GST-UNet: Spatiotemporal Causal Inference with Time-Varying Confounders

Space X Time X Causality Reading Group

Gerrit Großmann, May 28

Reminder

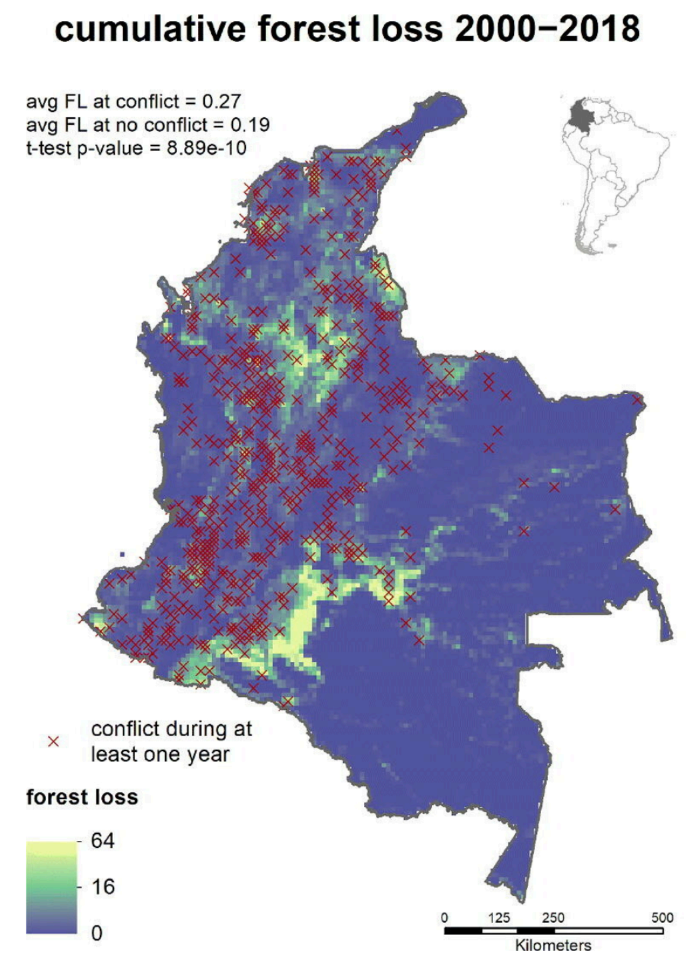

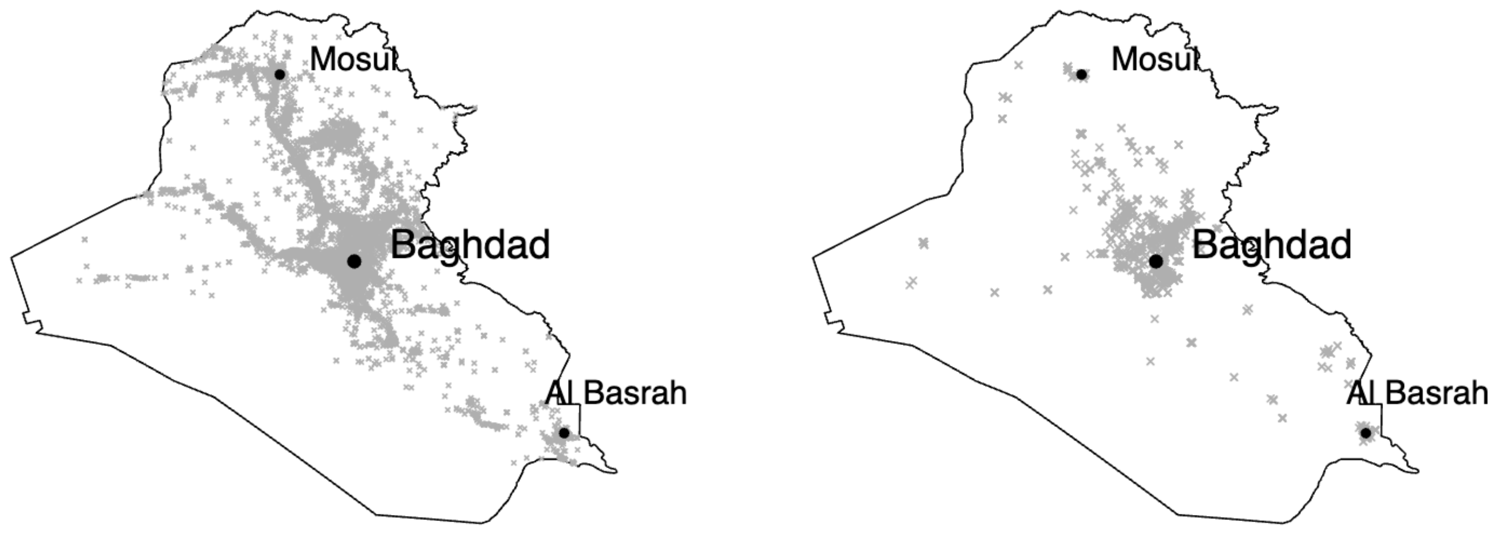

(a) Insurgent violence

(b) Airstrikes

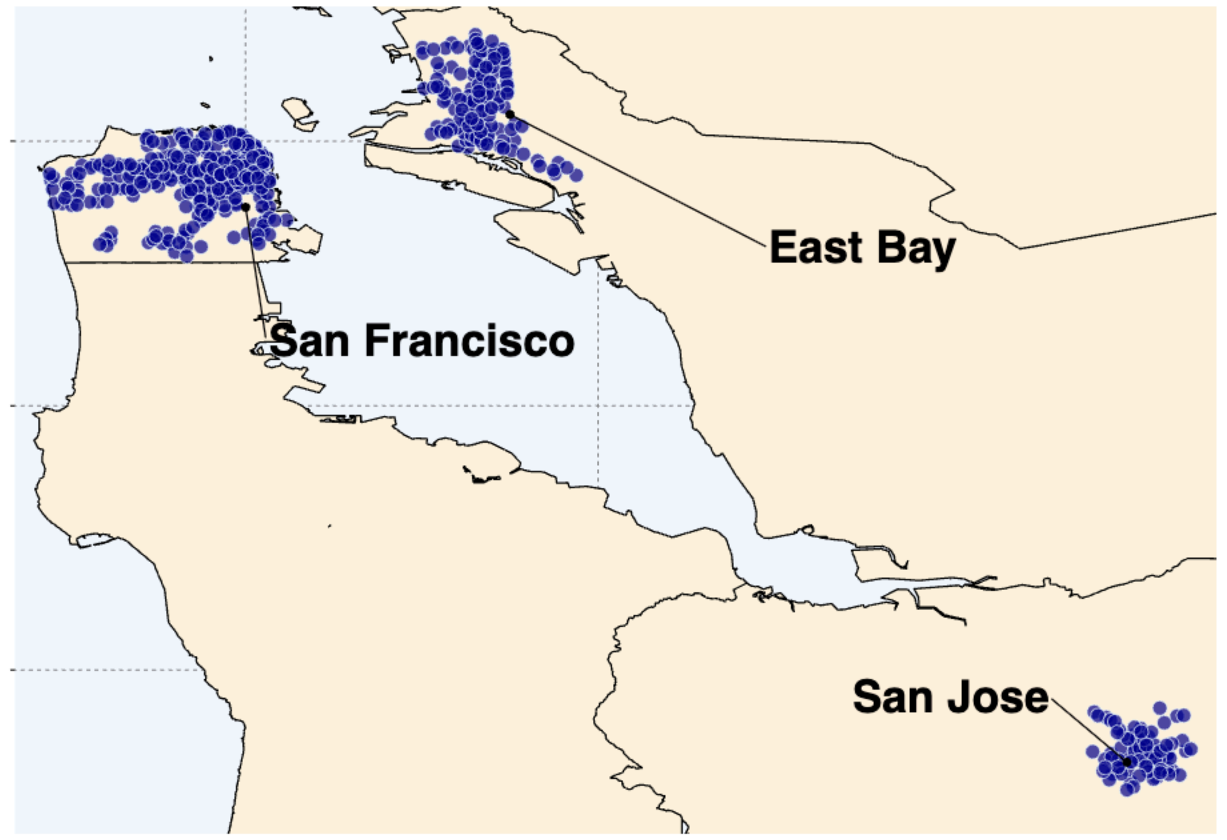

Conflict
Deforestation
Road Network

Airstrikes
Insurgency Violence
Aid

Wildfire
Bike Rental Hours
Temperature + Wind

Challanges

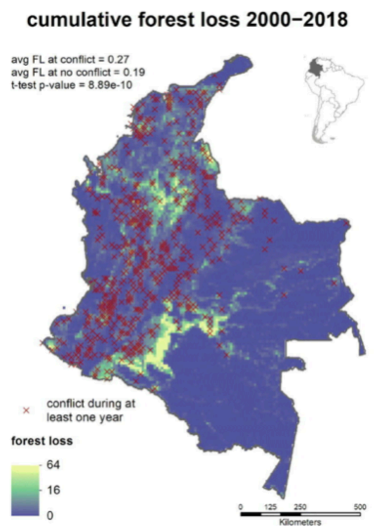

Conflict
Deforestation
Road Network

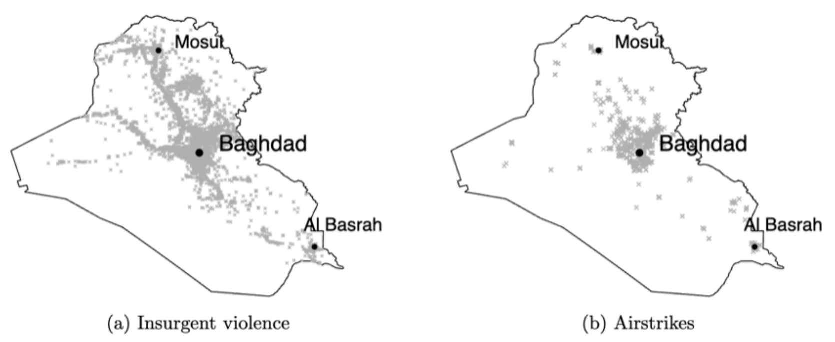

Airstrikes
Insurgency Violence
Aid

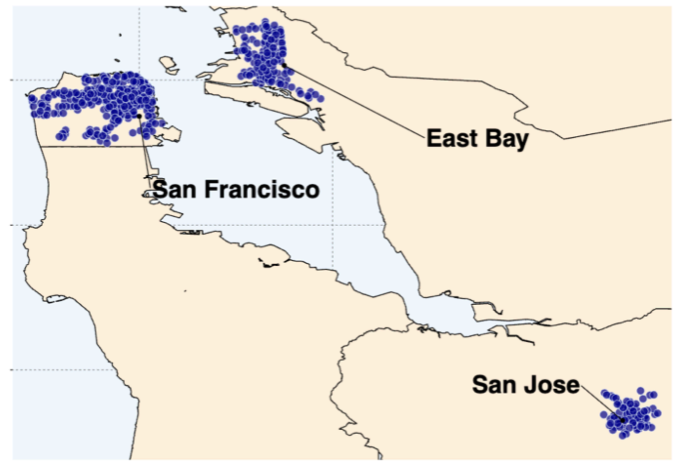

Wildfire
Bike Rental Hours
Temperature + Wind

(Temporal) **Carryover:** A wildfire today increases bike rentals tomorrow.

(Spatial) **Spillover:** A wildfire in one county drifts into neighboring counties.

Time-Varying Confounding: Weather affects wildfires and bike rentals.

Treatment–Covariate Feedback: A Covid lockdown reduces mobility, which in turn affects lockdown necessity.

GST-UNet: Spatiotemporal Causal Inference with Time-Varying Confounders

Miruna Oprescu [1] David K. Park [2] Xihaier Luo [2] Shinjae Yoo [2] Nathan Kallus [1]

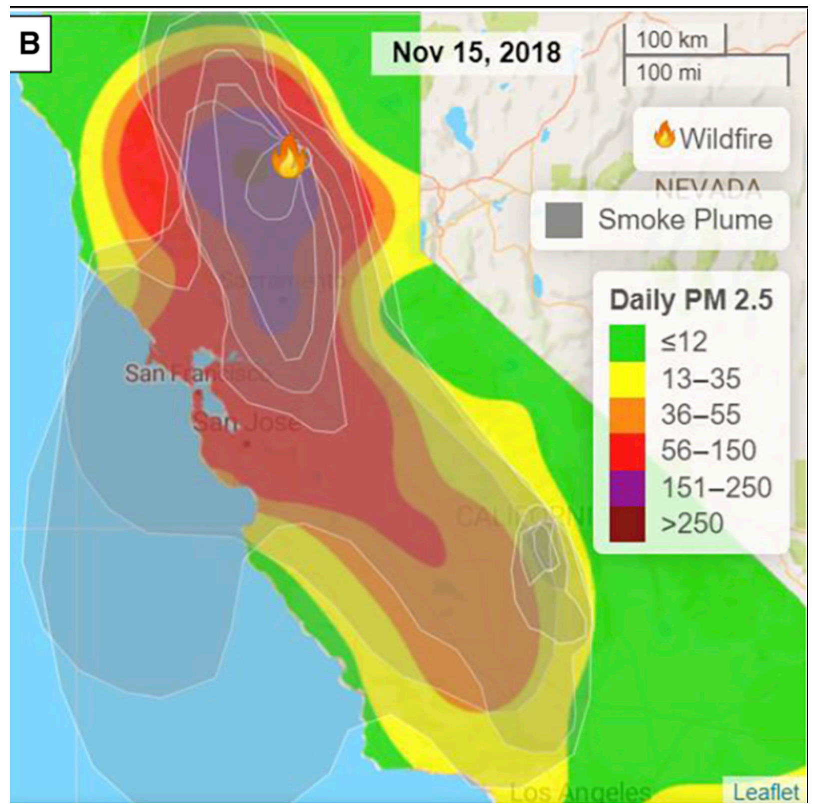

Wildfire Air Pollution and Rates of Cardiovascular Events and Mortality in Northern California in 2018 (Alexeeff et al., 2025)

?

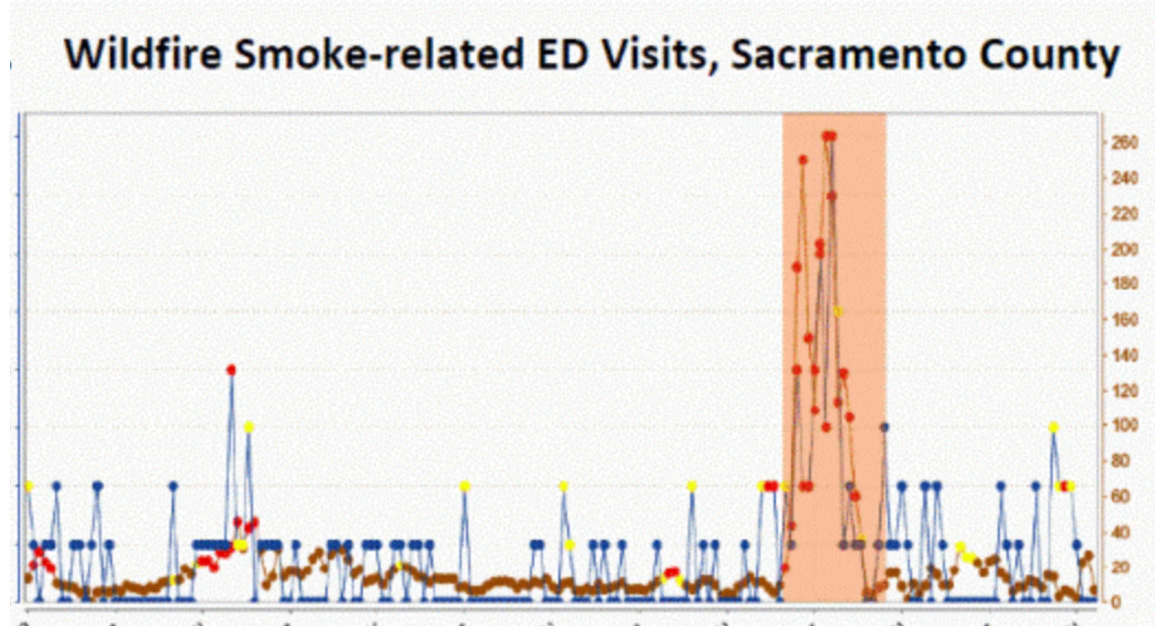

https://www.cdc.gov/nssp/php/story/regional-collaboration-during-california-wildfires.html

Wildfire smoke exposure during the 2018 Camp Fire

Respiratory hospitalizations

GST-UNet: Spatiotemporal Causal Inference with Time-Varying Confounders

Miruna Oprescu [1] David K. Park [2] Xihaier Luo [2] Shinjae Yoo [2] Nathan Kallus [1]

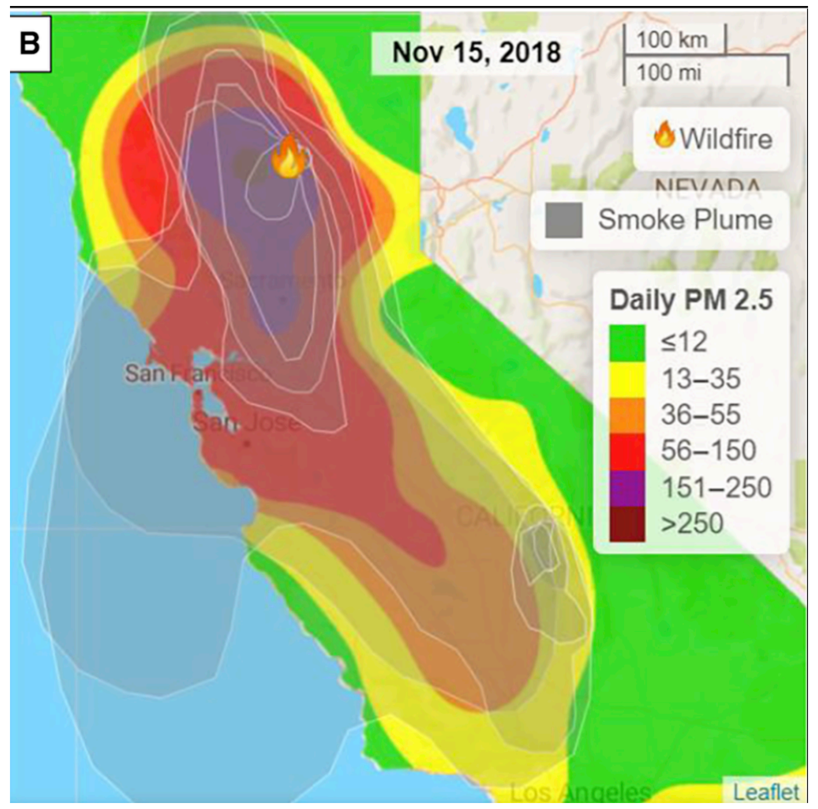

Wildfire Air Pollution and Rates of Cardiovascular Events and Mortality in Northern California in 2018 (Alexeeff et al., 2025)

?

- G-computation
- NN-based spatio-temporal prediction
- Single realization of the process

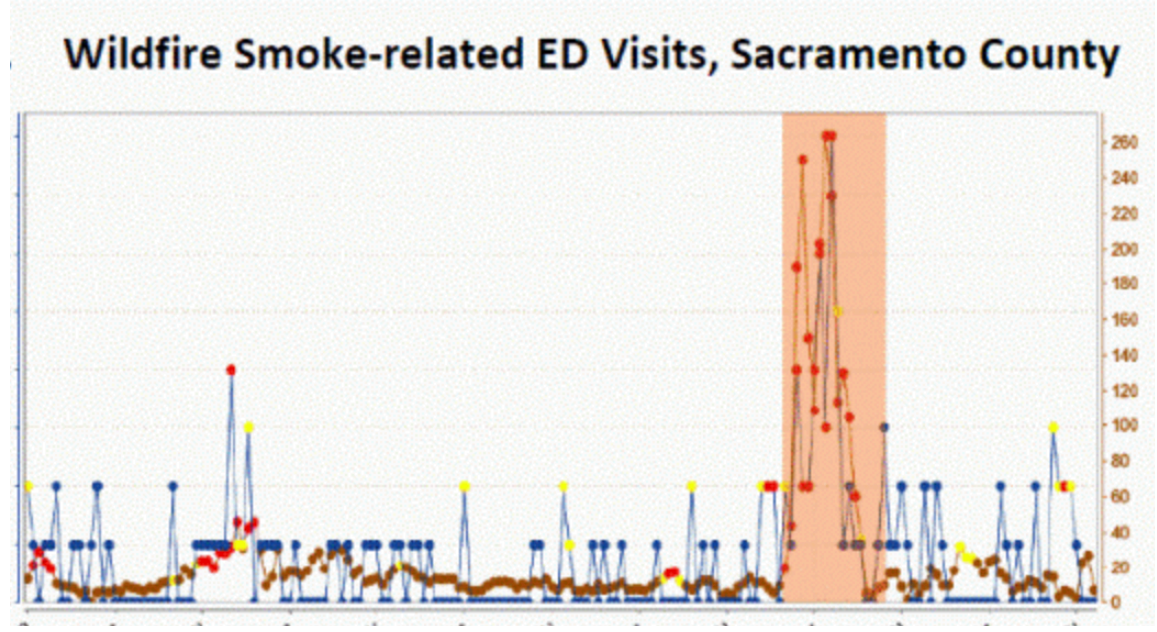

https://www.cdc.gov/nssp/php/story/regional-collaboration-during-california-wildfires.html

