Causal Inference for Spatiotemporal Interventions

Miruna Oprescu

Joint work with David K. Park[2], Xihaier Luo[2], Shinjae Yoo[2], and Nathan Kallus[1]

[1] Cornell University, Cornell Tech

[2] Computing and Data Sciences, Brookhaven National Laboratory

Decision-Making in Spatiotemporal Contexts

- **Spatiotemporal Data**
  - Observations that vary across both spatial and temporal dimensions. E.g.: $PM_{2.5}$ levels during the 2018 California wildfires.

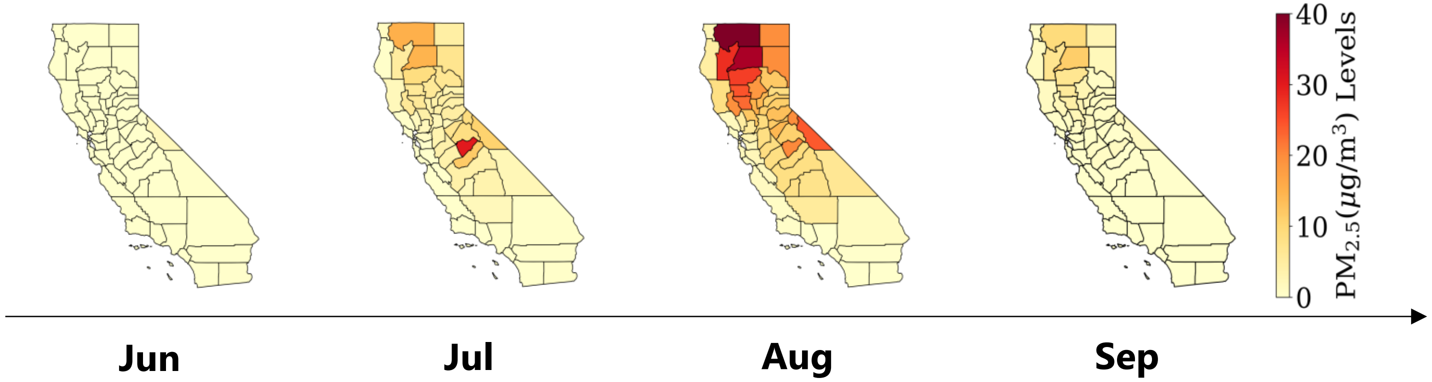

  - Often sourced from satellites, ground sensors, and weather stations, capturing how conditions evolve day by day and region by region.
- **Spatiotemporal Interventions**
  - Real-world actions or policies applied across space and time—such as wildfire prevention or pollution control measures—that shape local and regional outcomes (e.g., $PM_{2.5}$ levels, public health).

Decision-Making in Spatiotemporal Contexts

- **Counterfactual / Policy-Relevant Questions**
  - “What if stricter wildfire prevention measures had been implemented 2 weeks earlier—how would PM2.5 and health outcomes change over τ time steps?”

$\downarrow$

Causal Inference Question!

Decision-Making in Spatiotemporal Contexts

- **Counterfactual / Policy-Relevant Questions**
  - "What if stricter wildfire prevention measures had been implemented 2 weeks earlier—how would PM2.5 and health outcomes change over τ time steps?"
- **Notation**
  - Time $t \in \{1, \dots, T\}$, horizon τ, spatial index $s \in \mathbb{G}$.
  - **Features (Covariates):** $X_{s,1}, X_{s,2}, \dots, X_{s,T}$.
  - **Interventions (Treatments):** $A_{s,1}, A_{s,2}, \dots, A_{s,T}$.
  - **Outcomes:** $Y_{s,1}, Y_{s,2}, \dots, Y_{s,T}$.
  - **History:** $H_{s,1:t} = \left(X_{s,1:t}, Y_{s,1:t}, A_{s,1:t-1}\right)$.
  - Shorthand:

    $W_{s,1:t} = \left\{W_{s,1}, W_{s,2}, \dots, W_{s,t}\right\}$, $\boldsymbol{W}_{1:t} = \{W_{s,1:t}\colon \forall s \in \mathbb{G}\}$

    for any $W \in \{X, A, Y, H\}$.

