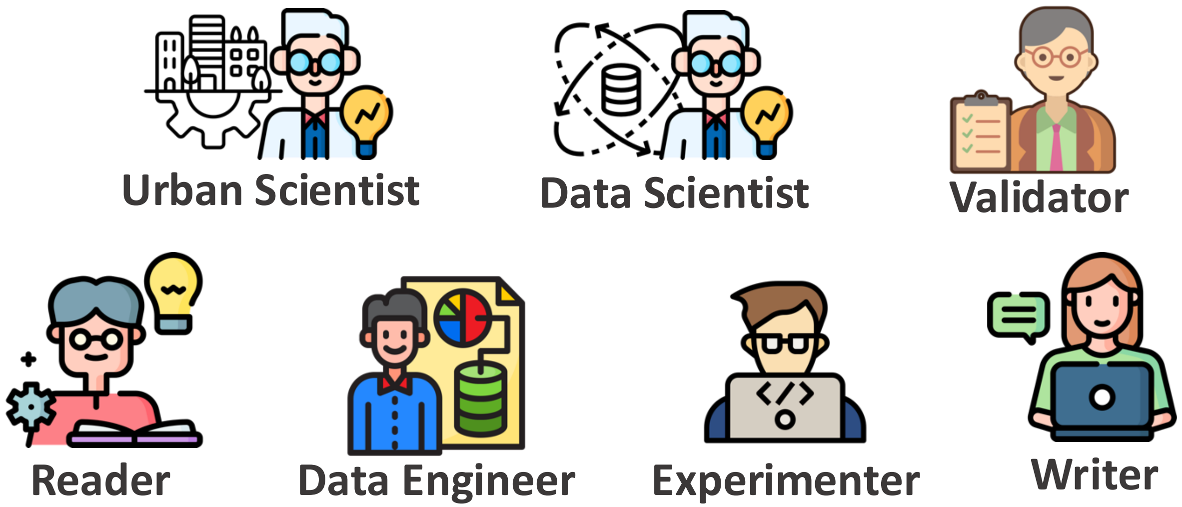

## Manual Workflow Limitzations

**Phase 1**
Generating Hypotheses
- Challenges in timely and novel question formulation
- Subjectivity in hypothesis formation

**Phase 2**
Assembling Urban Data
- Multimodal, multisource, and multiresolution data complexity
- Repetitive manual data integration and cleaning

**Phase 3**
Executing CI Experiments
- Difficulty selecting the most suitable causal inferece methods
- Ensuring methodological rigor

**Phase 4**
Drawing Policy Insights
- Translating findings into generalizable, policy-relevant insights



Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

## Manual Workflow Limitzations

**Phase 1**
Generating Hypotheses

- Challenges in timely and novel question formulation
- Subjectivity in hypothesis formation

**Phase 2**
Assembling Urban Data

- Multimodal, multisource, and multiresolution data complexity
- Repetitive manual data integration and cleaning

**Phase 3**
Executing CI Experiments

- Difficulty selecting the most suitable causal inferece methods
- Ensuring methodological rigor