Wildfire smoke exposure during the 2018 Camp Fire

Respiratory hospitalizations

Running Example

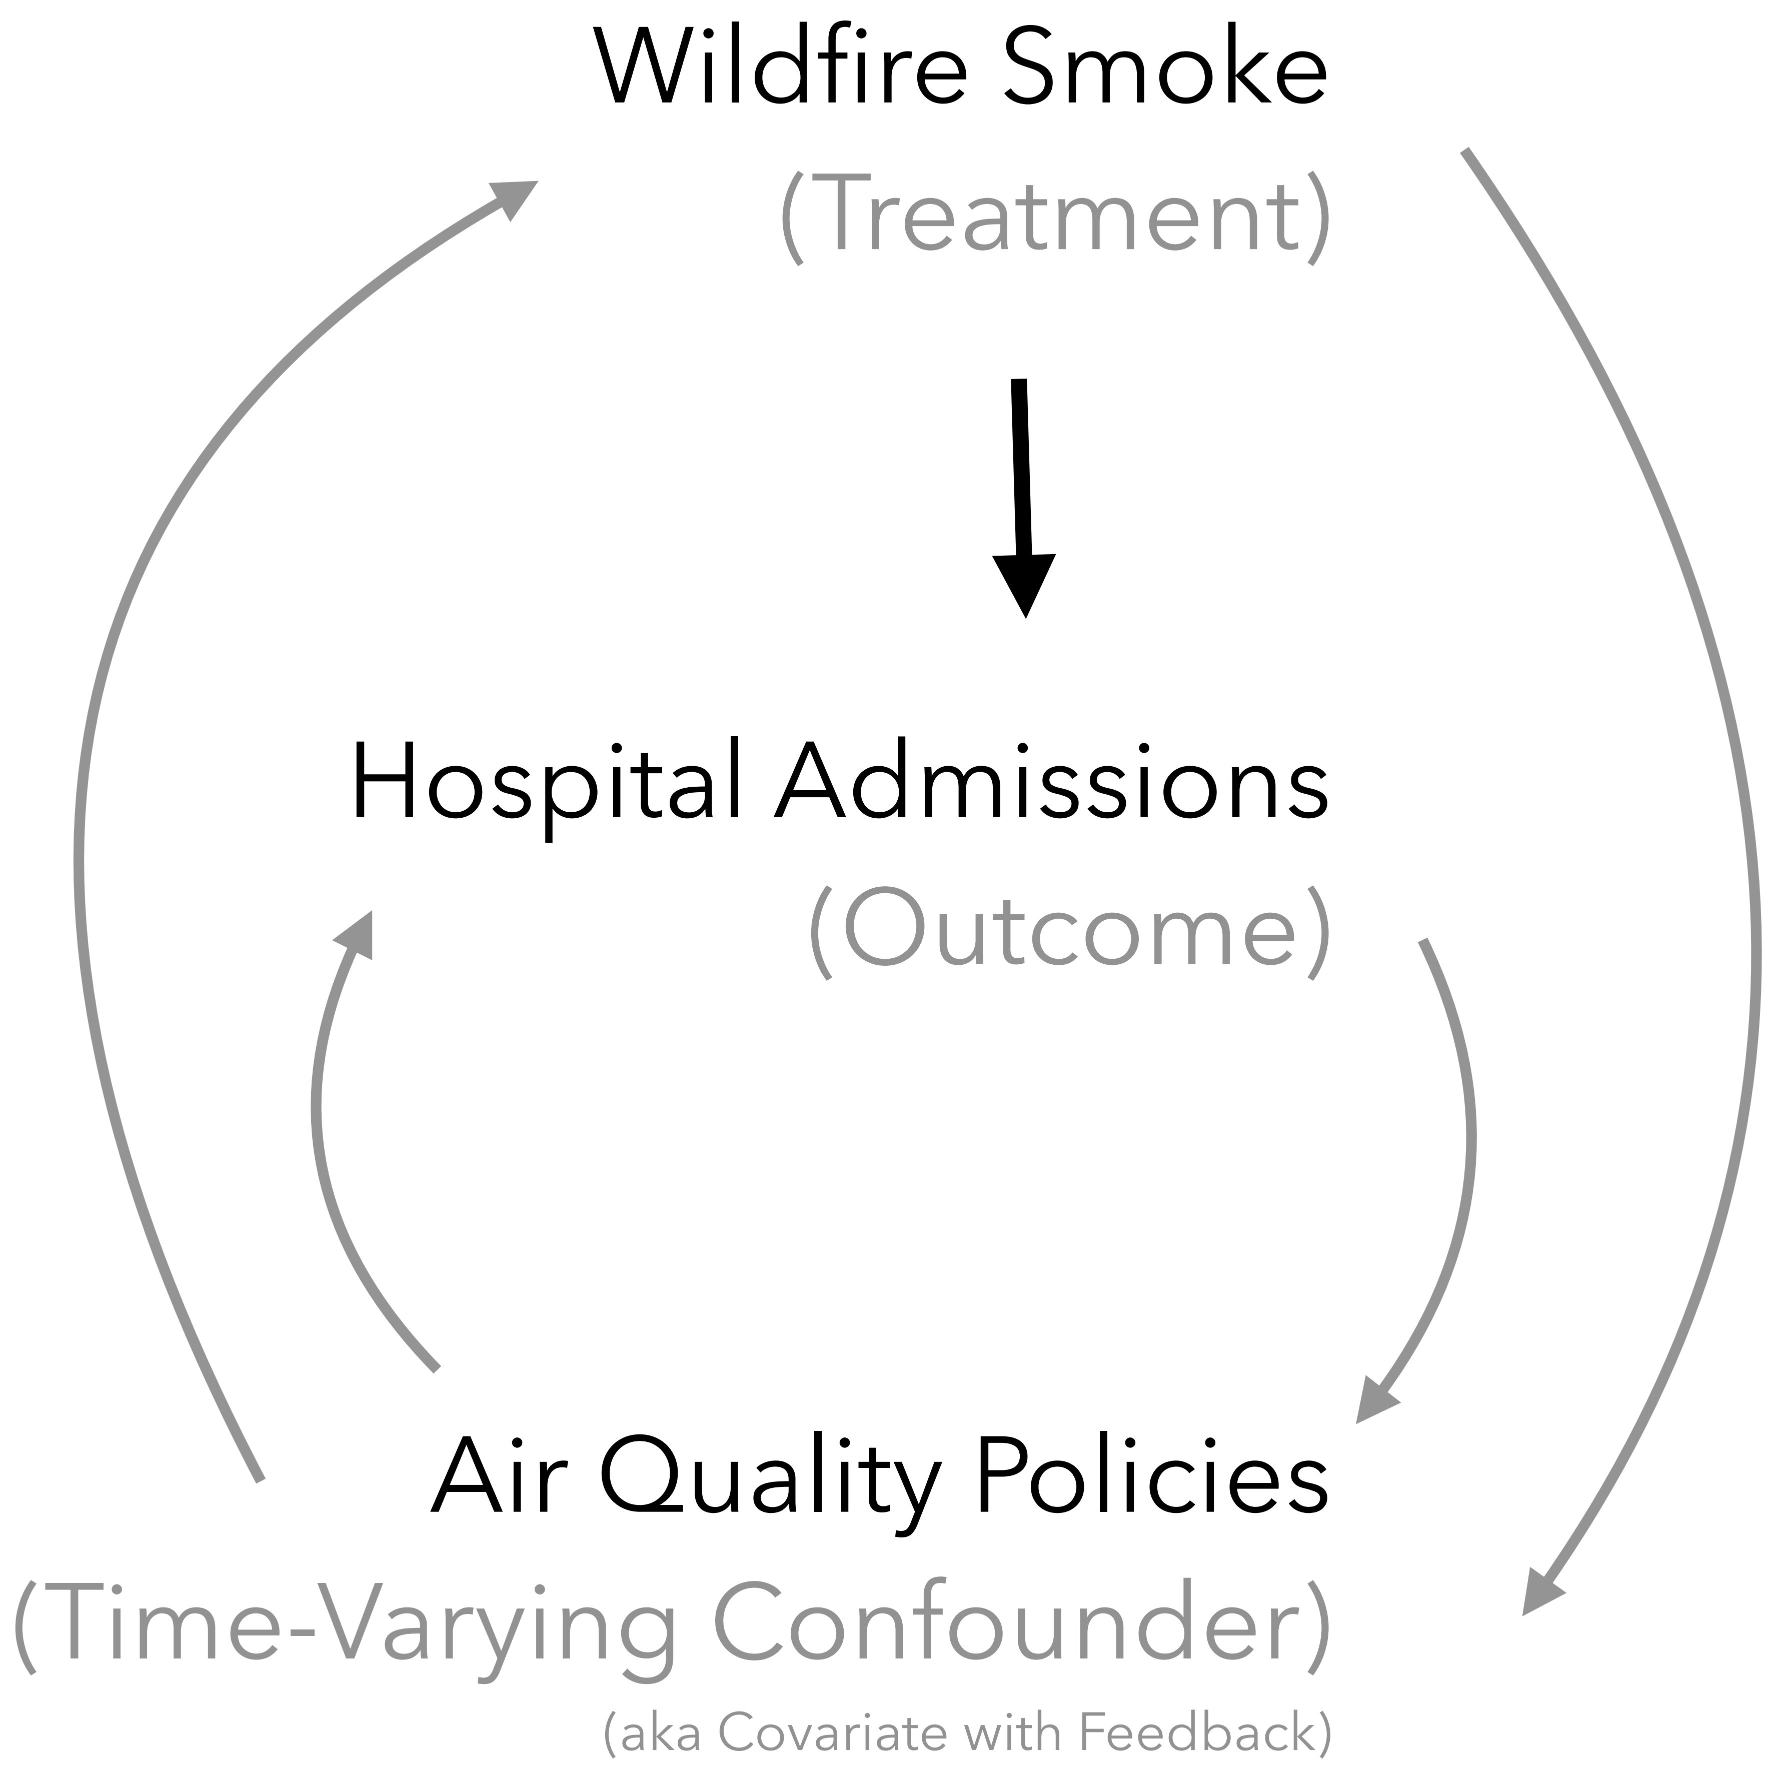

Running Example

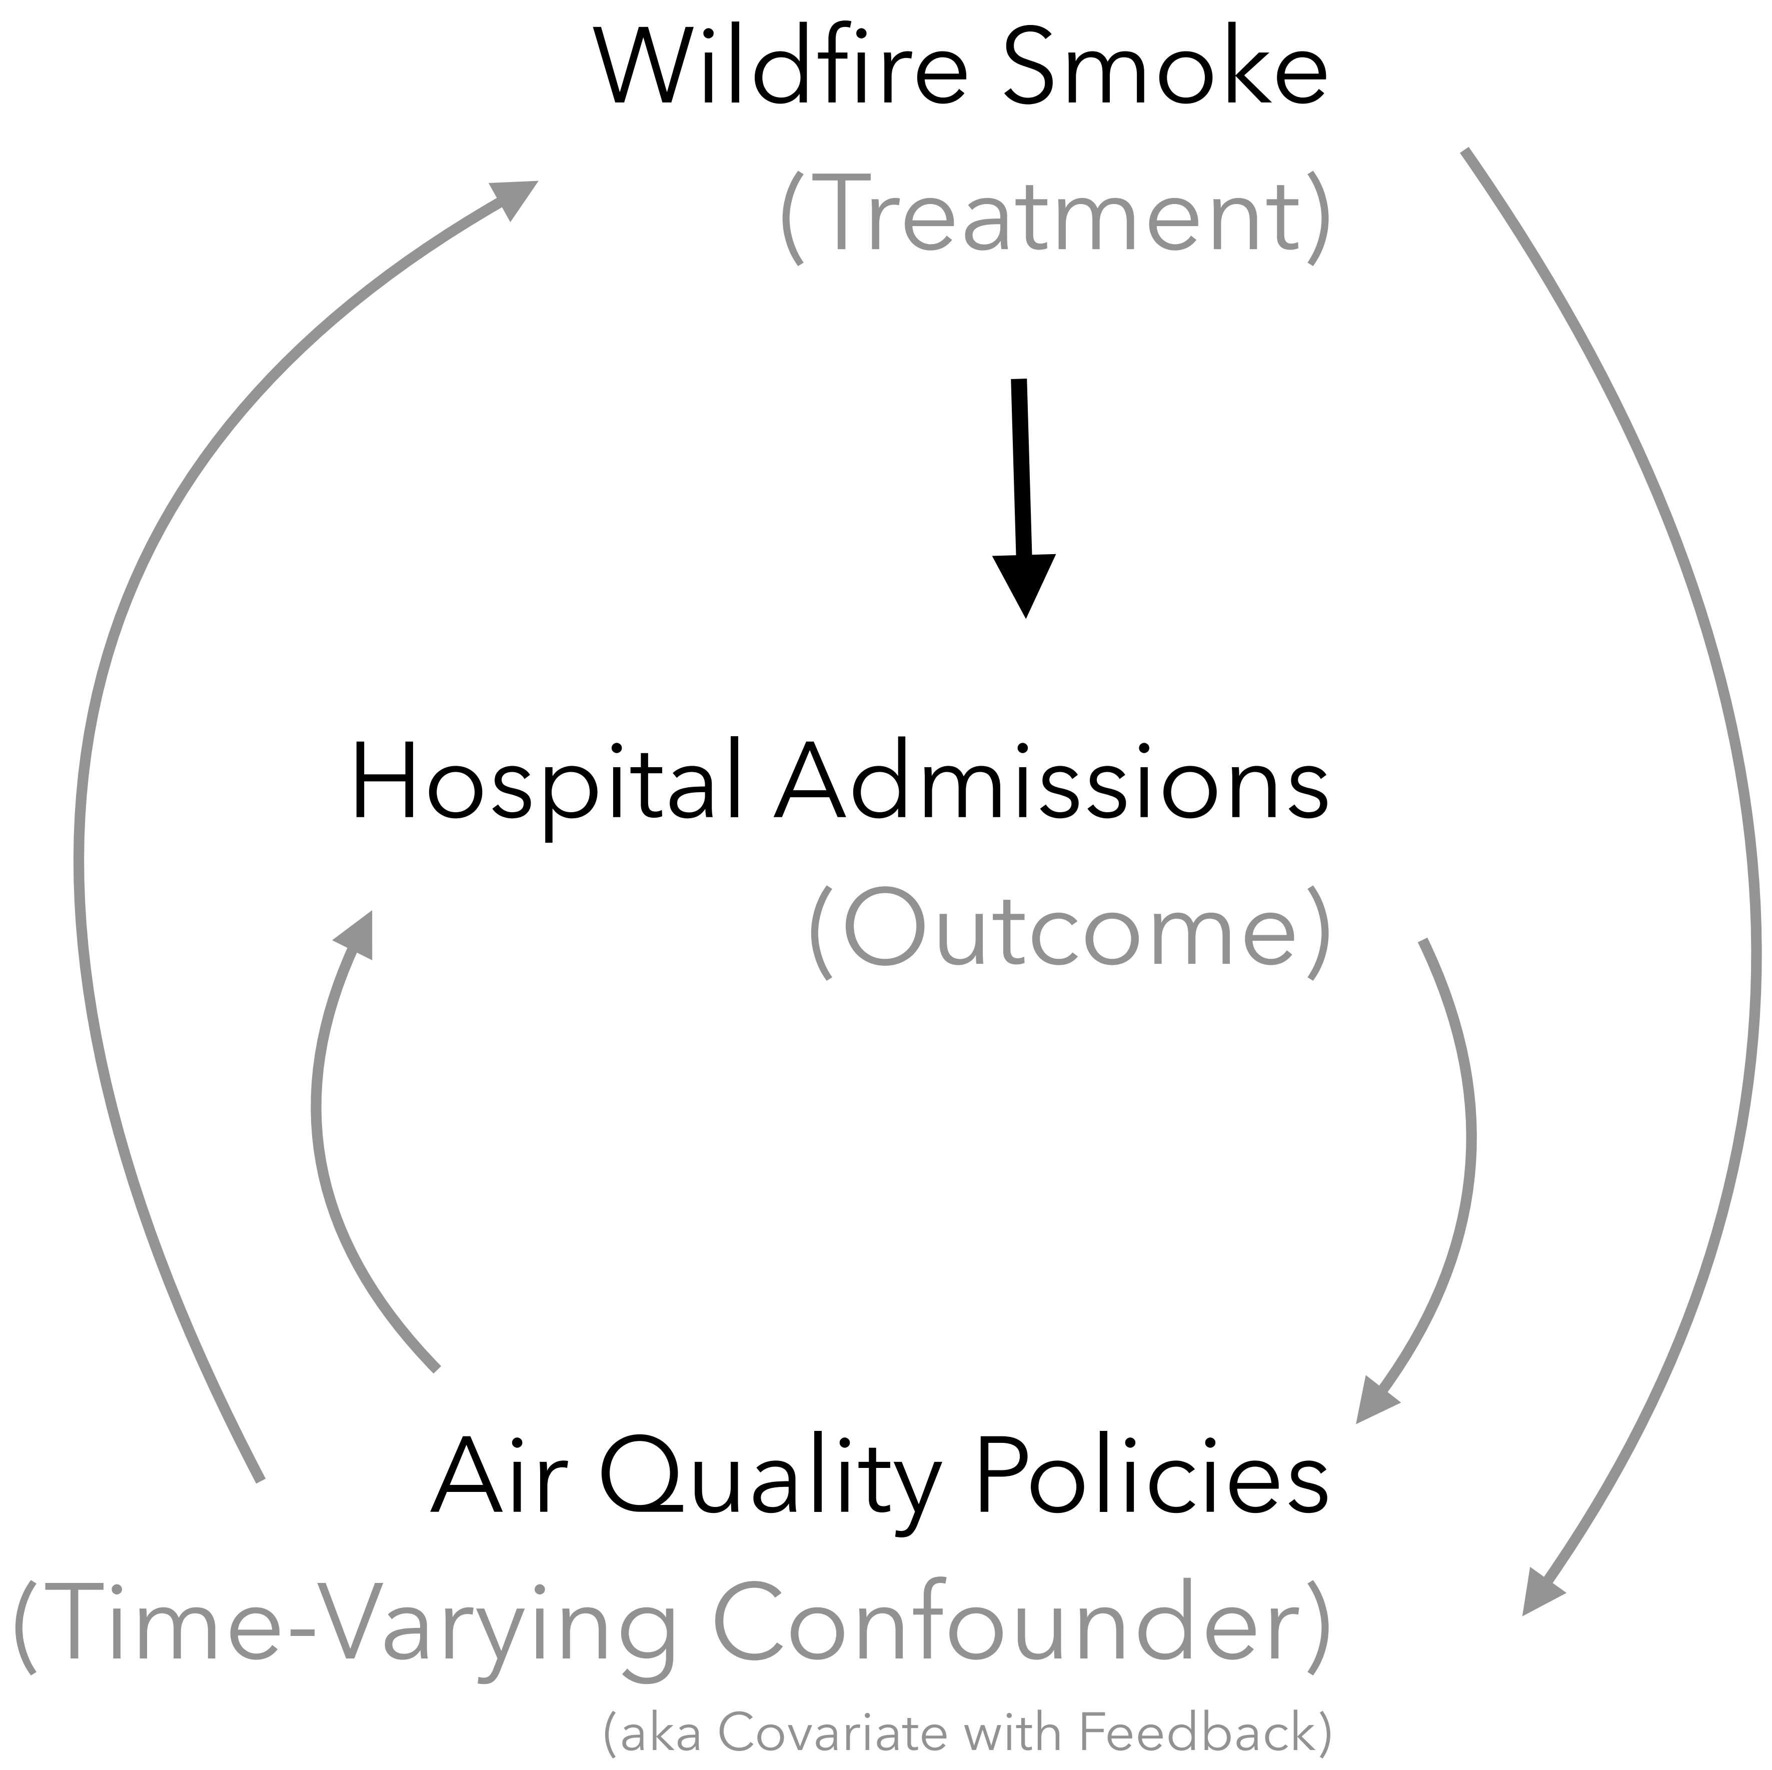

Time-varying confounders:
Covariates that both influence and are influenced by past treatments and outcomes (creates feedback loops).

Running Example

Wildfire Smoke
(Treatment)

Hospital Admissions
(Outcome)

Air Quality Policies
(Time-Varying Confounder)
(aka Covariate with Feedback)

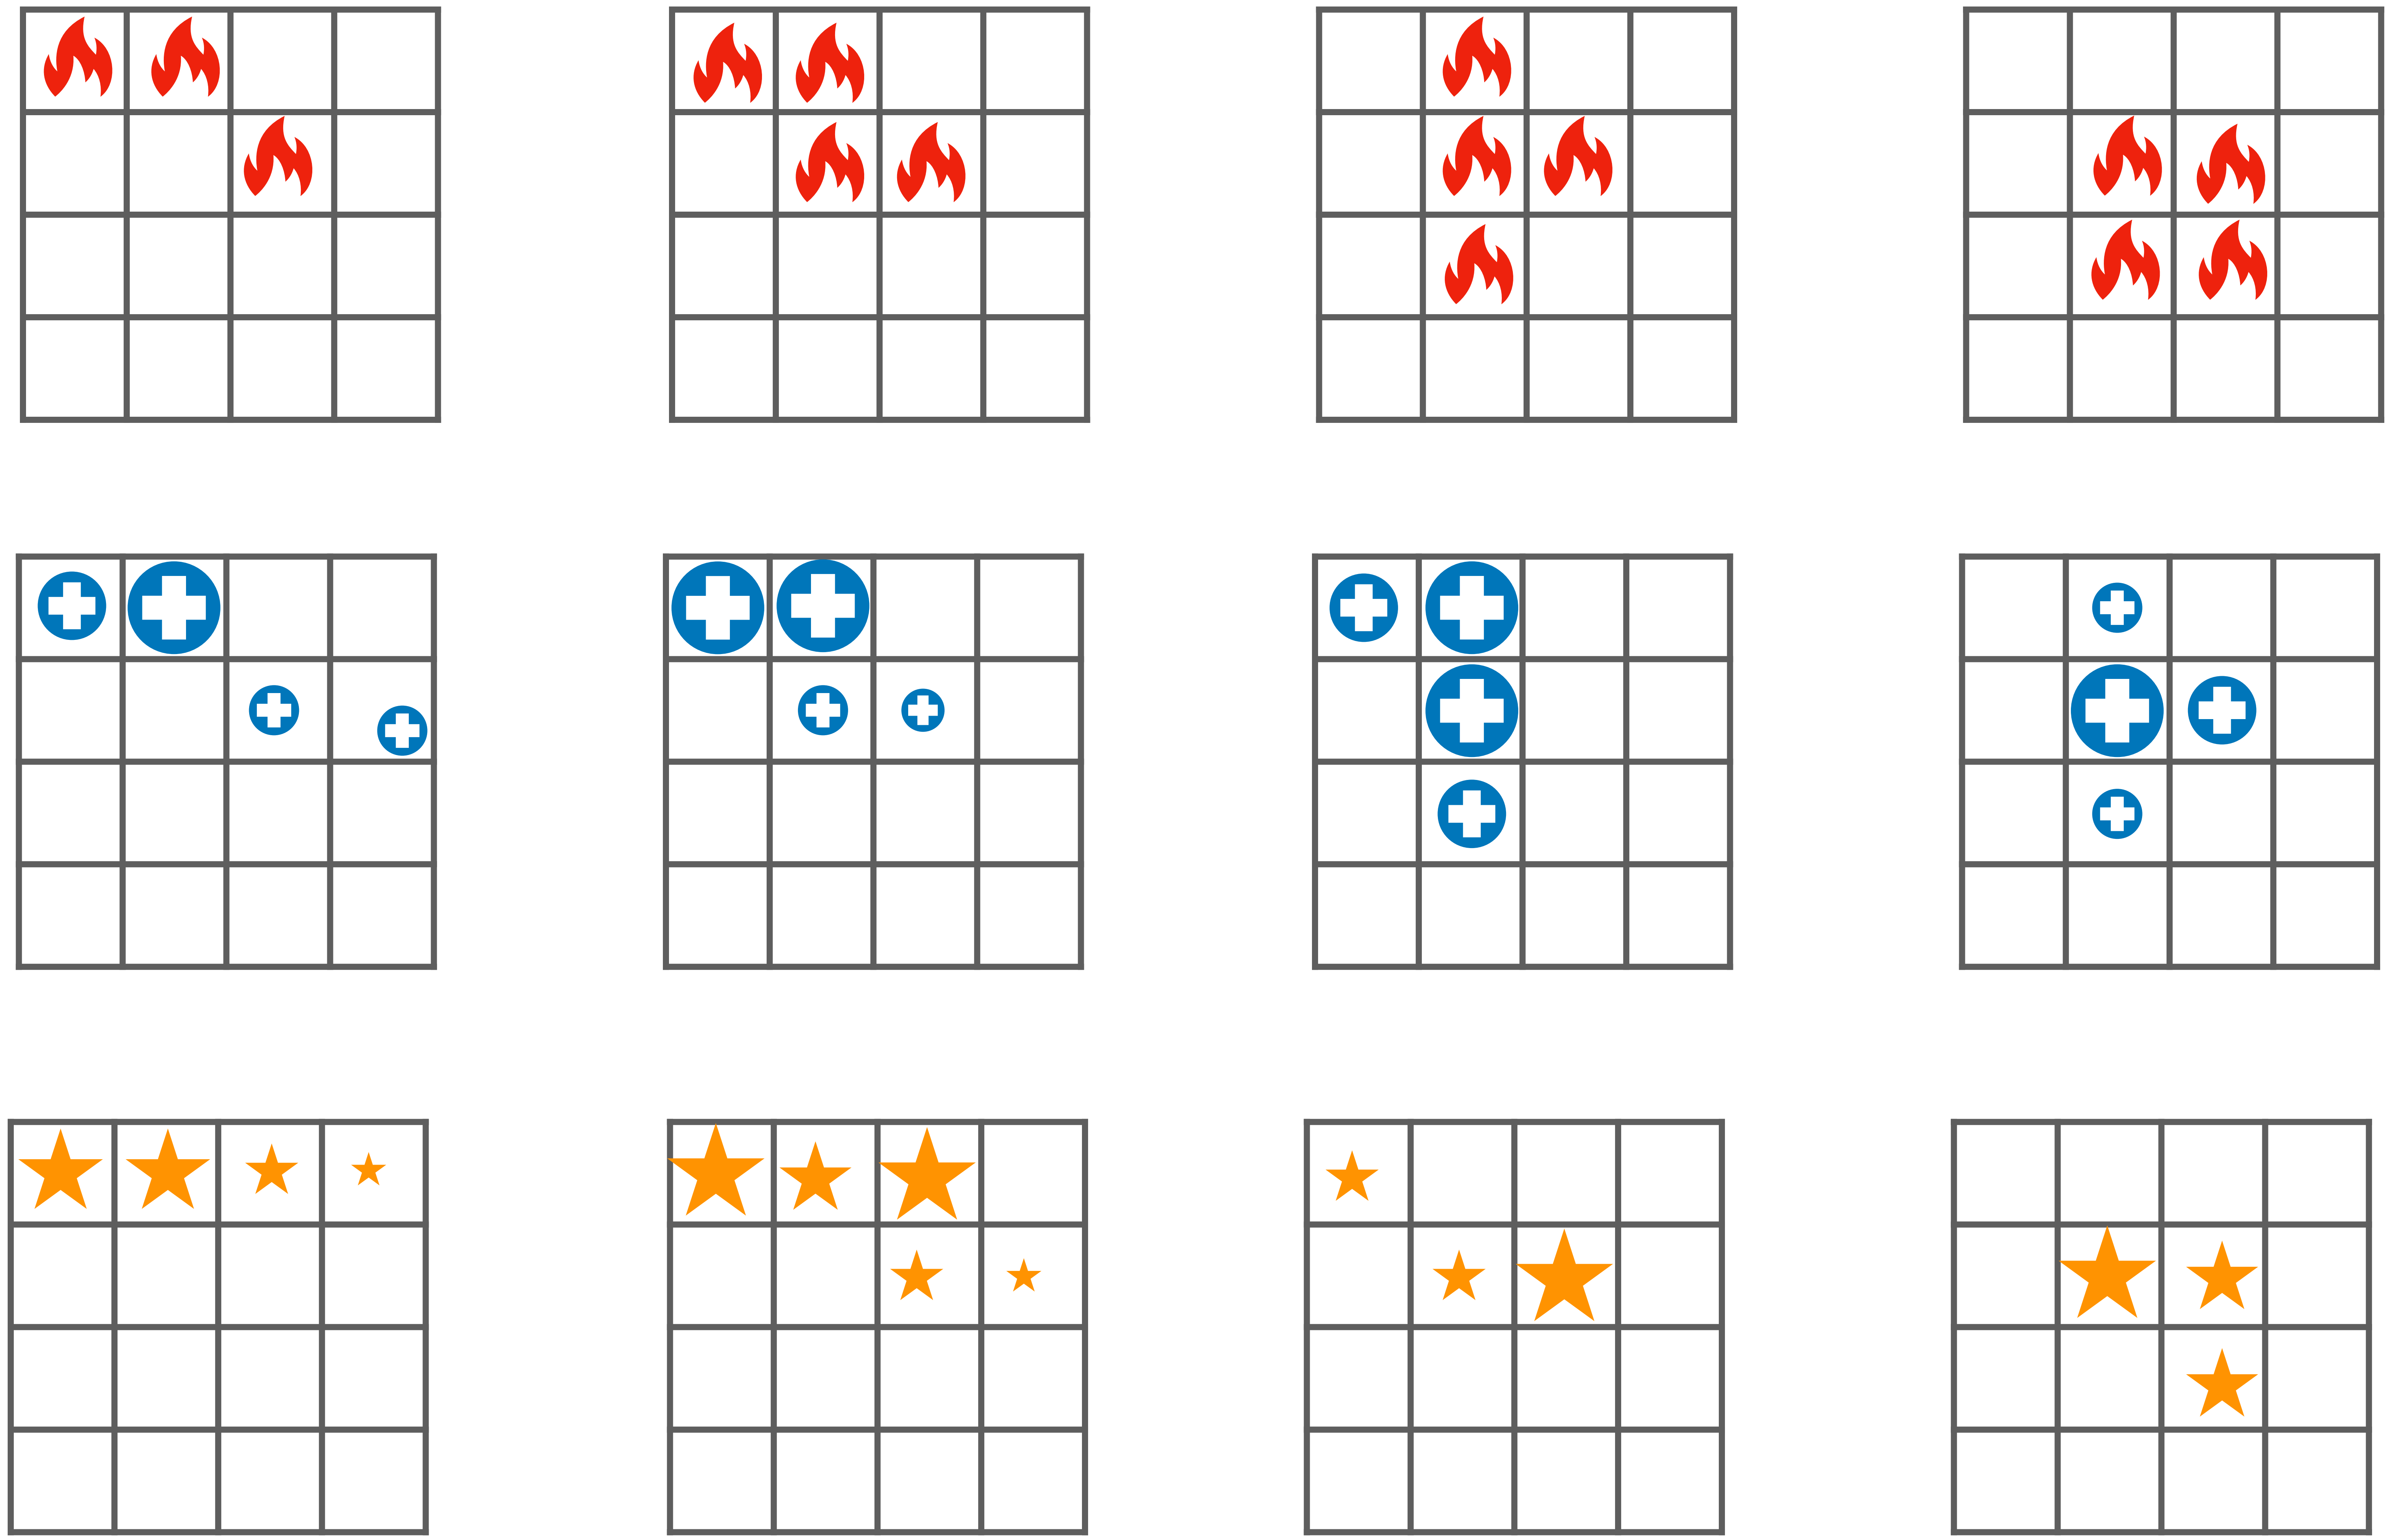

Running Example

They also have fixed covariates (like elevation) that we ignore for now.

Running Example

Wildfire Smoke
(Treatment)

Hospital Admissions
(Outcome)

Air Quality Policies
(Time-Varying Confounder)
(aka Covariate with Feedback)

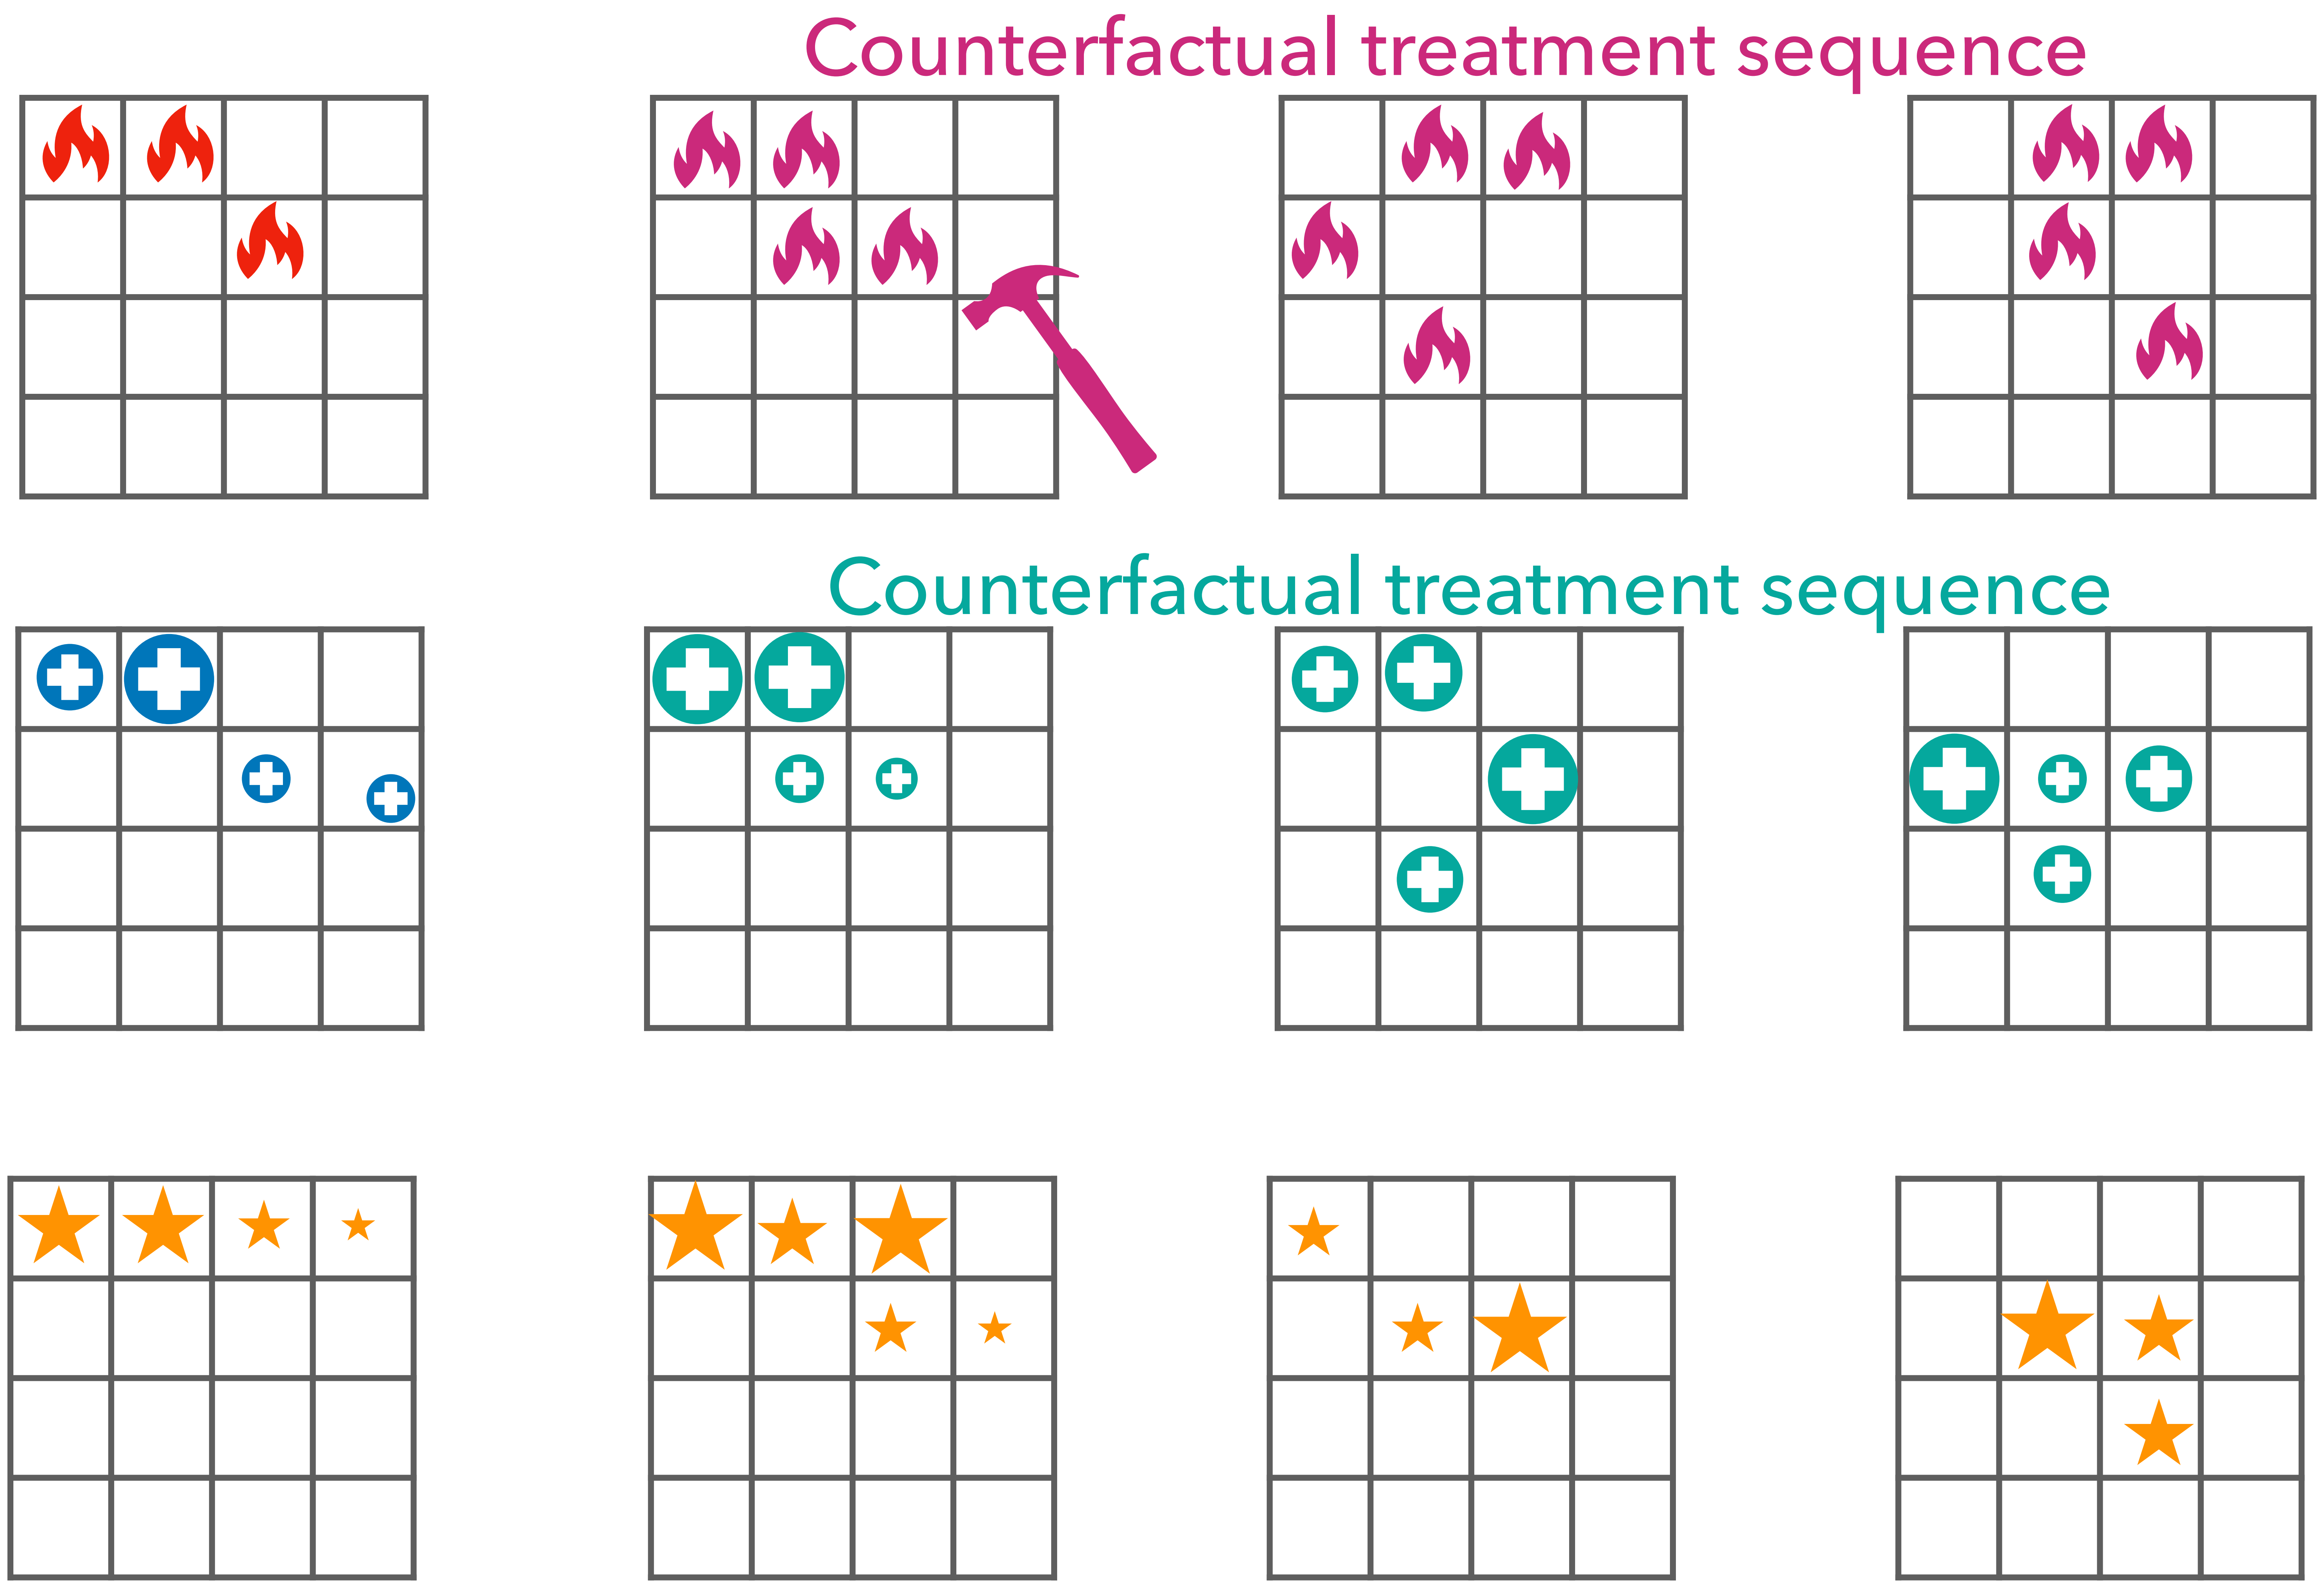

Running Example

Wildfire Smoke
(Treatment)

Hospital Admissions
(Outcome)

Air Quality Policies
(Time-Varying Confounder)
(aka Covariate with Feedback)