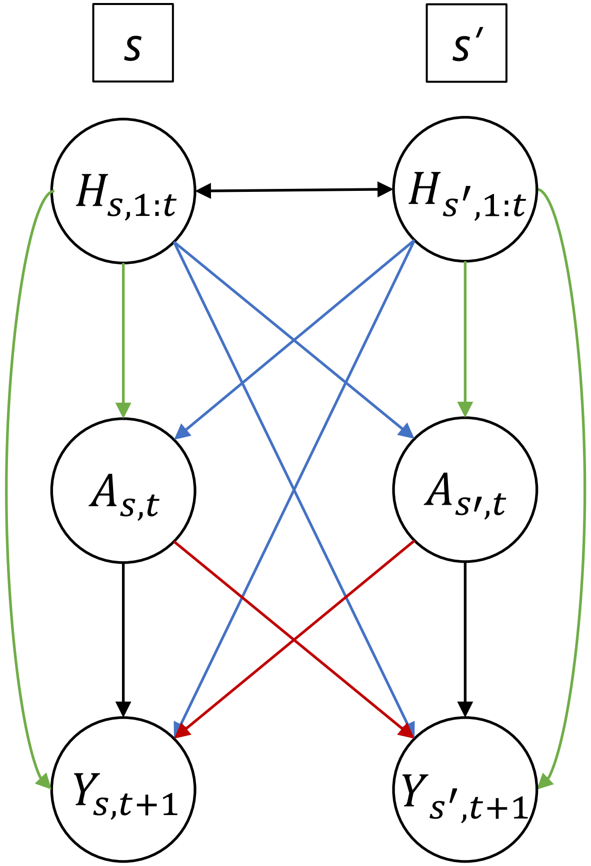

Schematic of the spatiotemporal data (X, A, Y, H) across time t and location s.

Challenges in Spatiotemporal Causal Inference

- **Complex Space-Time Dependencies**
  - Observations at different locations and times can strongly influence one another, complicating standard causal analyses.
- **Observational vs. Interventional Data**
  - We need to learn features of an interventional distribution (new policy scenarios) from observational data, where interventions were applied differently (or non-randomly).

