GST-UNet: A Neural Framework for Spatiotemporal Causal Inference with Time-Varying Confounding

Miruna Oprescu[1], 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

Causal Inference in Spatiotemporal Contexts

- **Notation**
  - Time $t \in \{1, \ldots, T\}$, spatial index $s \in \mathbb{G}$.
  - **Features (Covariates):** $X_{s,1}, X_{s,2}, \ldots, X_{s,T}$.
  - **Interventions (Treatments):** $A_{s,1}, A_{s,2}, \ldots, A_{s,T} \in \{0,1\}$.
  - **Outcomes:** $Y_{s,1}, Y_{s,2}, \ldots, Y_{s,T}$.
  - **History:** $H_{s,1:t} = \left(X_{s,1:t}, Y_{s,1:t}, A_{s,1:t-1}\right)$.
- **Counterfactuals**

$$\mathbb{E}[\boldsymbol{Y}_{t+\tau}[\boldsymbol{a}_{t:t+\tau-1}] \mid \boldsymbol{H}_{1:t} = \boldsymbol{h}_{1:t}]$$

  - Average potential outcome after τ time steps under a series of fixed interventions, $\boldsymbol{a}_{t:t+\tau-1}$, given history $\boldsymbol{h}_{1:t}$.
  - “If factory filters had been installed earlier, how would health outcomes have changed over τ time steps?”

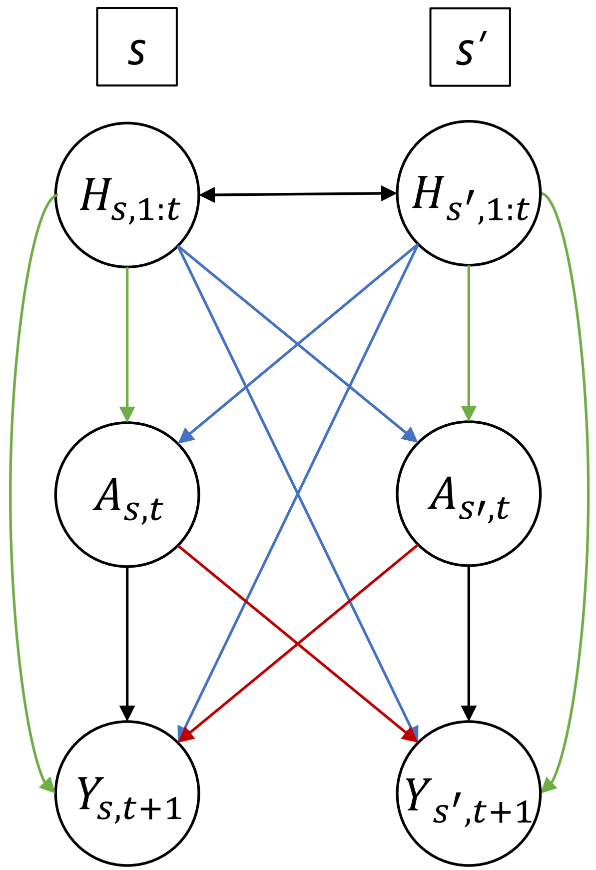

Schematic of the spatiotemporal data (X, A, Y, H) across time t and location s.