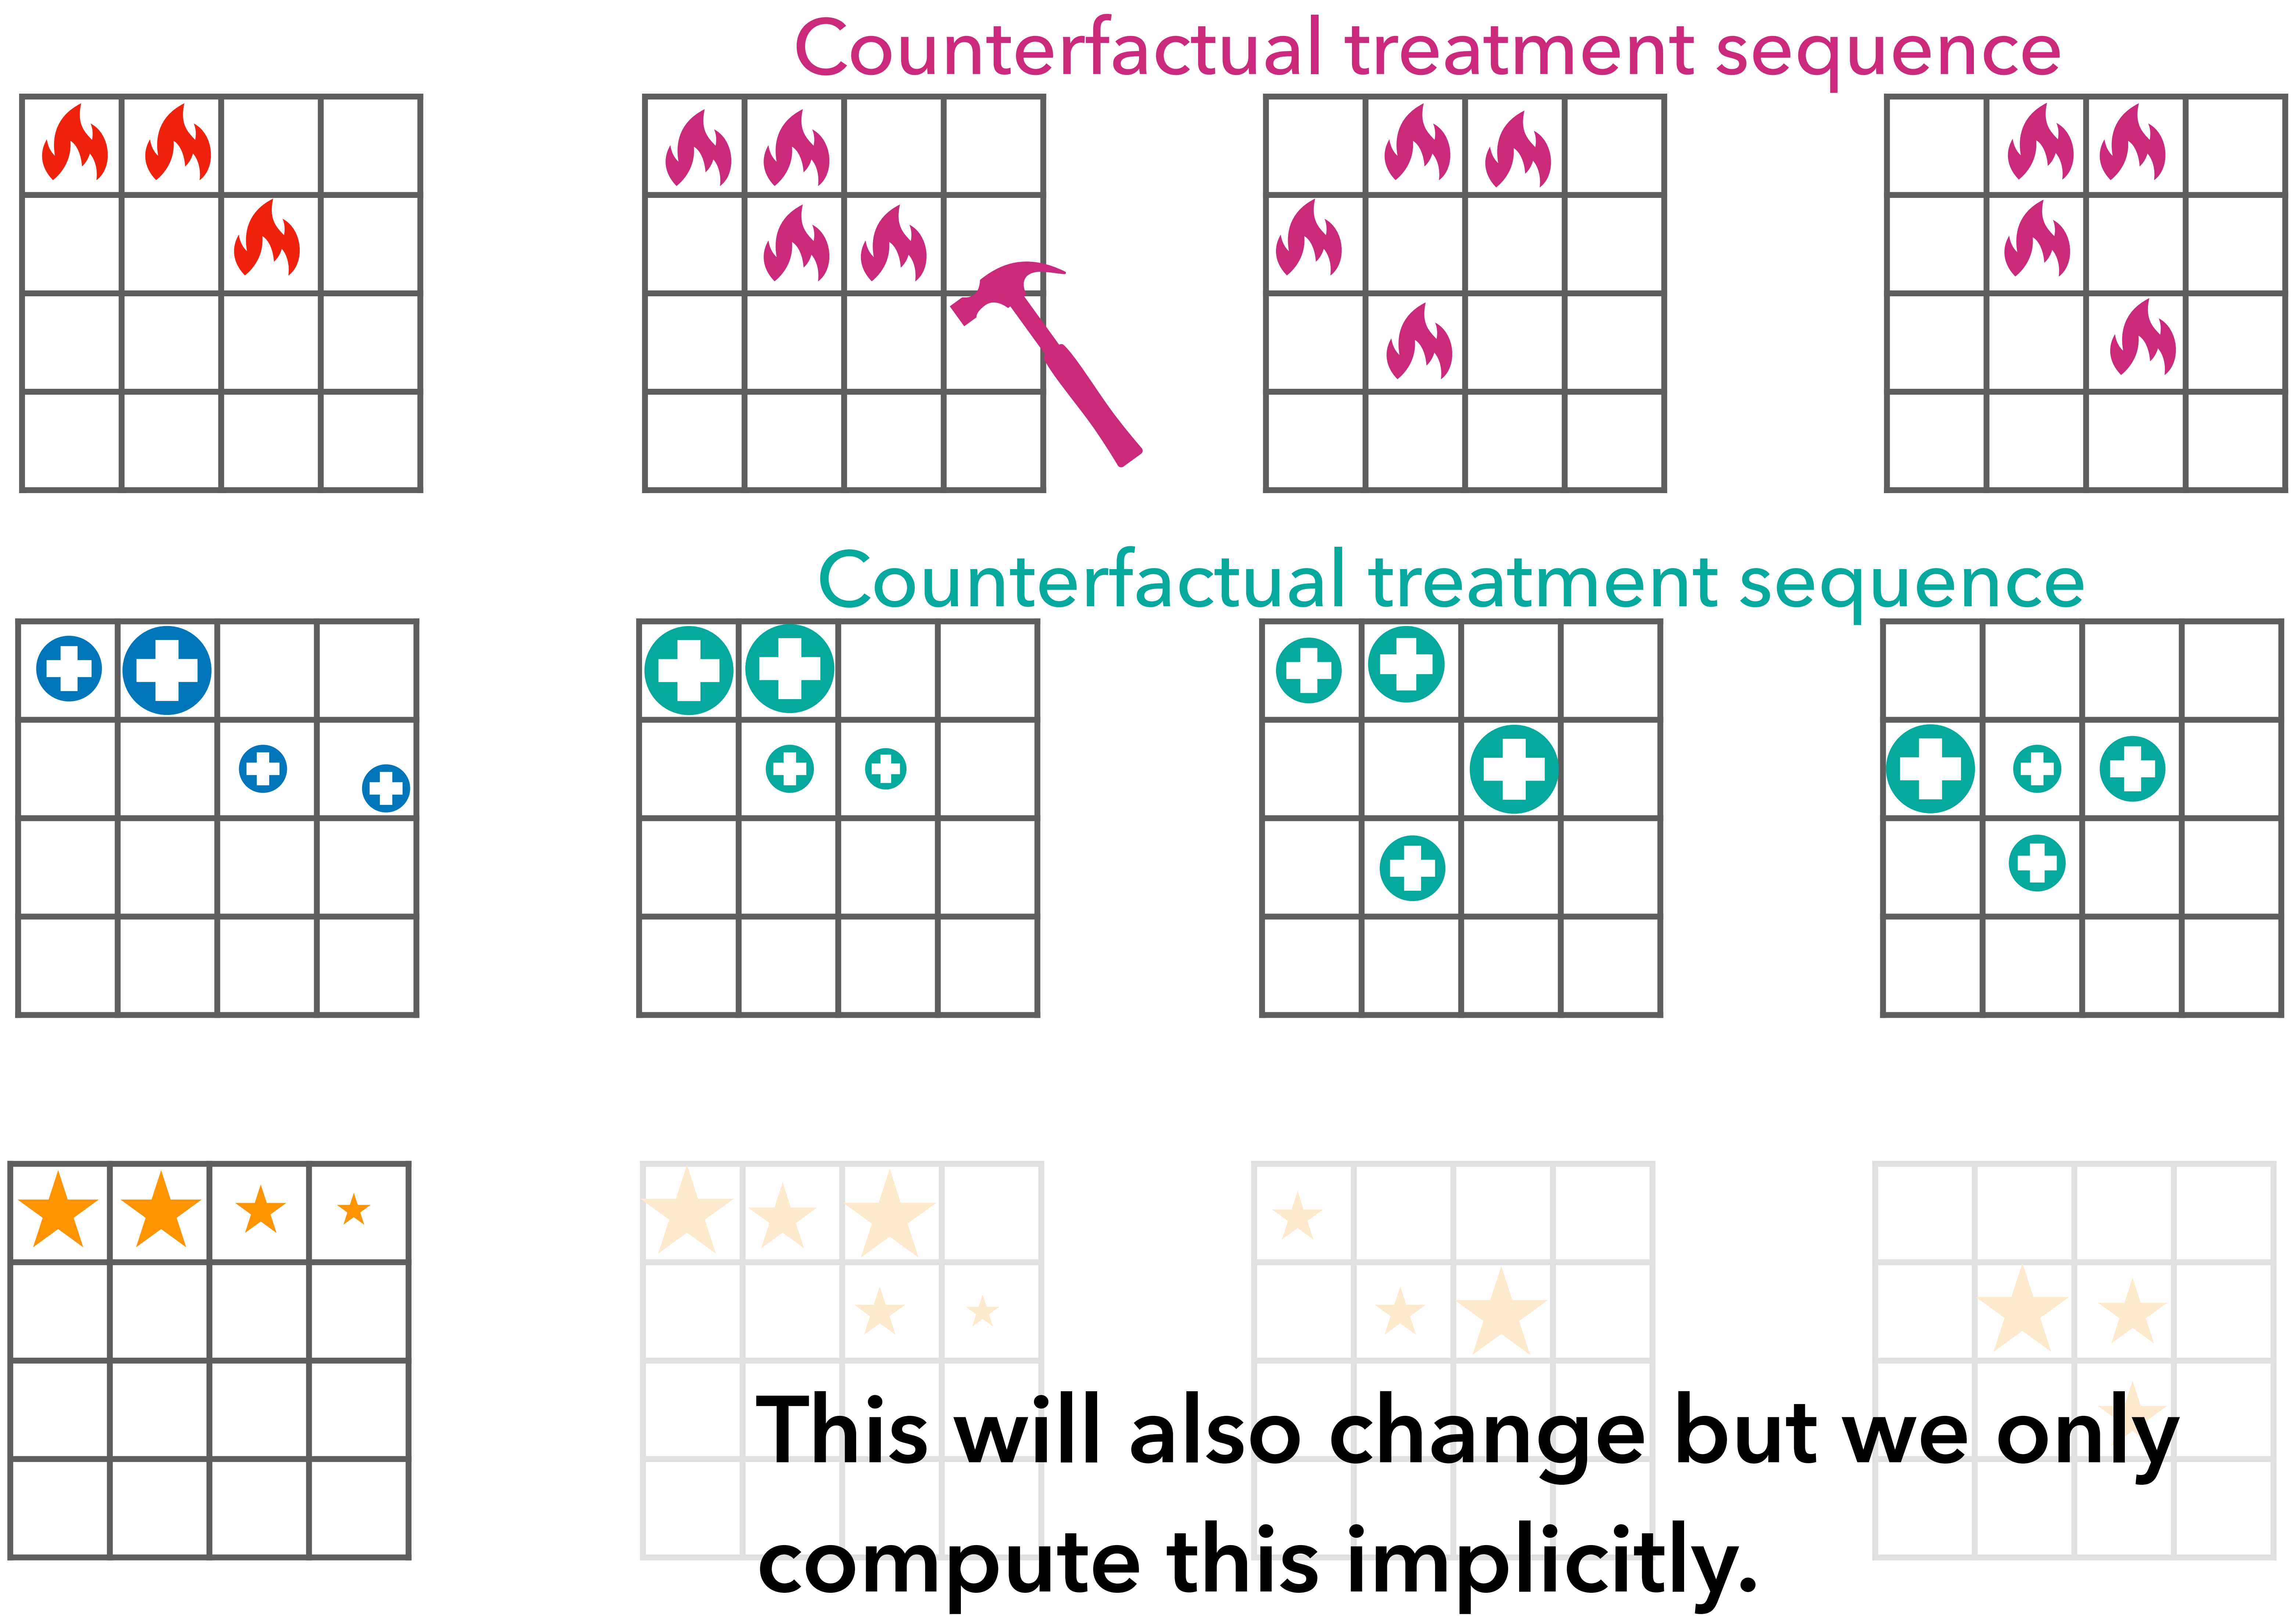

CAPO

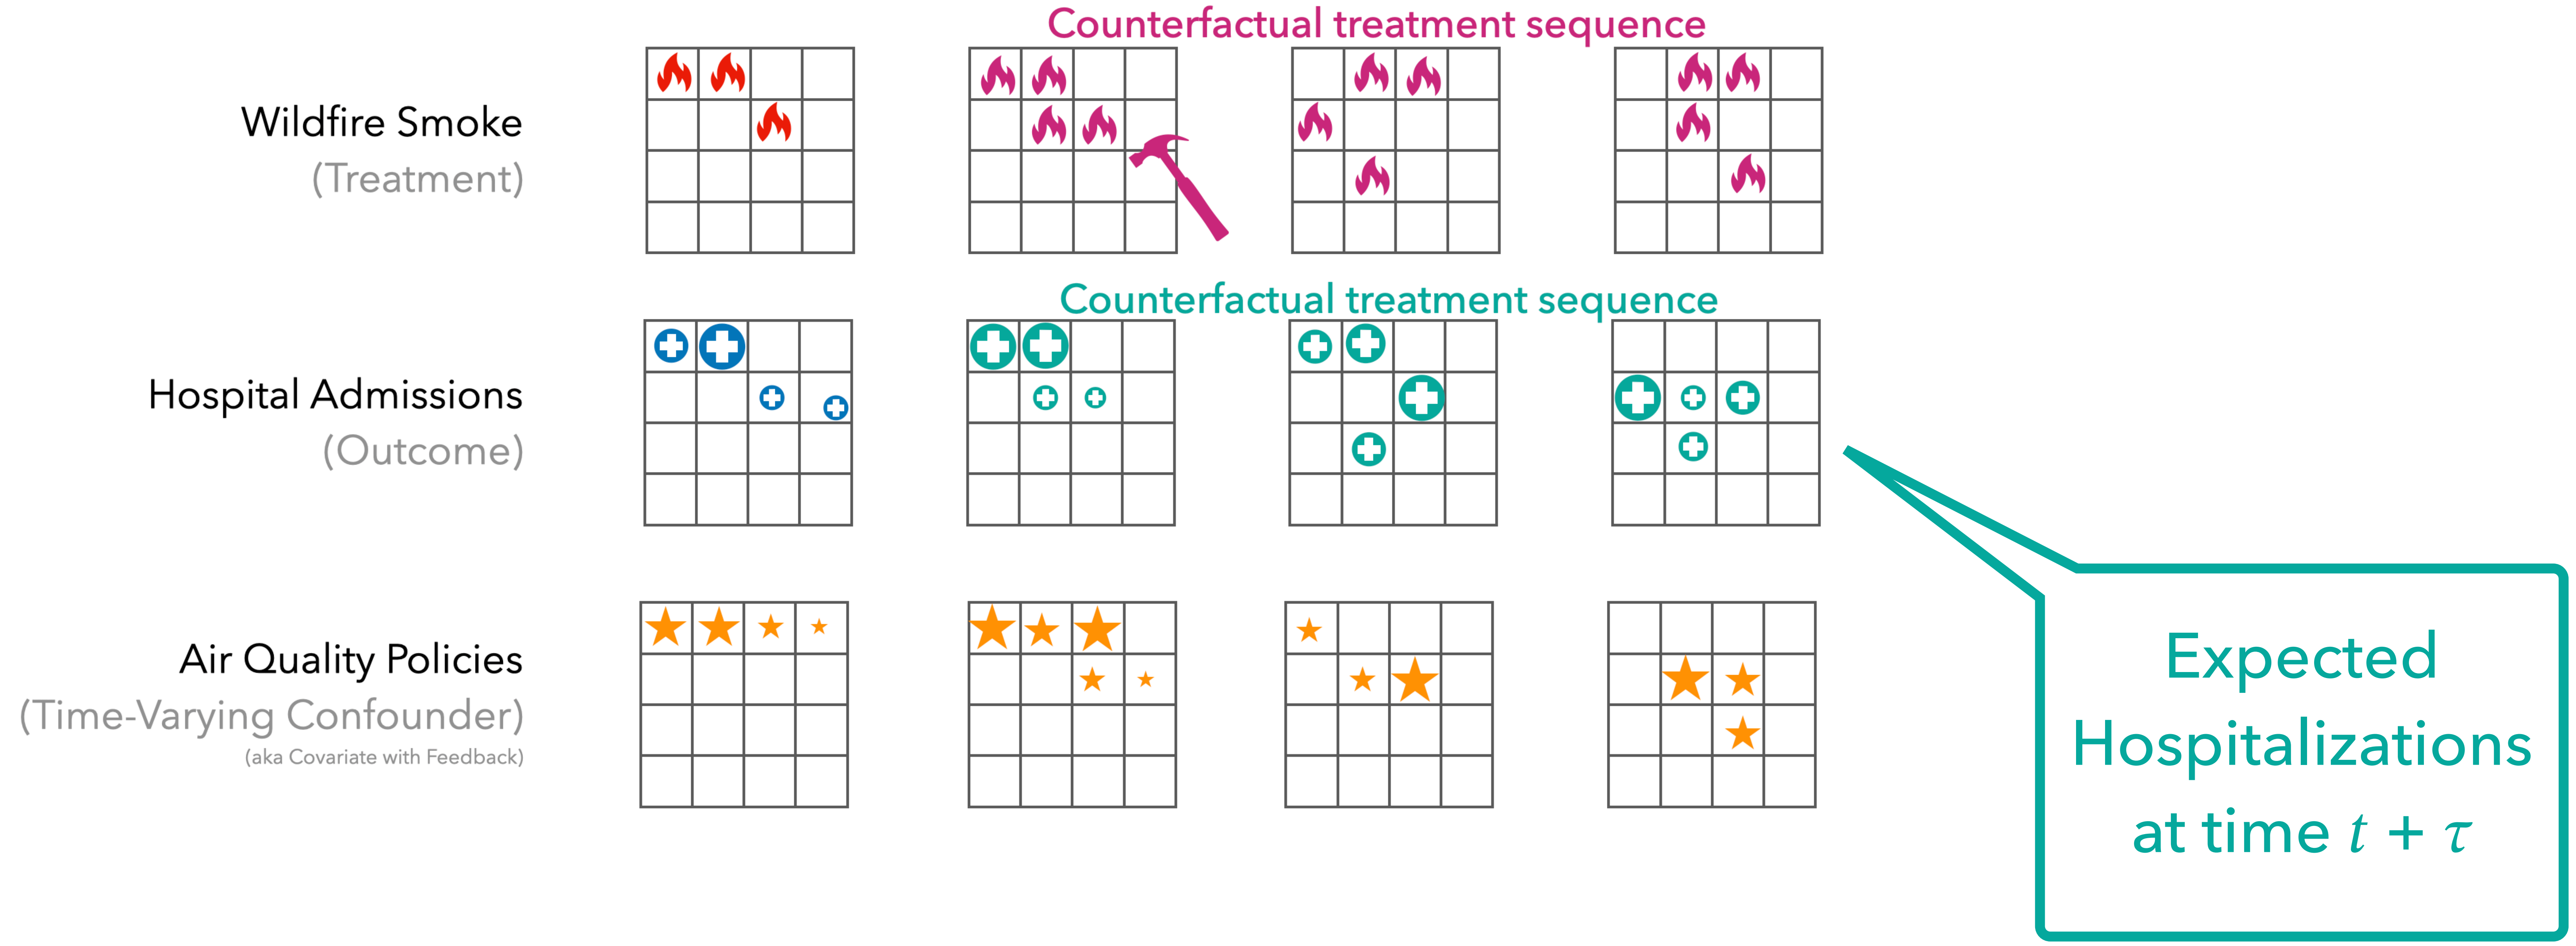

CAPO = Conditional Average Potential Outcomes

(w.r.t. history, counterfactual sequence, horizon τ)

=

Expected number of hospitalizations at time $t + \tau$ if we were to apply a specific sequence of treatments from time t t to $t + \tau - 1$.

Assumptions

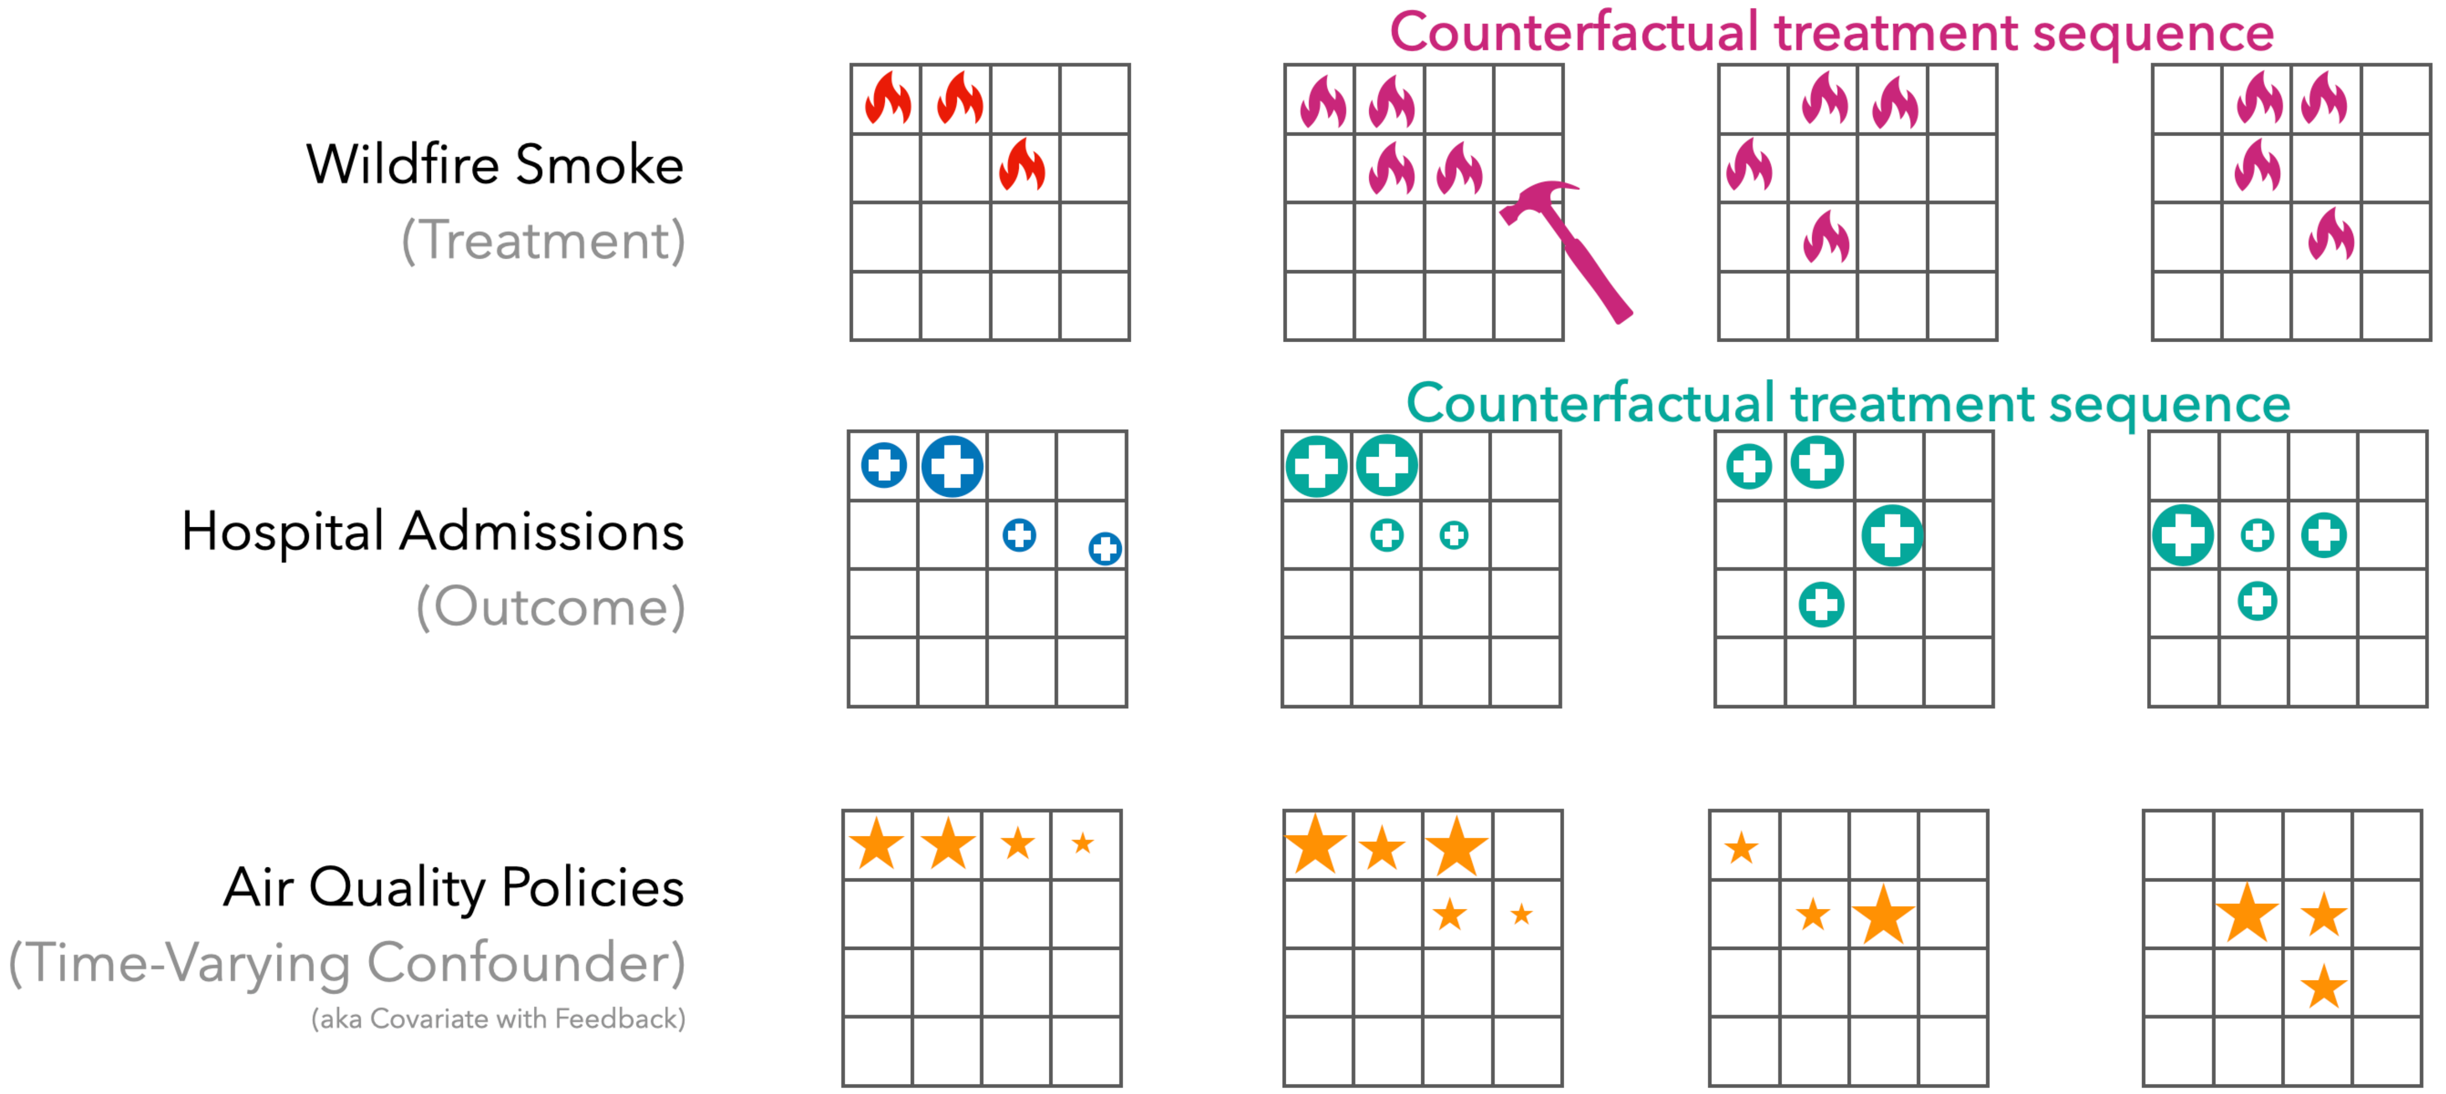

Assumptions	Where it breaks:
Causal Sufficiency:	Political pressure from activist groups that is not captured.
Temporal Ordering:	Correct the air quality sensor readings retrospectively.
Time-Invariant Dynamics:	Behavior changes after 5 month.

G-Computation (Robins, 1986)

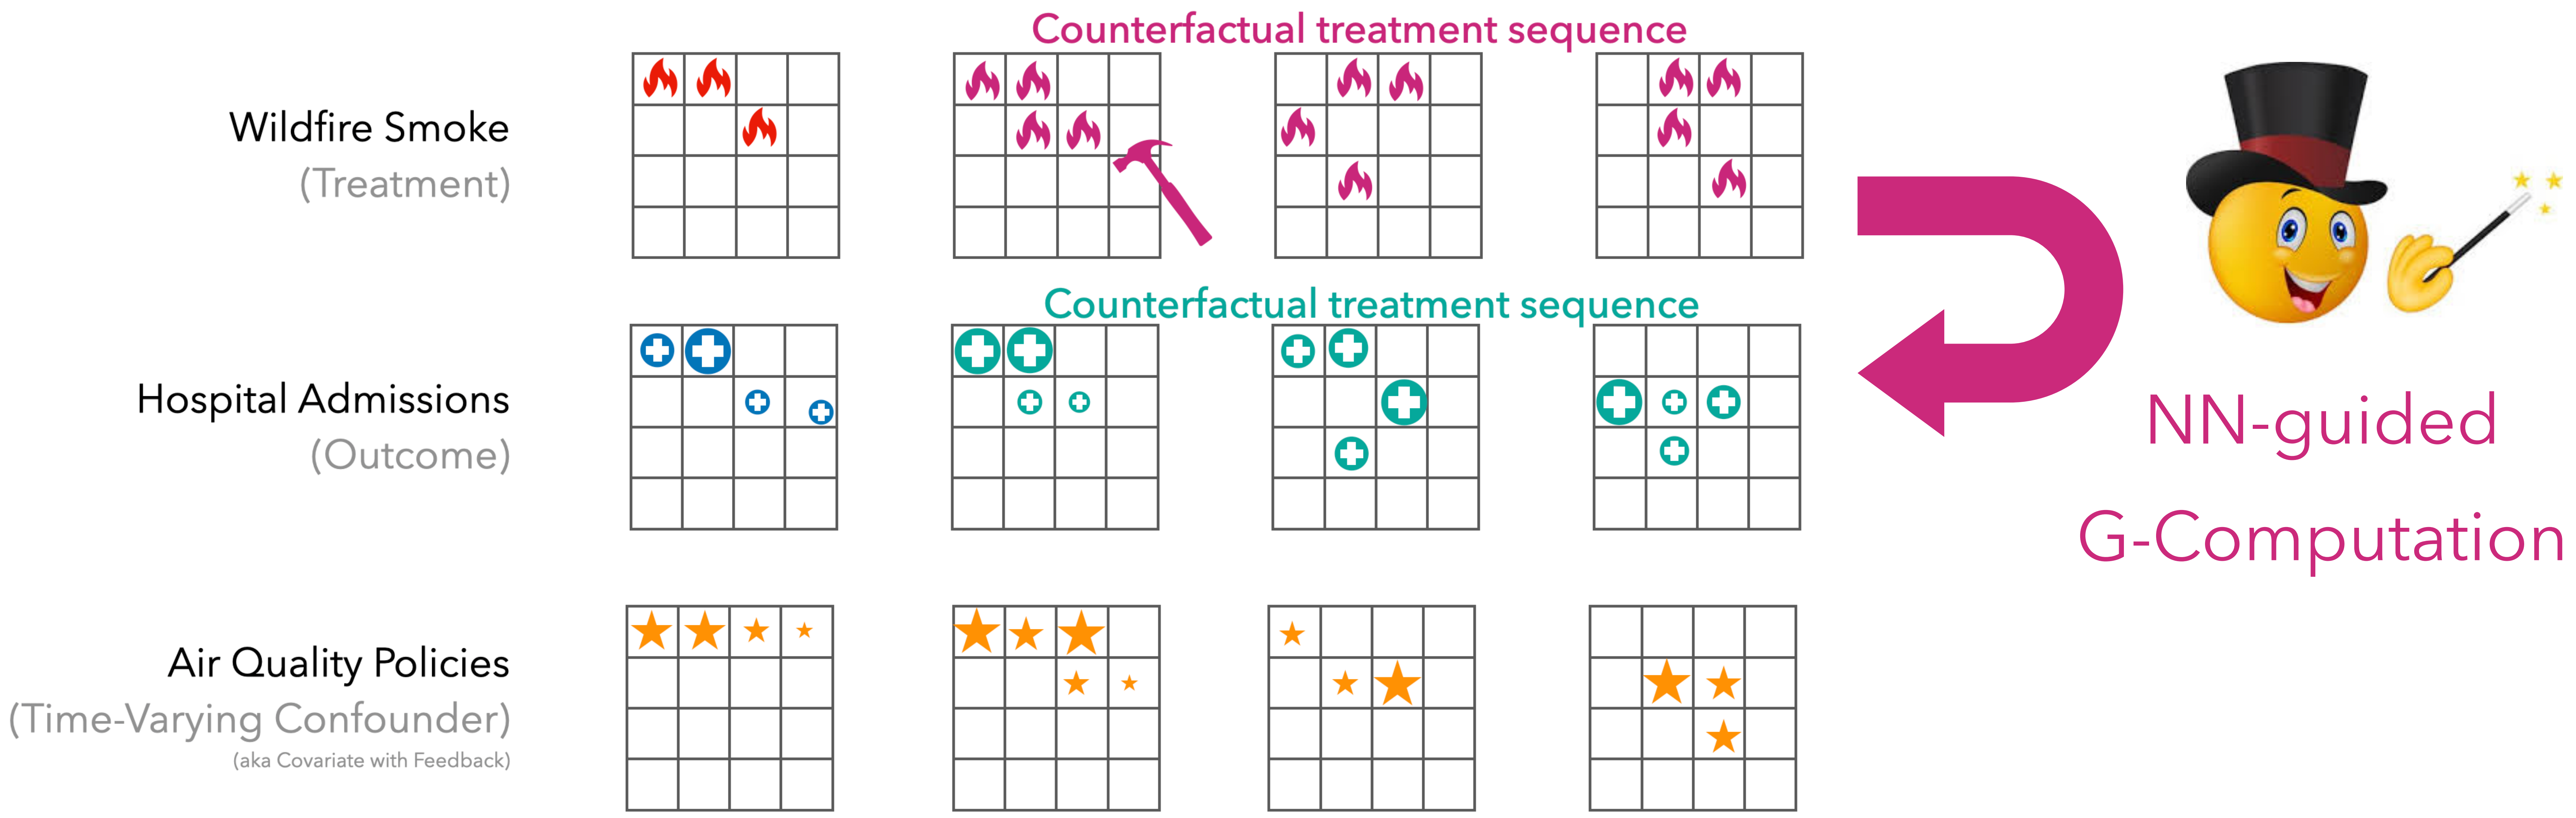

G-Computation (Robins, 1986)

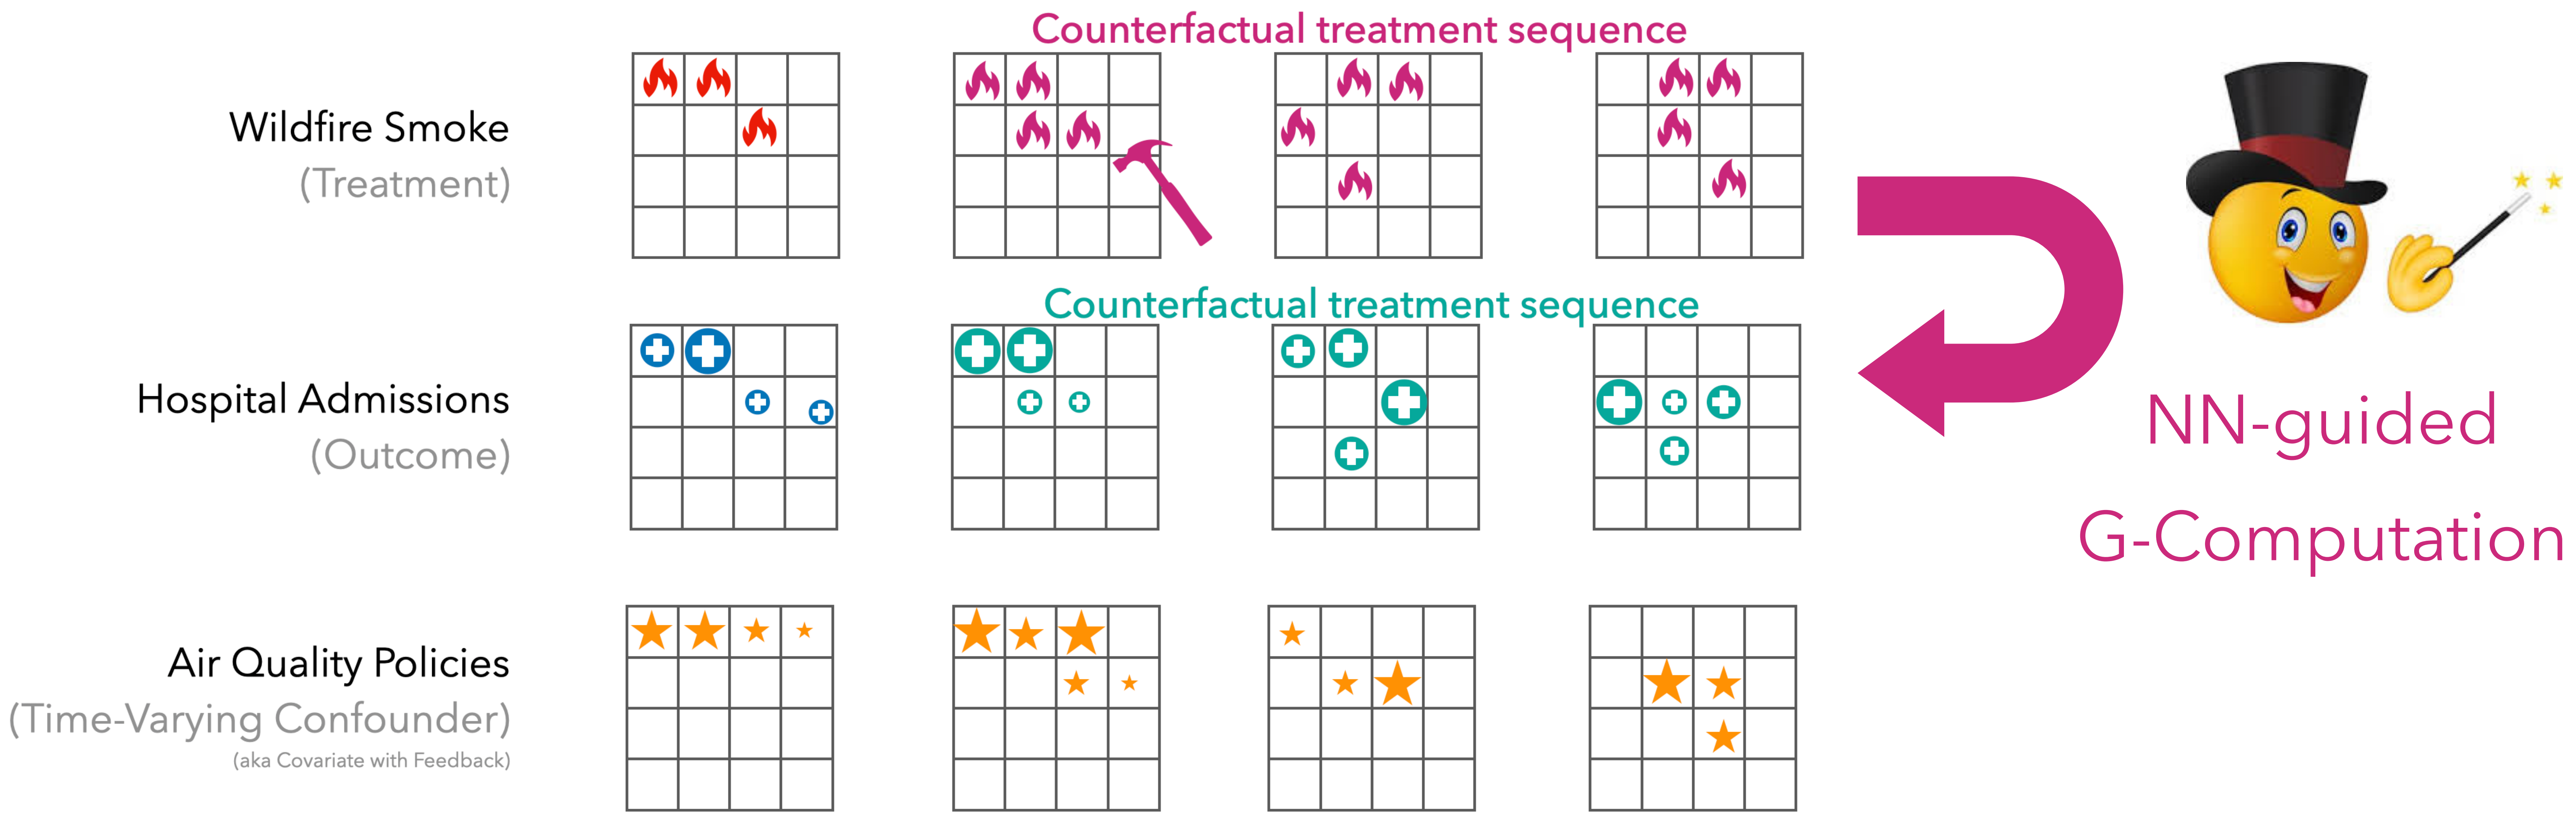

- Estimate counterfactual outcomes with regression + recursion
- Good understanding of how outcomes depend on covariates and treatments.
- You cannot estimate effects of treatments that never occur in some part of the covariate space.

Setup

Intervention (binary, *good* vs *bad* air) at time t at City S.