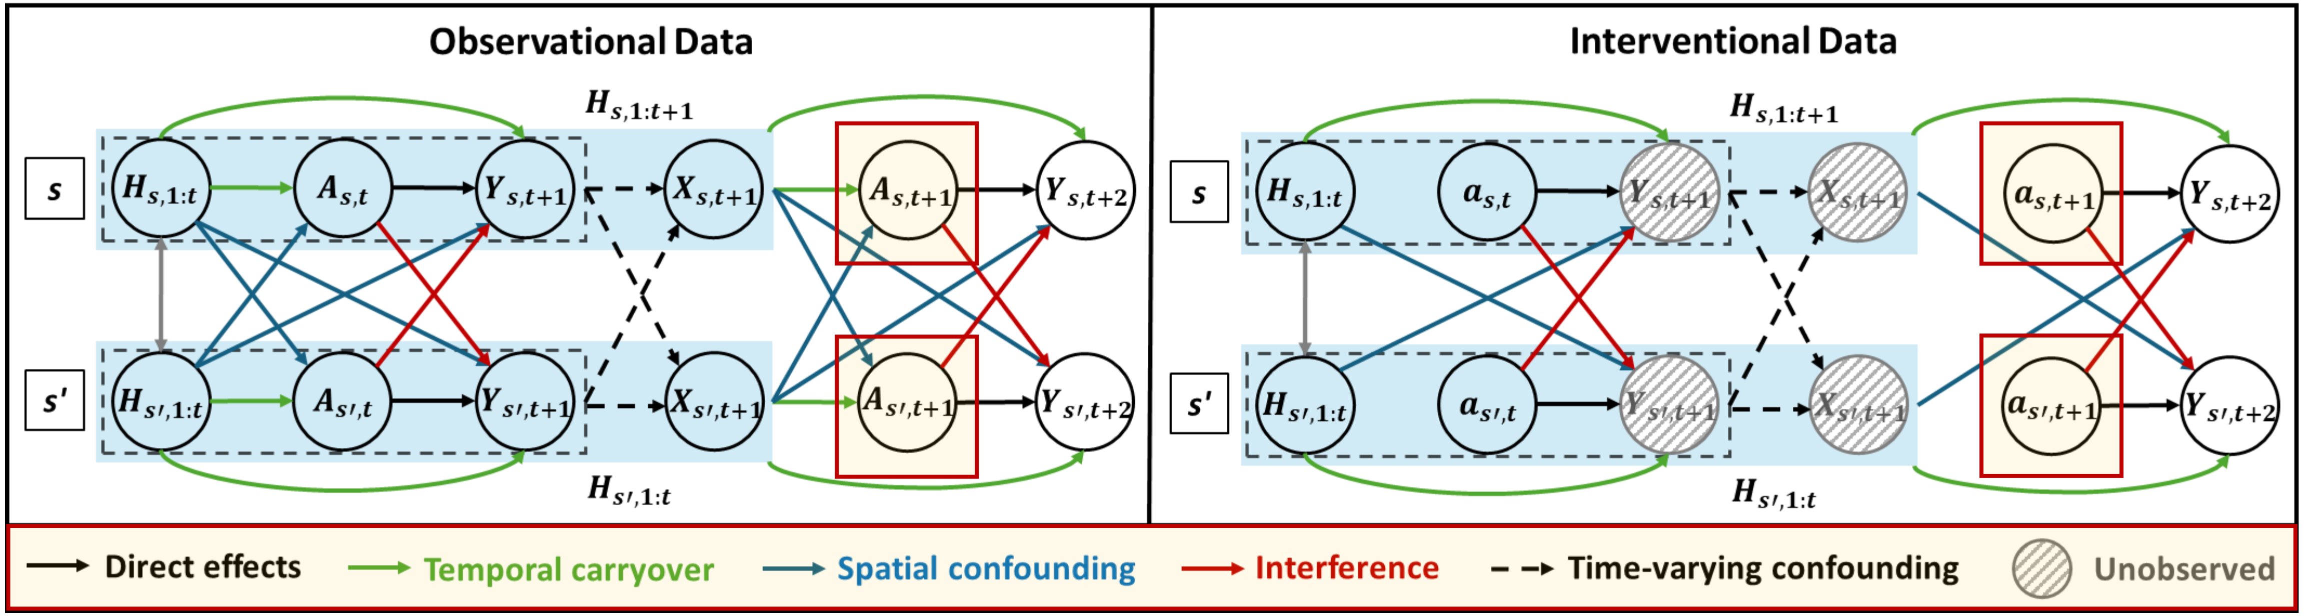

Challenges in Spatiotemporal Causal Inference

- **Time-Varying Confounders**
  - *A confounder* is any variable that affects both treatments and outcomes, and must be controlled to avoid biased causal estimates.
  - A *time-varying confounder* is a variable that affects both future treatments and outcomes, creating feedback loops (e.g. past interventions shape future covariates, which in turn drive subsequent interventions and outcomes).