**Phase 4**
Drawing Policy Insights

- Translating findings into generalizable, policy-relevant insights

*Motivate*

*Support*

## LLMs/MLLMs' Abilities

*Knowledge reasoning*

- Contextual Understanding
- Creativity
- Writing

*Procedural execution*

- Data Integration
- Programming
- Tool Use



Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

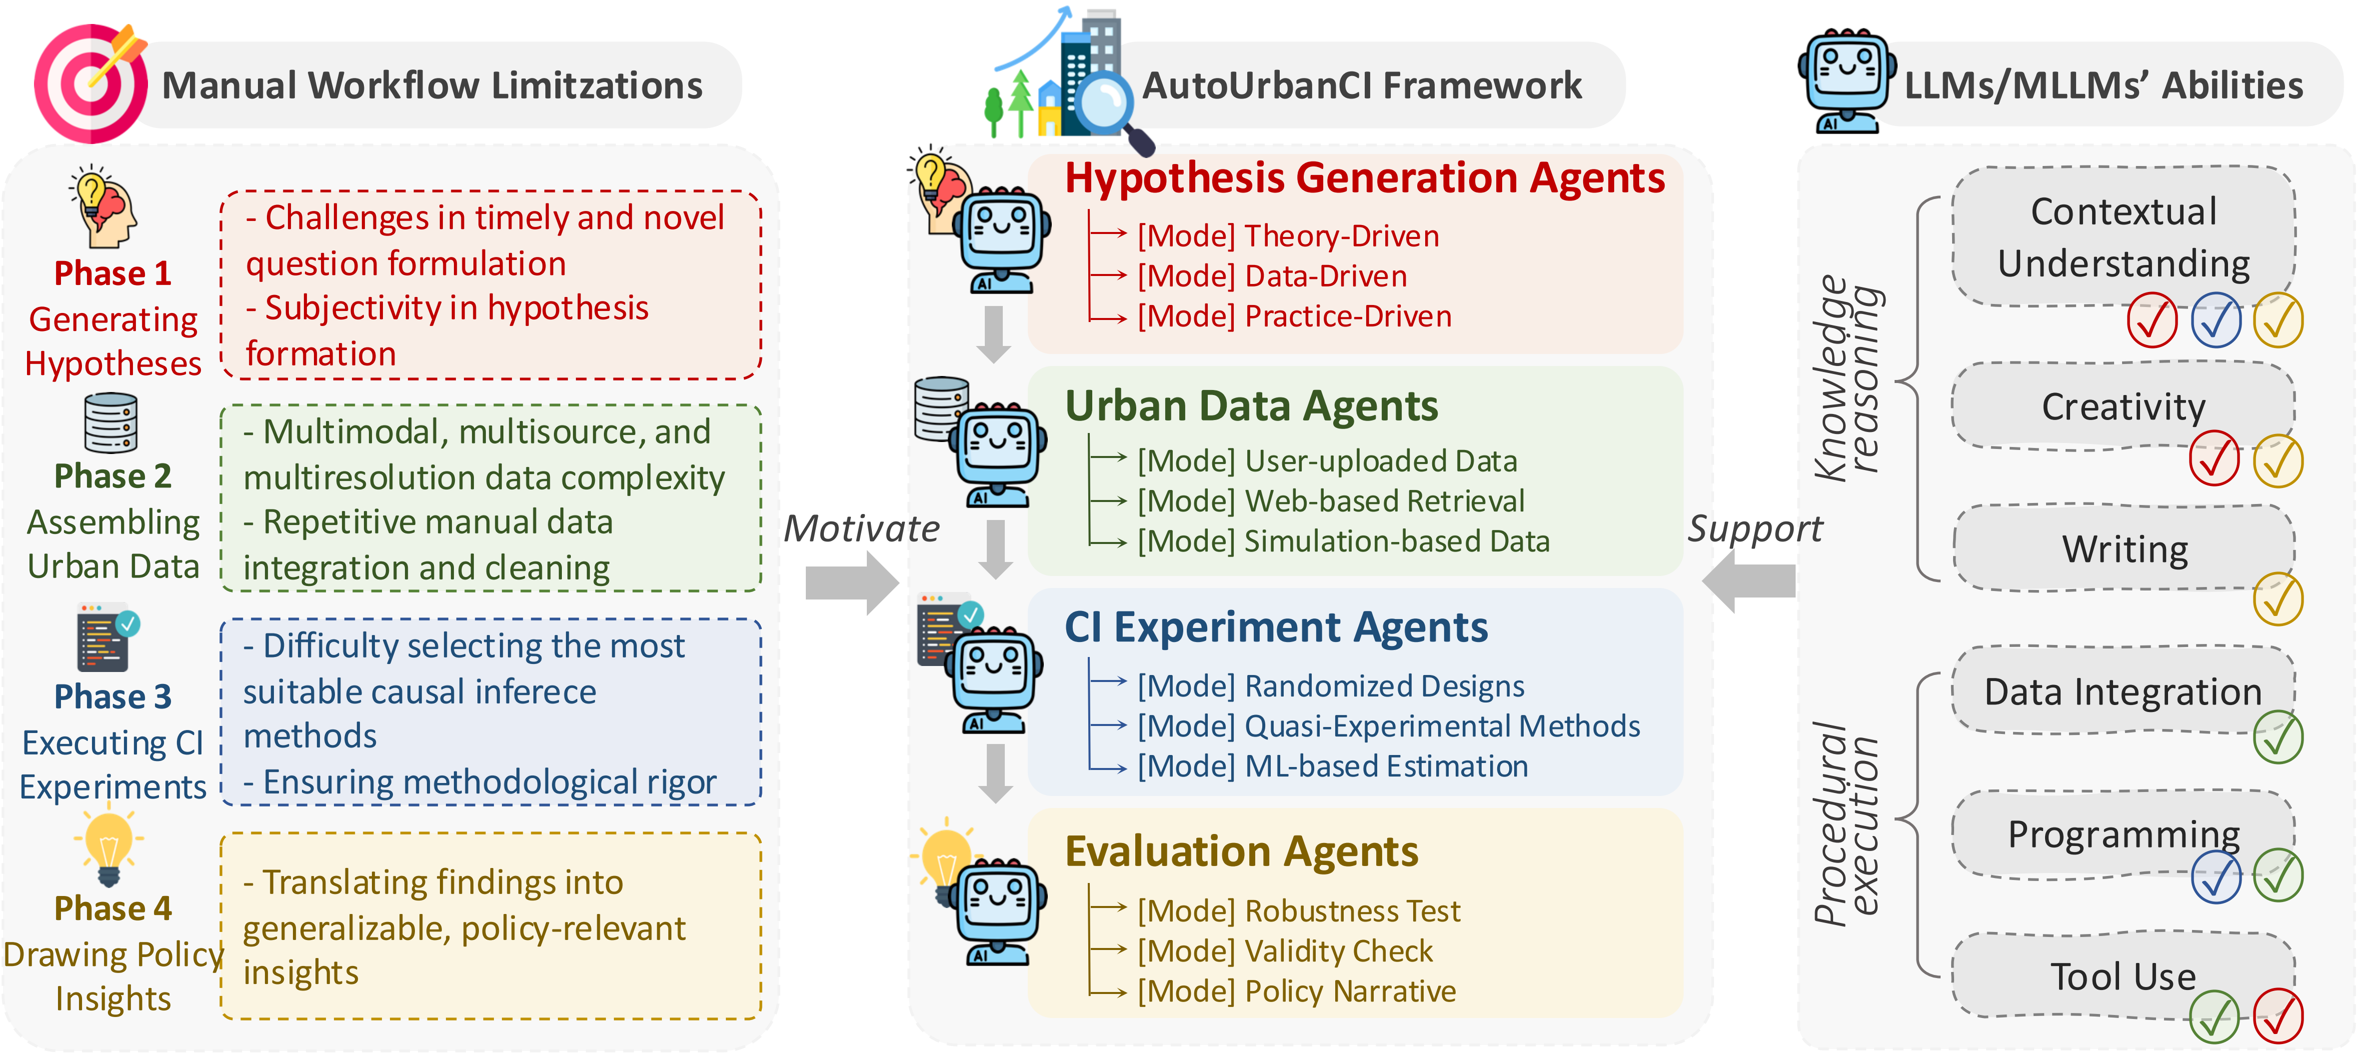