City S

City S′

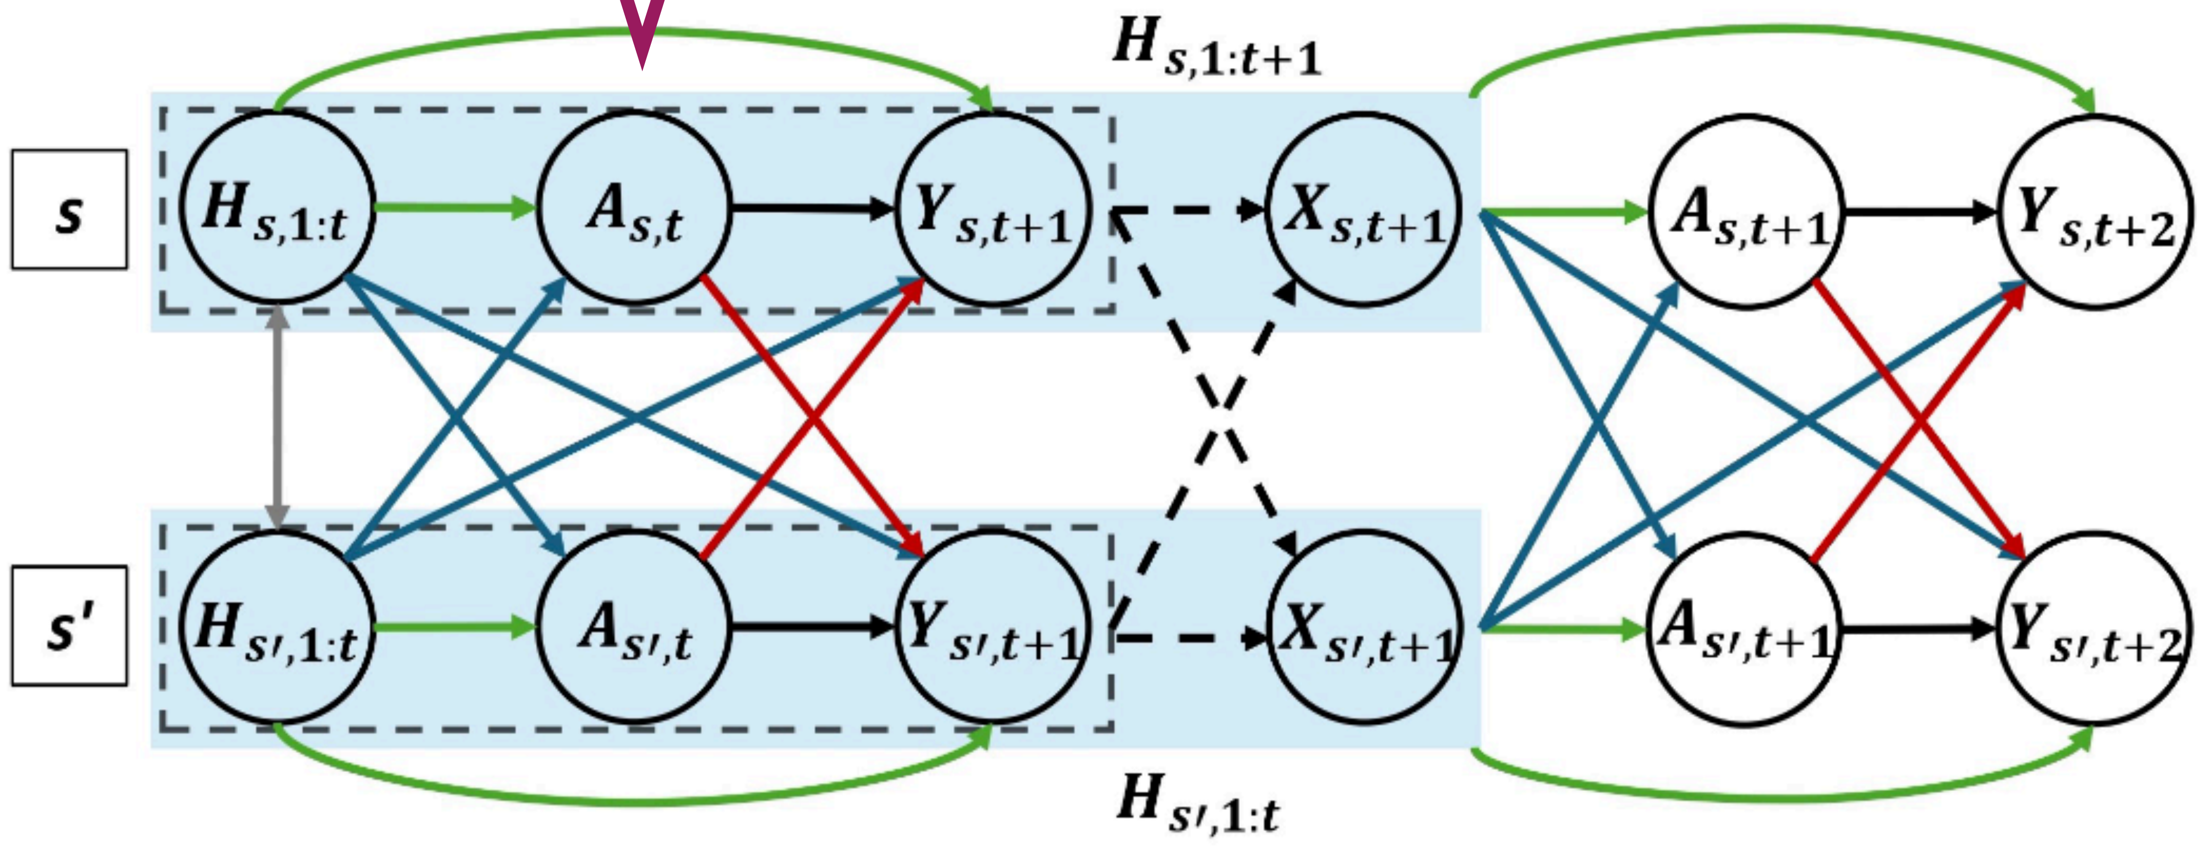

Setup

Outcome (scalar, hospitalization level) at time t+1 at City S.

City S

City S′

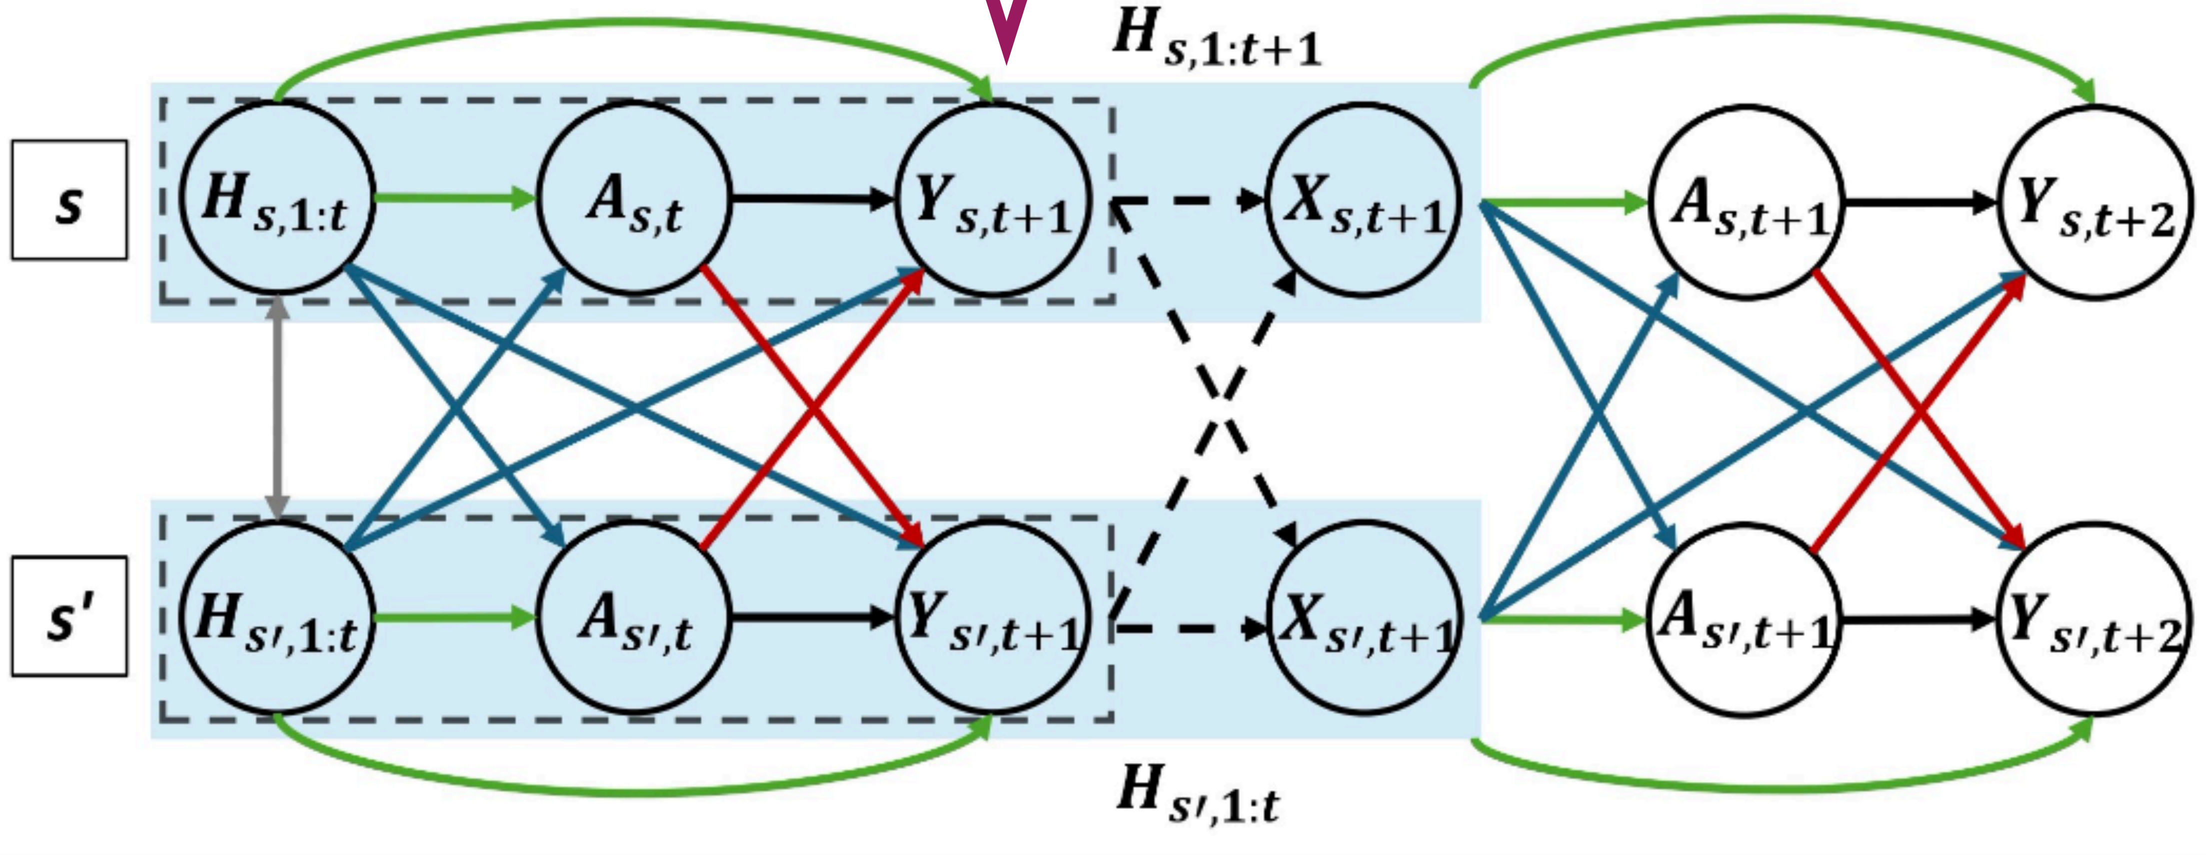

Setup

Covariates (temperature) at time t+1 at City S.

City S

City S'

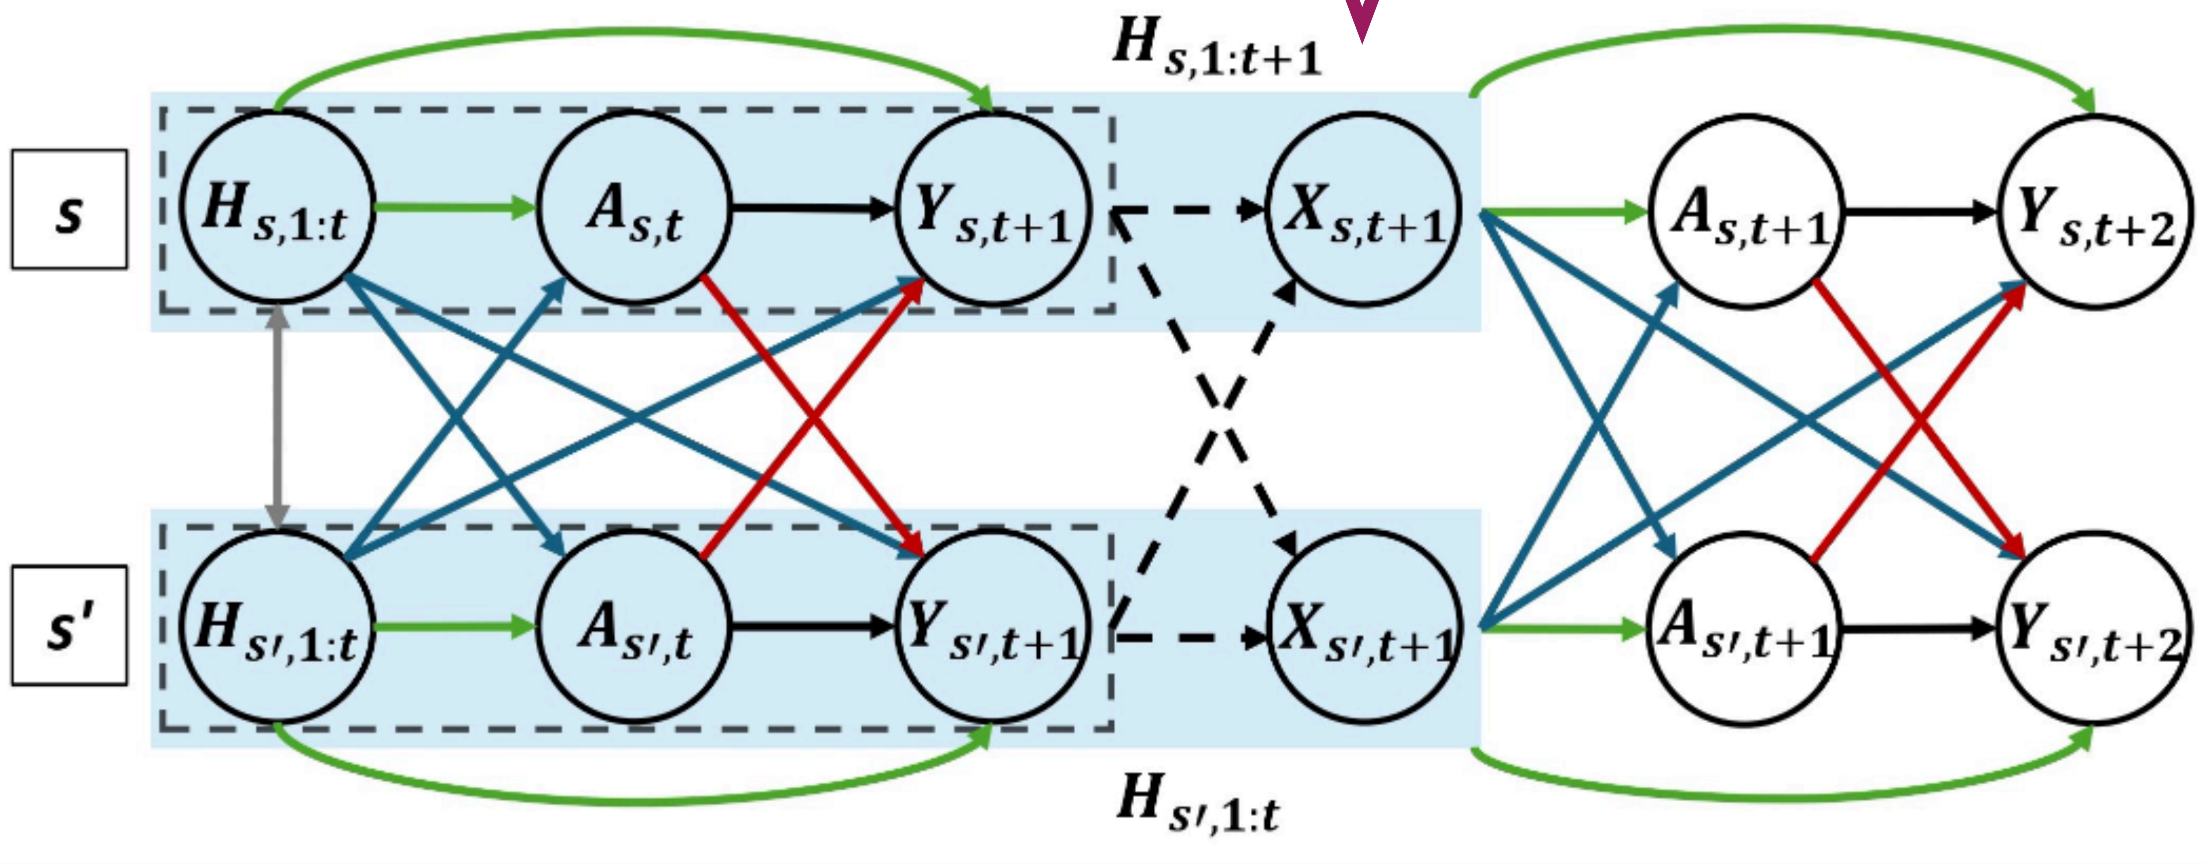

Setup

History of all variables until time t in City S

City S

City S′

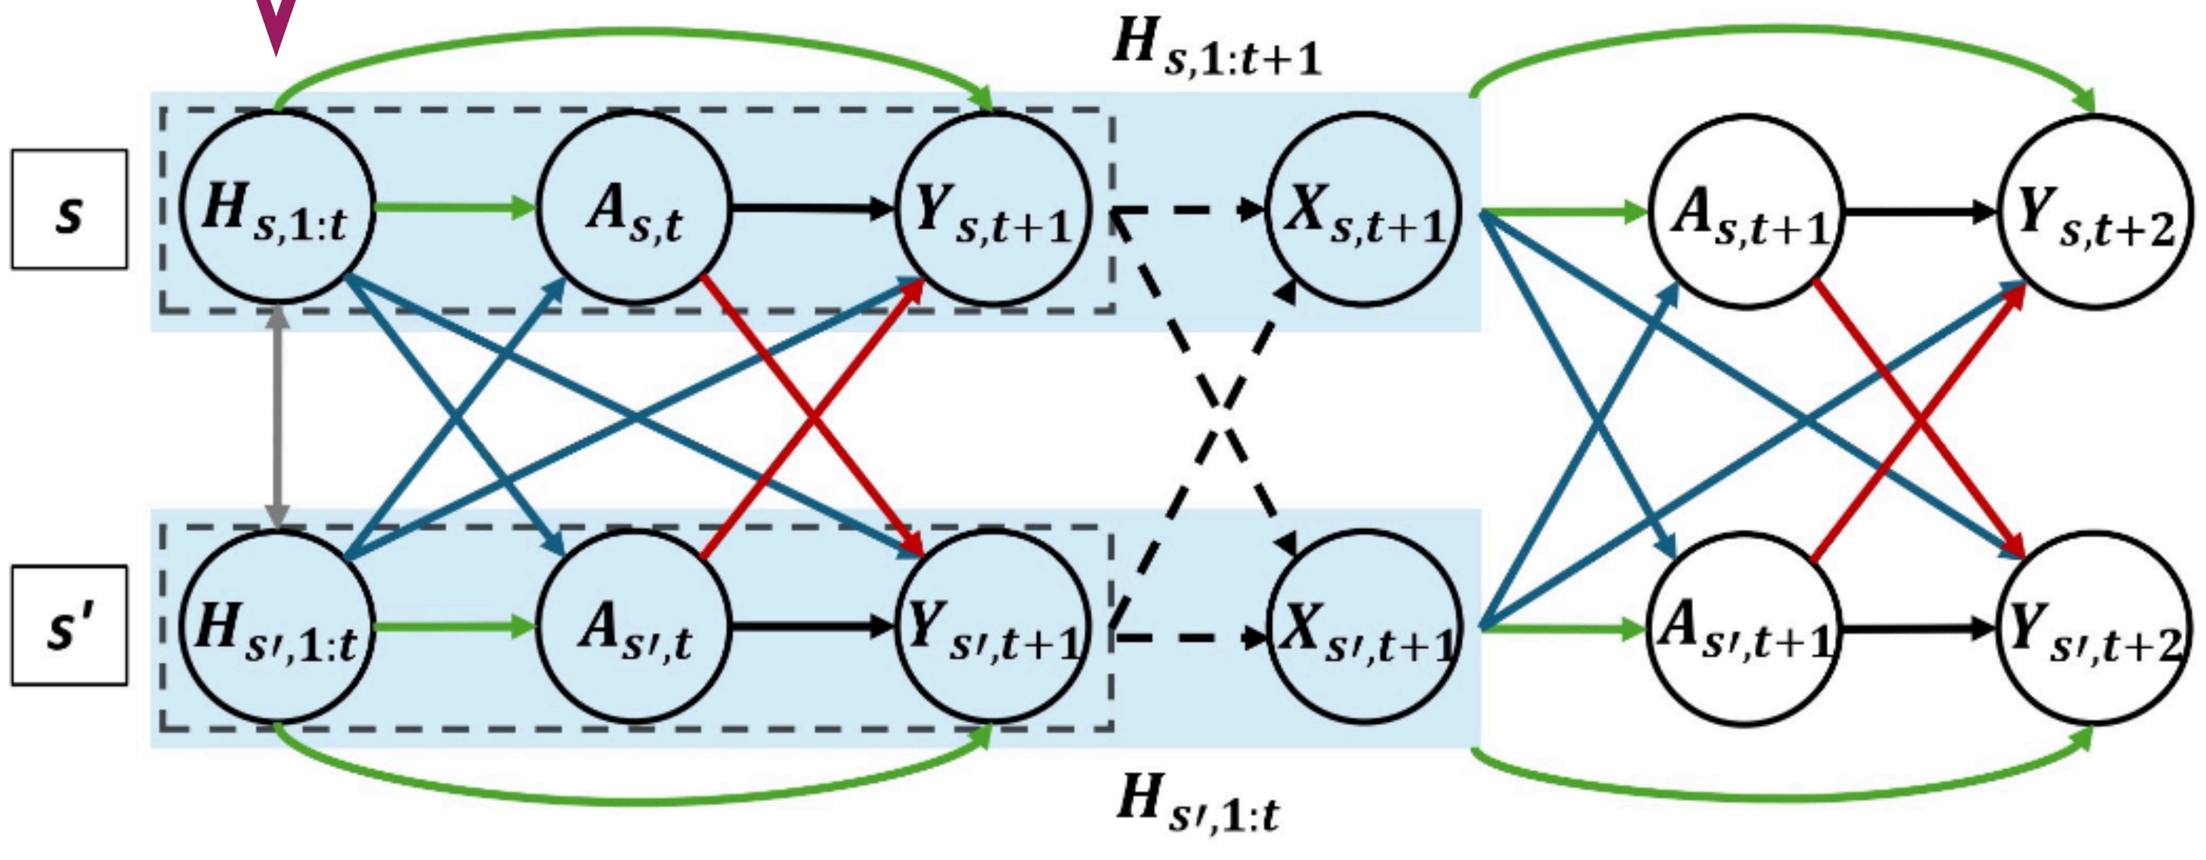

Setup

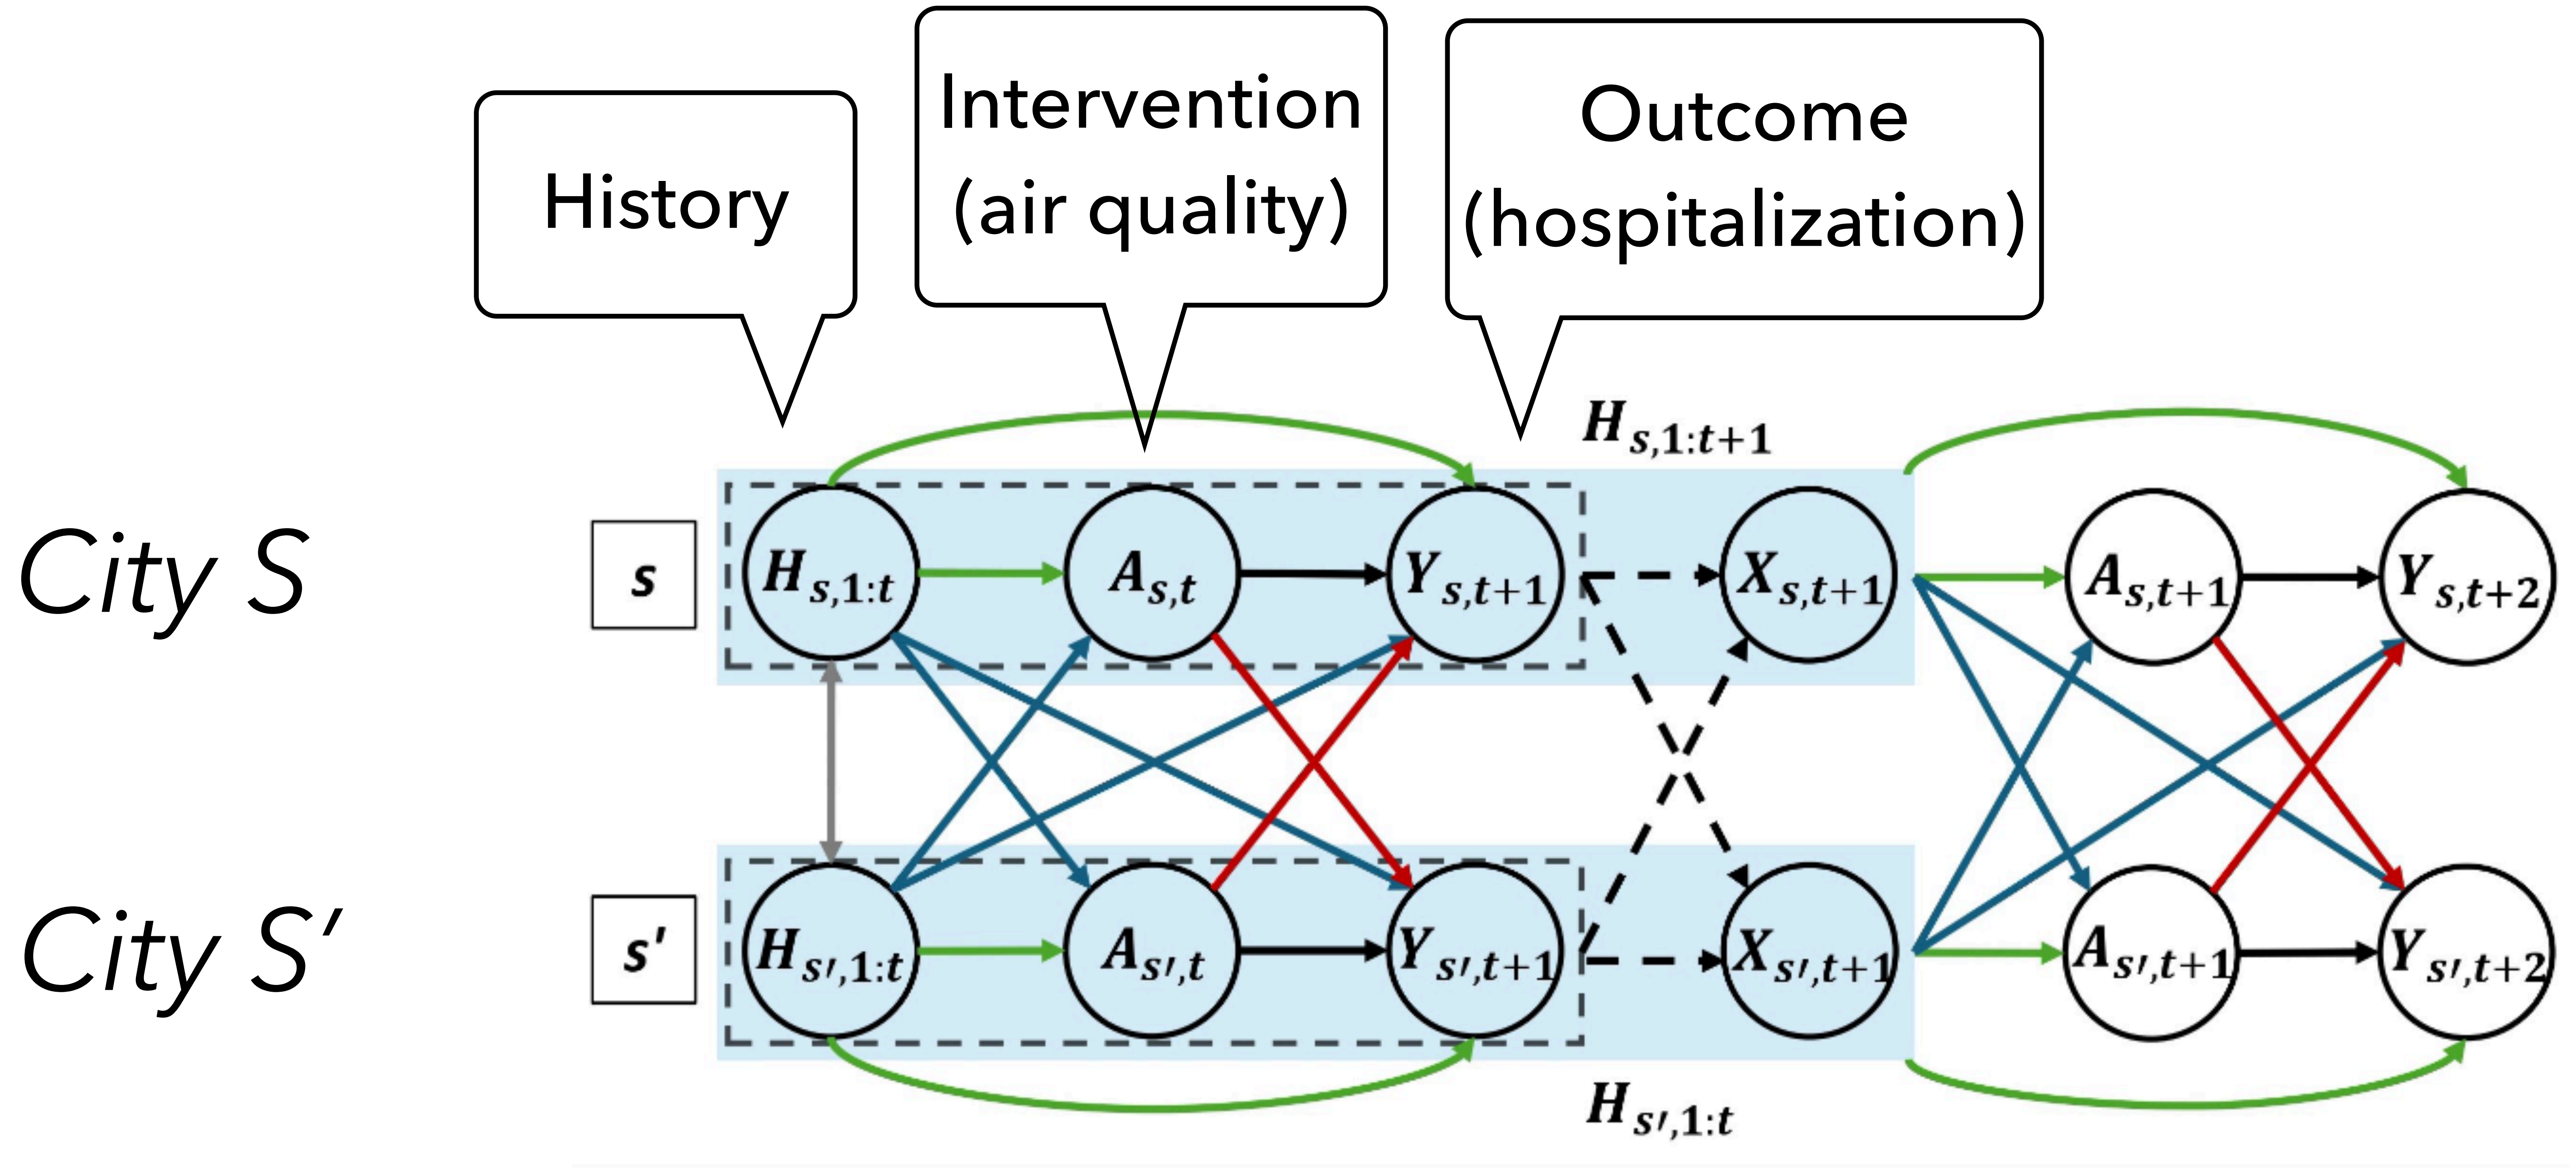

Direct effects **Temporal carryover** **Spatial confounding** **Interference** **Time-varying confounding**

Air pollution today causes more hospitalizations today.