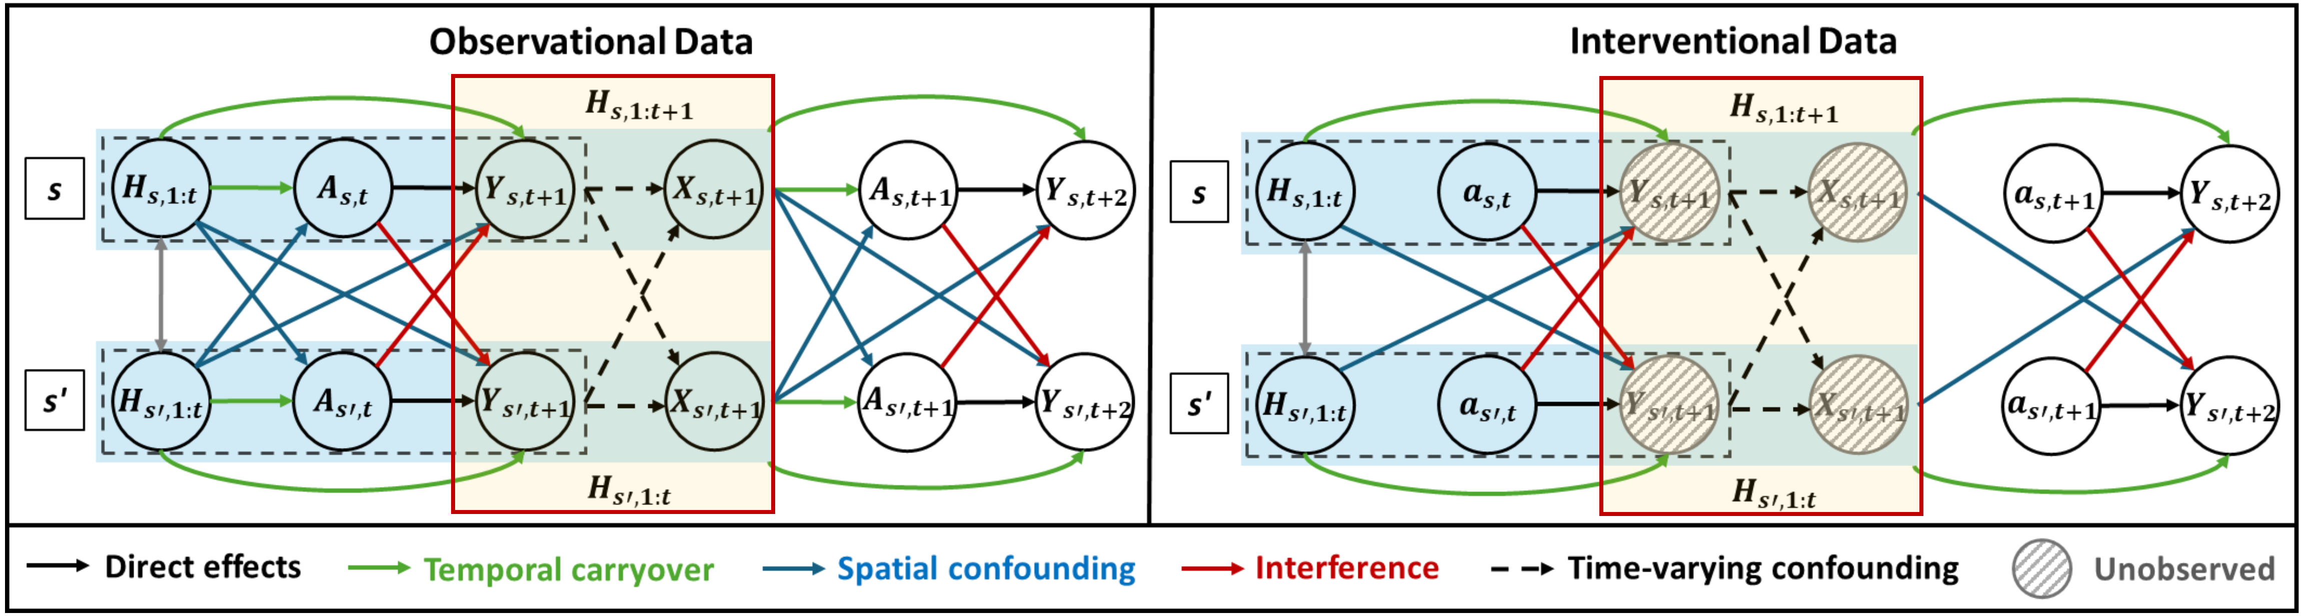

Machine Learning for Spatiotemporal Modeling

Approach: Use neural networks to capture spatiotemporal patterns

- **U-Net** for spatial dependencies [1]
  - Encoder-decoder architecture that captures multi-scale spatial features.
- **ConvLSTM** for temporal dynamics
  - Merges convolution and LSTM to model temporal dynamics within a single series.
- **Attention** to highlight key spatial regions and time steps [2].