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

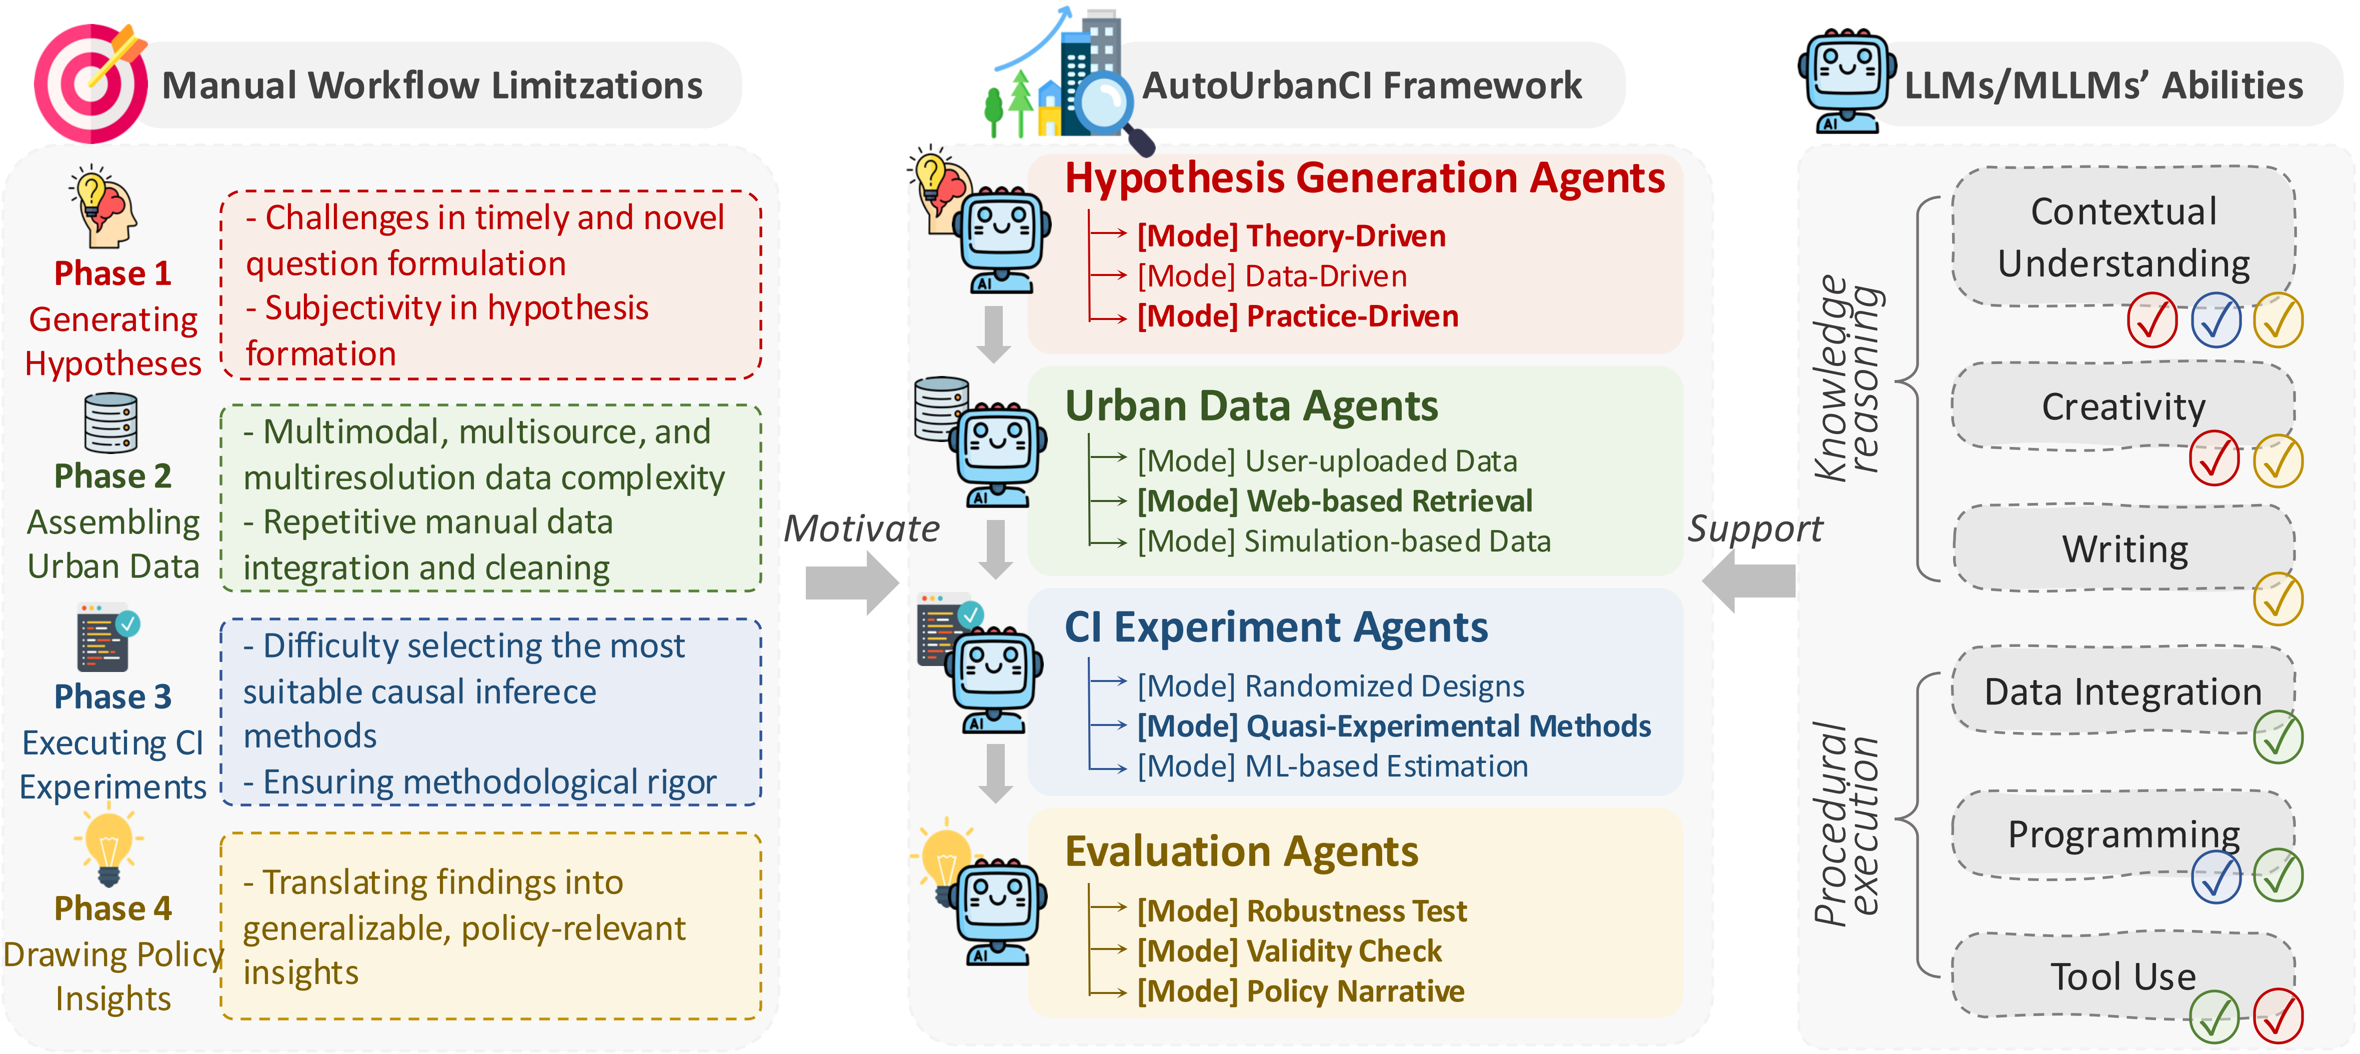




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# A Example Use Case

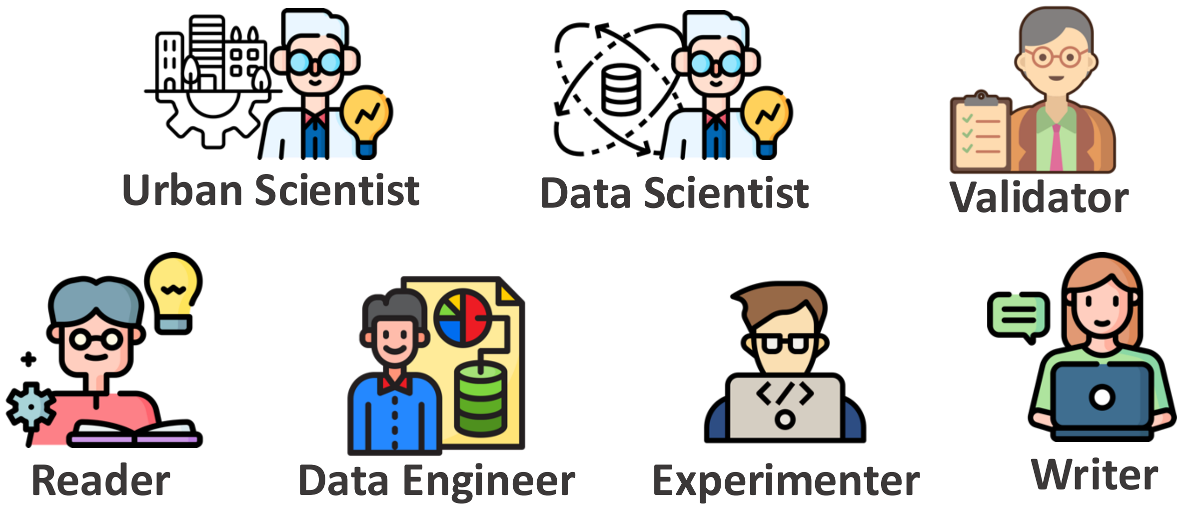




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# A Example Use Case

What if we created **a team of agents** to simulate an **urban research lab**?