Smoke exposure yesterday increases hospitalizations tomorrow

Temperature in City S causes air quality changes in City S′

Pollution in *City S* causes hospitalizations in *City S′*

Weather influences pollution levels and respiratory health

Setup

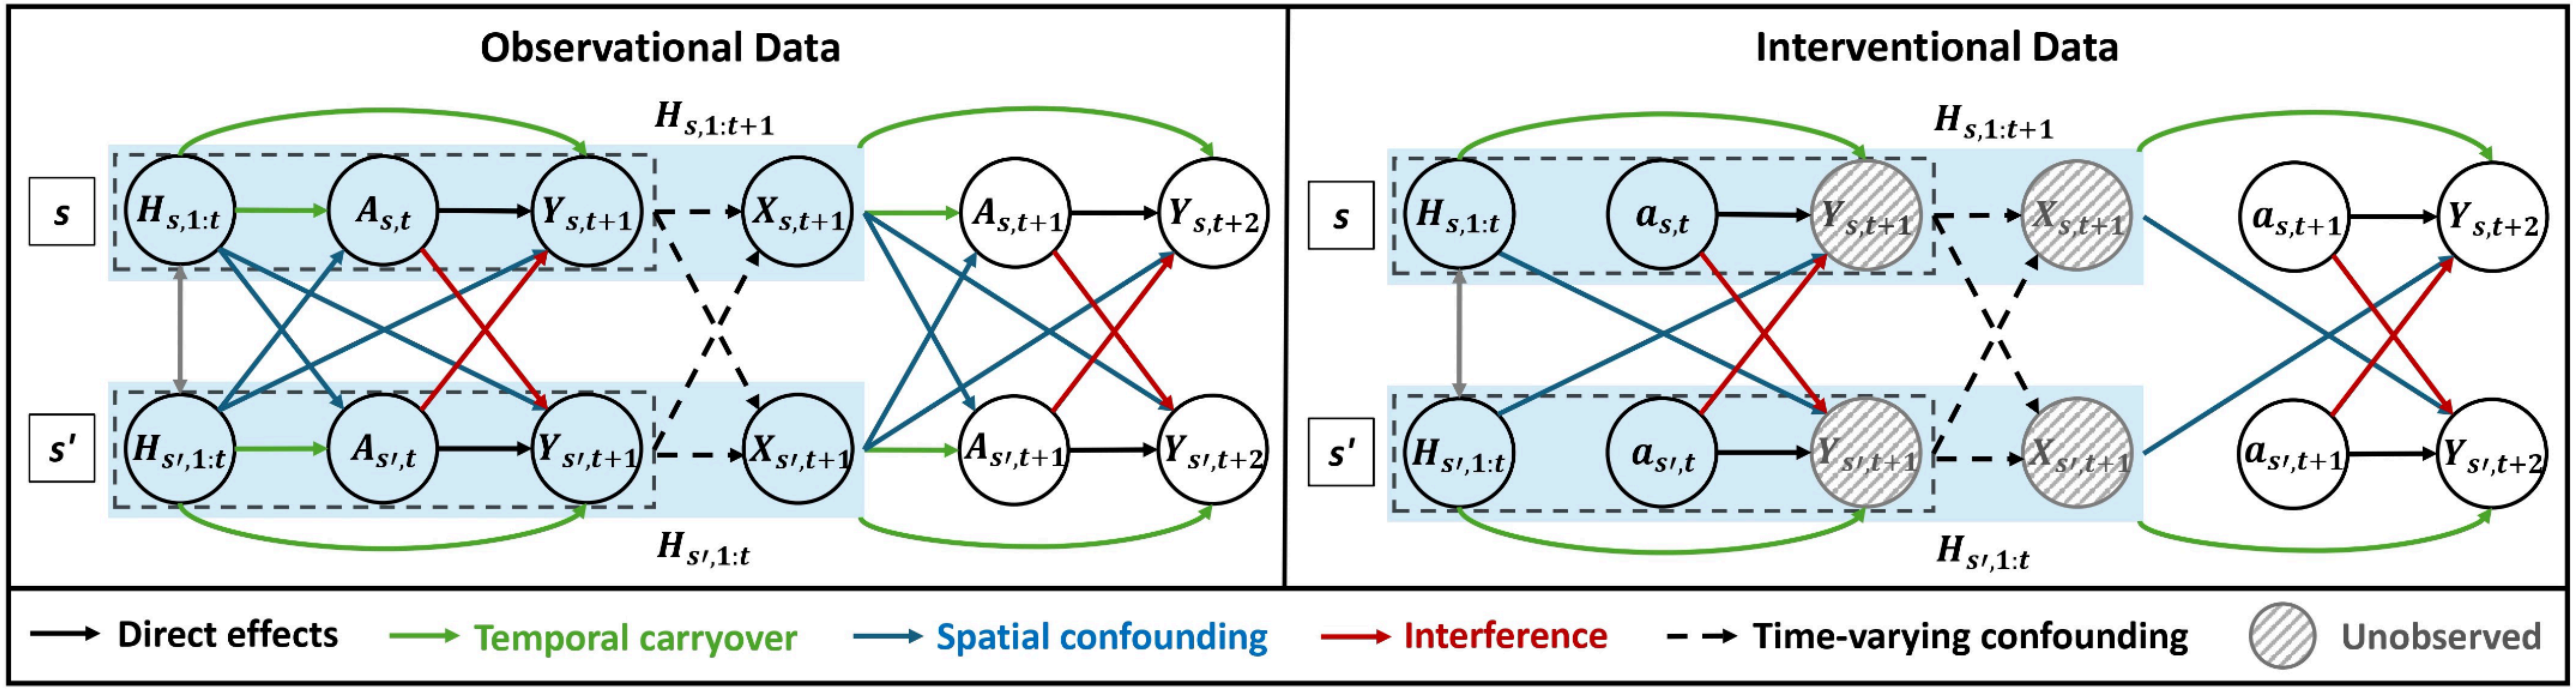

Figure 1. Observational data (left) versus interventional data (right) for a horizon $\tau = 2$ across multiple locations (s, s'). Green arrows indicate temporal carryover, blue arrows show spatial confounding, and red arrows depict interference; dashed arrows denote time-varying confounding, and dashed circles represent unobserved variables at inference time. Under the intervention (right), treatments are set independently of confounders, and the full history is not observed for the entire horizon.

G-Computation (Robins, 1986)

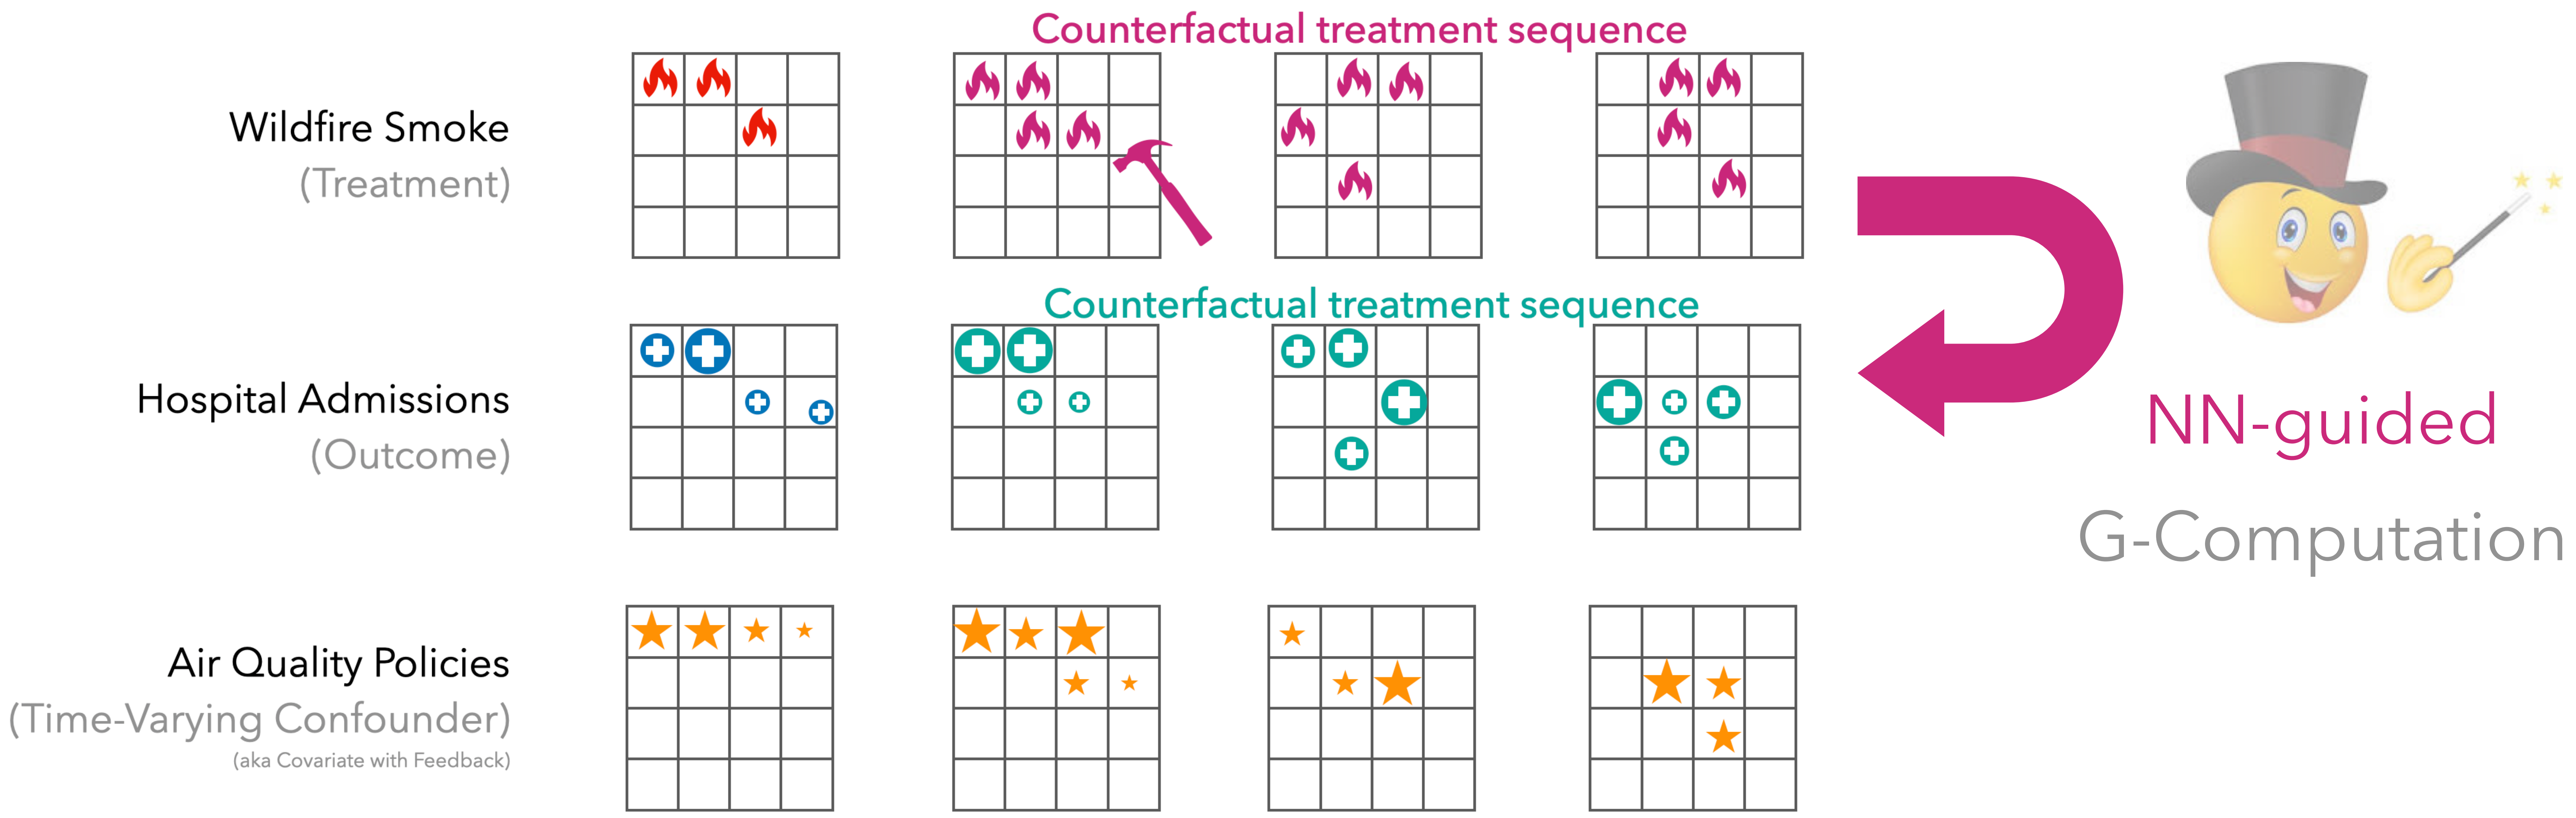

G-Computation (Robins, 1986)

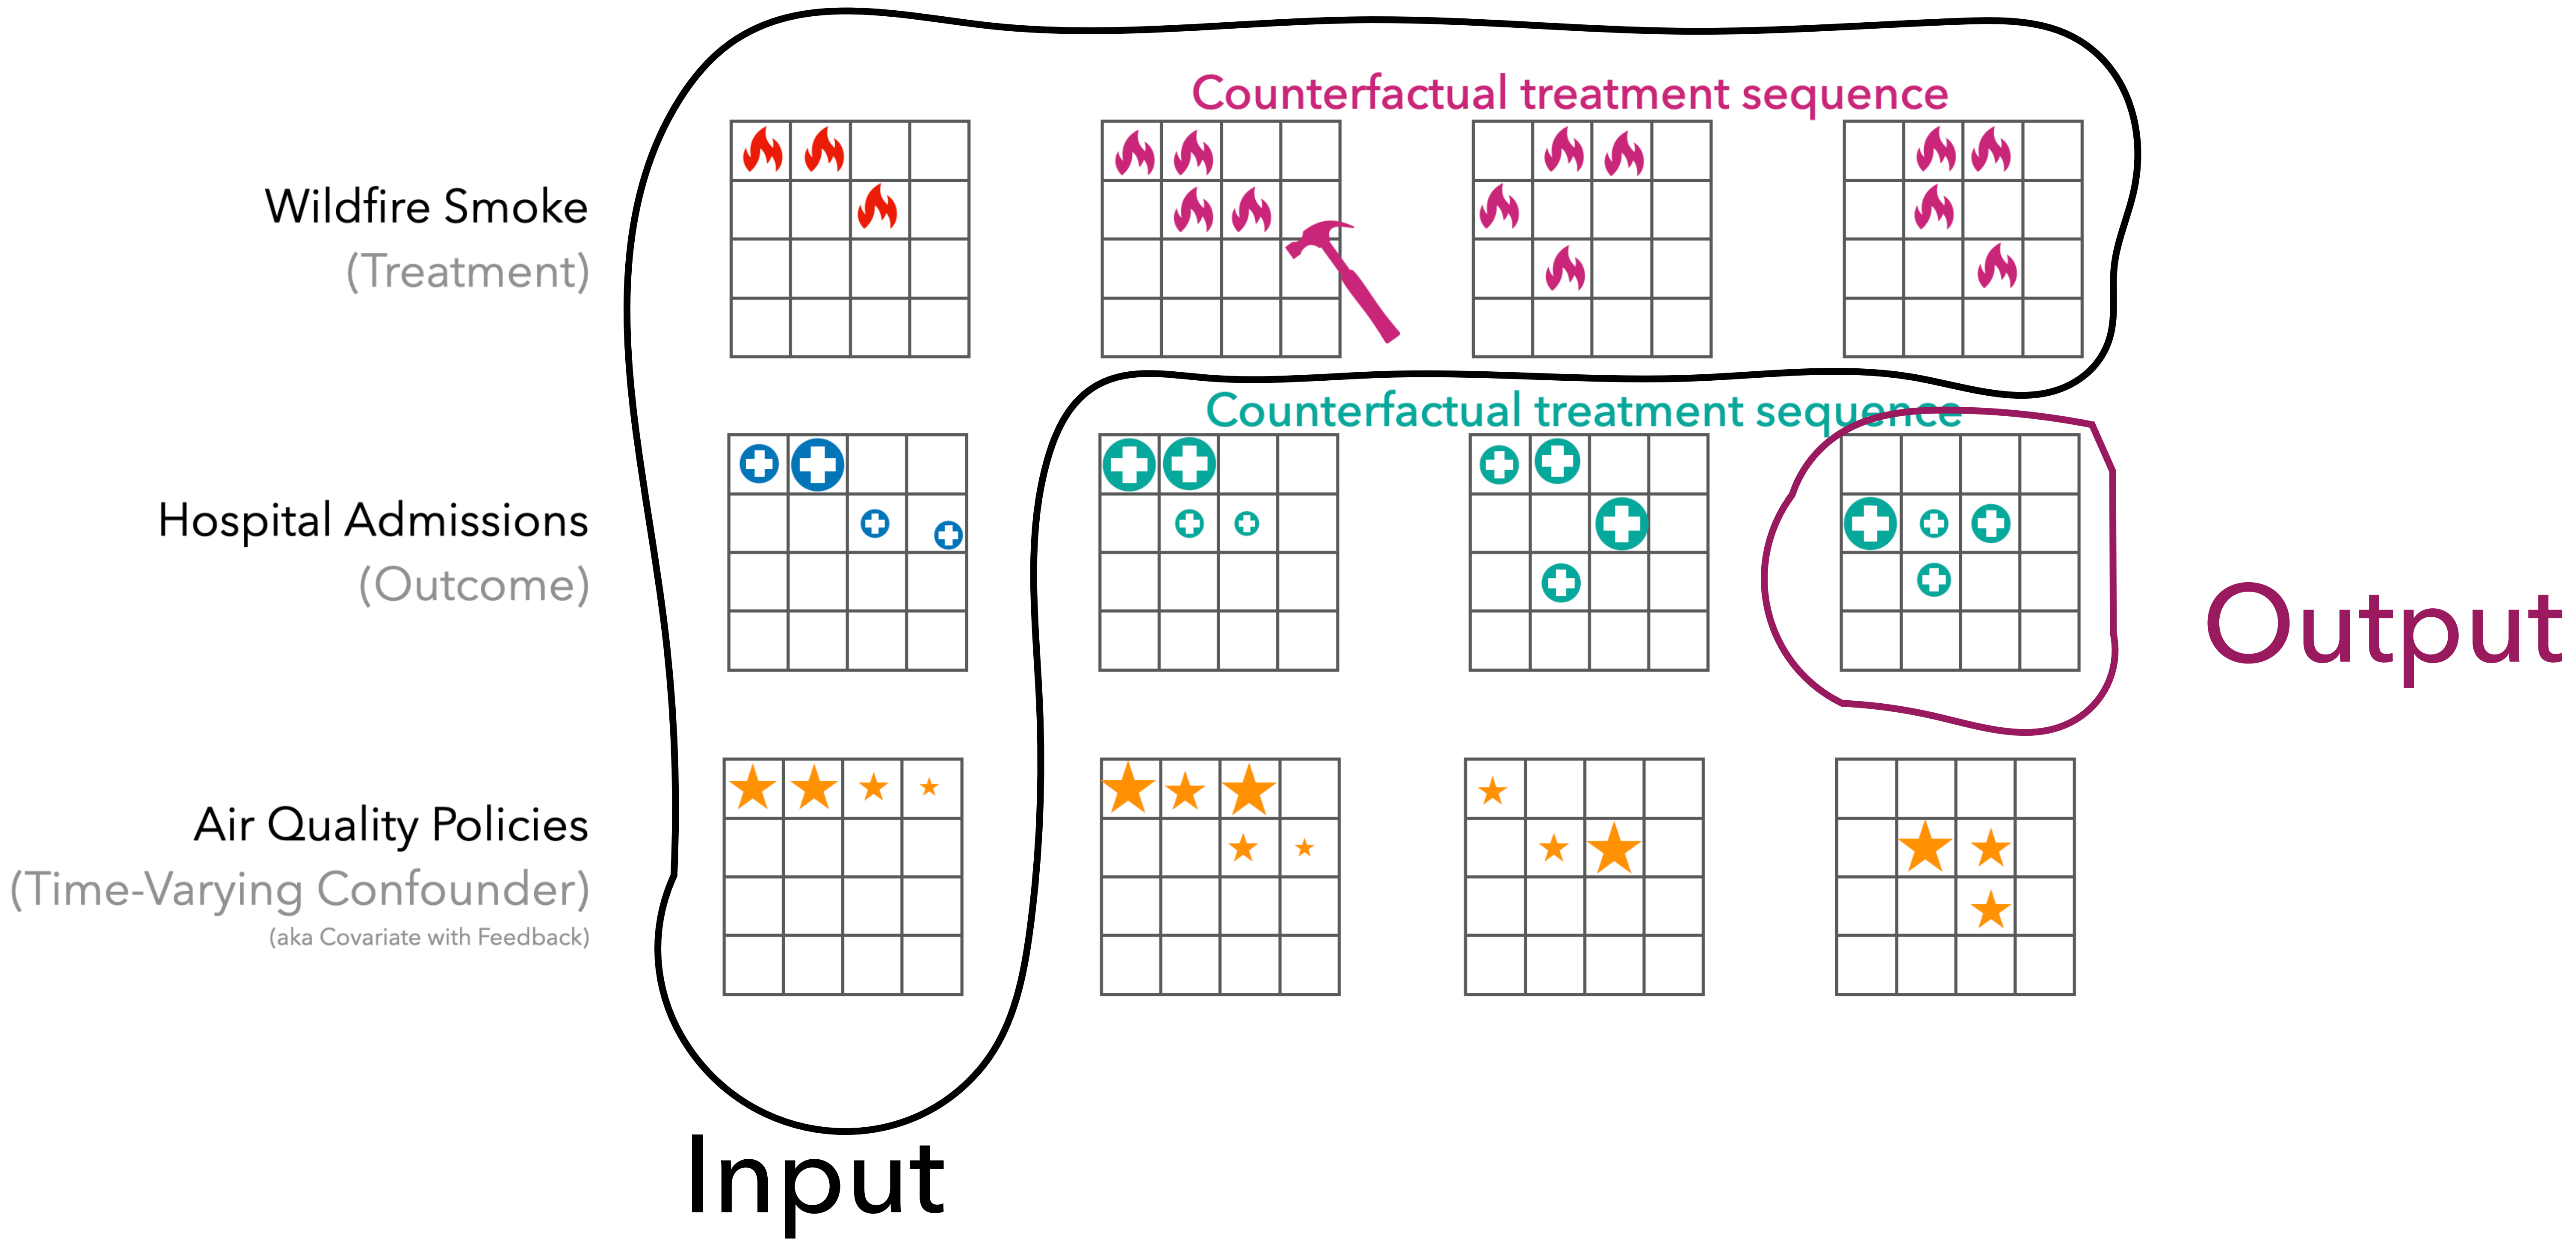

G-Computation (Robins, 1986)

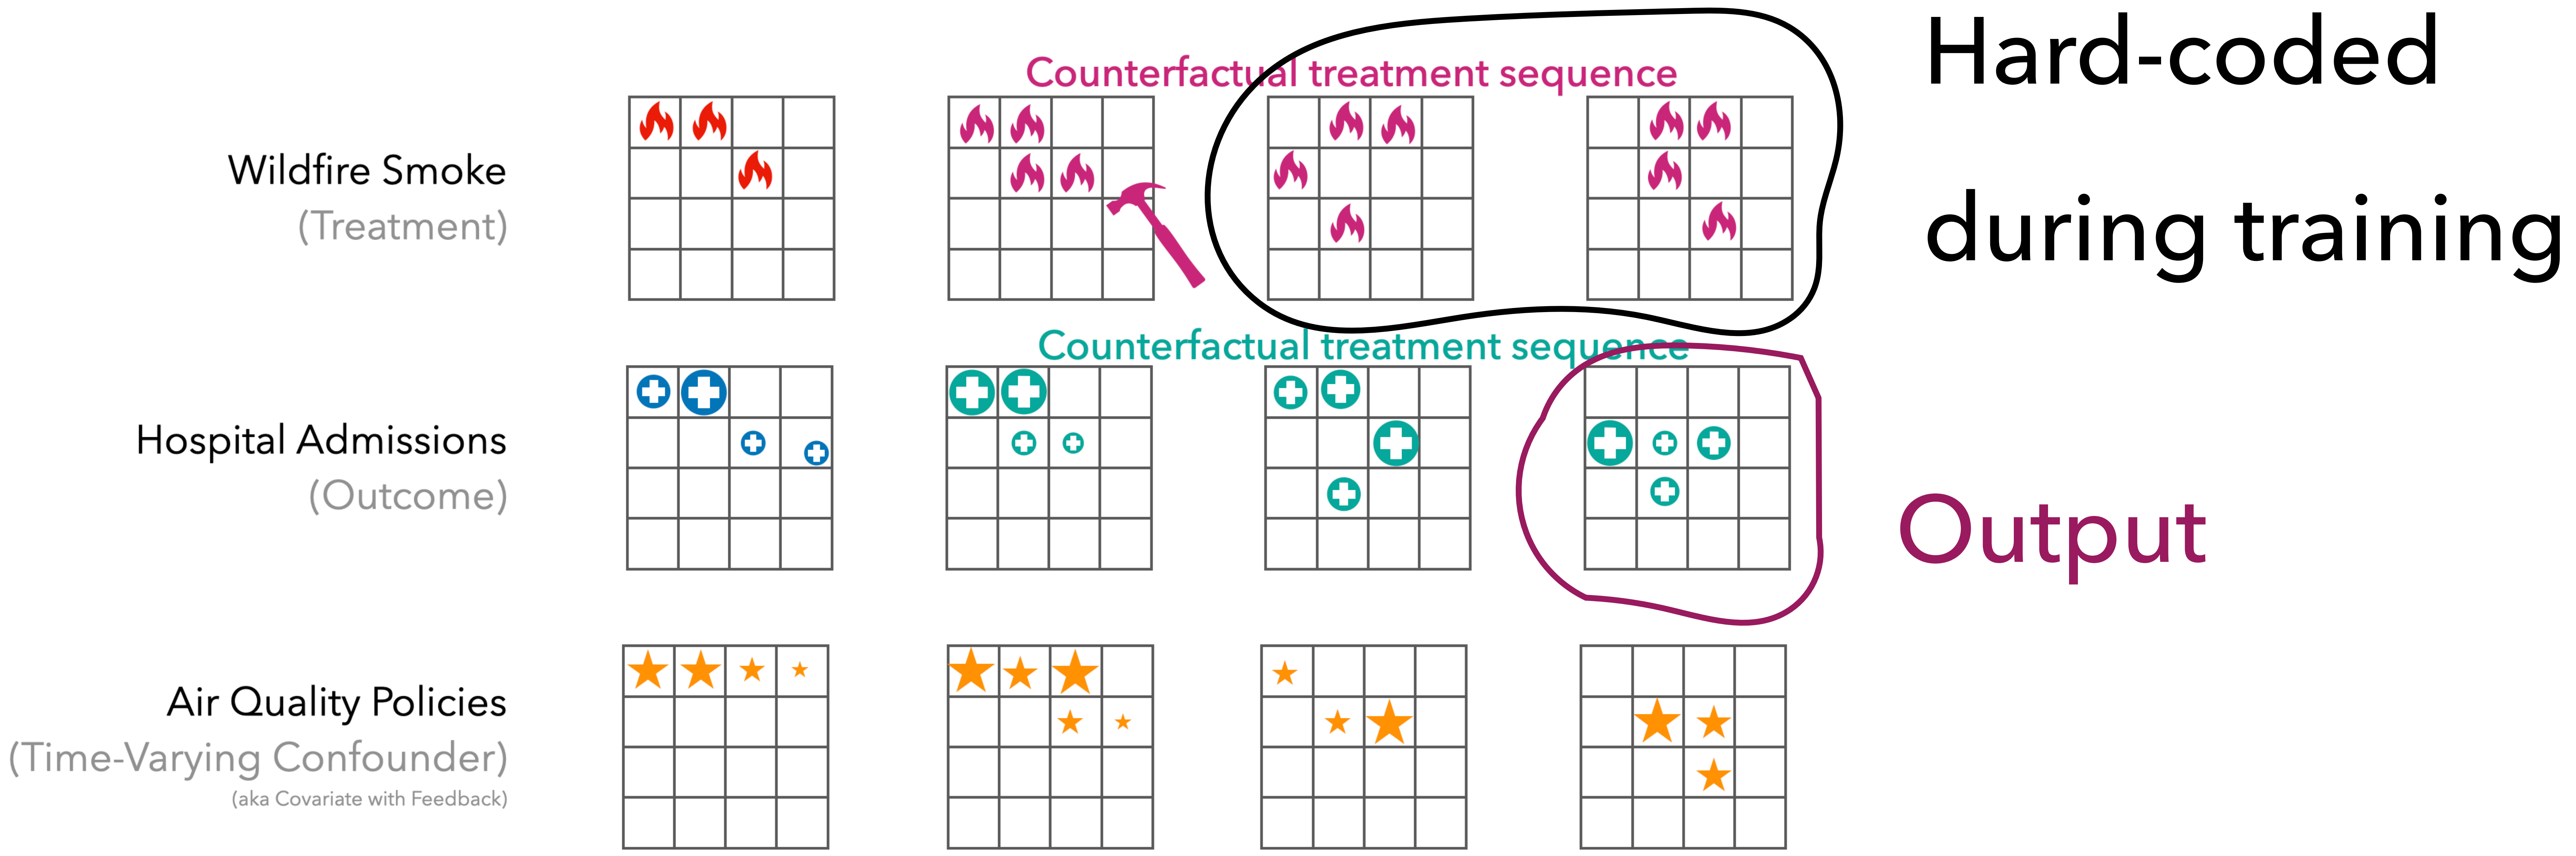

G-Computation (Robins, 1986)

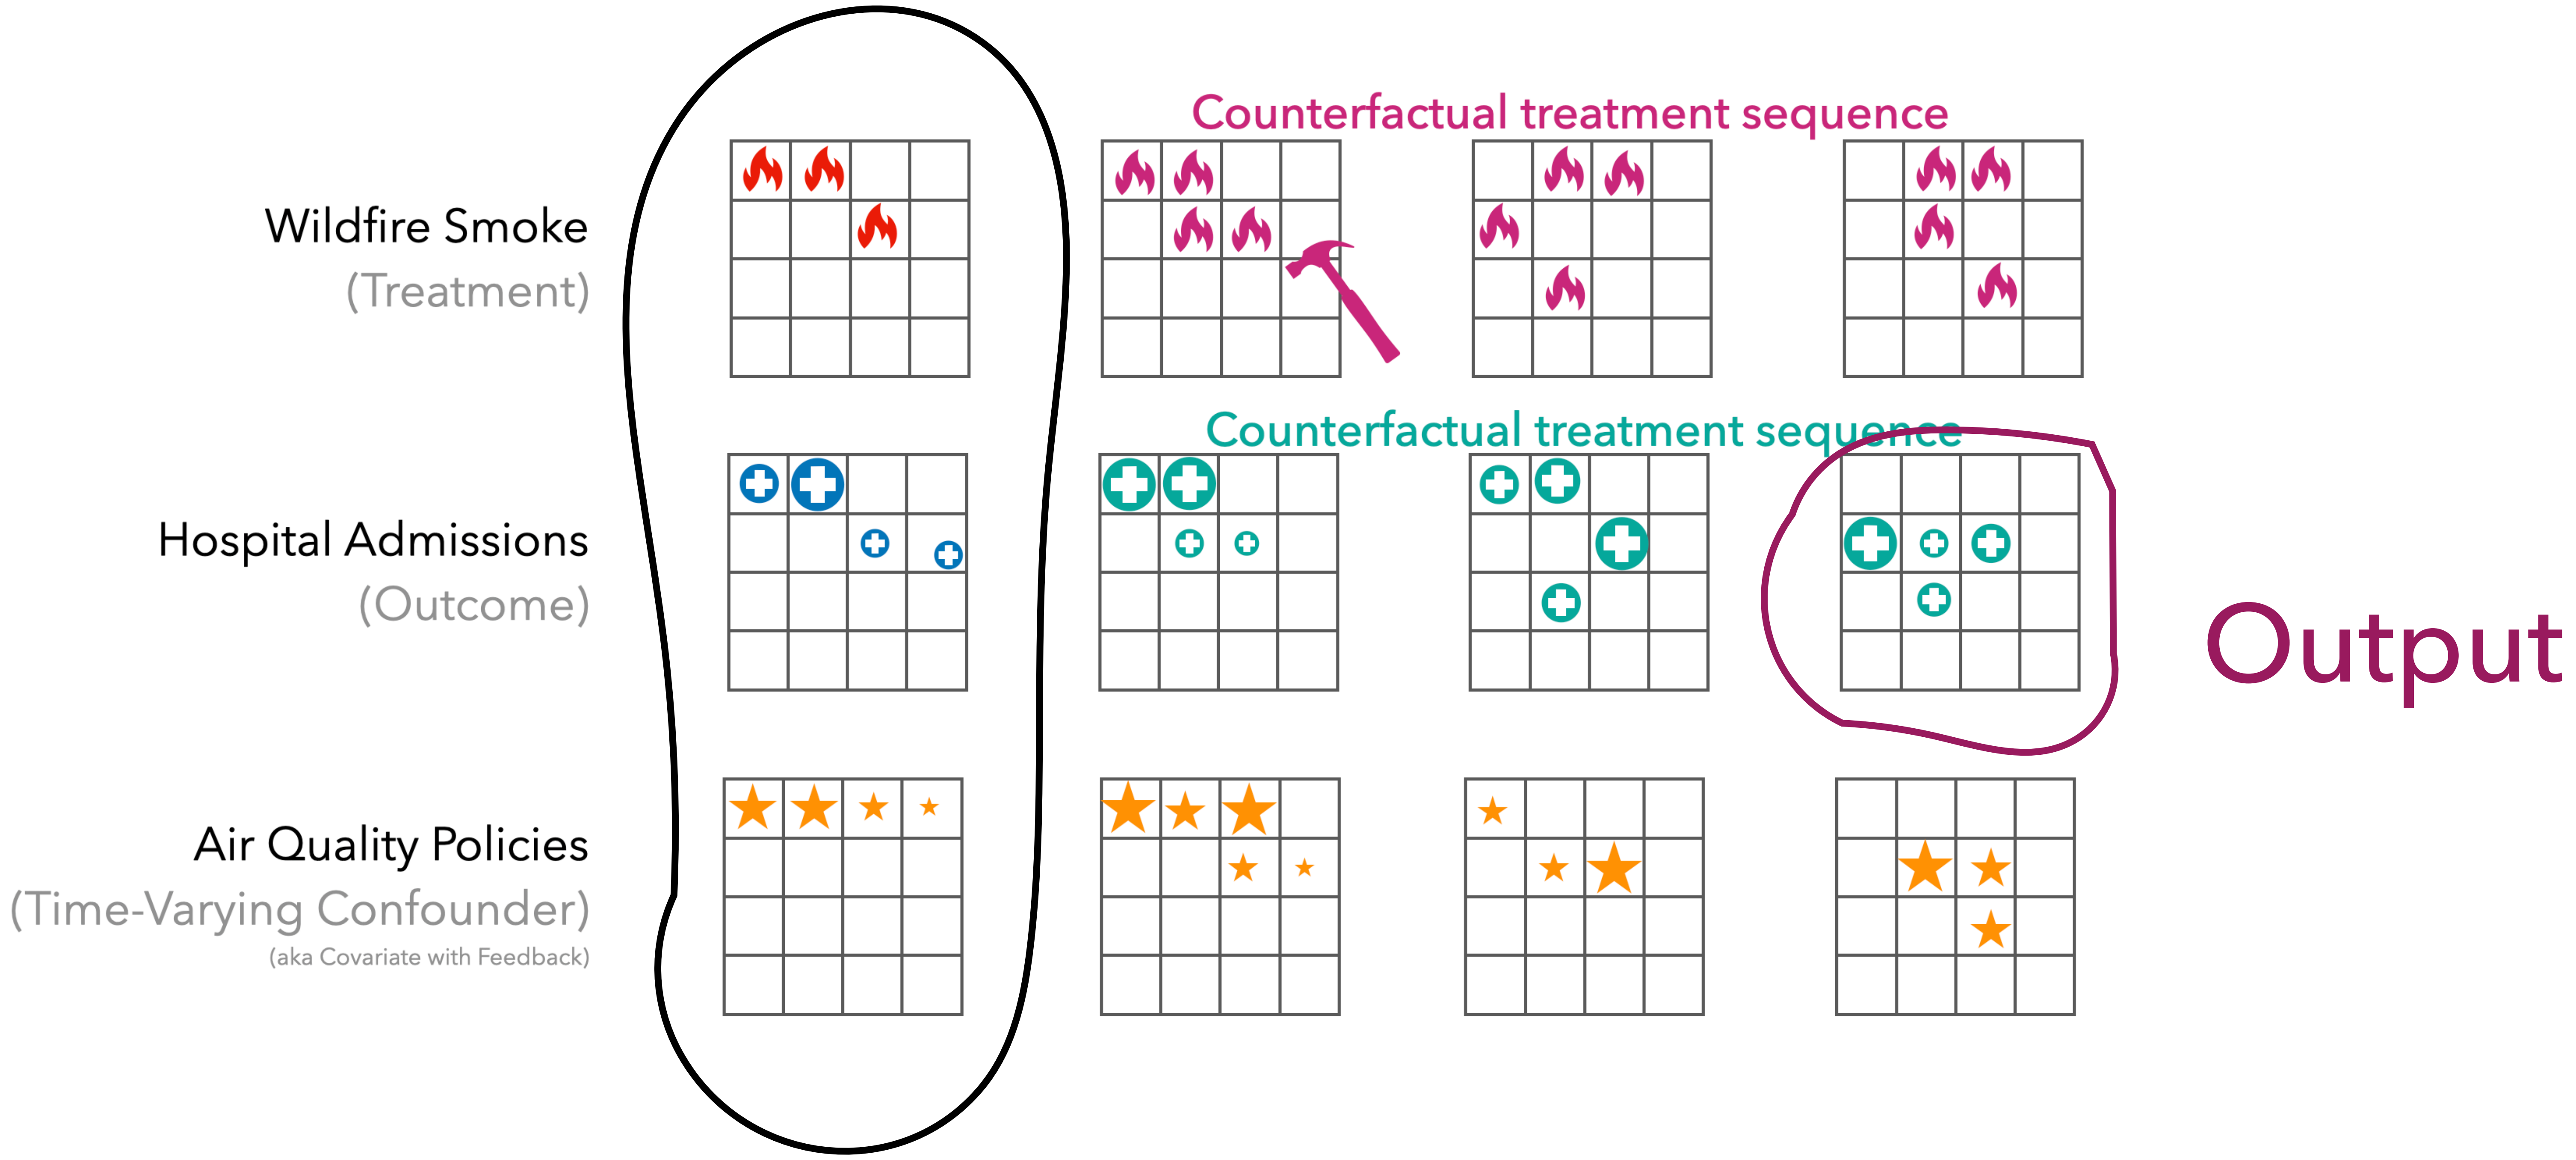

History Encoder

(Makes everything Markovian)