Challenges in Spatiotemporal Causal Inference

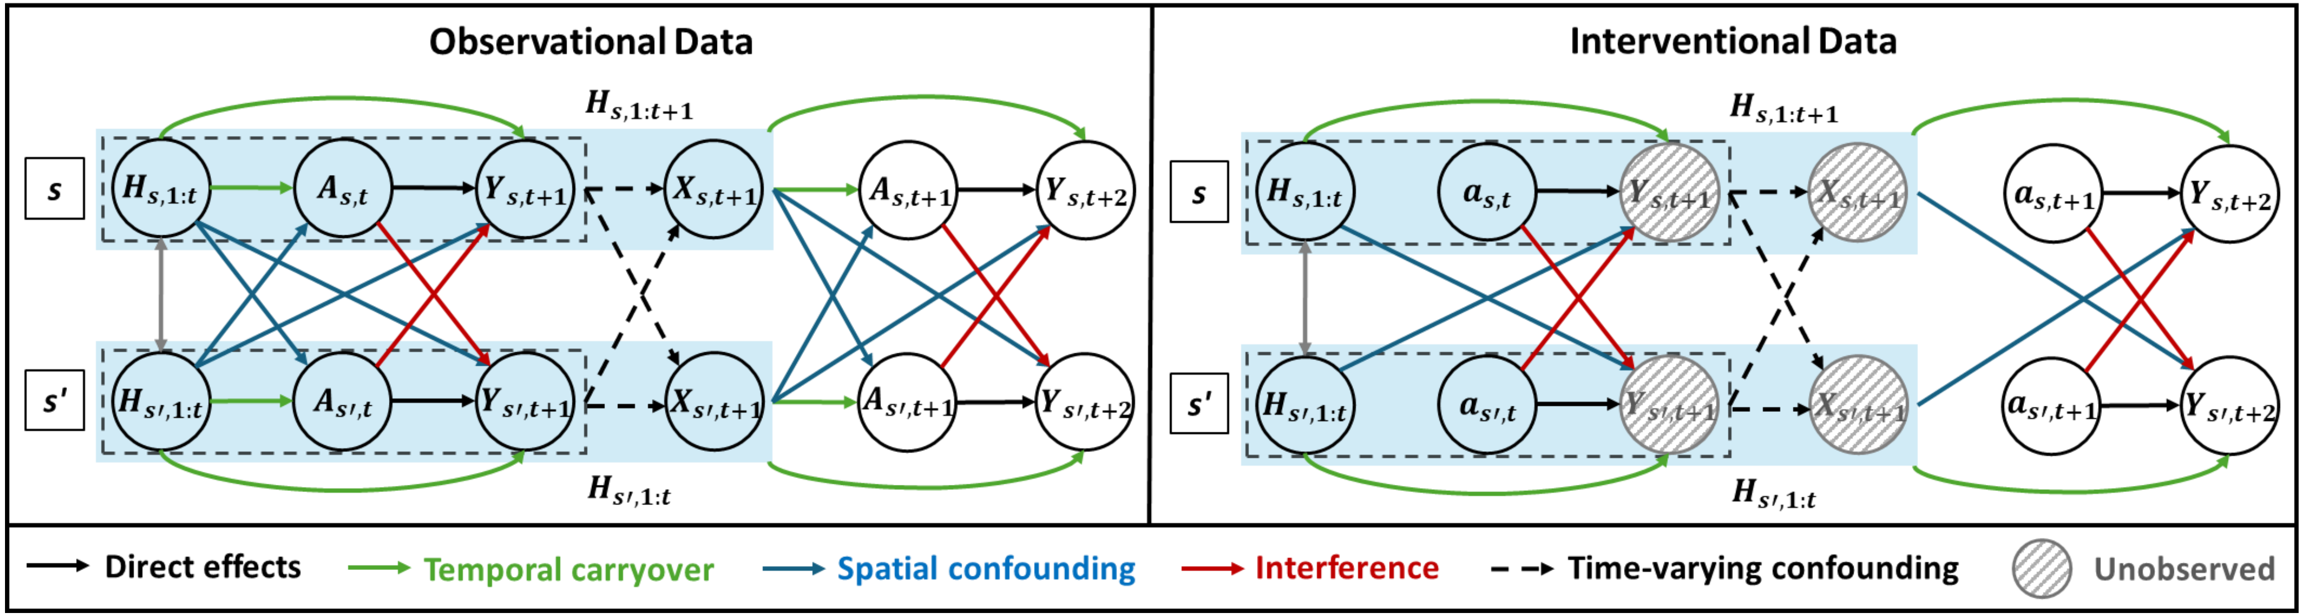

1. **Single Spatiotemporal Chain**
2. **Complex Space-Time Dependencies**
3. **Time-Varying Confounders**
   - Variables that affect both future treatments and outcomes, creating feedback loops (e.g. past interventions shape covariates, which in turn affect future interventions and outcomes).