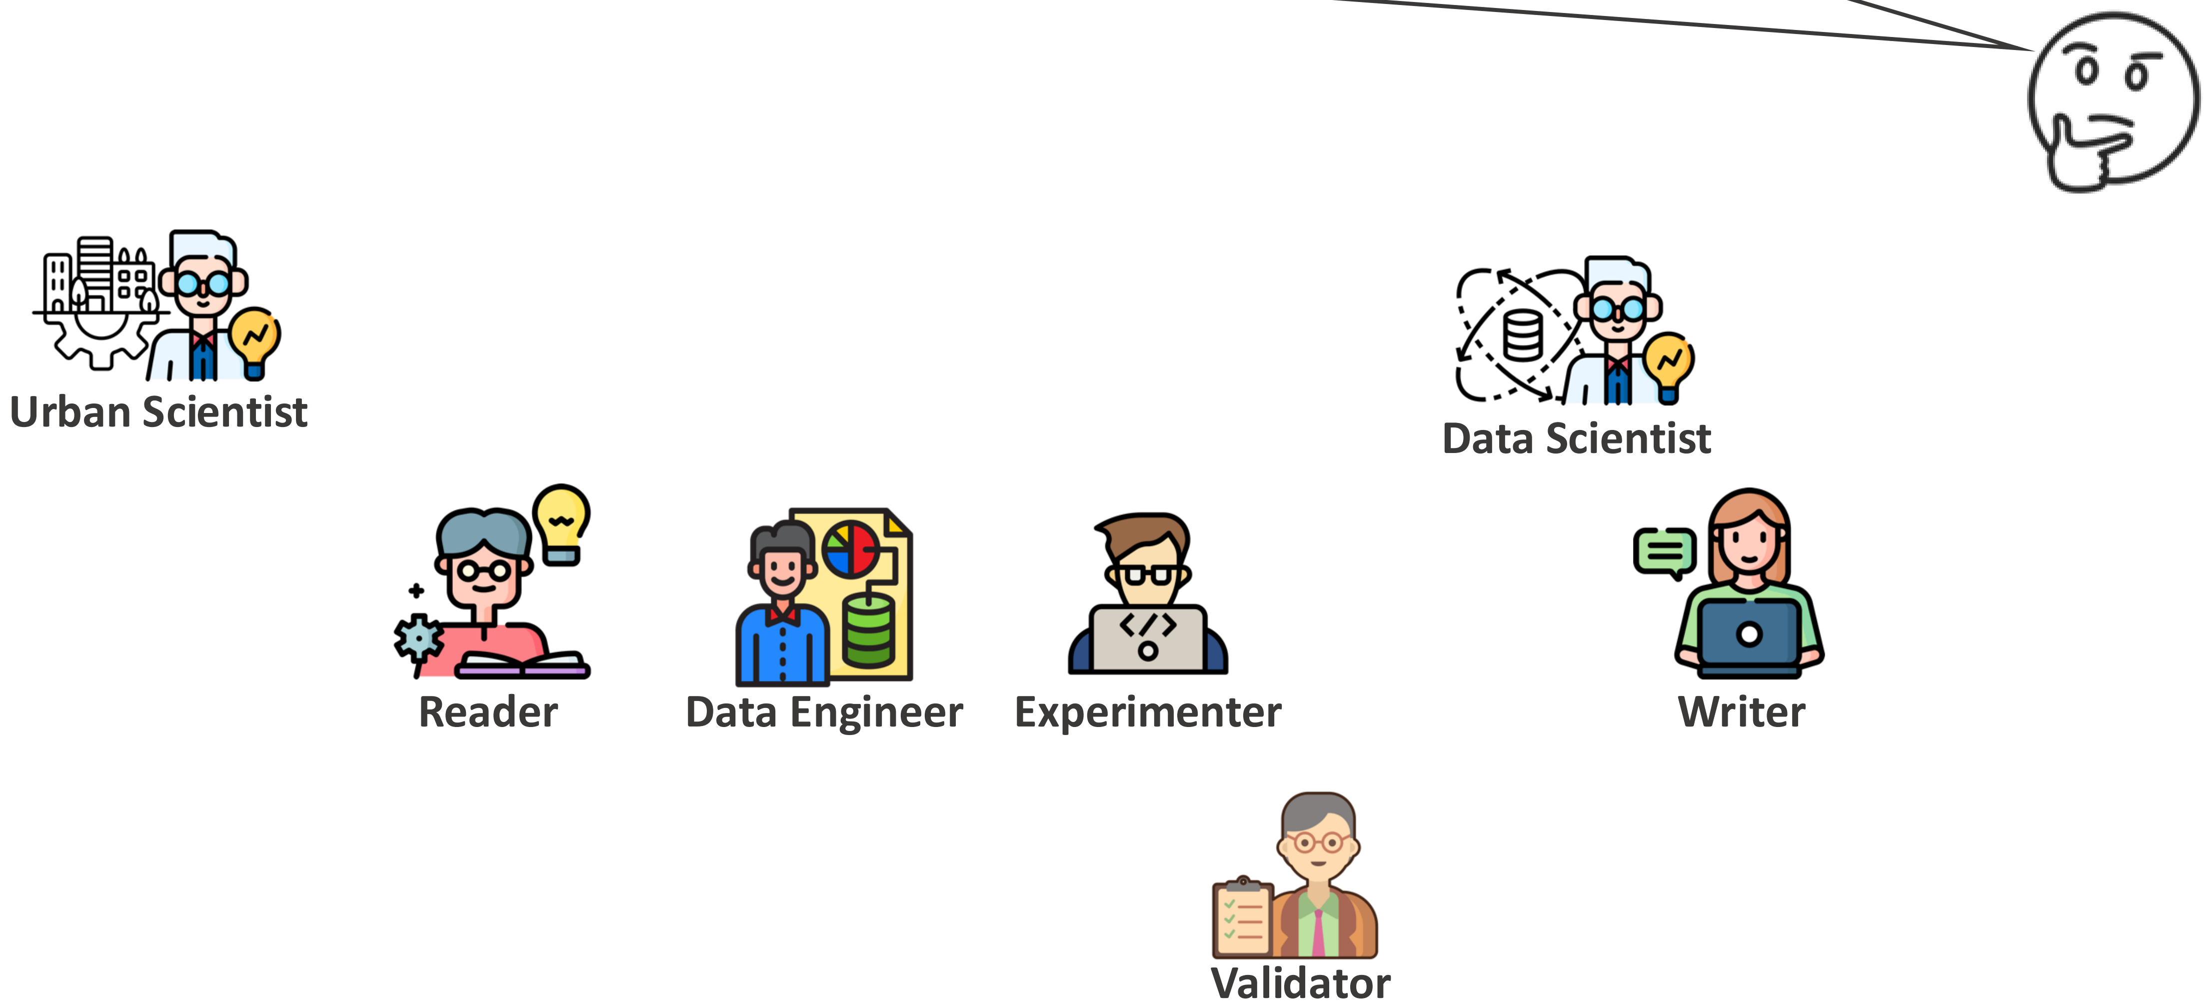




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# A Example Use Case

**Hypothesis Generation**

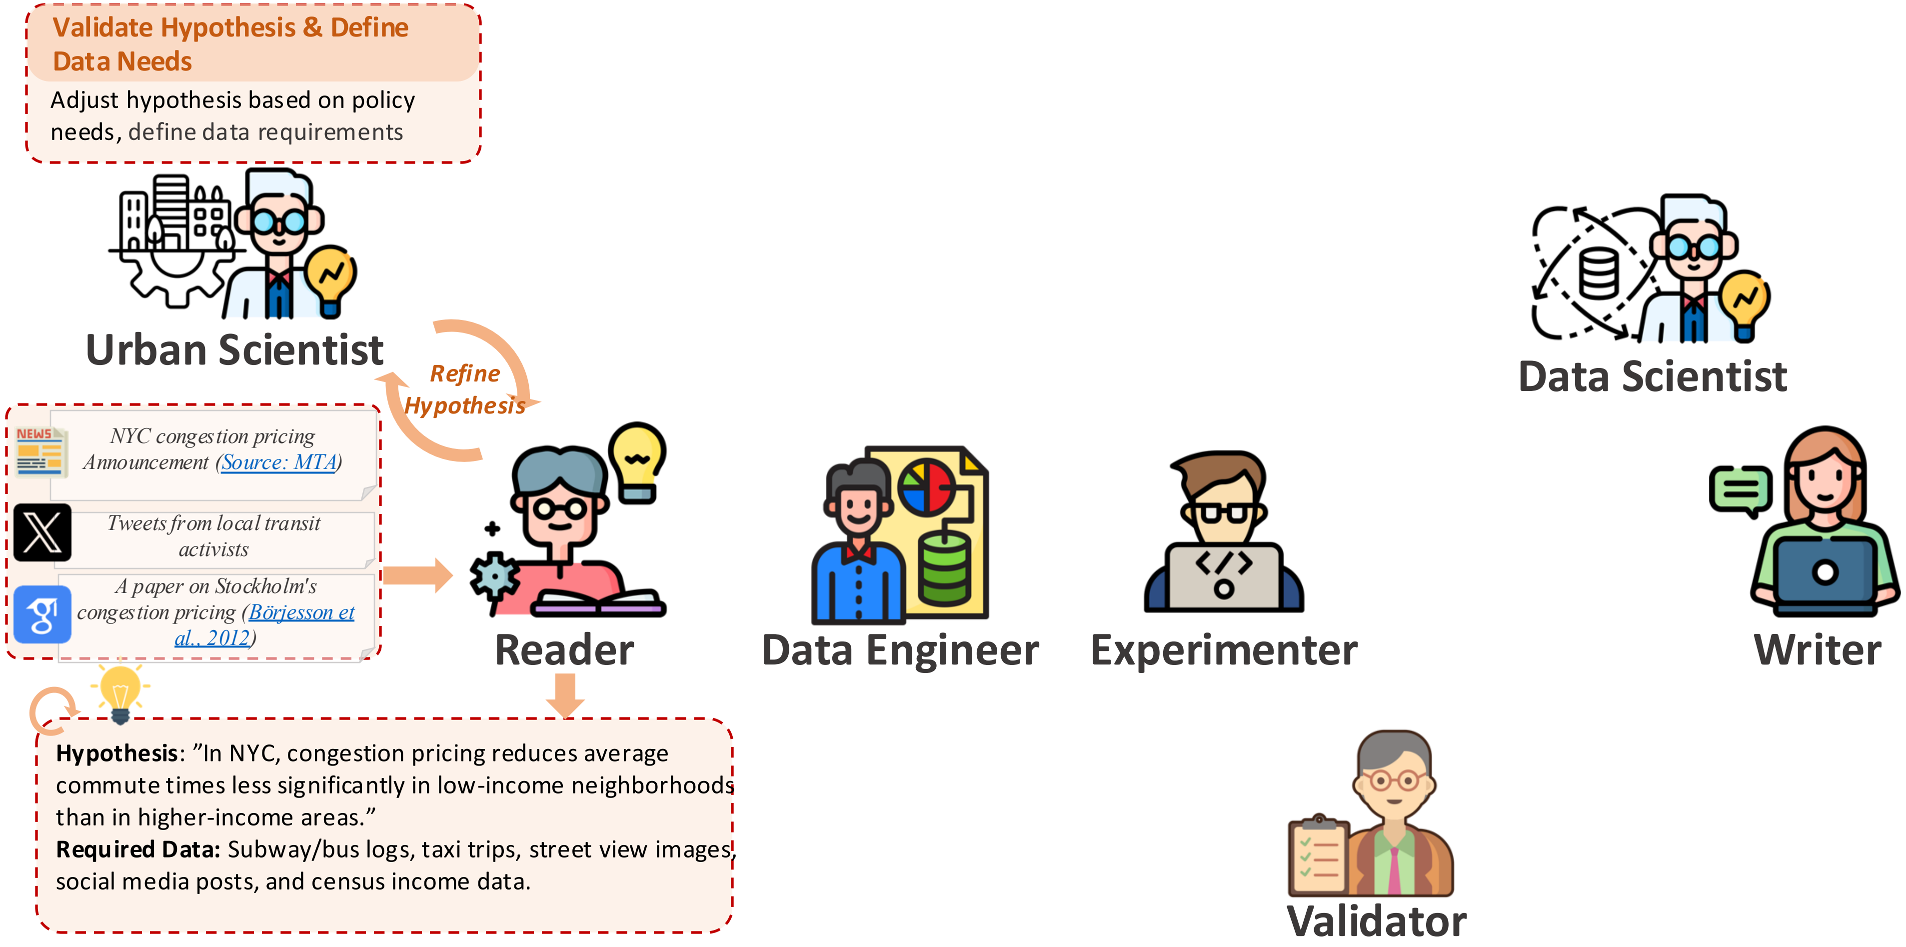




Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# A Example Use Case

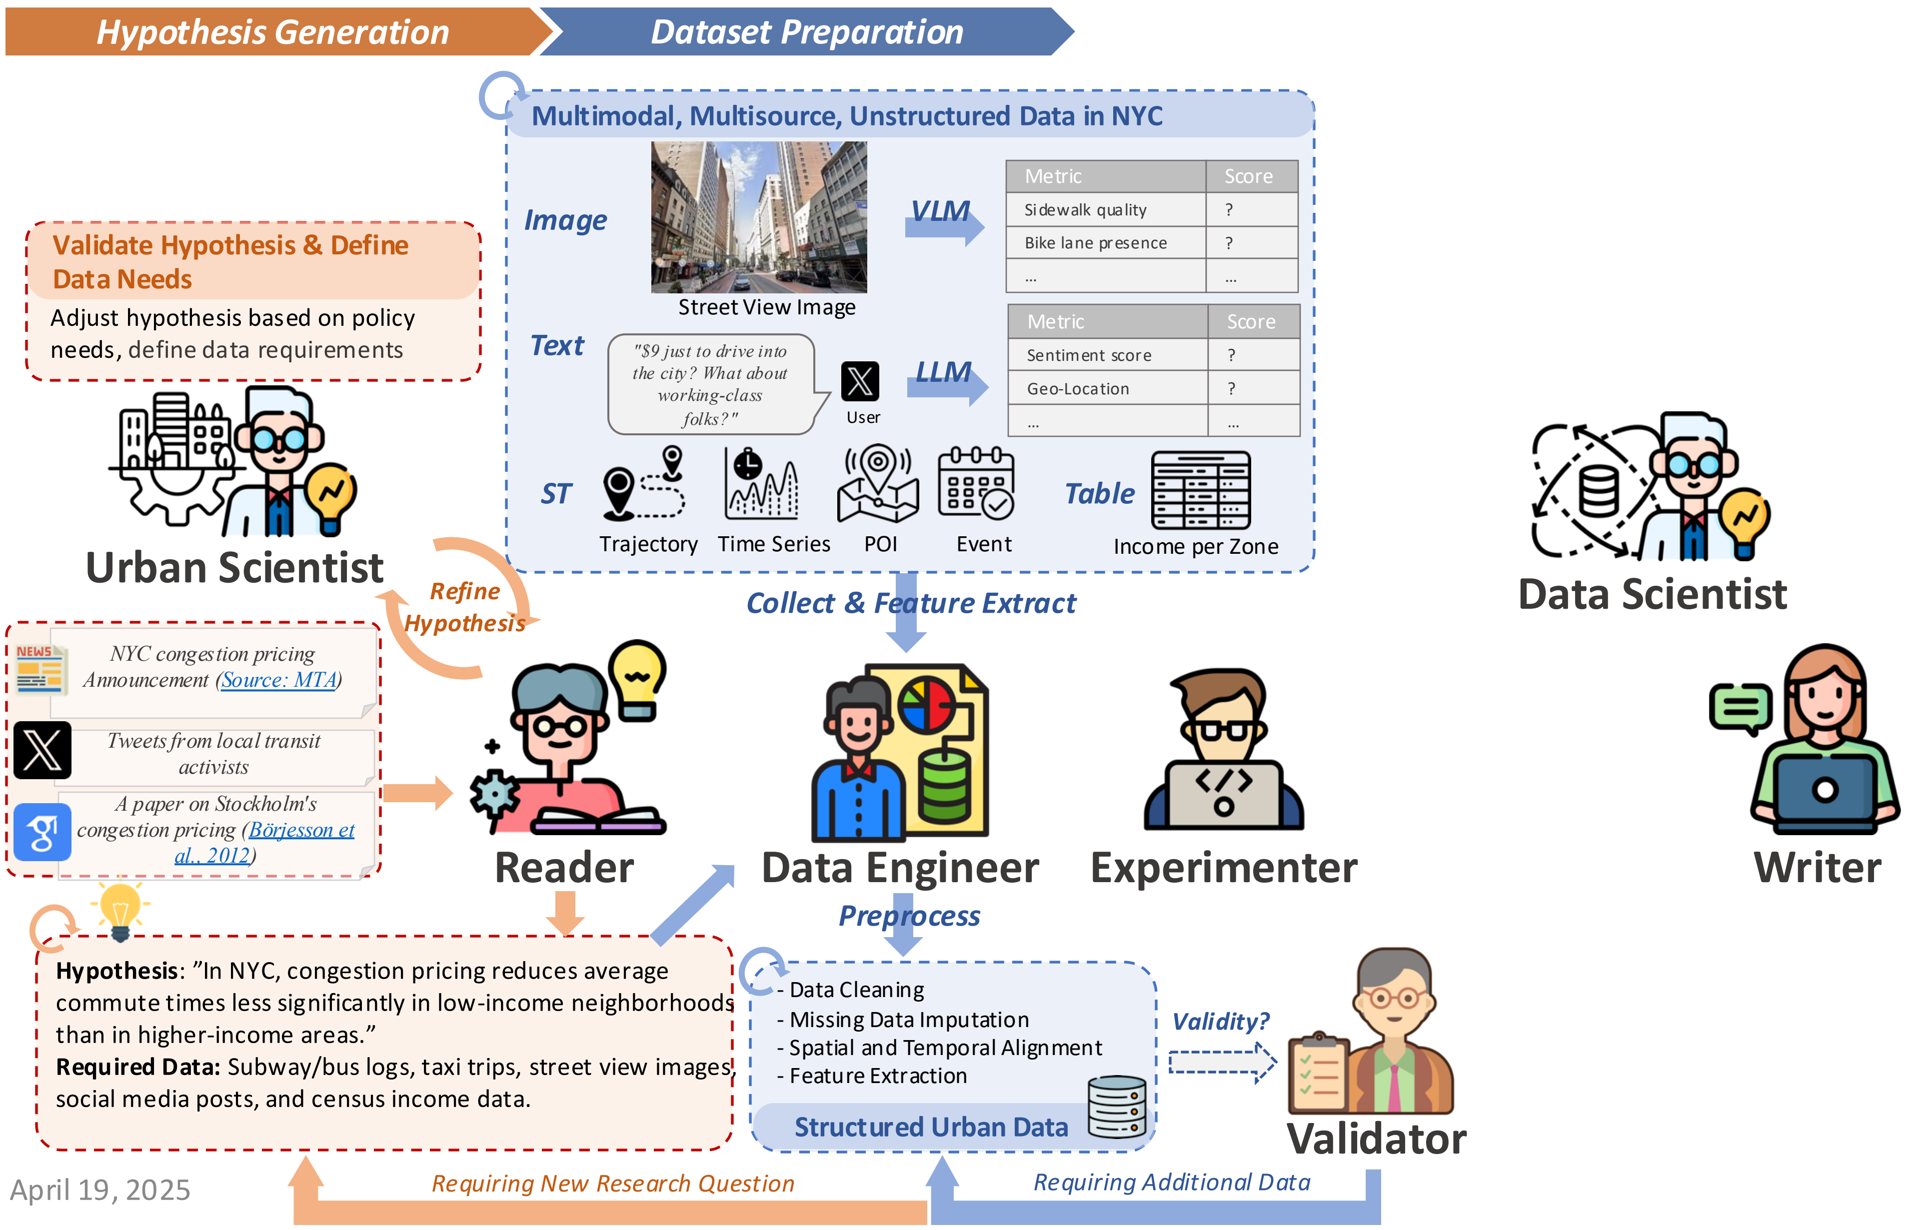

# A Example Use Case

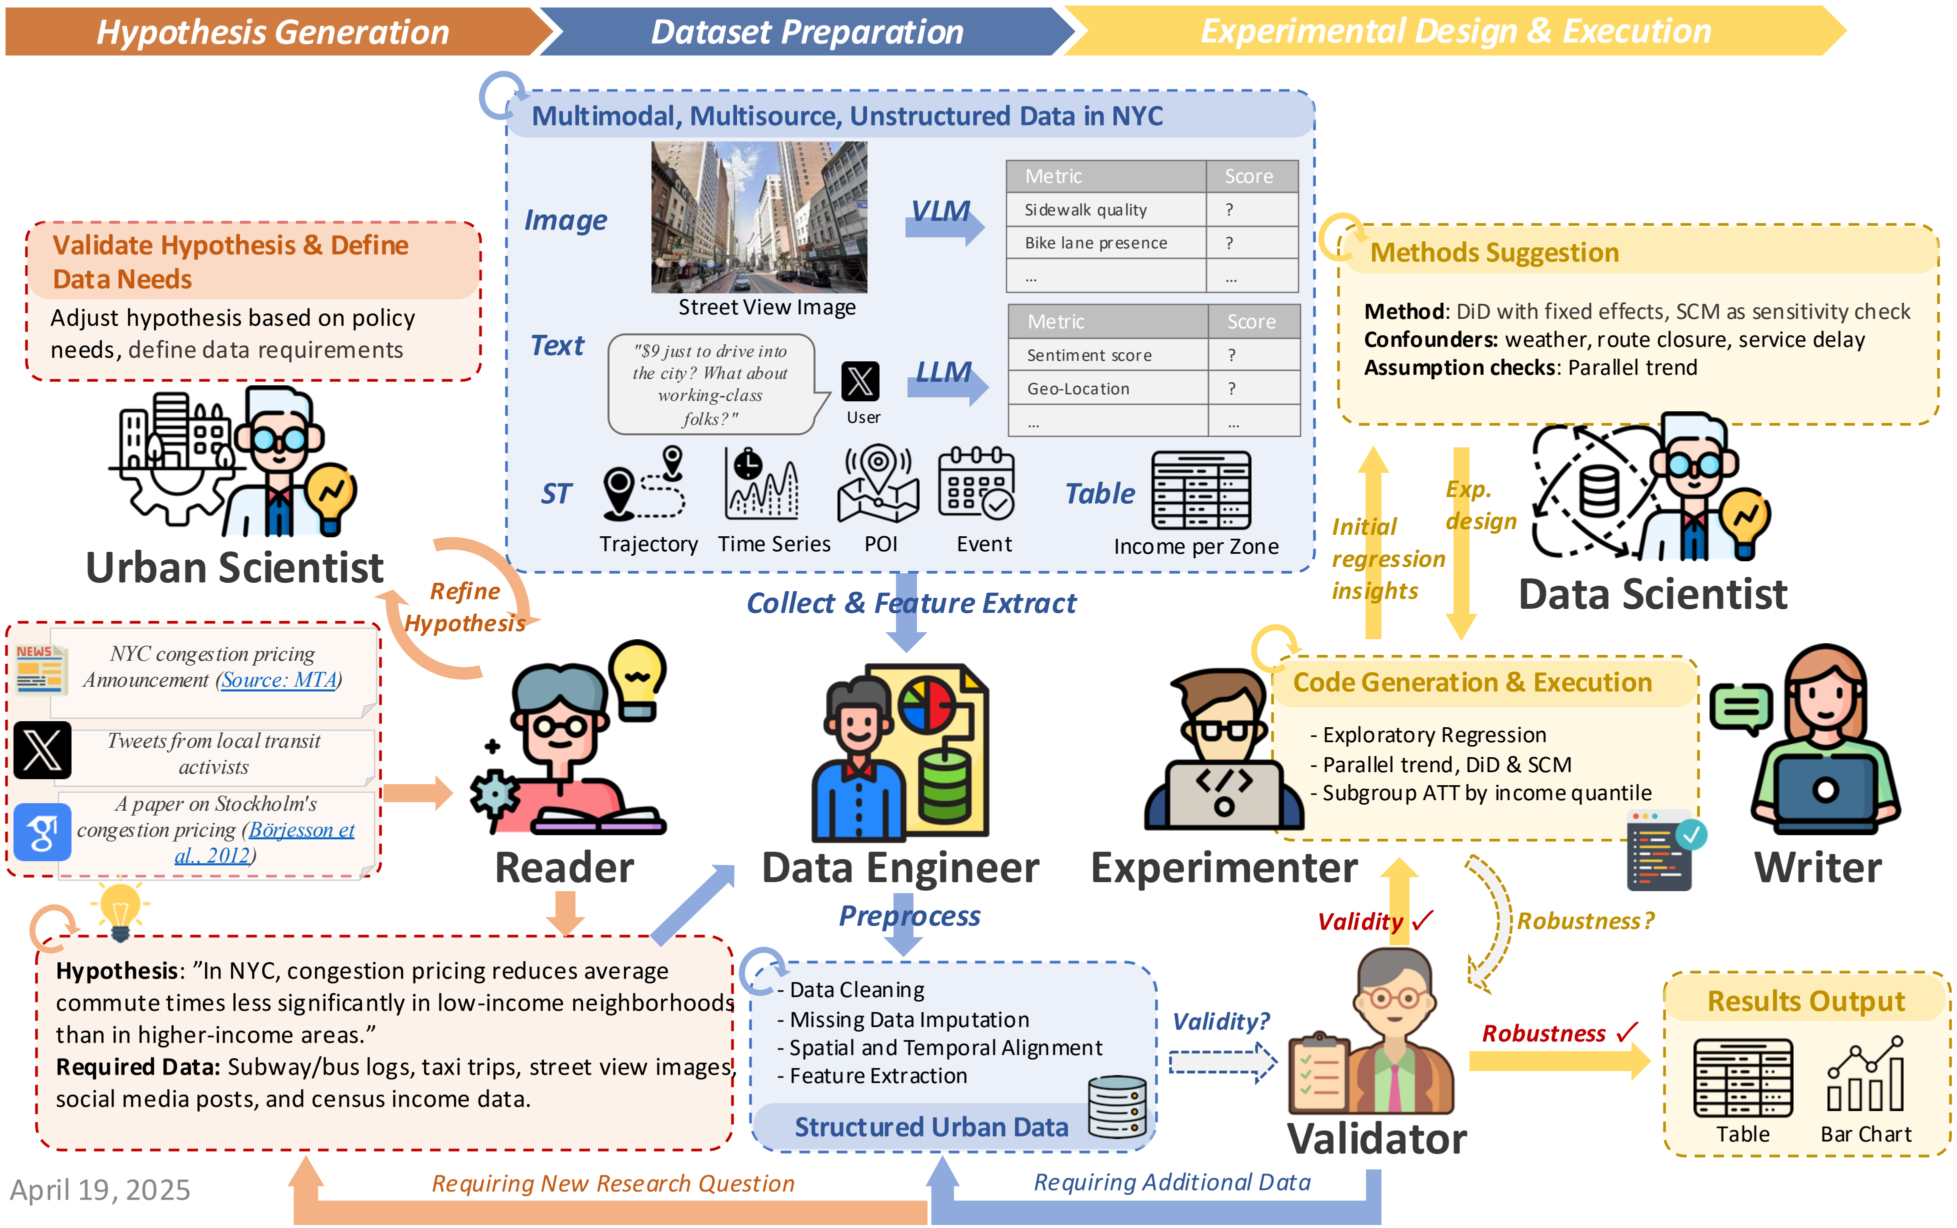

# A Example Use Case

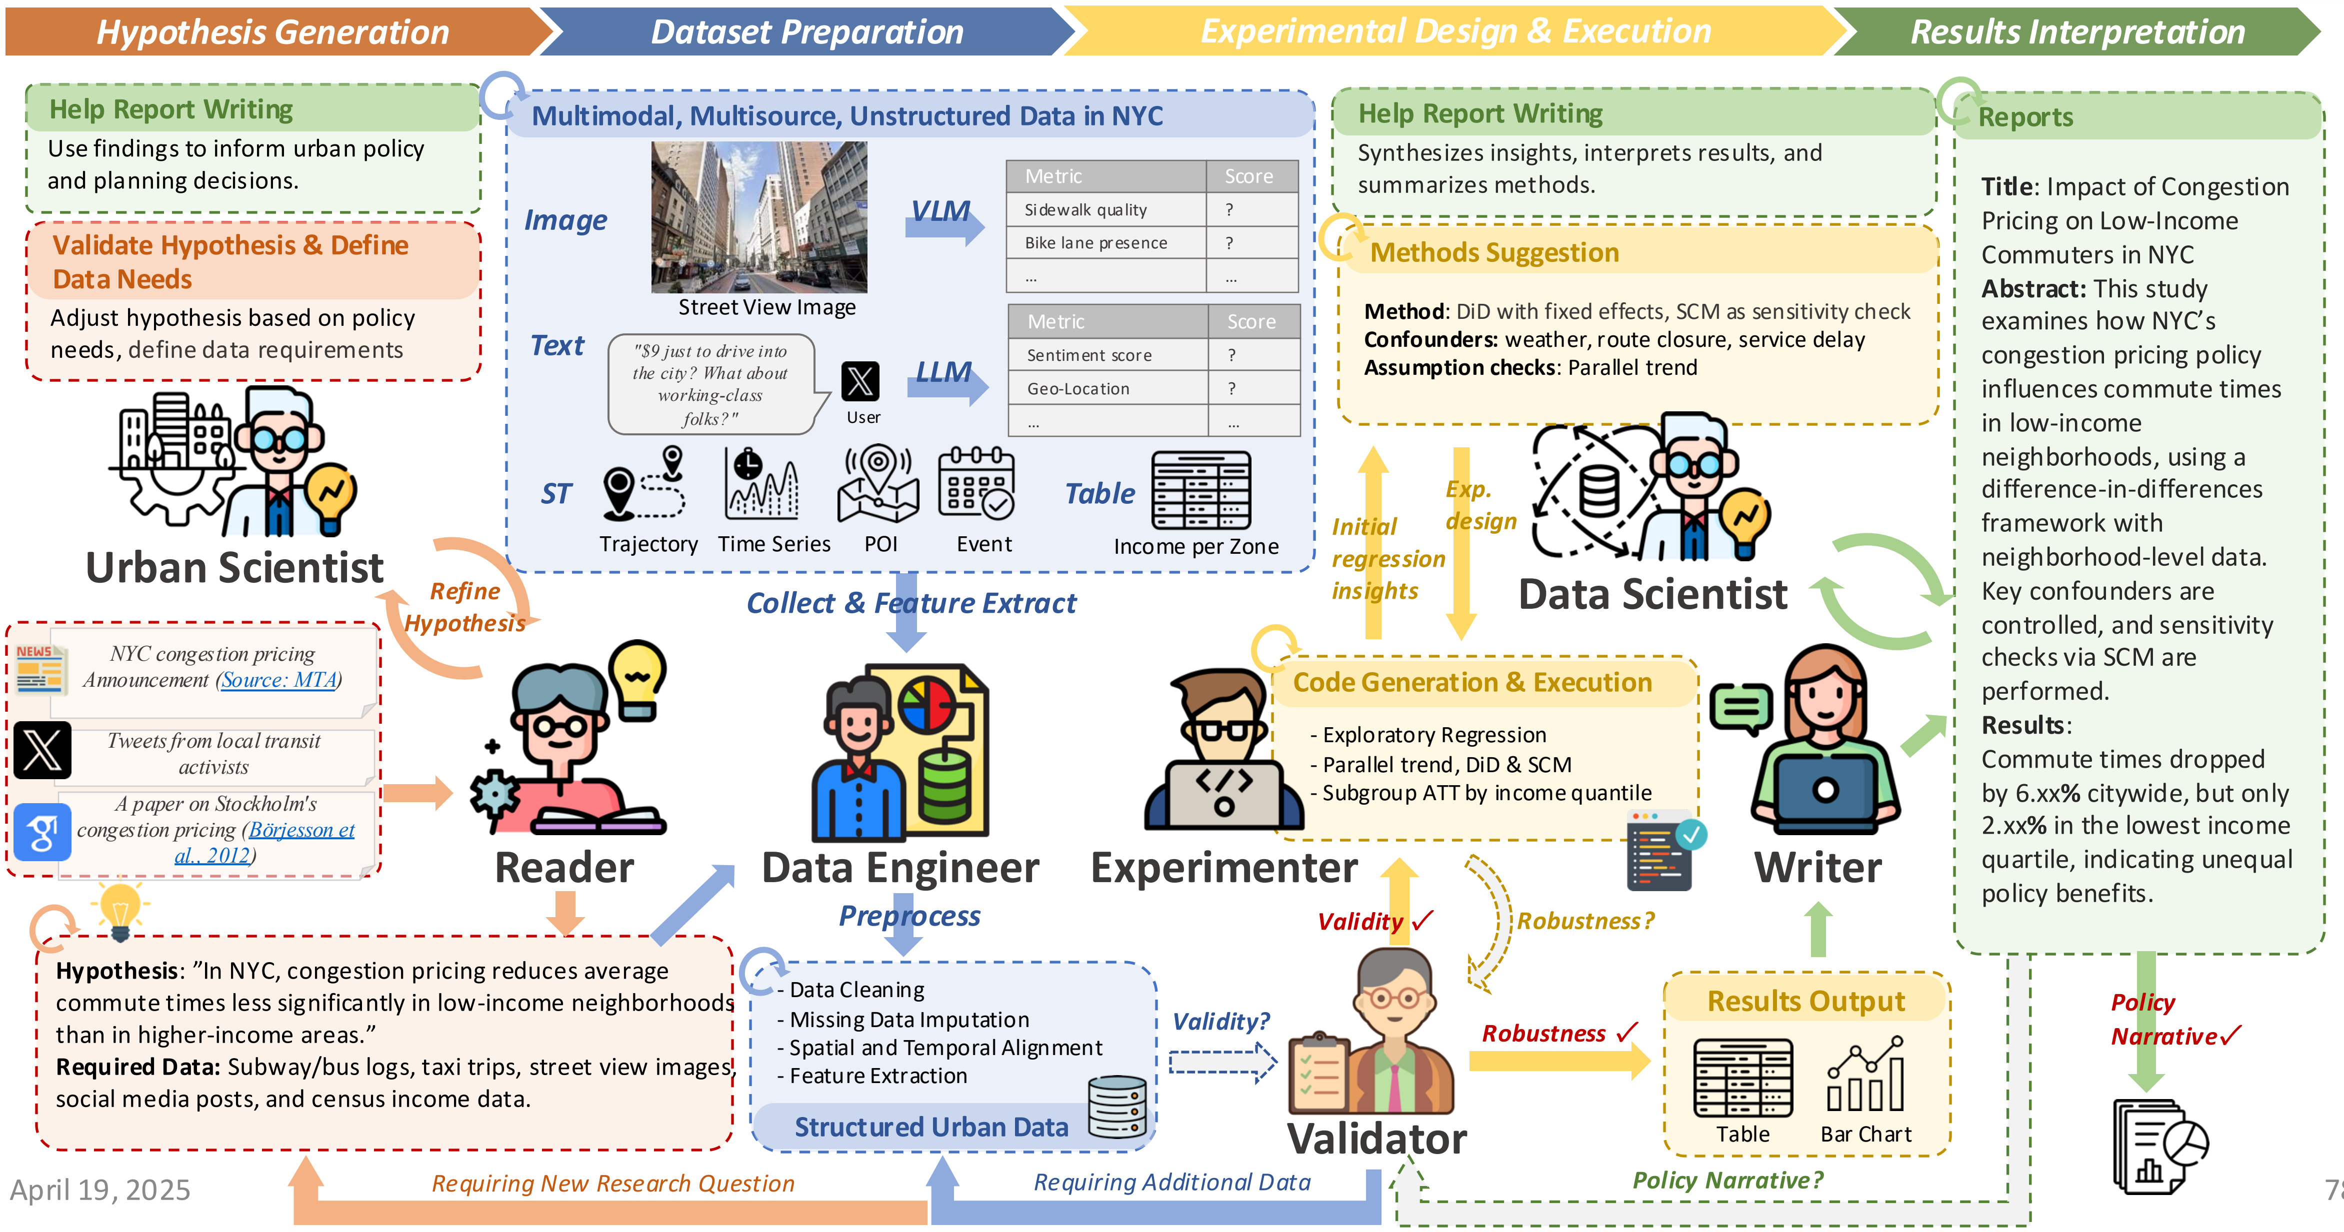

# Social Impact

For ***Urban Researchers***:

- Assists and accelerates causal research

For ***Policy Makers***

- Enhances evidence-based urban policy

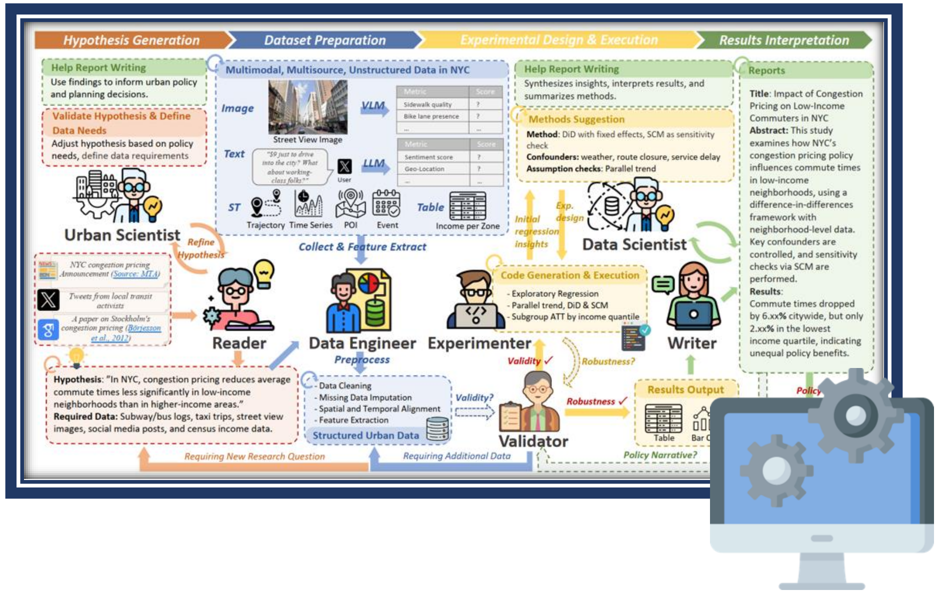

# Social Impact

For ***Urban Researchers***:

- Assists and accelerates causal research