G-Computation (Robins, 1986)

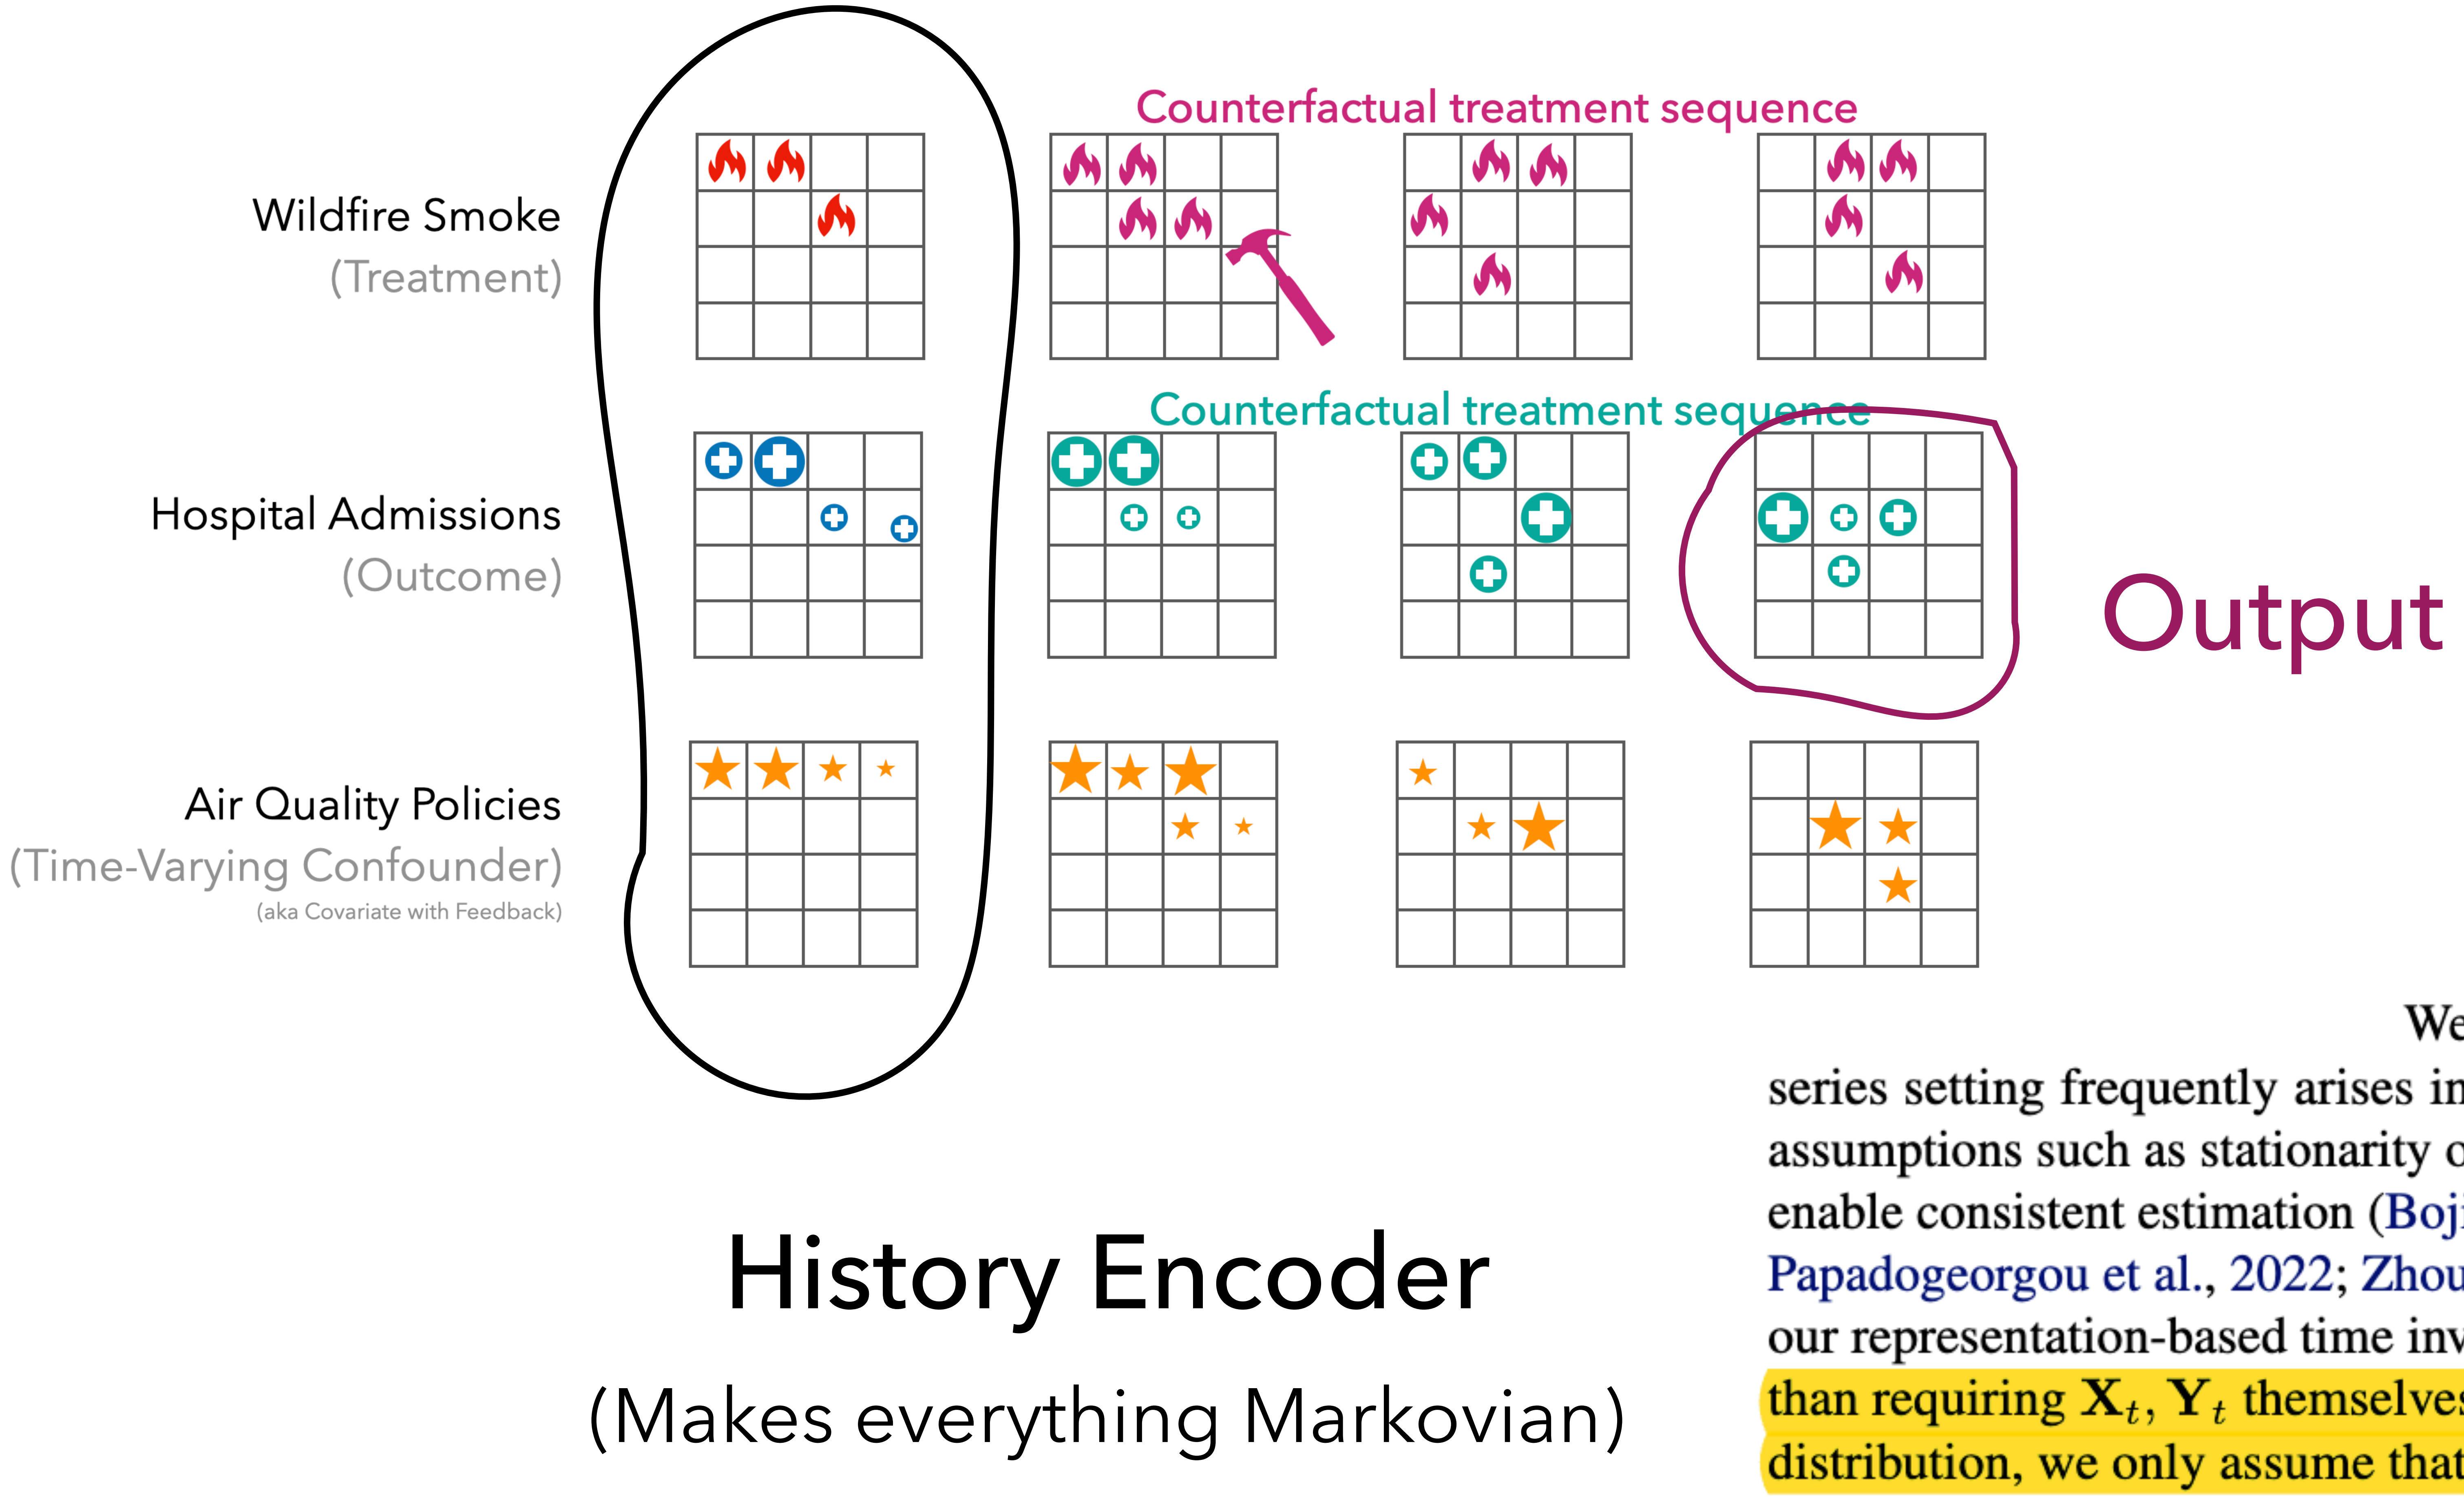

We note that the single time-series setting frequently arises in causal inference, where assumptions such as stationarity or strict time homogeneity enable consistent estimation (Bojinov and Shephard, 2019; Papadogeorgou et al., 2022; Zhou et al., 2024). In contrast, our representation-based time invariance is *weaker*: rather than requiring $\mathbf{X}_t, \mathbf{Y}_t$ themselves to have a time-invariant distribution, we only assume that, once the history is summarized by $\phi(\mathbf{H}_{1:t}, \mathbf{A}_t)$, the transition to $(\mathbf{X}_{t+1}, \mathbf{Y}_{t+1})$ follows a single shared mechanism.

NN

History until t

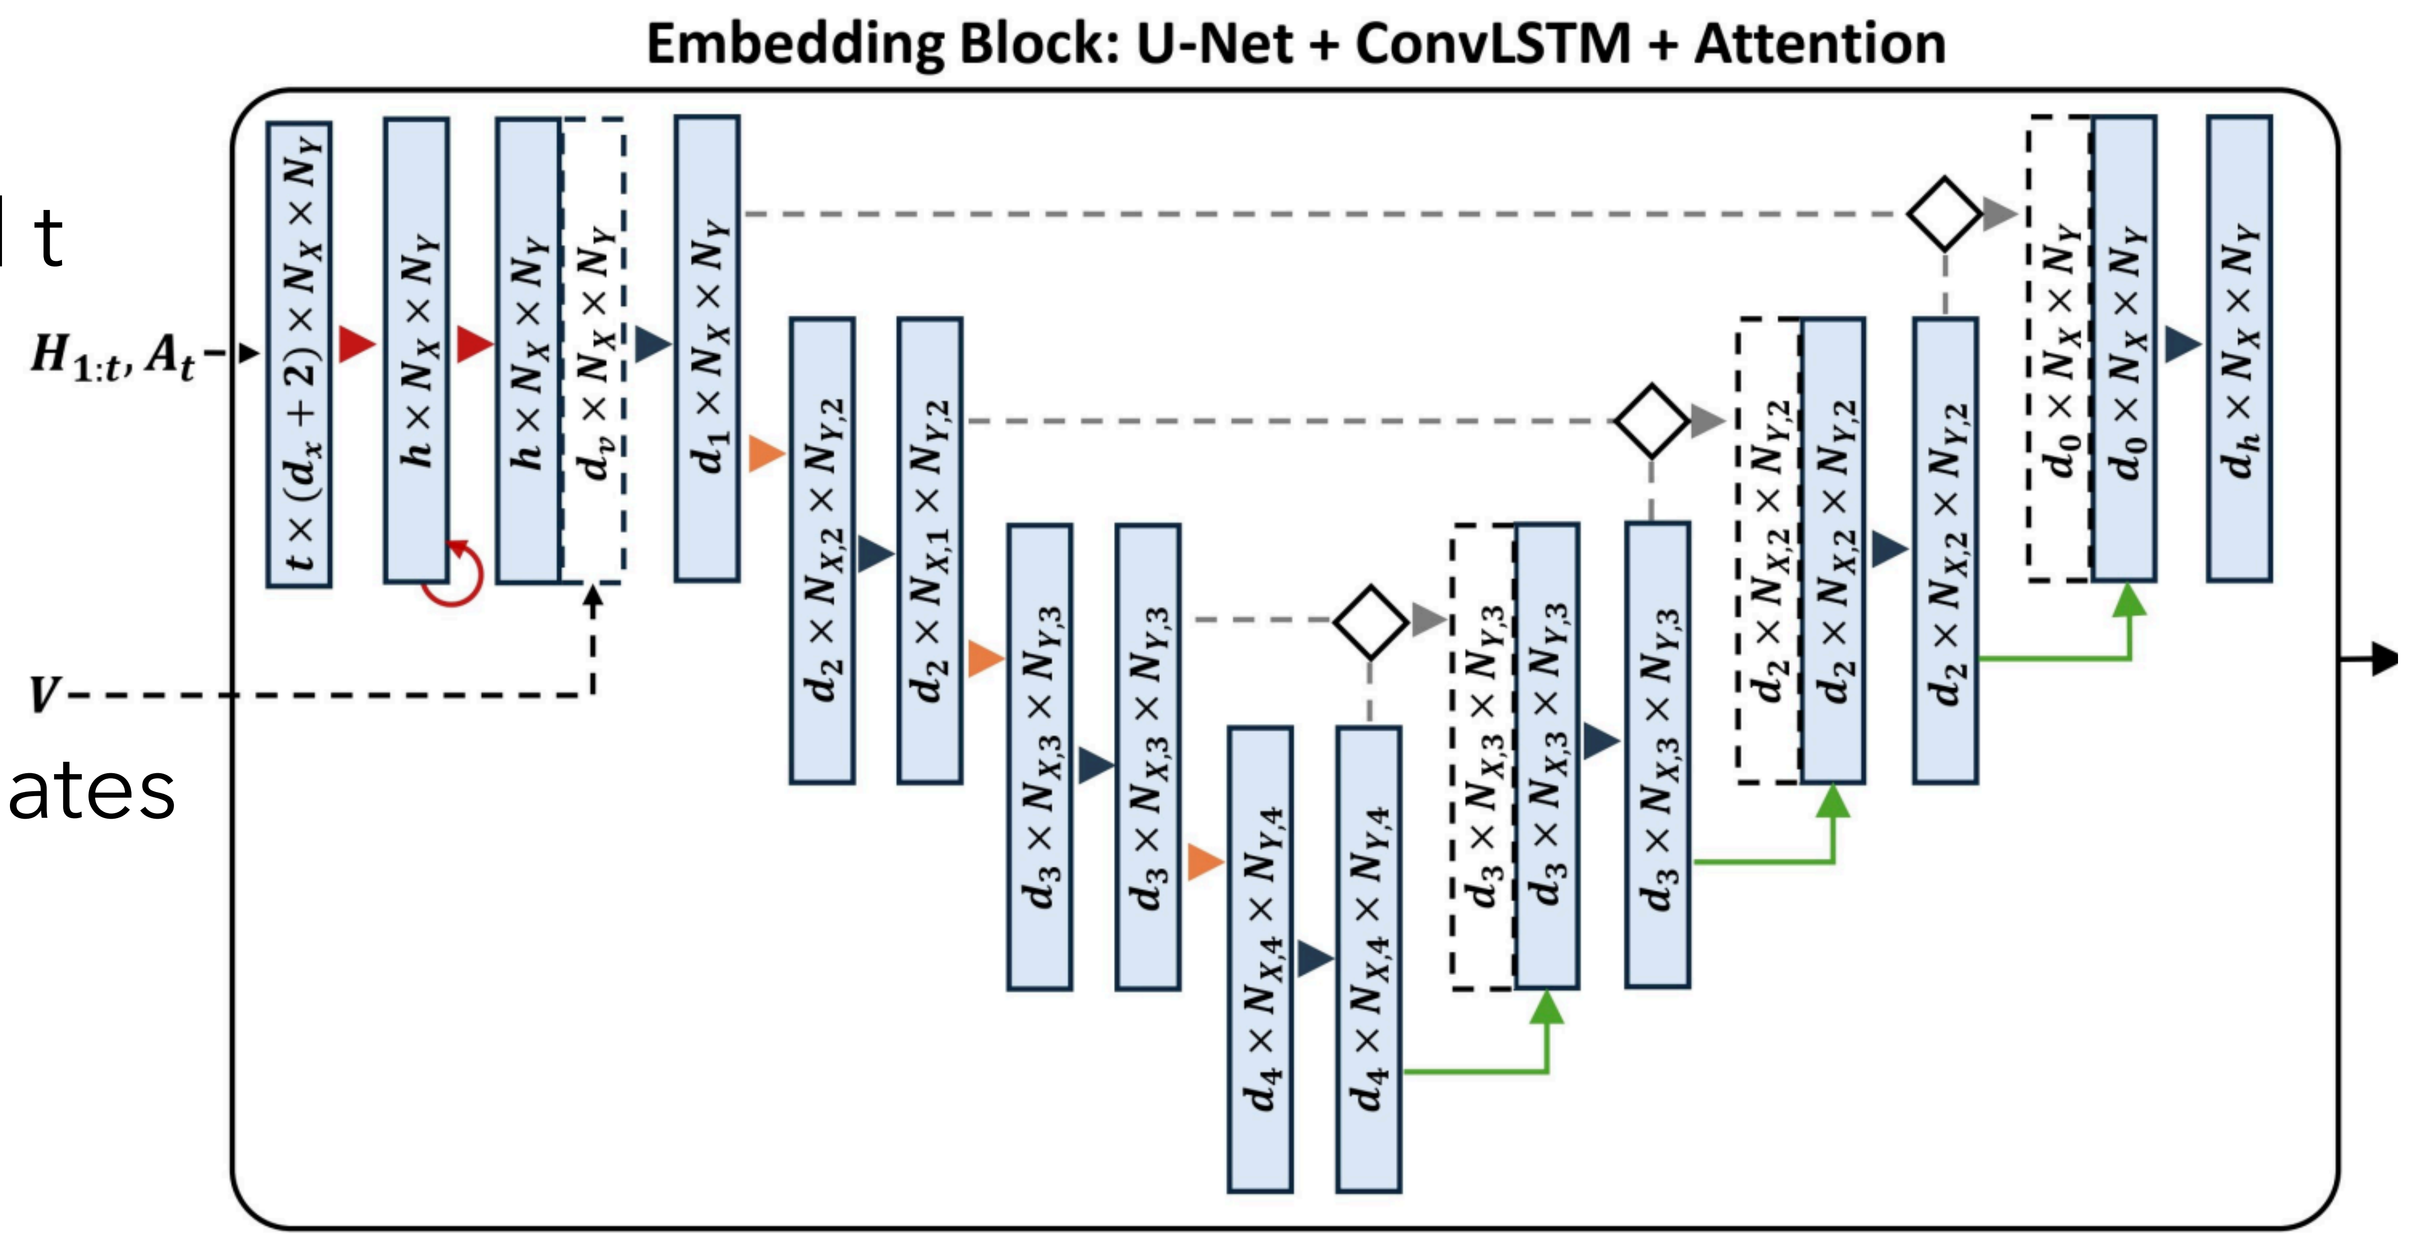

Static Covariates

History Encoding

(Makes everything Markovian)

NN

History until t

Static Covariates

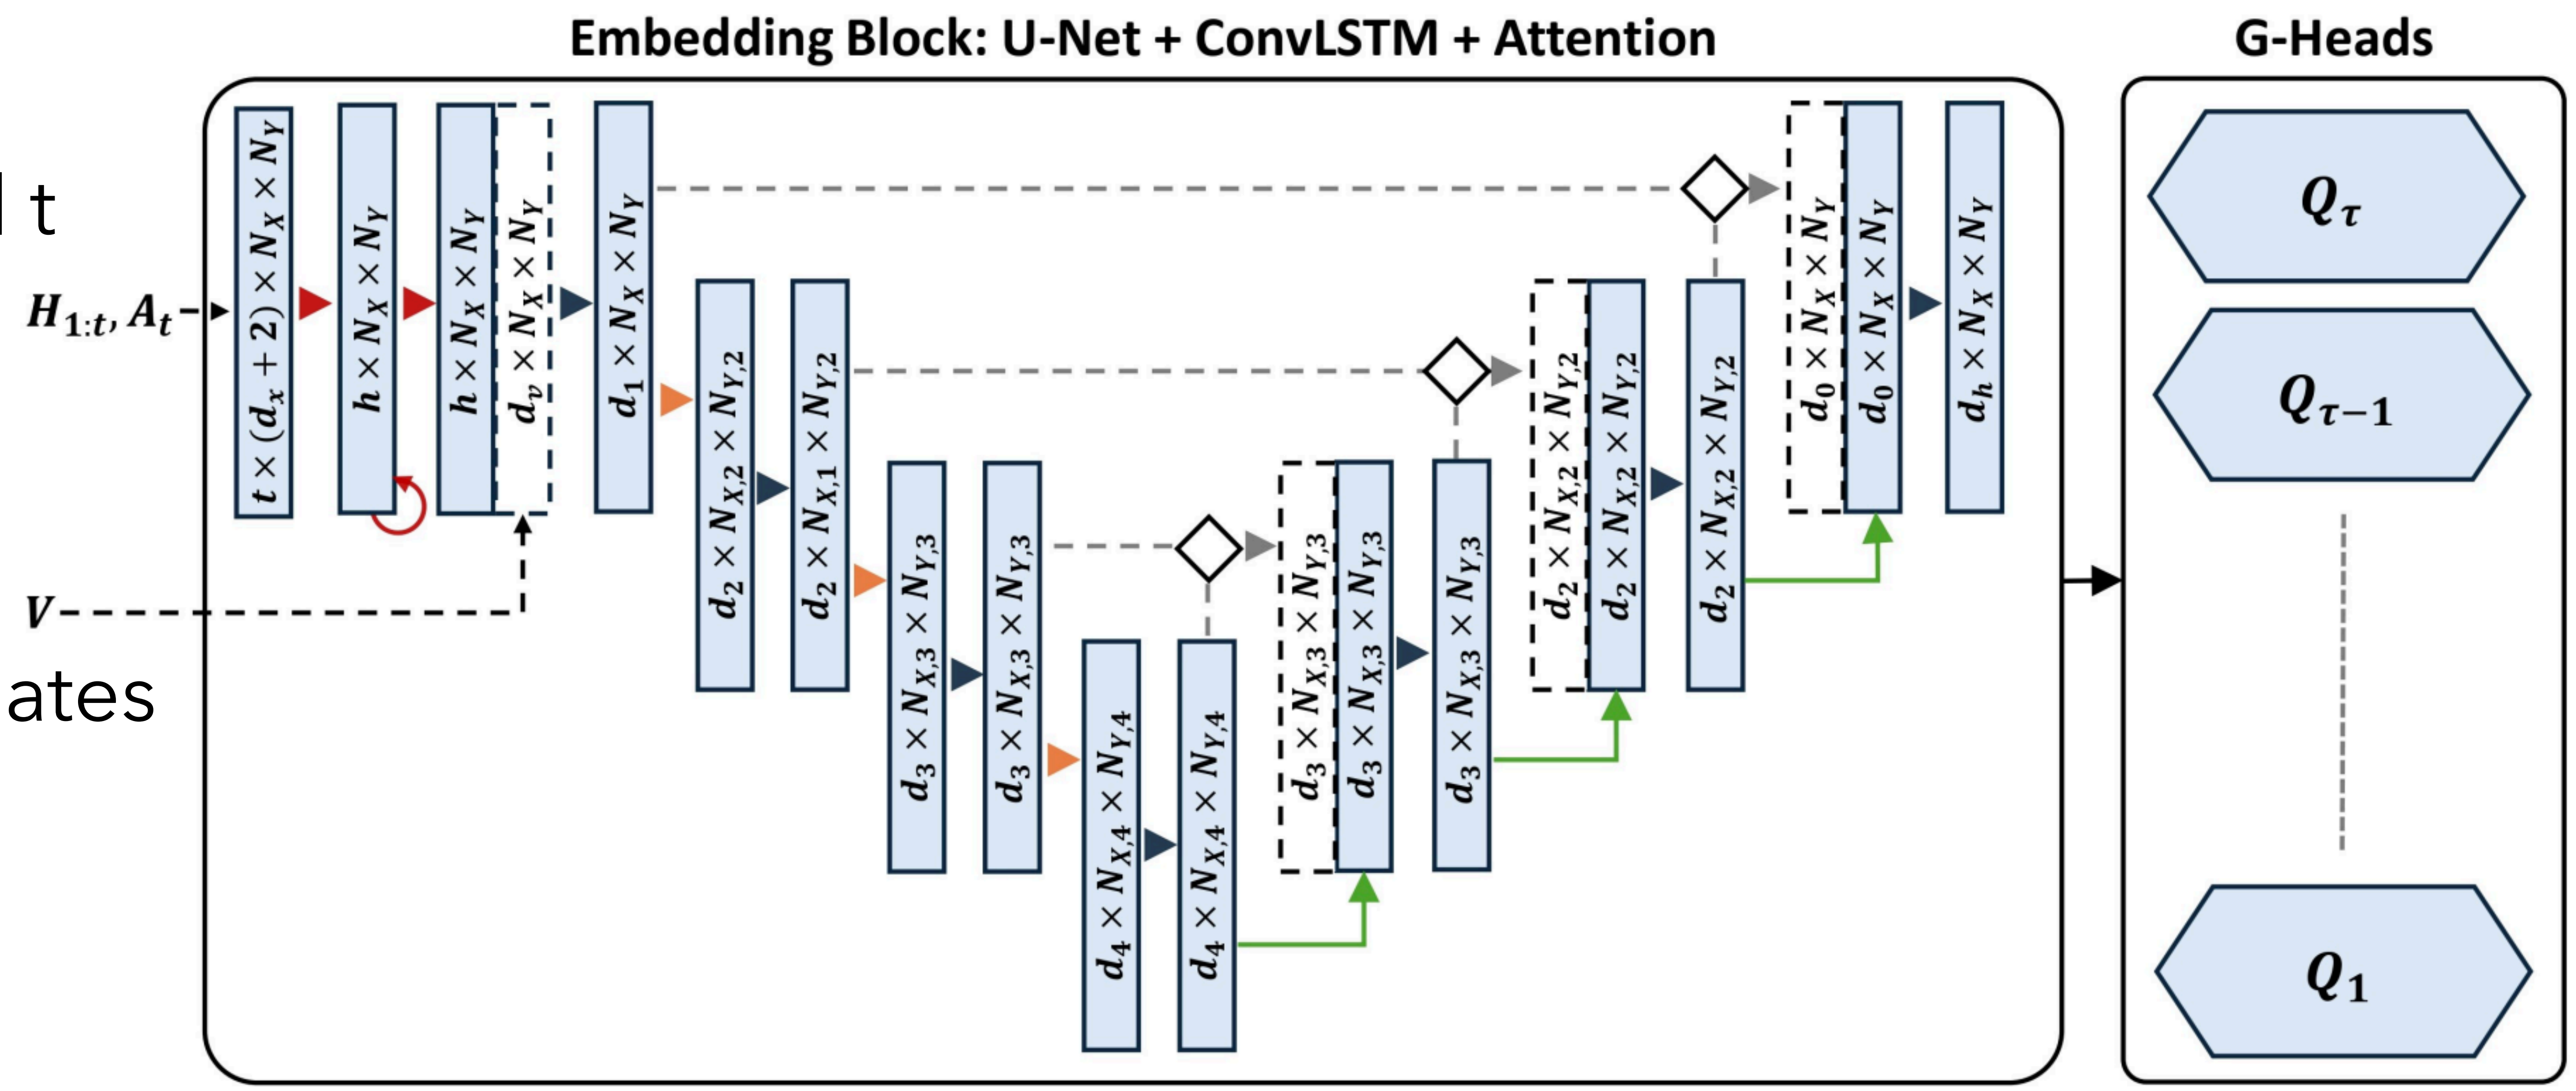

Each head Q_k estimates the CAPO: The expected hospitalisations on the final day $t + \tau$.

NN

History until t

Static Covariates

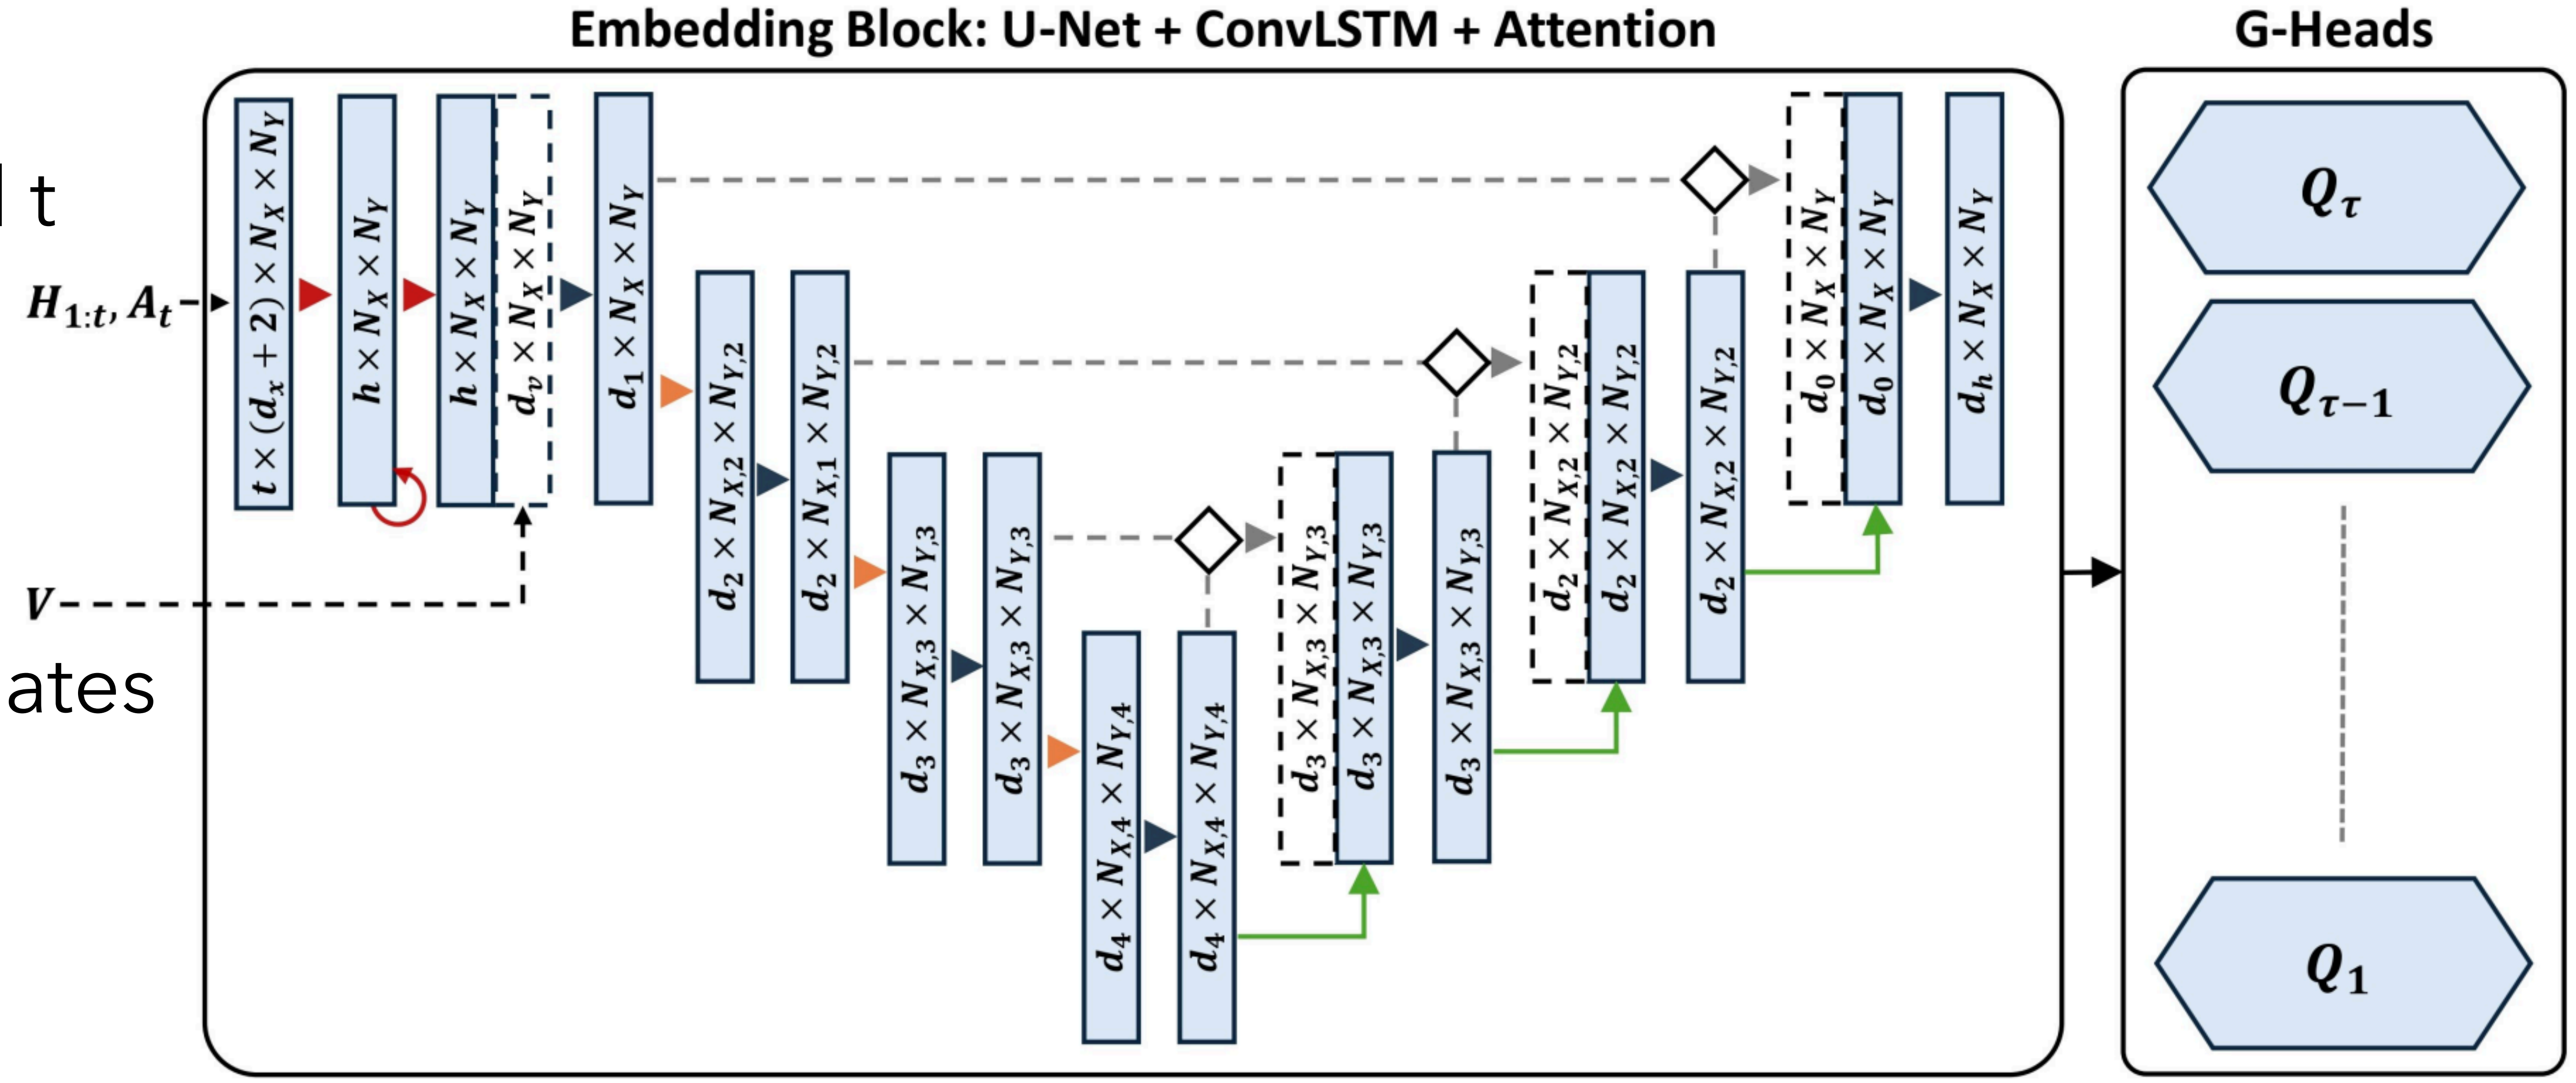

Each head Q_k estimates the CAPO: The expected hospitalisations on the final day $t + \tau$.