Using a Single Spatiotemporal Chain

Assumption 1: Representation-Based Time Invariance

- There exists an embedding $\phi: \mathcal{H} \times \mathcal{A} \to Z \subset \mathbb{R}^h$ such that, once we condition on $z = \phi(\boldsymbol{H}_{1:t}, \boldsymbol{A}_t)$ the distribution of $(\boldsymbol{X}_{t+1}, \boldsymbol{Y}_{t+1})$ does not explicitly depend on t. Formally:

$$p(\boldsymbol{X}_{t+1}, \boldsymbol{Y}_{t+1} \mid \phi(\boldsymbol{H}_{1:t}, \boldsymbol{A}_t) = z) = p(\boldsymbol{X}_{t'+1}, \boldsymbol{Y}_{t'+1} \mid \phi(\boldsymbol{H}_{1:t'}, \boldsymbol{A}_{t'}) = z)$$

Splicing the Single Time Series

- For each $t \in \{1, \ldots, T - \tau\}$, define a "prefix"

$$\boldsymbol{P}_t^{\tau} = (\boldsymbol{X}_{1:t+\tau}, \boldsymbol{A}_{1:t+\tau}, \boldsymbol{Y}_{1:t+\tau})$$

- Under representation-based time invariance, conditioning on $\phi(\boldsymbol{H}_{1:t}, \boldsymbol{A}_t)$ renders the distribution of $\boldsymbol{Y}_{t+\tau}$ independent of t.
- We can then write expectations over these prefixes as

$$\mathbb{E}_{\boldsymbol{P}}[\boldsymbol{Y}_{t+\tau} \mid \phi(\boldsymbol{H}_{1:t}, \boldsymbol{A}_t)]$$

Causal Inference with Time-Varying Confounders

Iterative G-Computation via Recursive Regression

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 $k = \tau - 1, \ldots, 1$):

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

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 $\boldsymbol{a}_{t:t+k+1}$.

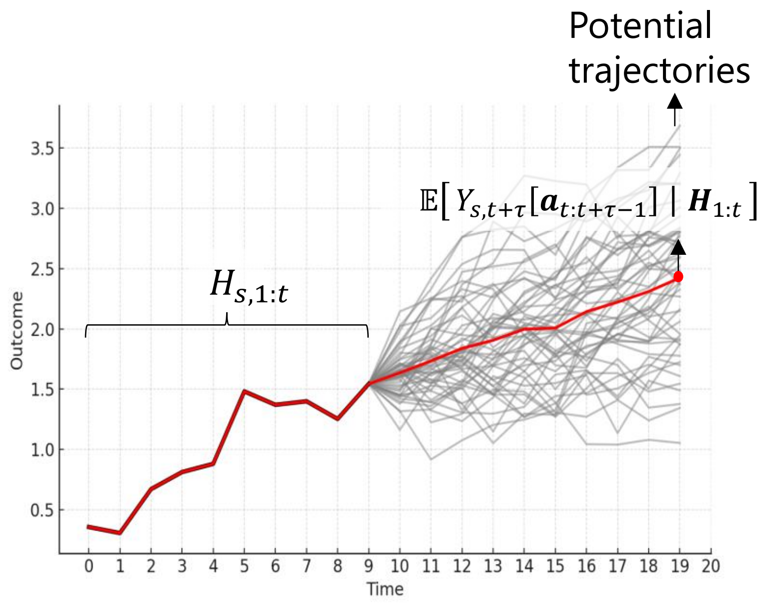

Introducing the GST-UNet

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.