For ***Policy Makers***

- Enhances evidence-based urban policy

For ***the Public***:

- Lowers the barriers for citizens, journalists, and grassroots organizations to explore urban issues

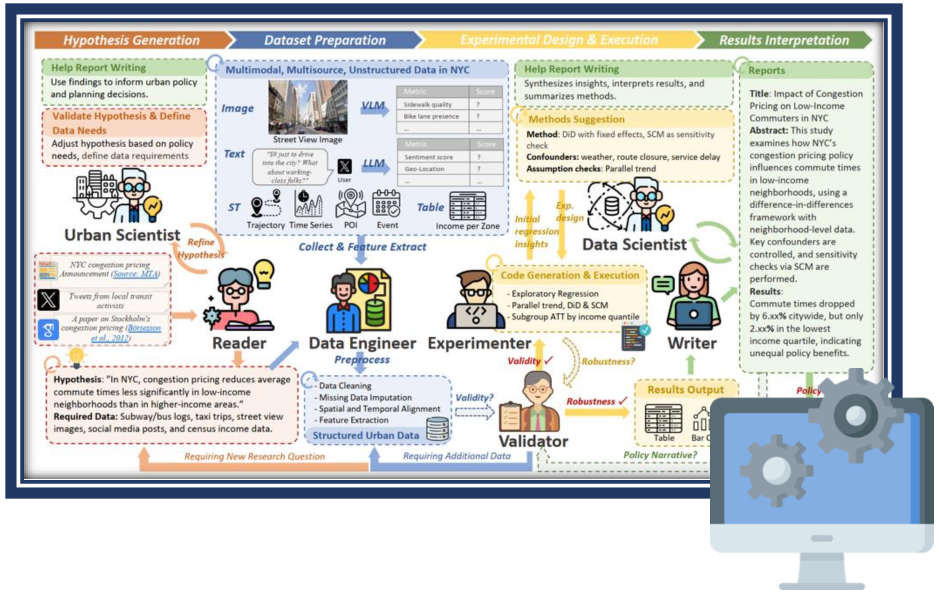


*"Cities have the capability of providing something for **everybody**, only because, and only when, they are created by **everybody**."*

*— Jane Jacobs*

# Our Perspetive


*AI* won't replace urban researchers,

but it can ***assist*** their thinking, ***accelerate*** their work, ***expand*** the horizons of urban science, and make urban insights a ***shared power***.

That's the ***future*** we're building toward.



Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.