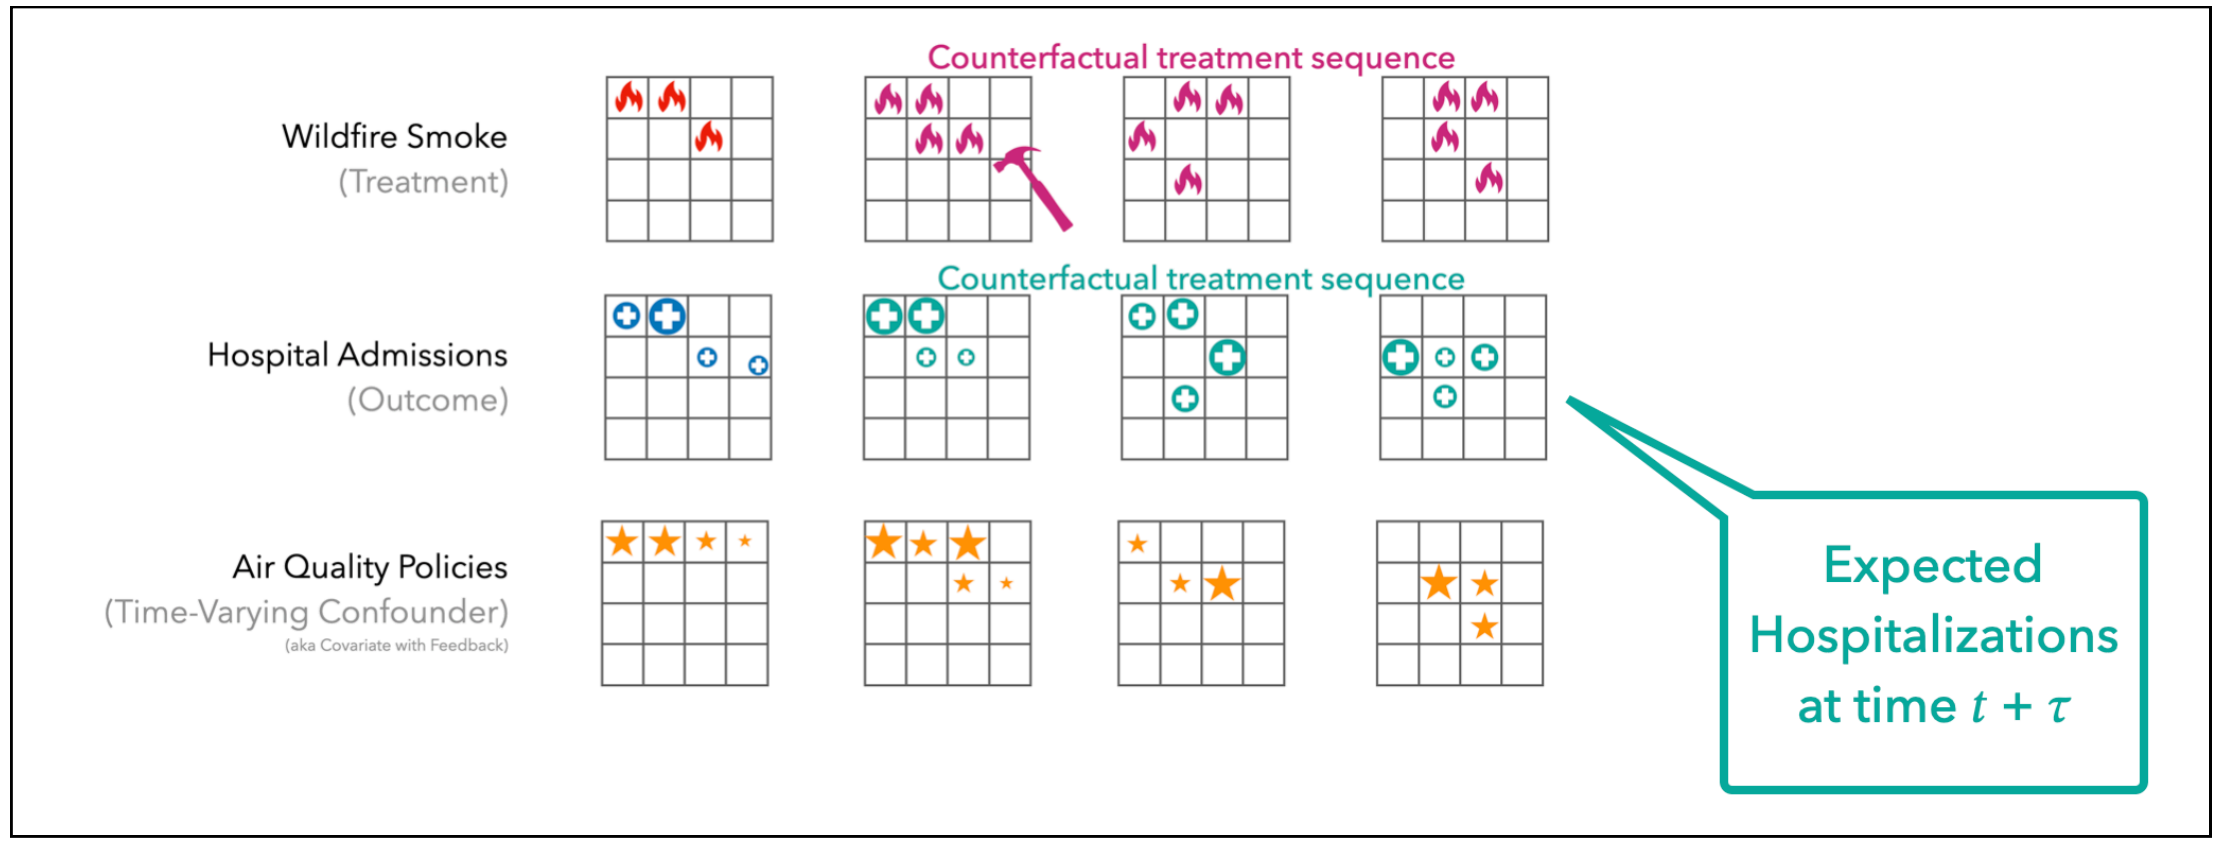

NN

History until t

Static Covariates

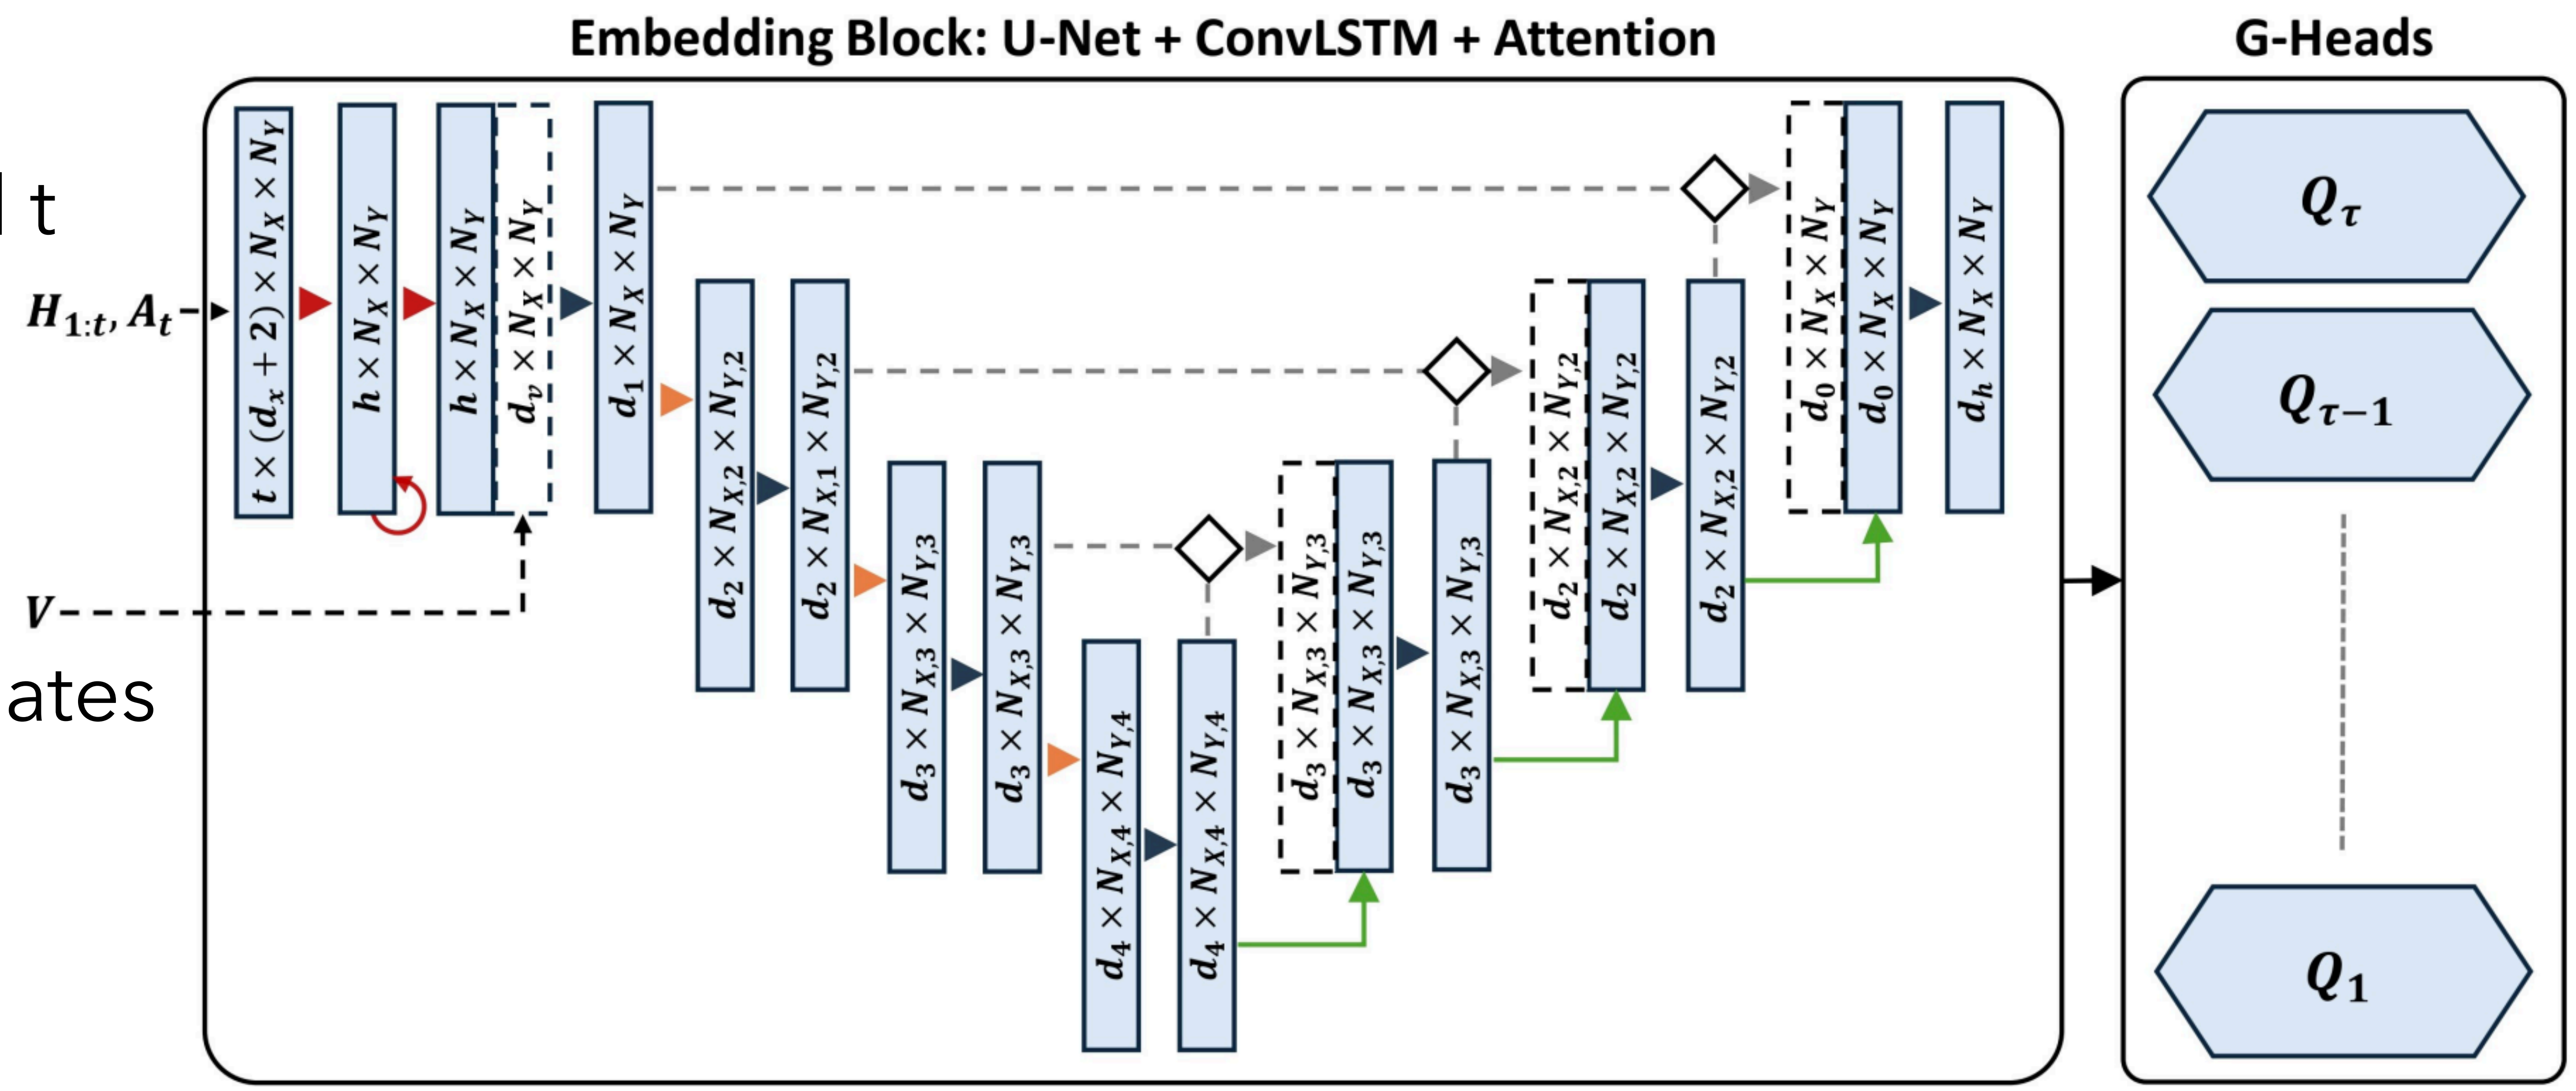

Each head Q_k estimates the CAPO: The expected hospitalisations on the final day $t + \tau$.

Q_k does so under the assumption that we know the history up to time $t + k - 1$.

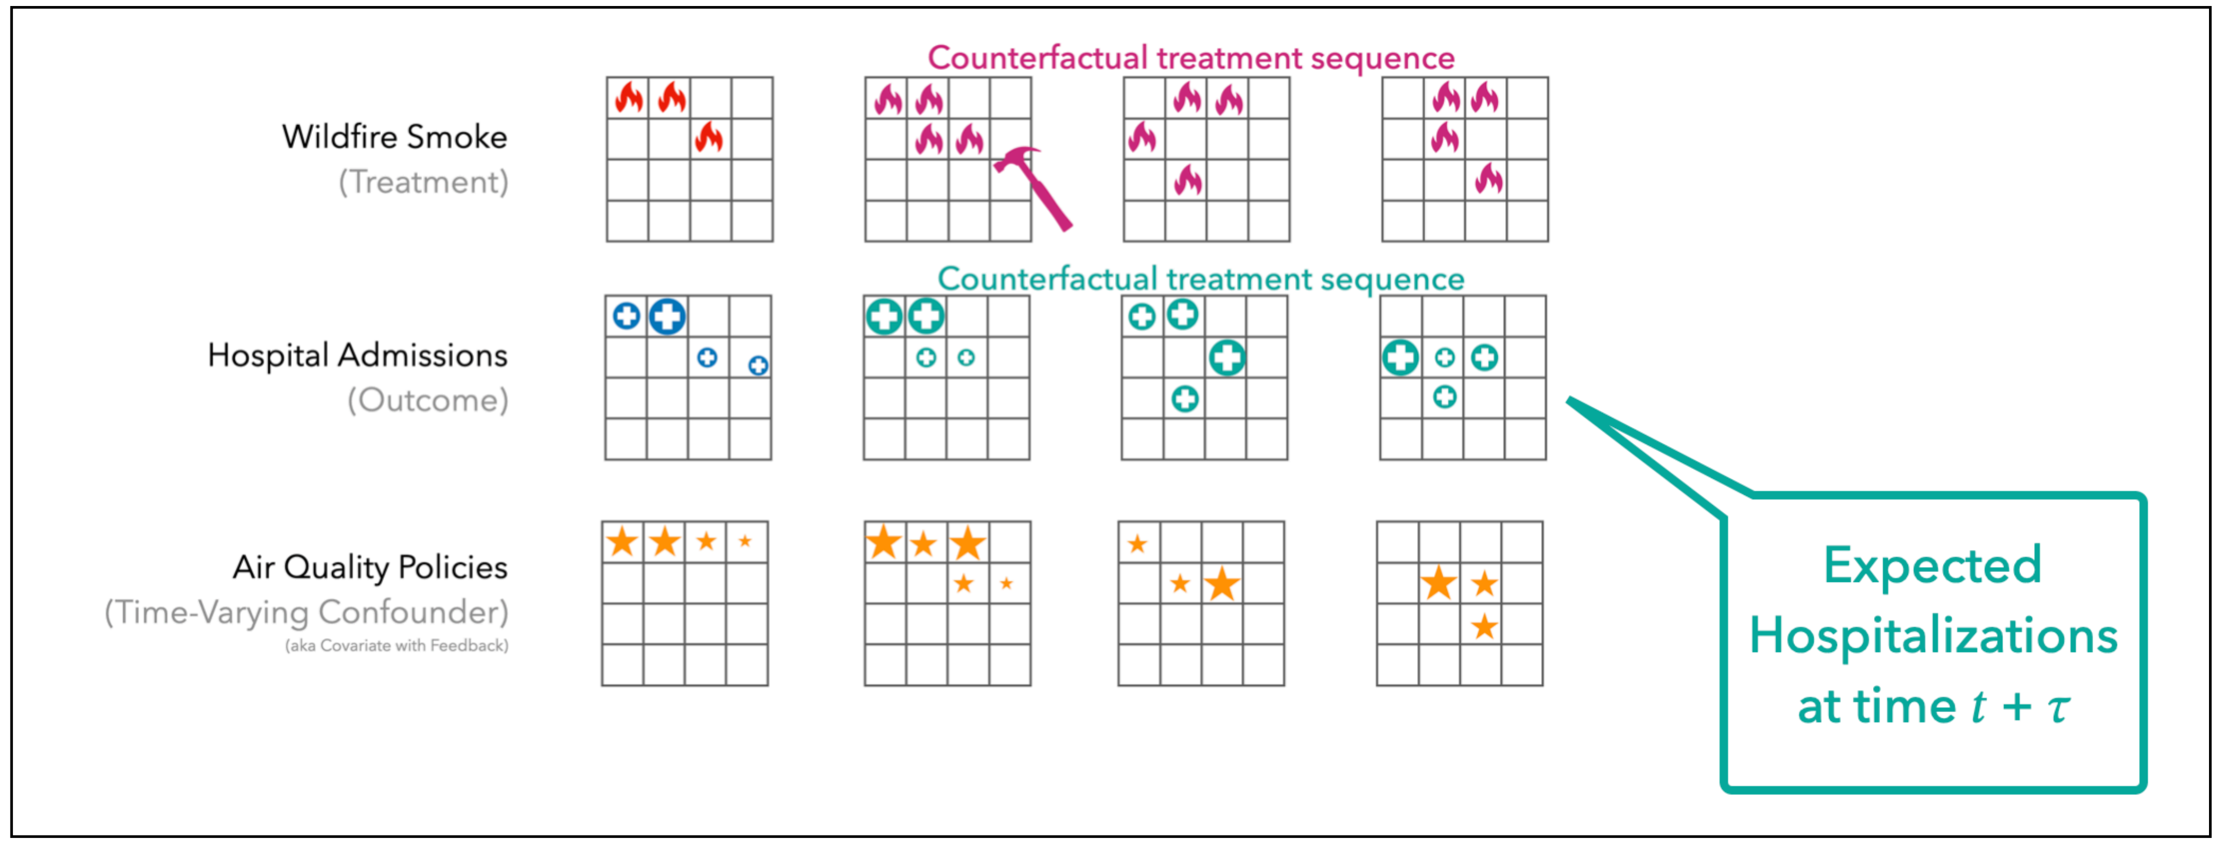

NN

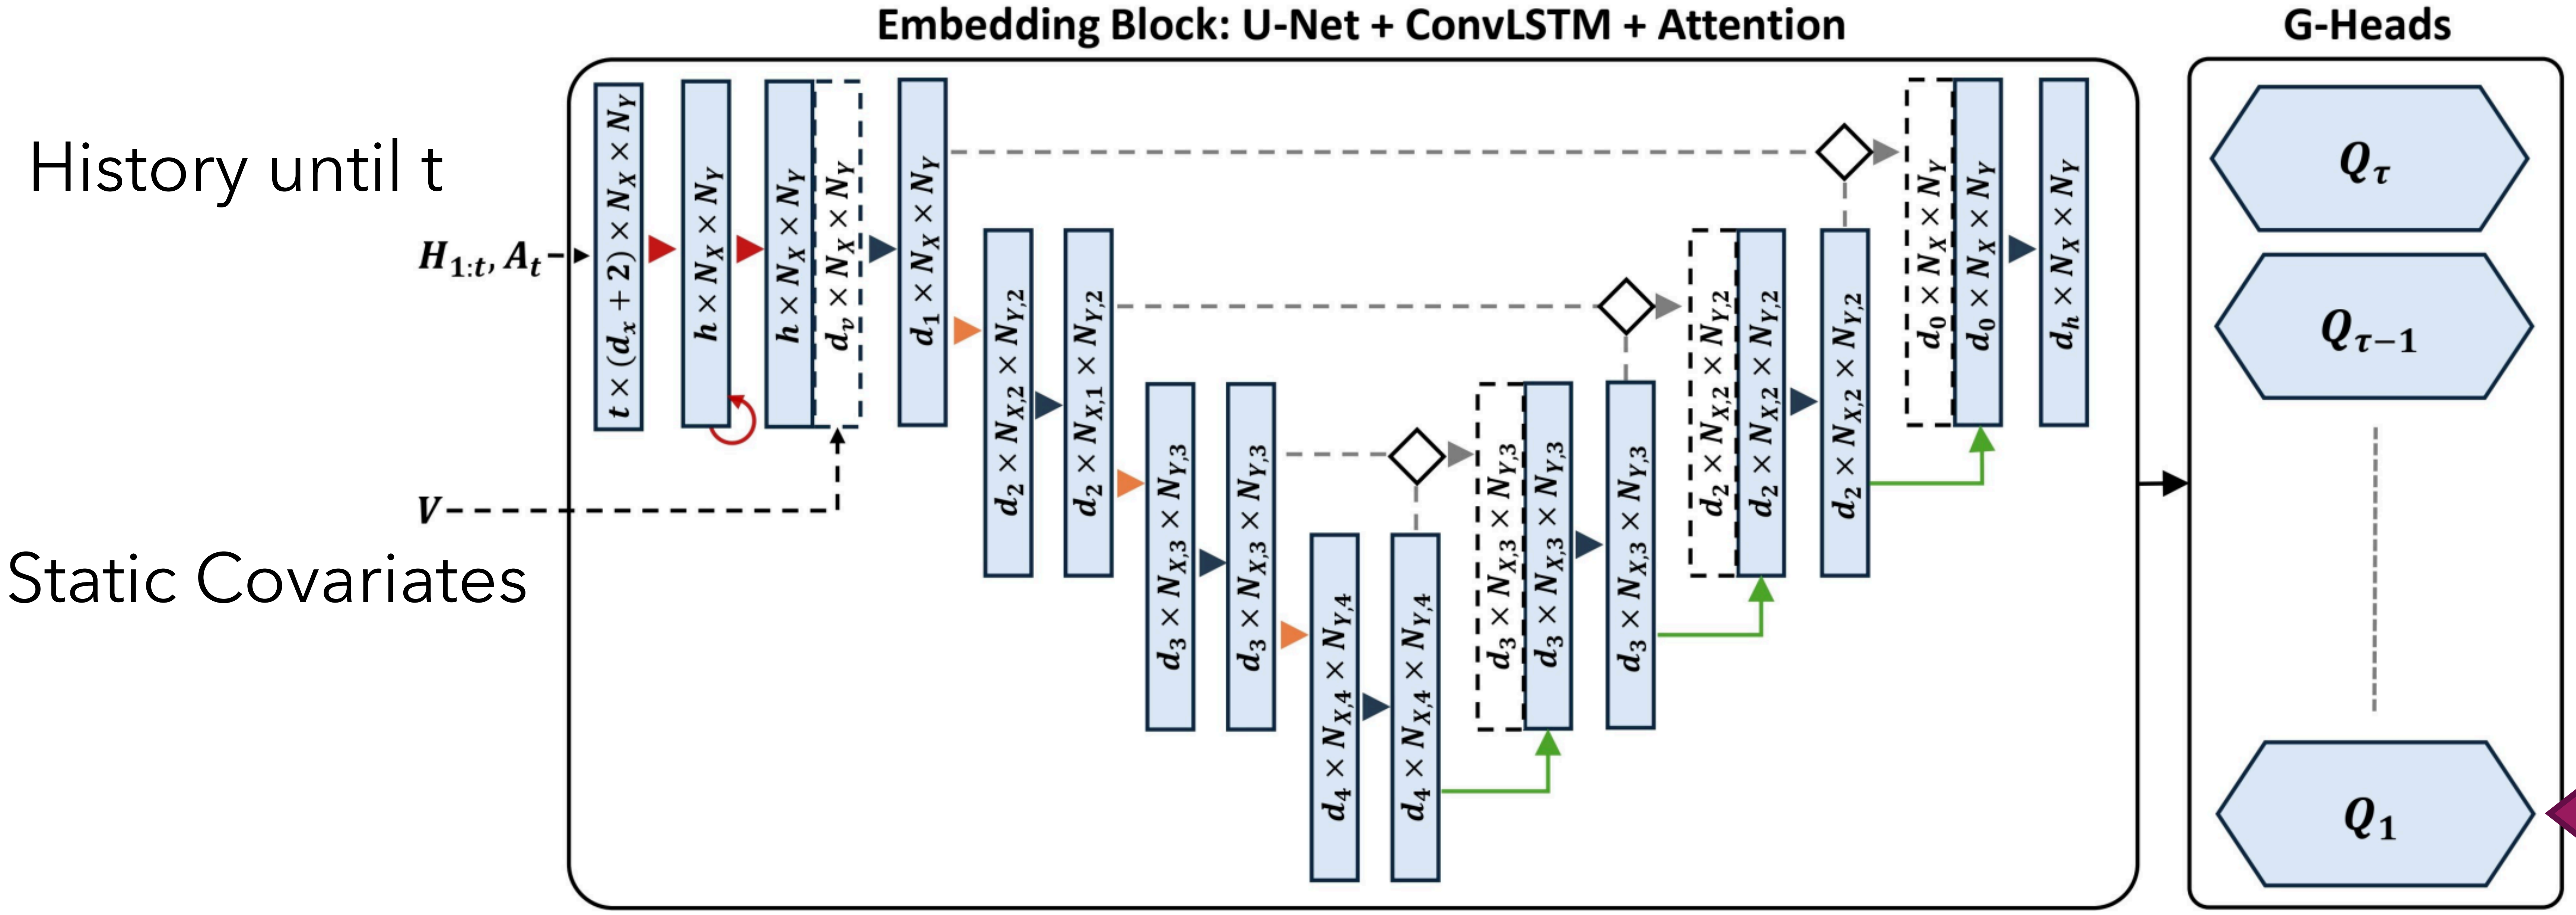

Each head Q_k estimates the CAPO: The expected hospitalisations on the final day $t + \tau$.

Q_k does so under the assumption that we know the history up to time $t + k - 1$.