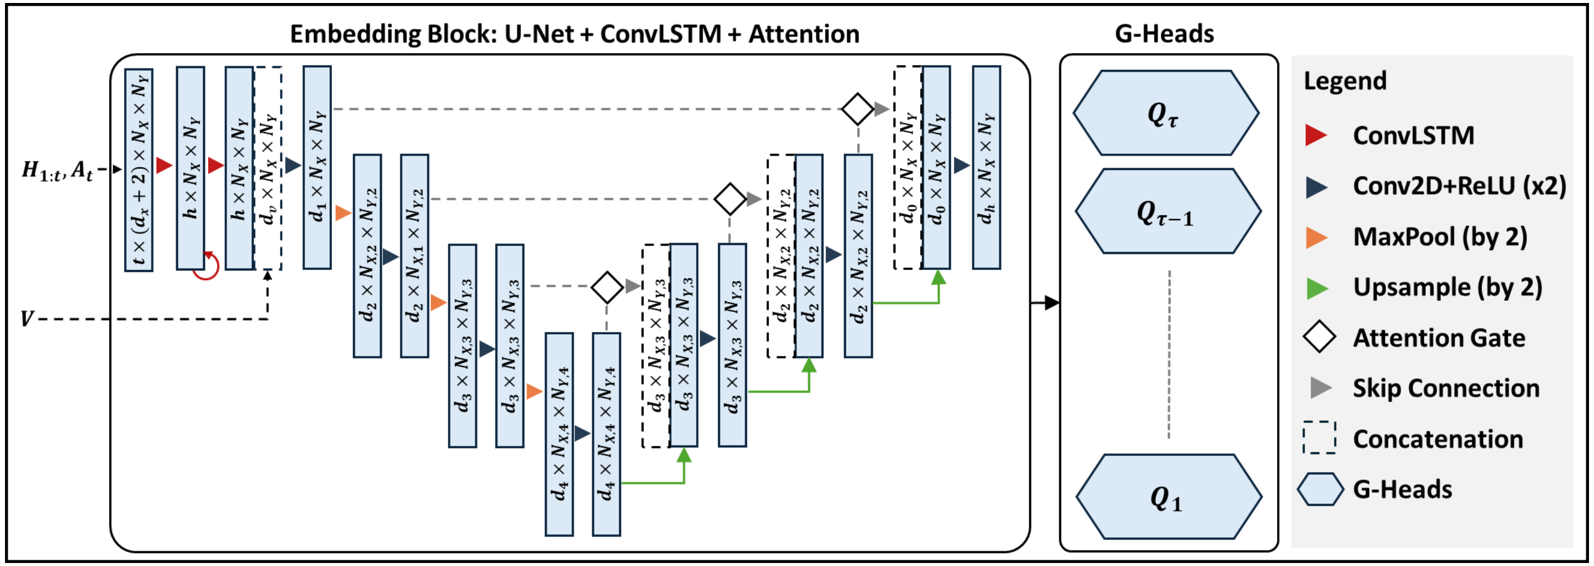

GST-UNet End-to-End Architecture

Simulation Results on Synthetic Data

Table 1: RMSE ± SD across test trajectories. Bold indicates lowest error per column; color shows improvement (RMSE **decrease** or **increase**) over best baseline (excluding ablations).