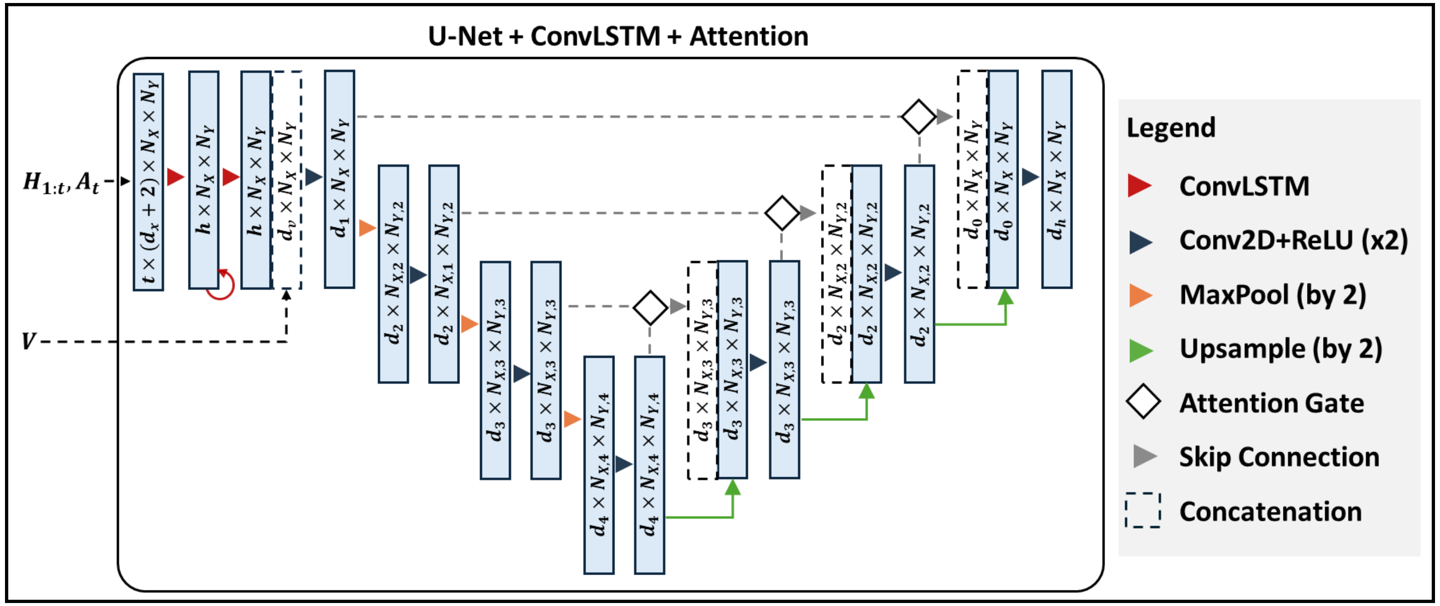

Causal Inference with Time-Varying Confounders

- **Estimand:** $\mathbb{E}[\,\boldsymbol{Y}_{t+\tau}[\boldsymbol{a}_{t:t+\tau-1}] \mid \boldsymbol{H}_{1:t}\,]$
  - Average potential outcome after τ time steps under a series of fixed τ interventions, $\boldsymbol{a}_{t:t+\tau-1}$, given an observed history $\boldsymbol{H}_{1:t}$.
- **Iterative G-Computation**
  - A standard causal inference technique to handle time-varying confounders.
  - Iteratively averages over potential trajectories from the observed history to the outcome over horizon τ, ensuring unbiased estimates.

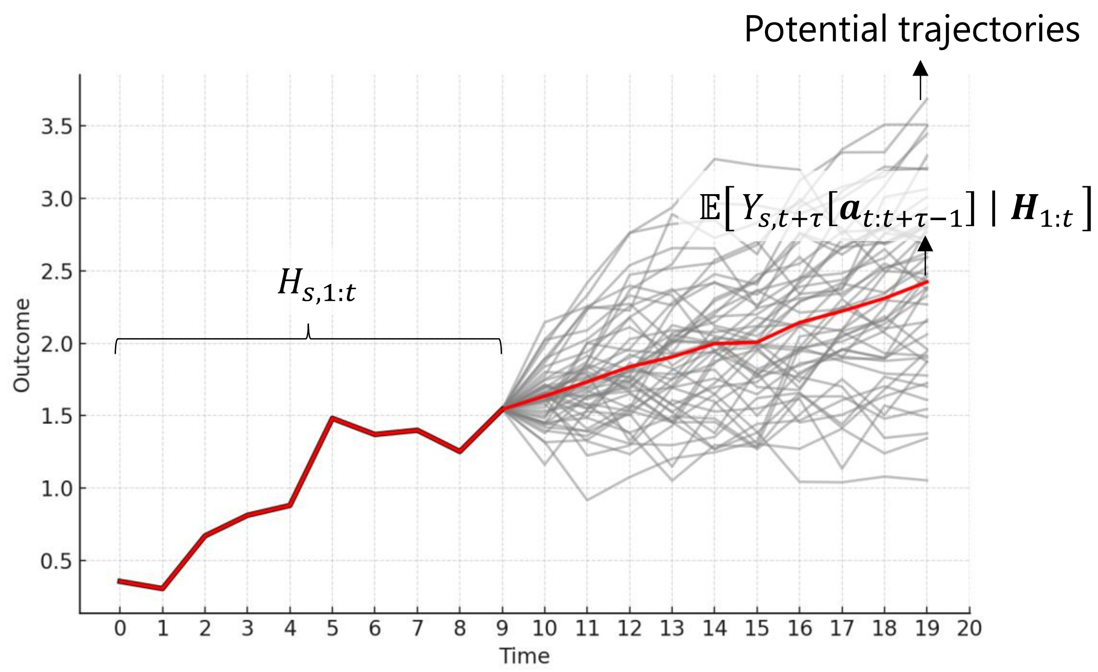

Causal Inference with Time-Varying Confounders

Iterative G-Computation via Recursive Regression [3]

1. Last Step:

$$Q_\tau(\boldsymbol{H}_{1:t+\tau-1}, \boldsymbol{A}_{t+\tau-1}) = \mathbb{E}[\boldsymbol{Y}_{t+\tau} \mid \boldsymbol{H}_{1:t+\tau-1}, \boldsymbol{A}_{t+\tau-1}]$$

2. Recursive Steps (for $k = \tau - 1, \ldots, 1$):

$$Q_k(\boldsymbol{H}_{1:t+k-1}, \boldsymbol{A}_{t+k-1}) = \mathbb{E}[Q_k(\boldsymbol{H}^{\boldsymbol{a}}_{1:t+k}, \boldsymbol{A}_{t+k}) \mid \boldsymbol{H}_{1:t+\tau-1}, \boldsymbol{A}_{t+\tau-1}]$$

3. Result:

$$\mathbb{E}[\boldsymbol{Y}_{t+\tau}[\boldsymbol{a}_{t:t+\tau-1}] \mid \boldsymbol{H}_{1:t}] = Q_1(\boldsymbol{H}_{1:t}, \boldsymbol{a}_t)$$

Here, $\boldsymbol{H}^{\boldsymbol{a}}_{1:t+k}$ denotes the history where treatments from time t onward are set to the intervention sequence $\boldsymbol{a}_{t:t+k+1}$.