# Takeaways

- **Cities are predictable.**
  - *STG forecasting* gives us powerful tools to model these signals in domains like air quality and mobility.
  - *Application-driven* modeling is tailored to the properties of specific data (e.g., air quality, traffic),
  - *Theory-driven* approaches incorporate causal lens and uncertainty modeling for deeper insight.
- **Towards prediction, what next?**
  - Toward a more *intelligent* and *accessible* urban causal analysis
- **Large Language Models unlock new possibilities.**
  - Not to replace urban scientists, but *assist* their thinking, *accelerate* their work, *expand* the horizons of urban science, and make urban insights a *shared power*.

# References


- Y. Zheng. Urban Computing. MIT Press, 2019.
- G. Jin et al. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. TKDE 2023.
- Y. Liang et al., AirFormer: Predicting Nationwide Air Quality in China with Transformers. AAAI 2023.
- X. Liu et al., LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. NeurIPS 2023.
- H. Zhang*, Y. Xia* et al., Predicting Carpark Availability in Singapore with Cross-Domain Data: A New Dataset and A Data-Driven Approach. IJCAI 2024.
- Q. Wang et al., AirRadar: Inferring Nationwide Air Quality in China with Deep Neural Networks. AAAI 2025.
- Y. Xia et al., Deciphering Spatio-Temporal Graph Forecasting: A Causal Lens and Treatment. NeurIPS 2023.
- J. Ho et al., Denoising diffusion probabilistic models. NeurIPS 2020.
- H. Wen, Lin Y, Y. Xia et al., DiffSTG: Probabilistic Spatio-Temporal Graph Forecasting with Denoising Diffusion Models. SIGSPATIAL 2023.
- Zou, et al., Deep Learning for Cross-Domain Data Fusion in Urban Computing: Taxonomy, Advances, and Outlook. Information Fusion 2025.
- J. Park et al., Generative Agents: Interactive Simulacra of Human Behavior. UIST 2023.
- Y. Xia*, A. Qu* et al., Reimagining Urban Science: Scaling Causal Inference with Large Language Models. arXiv 2025.
- Y. Liang, H. Wen, Y. Xia et al., Foundation Models for Spatio-Temporal Data Science: A Tutorial and Survey. arXiv 2025.

# Thanks!

**Slides for this Talk**

**Personal Homepage**