τ	Model	$\beta_1 = 0.0$	$\beta_1 = 0.5$	$\beta_1 = 1.0$	$\beta_1 = 1.5$	$\beta_1 = 2.0$
5	UNet+	**0.28 ± 0.00**	0.36 ± 0.00	0.54 ± 0.01	0.71 ± 0.01	0.81 ± 0.01
	STCINet	0.29 ± 0.00	0.38 ± 0.01	0.62 ± 0.01	0.80 ± 0.01	0.90 ± 0.01
	IPWUNet	0.60 ± 0.01	0.58 ± 0.01	0.58 ± 0.01	0.59 ± 0.01	0.59 ± 0.01
	GST-UNet w/o Attention	0.50 ± 0.00	0.46 ± 0.00	0.51 ± 0.00	0.45 ± 0.01	0.47 ± 0.01
	GST-UNet w/o Curriculum	0.69 ± 0.00	0.64 ± 0.00	0.63 ± 0.00	0.61 ± 0.01	0.61 ± 0.01
	GST-UNet	0.33 ± 0.00	**0.35 ± 0.00**	**0.40 ± 0.00**	**0.44 ± 0.00**	**0.40 ± 0.01**
		(**+17.9%**)	(**-2.7%**)	(**-21.6%**)	(**-25.4%**)	(**-32.2%**)
10	UNet+	**0.28 ± 0.00**	0.61 ± 0.00	1.18 ± 0.00	1.45 ± 0.00	1.71 ± 0.01
	STCINet	0.31 ± 0.00	0.68 ± 0.00	1.25 ± 0.00	1.47 ± 0.01	1.60 ± 0.01
	IPWUNet	0.78 ± 0.01	0.80 ± 0.01	0.96 ± 0.01	1.19 ± 0.02	1.08 ± 0.01
	GST-UNet w/o Attention	0.42 ± 0.00	0.60 ± 0.00	0.61 ± 0.00	0.79 ± 0.01	1.07 ± 0.01
	GST-UNet w/o Curriculum	0.62 ± 0.00	0.88 ± 0.00	1.02 ± 0.00	1.08 ± 0.01	1.12 ± 0.01
	GST-UNet	0.38 ± 0.00	**0.55 ± 0.00**	**0.68 ± 0.00**	**0.73 ± 0.01**	**0.85 ± 0.01**
		(**+35.7%**)	(**-9.8%**)	(**-29.2%**)	(**-38.7%**)	(**-21.3%**)

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?

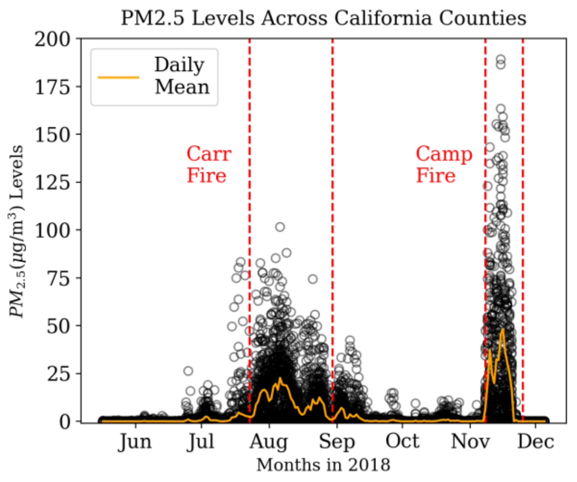

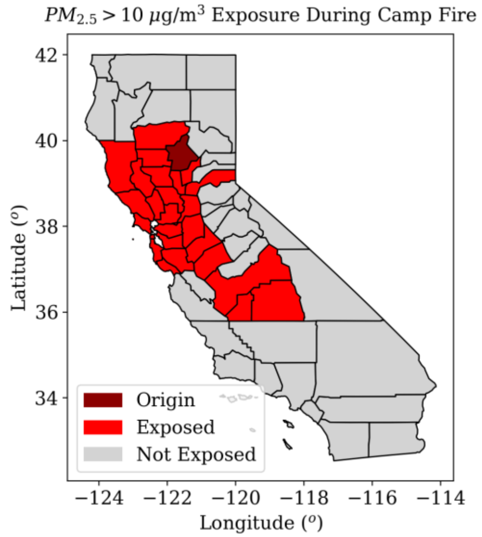

Case Study: Effect of Wildfire Smoke on Respiratory Illness during the 2018 California Camp Fire

Results

The GST-UNet estimates that the peak period of the Camp Fire (November 8–17, 2018) contributed to an excess 4650 ([1890, 6535] 95% CI) (465 per day)[1] respiratory-related hospitalizations in the affected counties.