Introducing the GST-UNet (Our Work)

G-computation **S**patio-**T**emporal **UNet** (**GST-UNet**):

- **Spatiotemporal Embedding:** U-Net + ConvLSTM + attention gates.
- **Neural Causal Modules:** G-computation heads (e.g. shallow feed-forward networks or convolutional layers) for iterative adjustment.
- **Key Innovation:** Flexible, end-to-end approach that avoids strong modeling assumptions and properly accounts for time-varying confounders.

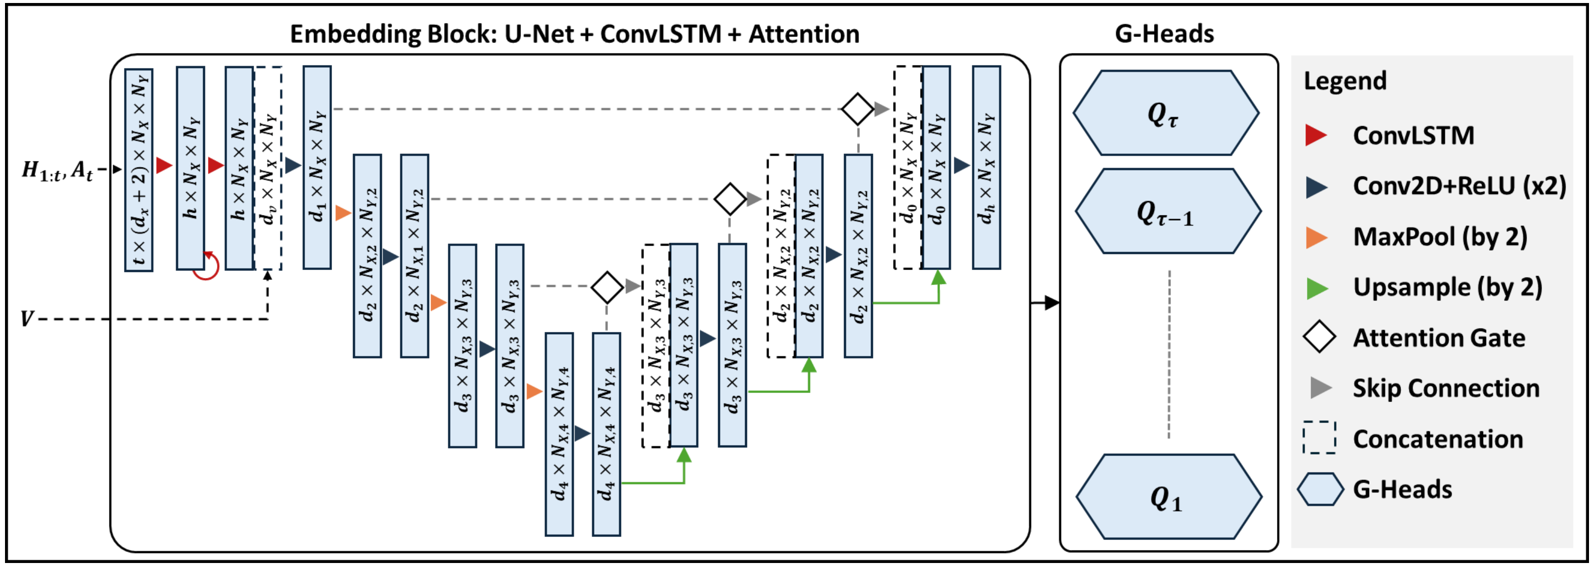

GST-UNet End-to-End Architecture

Simulation Results on Synthetic Data

- **Data:** We generate $T = 200$ steps of a 64×64 grid of observational data from:

$$\boldsymbol{X}_t = \alpha_0 + \alpha_1 \boldsymbol{X}_{t-1} + \alpha_2 \boldsymbol{A}_{t-1} + \alpha_3 K_X * \boldsymbol{X}_{t-1} + \epsilon_X$$

$$\boldsymbol{A}_t \sim Bern\left(\sigma\left(\beta_1\left(\beta_0 + \frac{1}{L}\sum_{l=0}^{L-1} K_A * \boldsymbol{X}_{t-l}\right)\right)\right)$$

$$\boldsymbol{Y}_t = \gamma_0 + \gamma_1 (K_{YA} * \boldsymbol{A}_{t-1}) + \gamma_2 \frac{1}{L}\sum_{l=1}^{L} K_{YX} * \boldsymbol{X}_{t-l} + \gamma_3 \boldsymbol{Y}_{t-1} + \epsilon_Y$$

Note: "*" is the convolution operation and β_1 controls the time-varying confounding.