Baseline Predictions

- **UNet+:** 3911 ([–899, 5202] 95% CI)
- **STCINet:** 343 ([–3077, 3281] 95% CI)

[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).

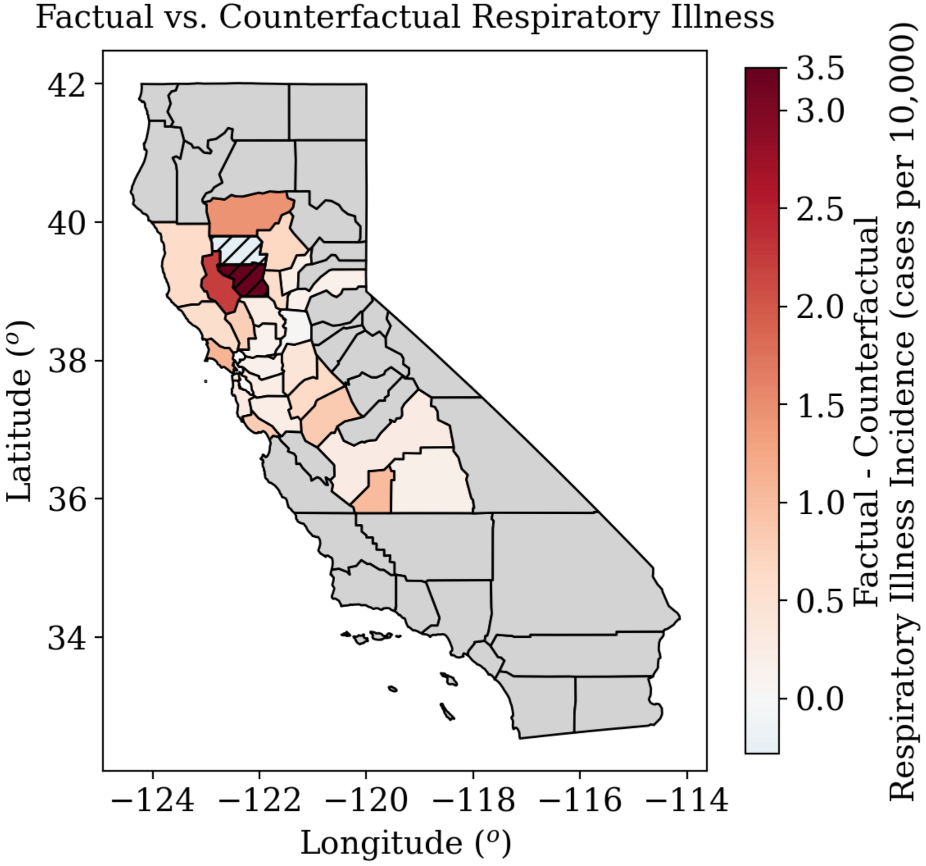

Observed minus predicted daily respiratory admissions at Camp Fire peak. Hashed areas mark small-population counties (<30,000).

Summary of Contributions and Impact

Key Contributions:

- **GST-UNet**: A **spatiotemporal encoder** (U-Net + ConvLSTM + attention) + **regression-based iterative G-computation** to estimate location-specific counterfactuals under complex treatment sequences.
- We establish **identification from a single observed trajectory** via a **representation-based time-invariance** assumption and prove **consistency** of the neural estimator.
- **Empirics:** In advection–diffusion simulations with increasing time-varying confounding, **GST-UNet remains accurate** while baselines degrade.

Broader Impact:

- Enables **credible effect estimation** from observational spatiotemporal data: e.g., the 2018 Camp Fire analysis estimates **≈4,650** excess respiratory ED visits (95% CI: 1,888–6,535)—supporting **public health** and **environmental policy**.