- **Results:**

τ	Model	$\beta_1 = 0.0$	$\beta_1 = 0.5$	$\beta_1 = 1.0$	$\beta_1 = 1.5$	$\beta_1 = 2.0$	$\beta_1 = 2.5$	$\beta_1 = 3.0$
5	UNet+	**0.51 ± 0.00**	0.62 ± 0.01	0.84 ± 0.01	1.03 ± 0.01	1.10 ± 0.01	1.16 ± 0.01	1.25 ± 0.01
	STCINet	0.52 ± 0.00	0.68 ± 0.01	0.93 ± 0.01	1.11 ± 0.01	1.20 ± 0.01	1.33 ± 0.01	1.32 ± 0.01
	GST-UNet	0.55 ± 0.01	**0.61 ± 0.01**	**0.60 ± 0.01**	**0.61 ± 0.01**	**0.64 ± 0.01**	**0.58 ± 0.01**	**0.64 ± 0.01**
		(+7.8%)	(-1.6%)	(-28.6%)	(-40.8%)	(-41.8%)	(-50.0%)	(-48.8%)
10	UNet+	0.56 ± 0.00	0.53 ± 0.00	0.68 ± 0.00	0.85 ± 0.00	1.00 ± 0.01	1.02 ± 0.01	1.01 ± 0.01
	STCINet	0.57 ± 0.00	0.59 ± 0.00	0.74 ± 0.00	0.86 ± 0.01	1.11 ± 0.01	1.15 ± 0.01	1.26 ± 0.01
	GST-UNet	**0.50 ± 0.00**	**0.44 ± 0.01**	**0.45 ± 0.01**	**0.57 ± 0.01**	**0.48 ± 0.01**	**0.53 ± 0.01**	**0.49 ± 0.01**
		(-10.7%)	(-16.7%)	(-33.8%)	(-32.9%)	(-52.1%)	(-48.0%)	(-51.5%)

Case Study: Effect of Wildfire Smoke on Respiratory Illness during the 2018 California Camp Fire

- **Data (2018 California, county-level data [4]):**
  - **Covariates:** wind, temperature, precipitation, humidity, shortwave radiation
  - **"Treatment":** $PM_{2.5} > 10\ \mu g/m^3$ (unhealthy)
  - **Outcome:** Respiratory hospitalizations.
- **Counterfactual/ Policy-Relevant Question:**
  - How did unhealthy $PM_{2.5}$ (Camp Fire smoke) affect respiratory hospitalization?
  - If Camp Fire never occurred (i.e. $PM_{2.5}$ never exceeded 10 $\mu g/m^3$), how would the daily respiratory hospitalizations differ during the same time period?

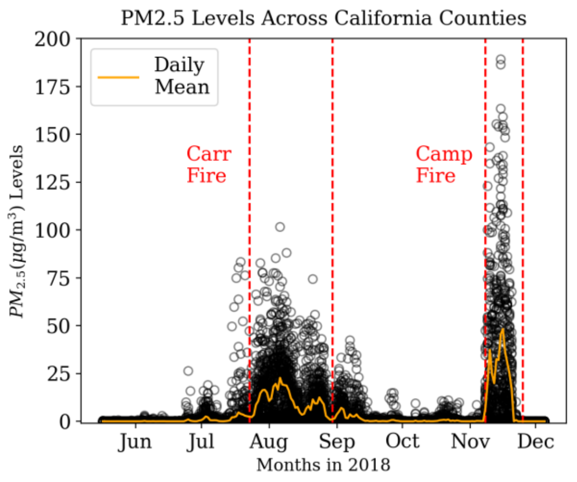

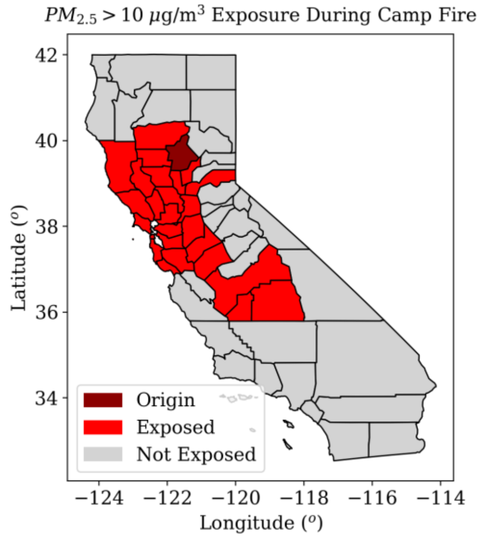

Case Study: Effect of Wildfire Smoke on Respiratory Illness during the 2018 California Camp Fire

Results

The GST-UNet estimates that the peak period of the Camp Fire (November 8–17, 2018) contributed to an excess 4,650 (465 per day)[1] respiratory-related hospitalizations in the affected counties.

[1] **Note:** This result aligns qualitatively with [4], who used a synthetic controls method and found about 259 excess daily cases from November 8–December 5 (including lower-intensity days, hence a smaller daily estimate).

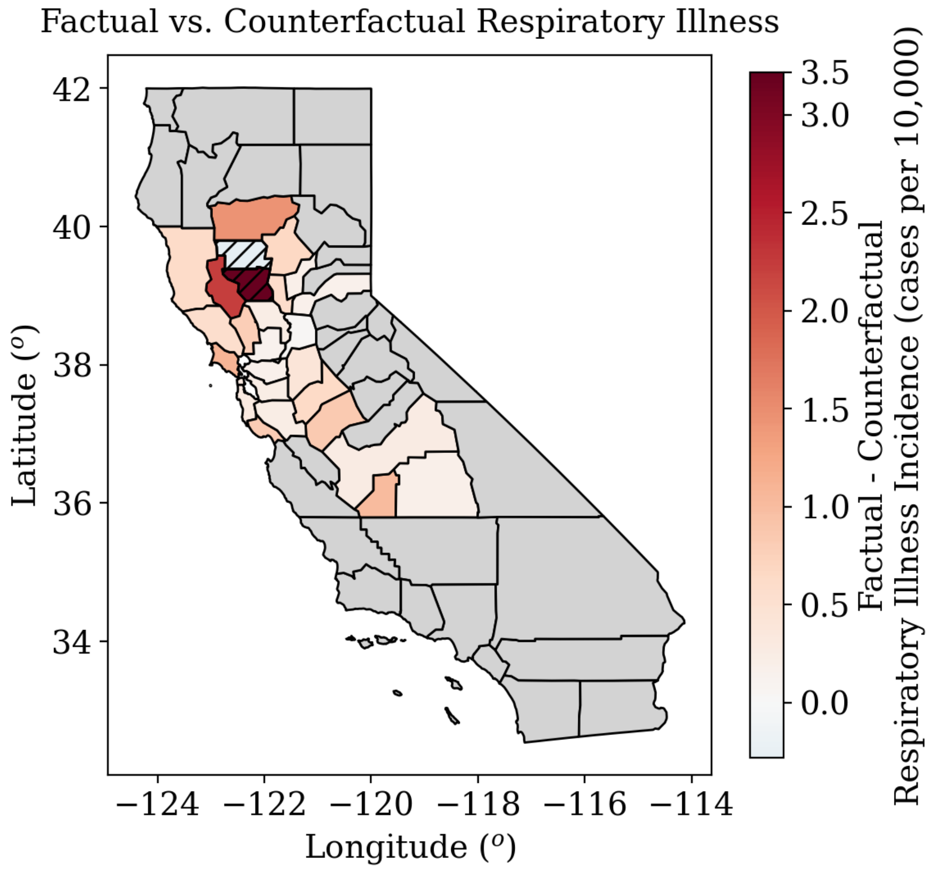

Observed minus predicted daily respiratory admissions at Camp Fire peak. Hashed areas mark small-population counties (<30,000).

Thank You!

Paper: *GST-UNet: Spatiotemporal Causal Inference with Time-Varying Confounders.*

Miruna Oprescu, David K. Park, Xihaier Luo, Shinjae Yoo, Nathan Kallus (Under Review, 2025).

Email: miruna@cs.cornell.edu.

References

[1] O. Ronneberger, P. Fischer, T. Brox. *U-Net: Convolutional Networks for Biomedical Image Segmentation*. MICCAI 2015.

[2] O. Oktay, J. Schlemper, L. Le Folgoc, et al.
Attention U-Net: Learning Where to Look for the Pancreas. MIDL 2018.

[3] J. Robins and M. Hernán.
Estimation of the causal effects of time-varying exposures.
In Chapman & Hall/CRC Handbooks of Modern Statistical Methods, 2008.

[4] N. Letellier, M. Hale, K. U. Salim, et al.
Applying a two-stage generalized synthetic control approach to quantify the heterogeneous health effects of extreme weather events: A 2018 large wildfire in California event as a case study.
Environmental Epidemiology, 2025.