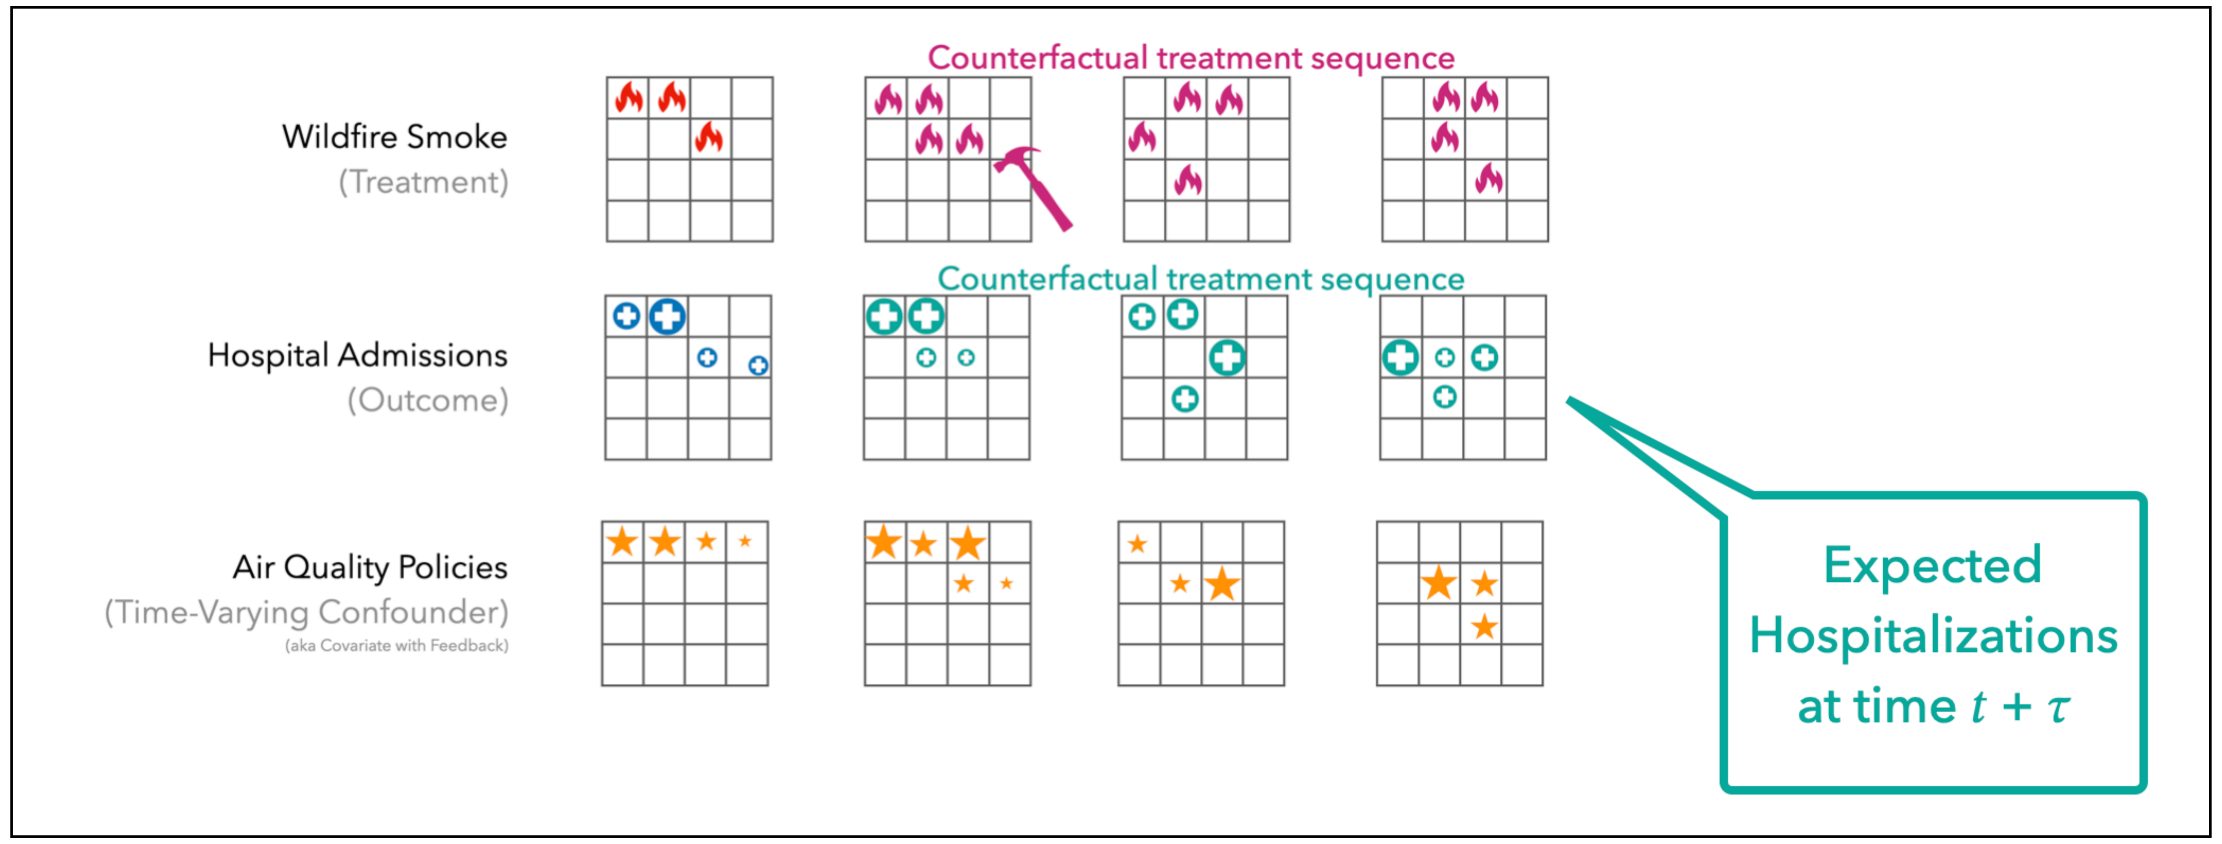

Splitting

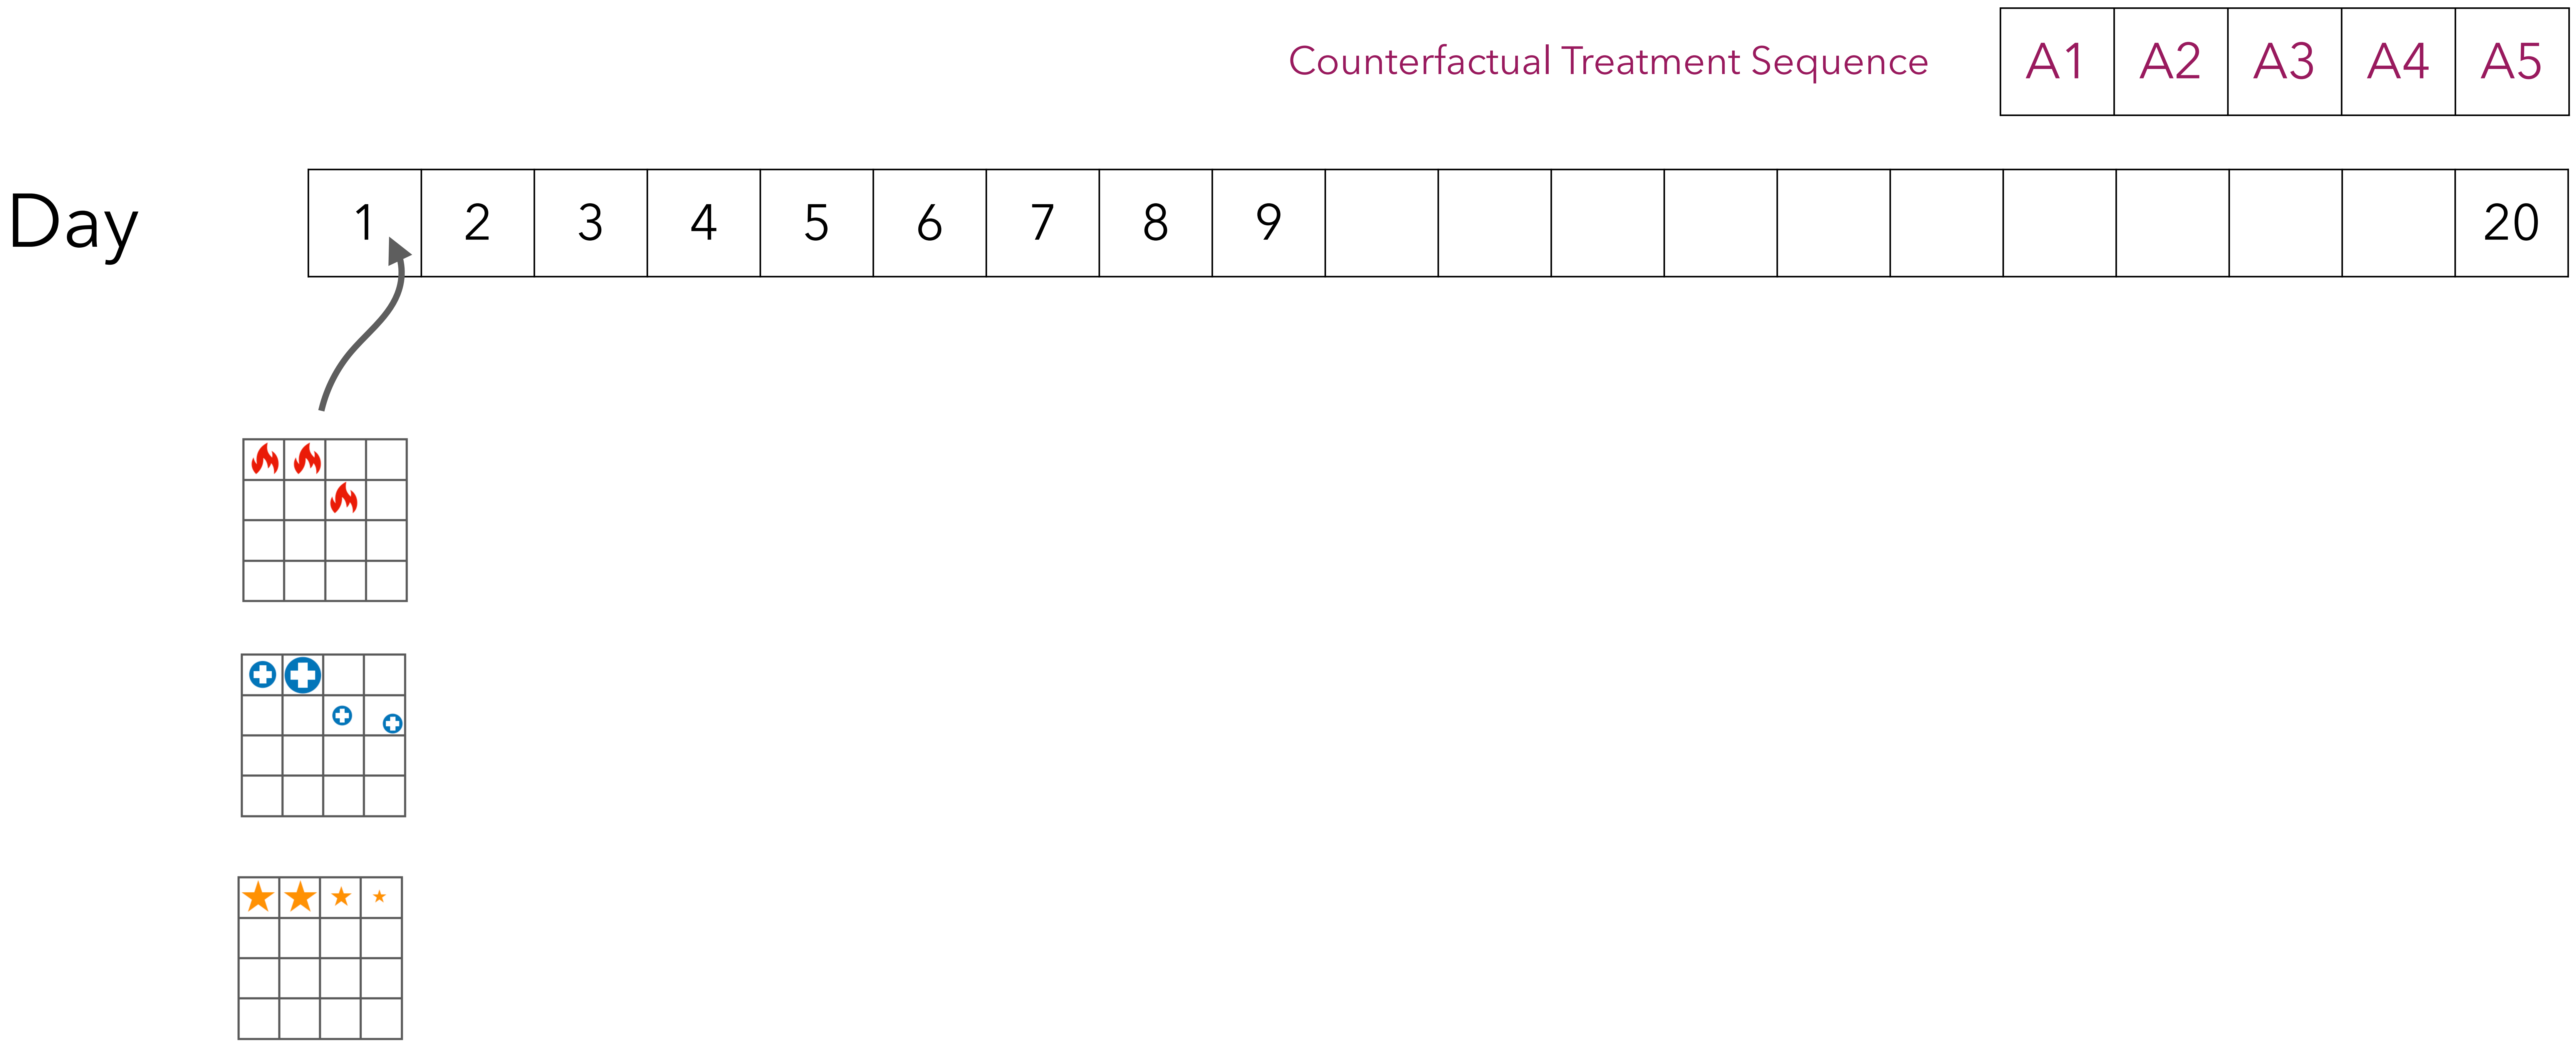
Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
20

Splitting


Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
NOW
20

Splitting

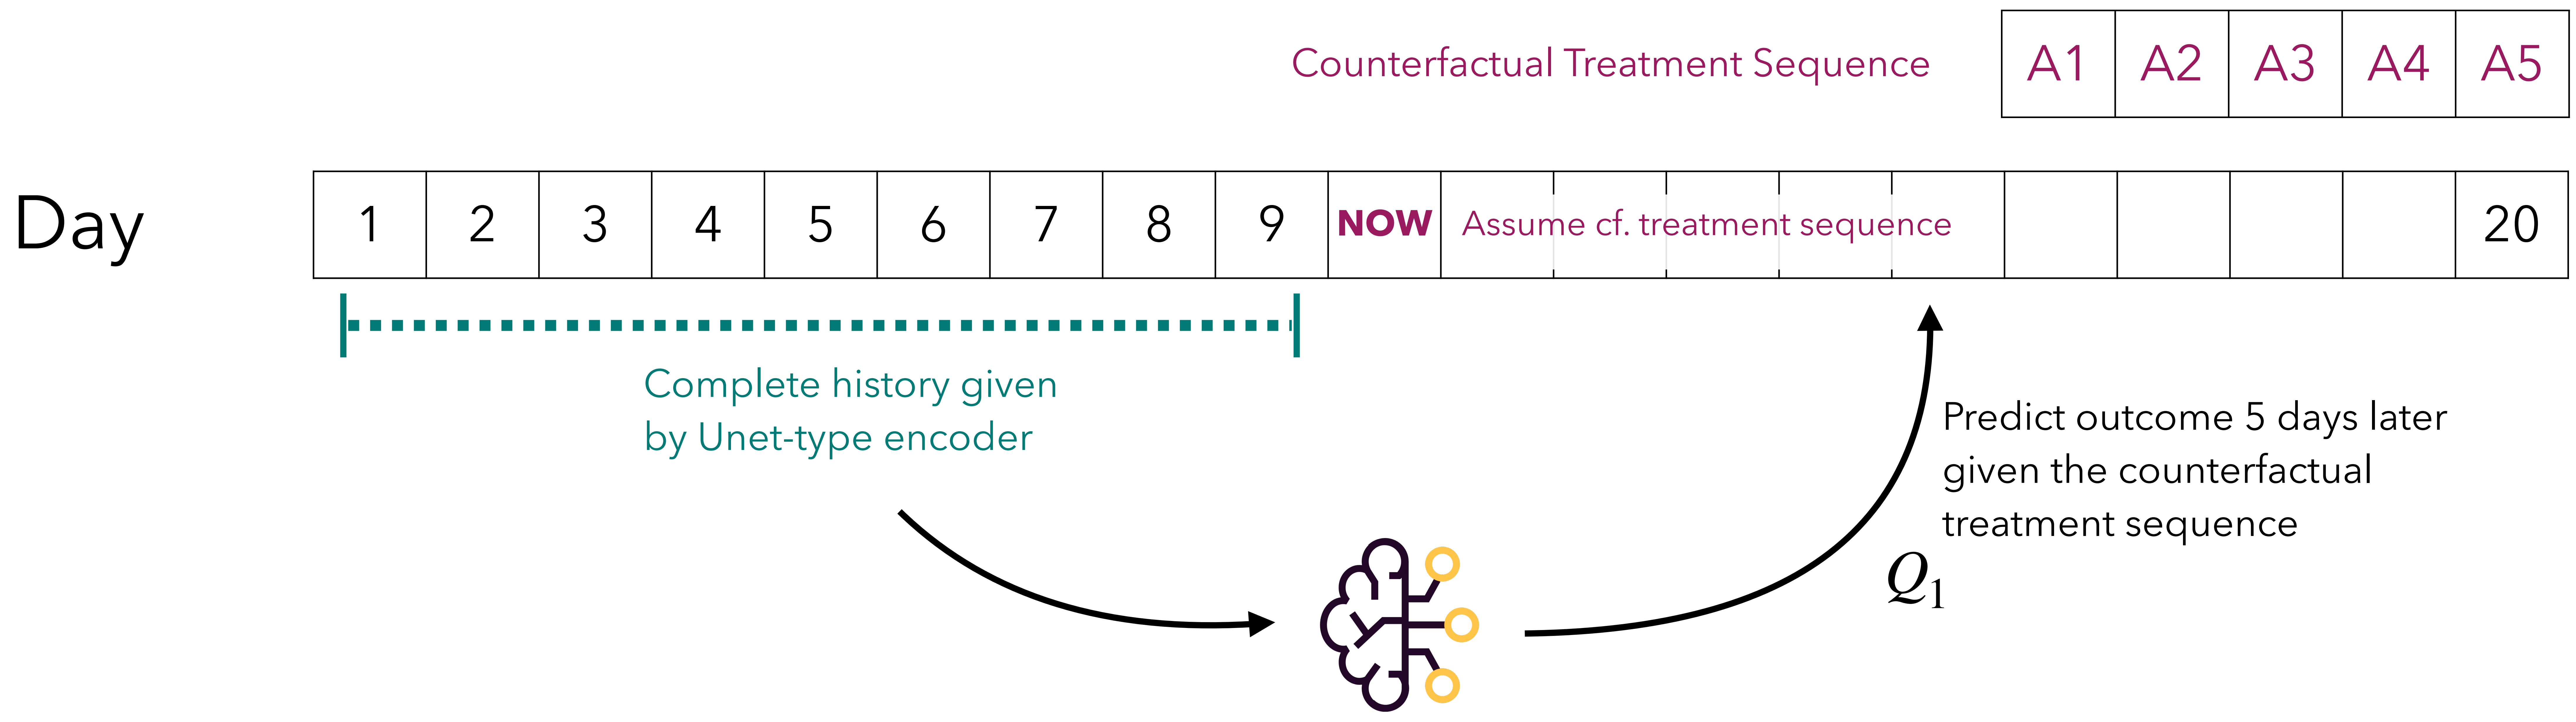
Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
NOW
Assume cf. treatment sequence
20
Complete history given
by Unet-type encoder
Predict outcome 5 days later
given the counterfactual
treatment sequence
Q_1

Splitting

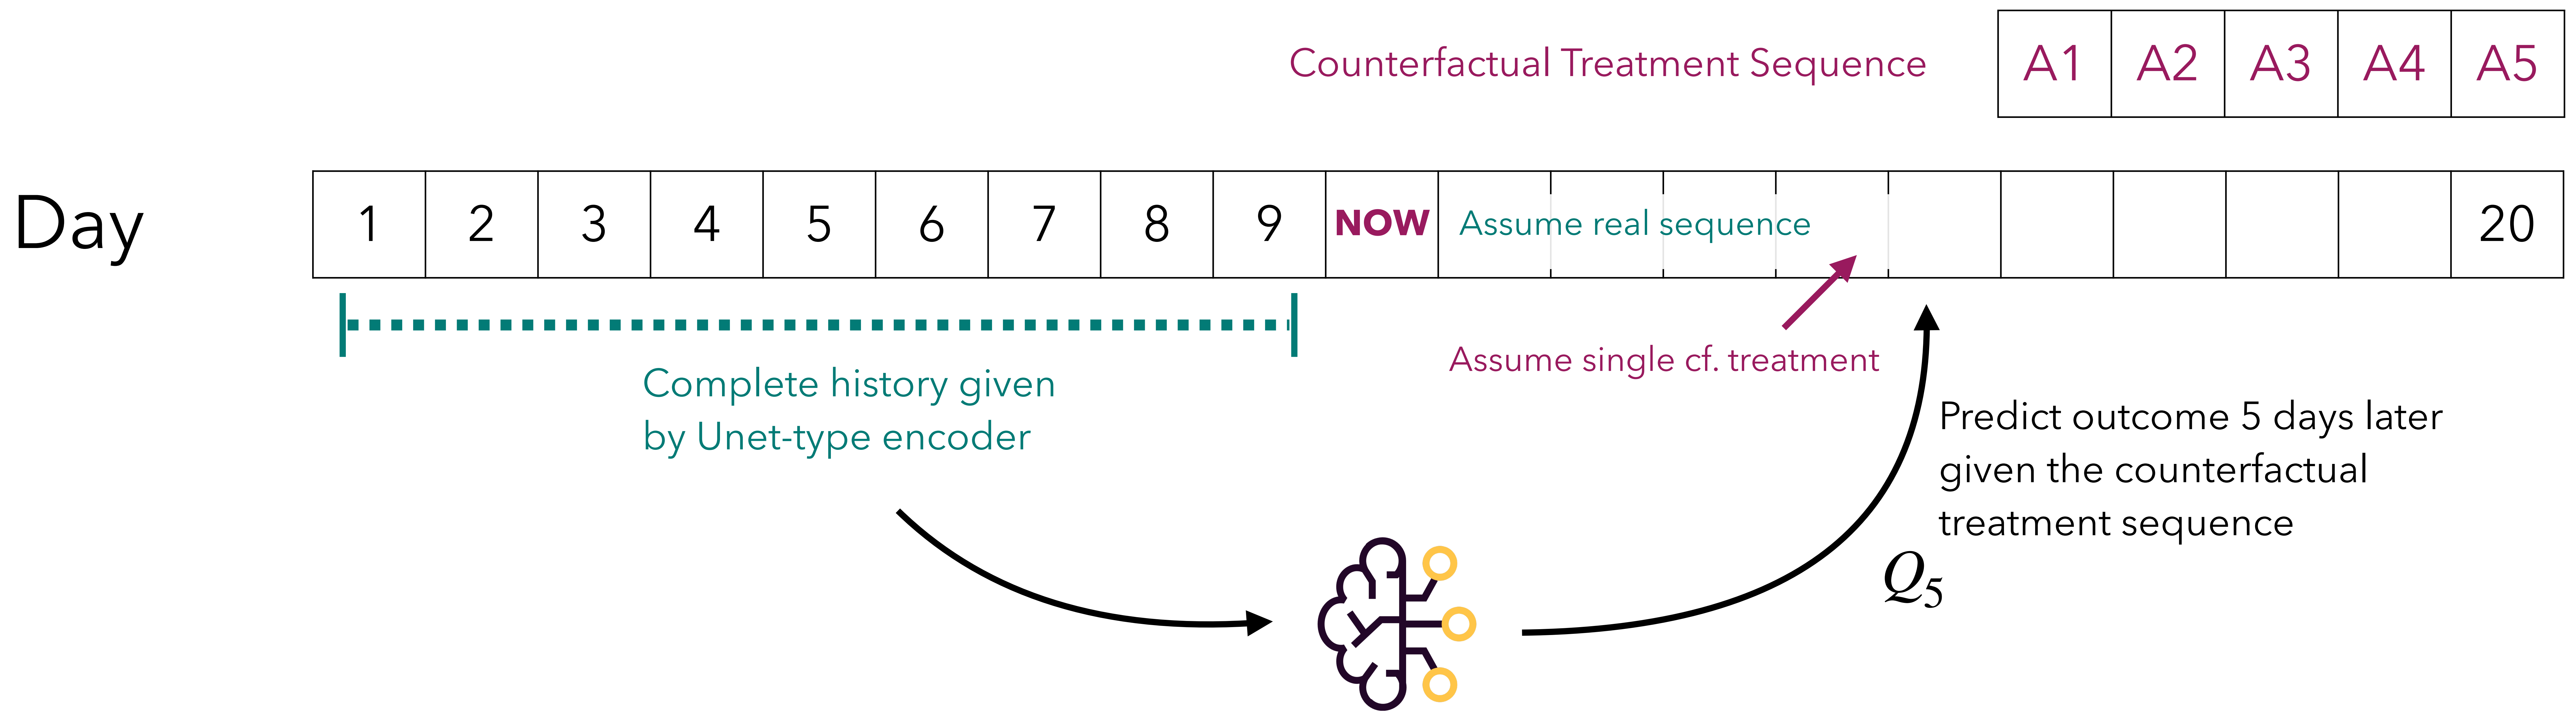
Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
NOW
Assume real sequence
20
Complete history given
by Unet-type encoder
Assume single cf. treatment
Predict outcome 5 days later
given the counterfactual
treatment sequence
Q_5

Splitting

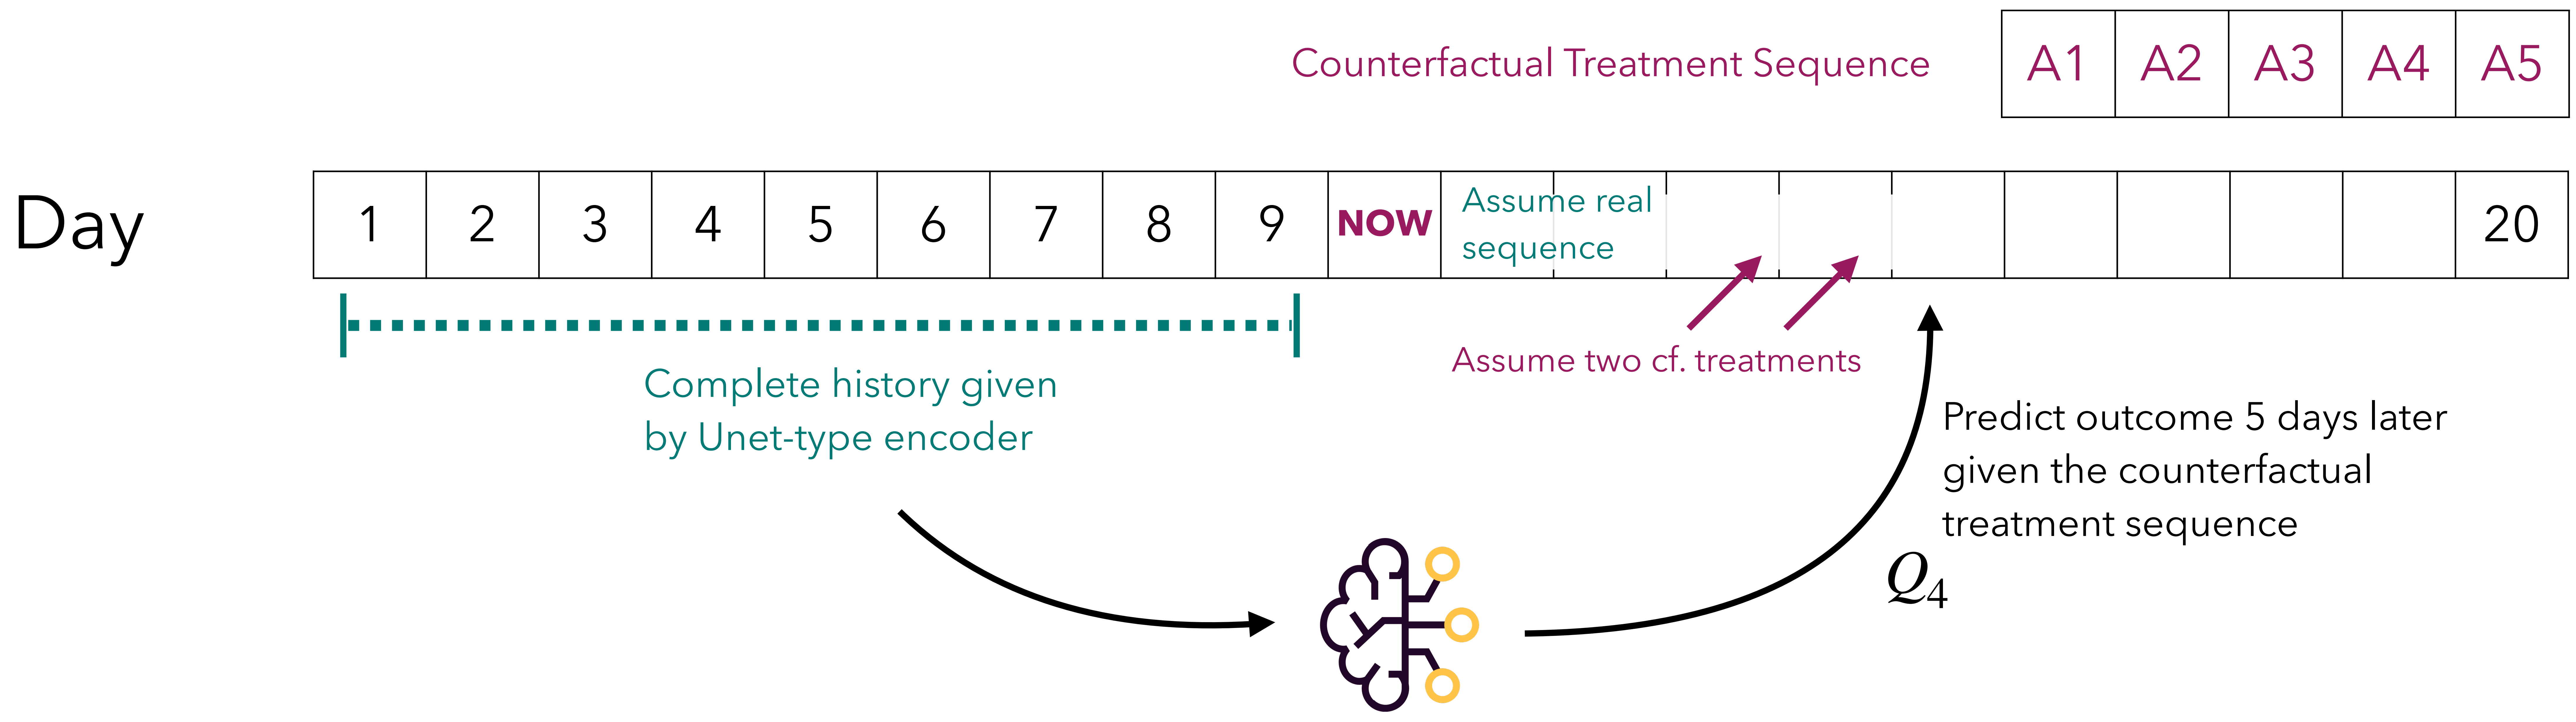
Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
NOW
Assume real sequence
20
Complete history given by Unet-type encoder
Assume two cf. treatments
Predict outcome 5 days later given the counterfactual treatment sequence
Q_4

Splitting

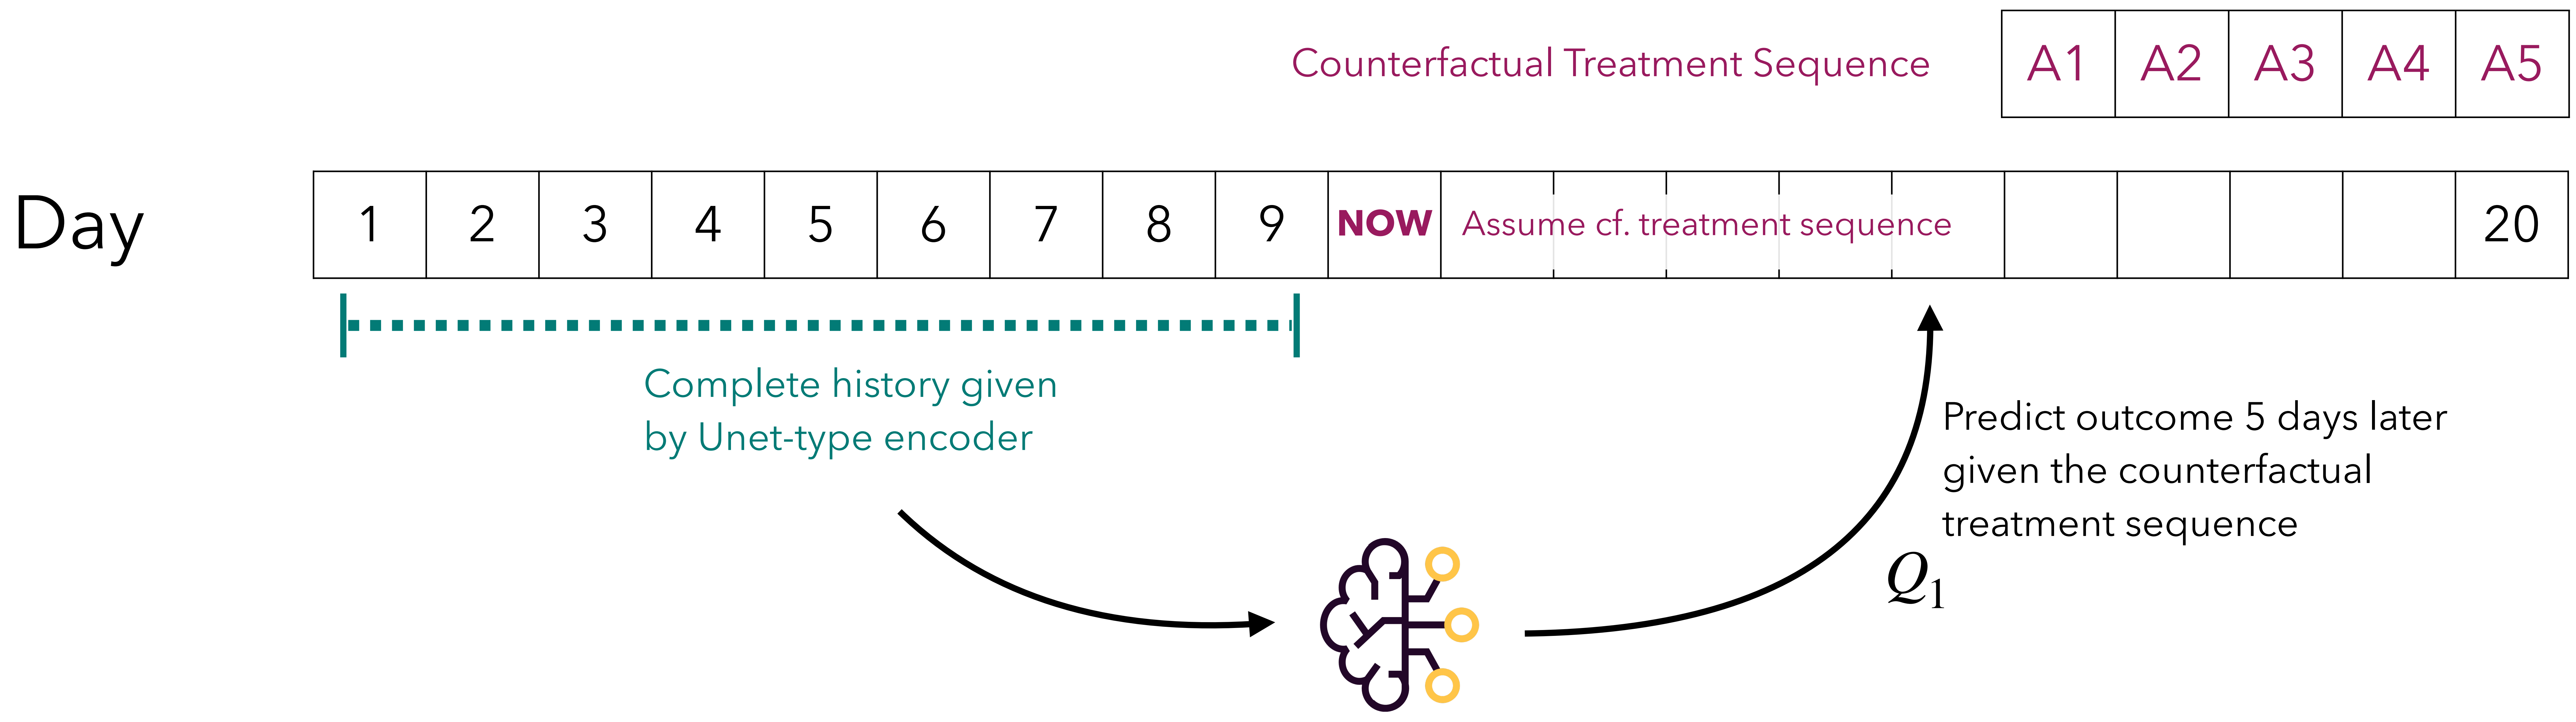
Counterfactual Treatment Sequence
A1
A2
A3
A4
A5
Day
1
2
3
4
5
6
7
8
9
NOW
Assume cf. treatment sequence
20
Complete history given
by Unet-type encoder
Predict outcome 5 days later
given the counterfactual
treatment sequence
Q_1

Simulation

Air Quality Policies
(Time-Varying Confounder)
(aka Covariate with Feedback)

Wildfire Smoke
(Treatment)

Hospital Admissions
(Outcome)

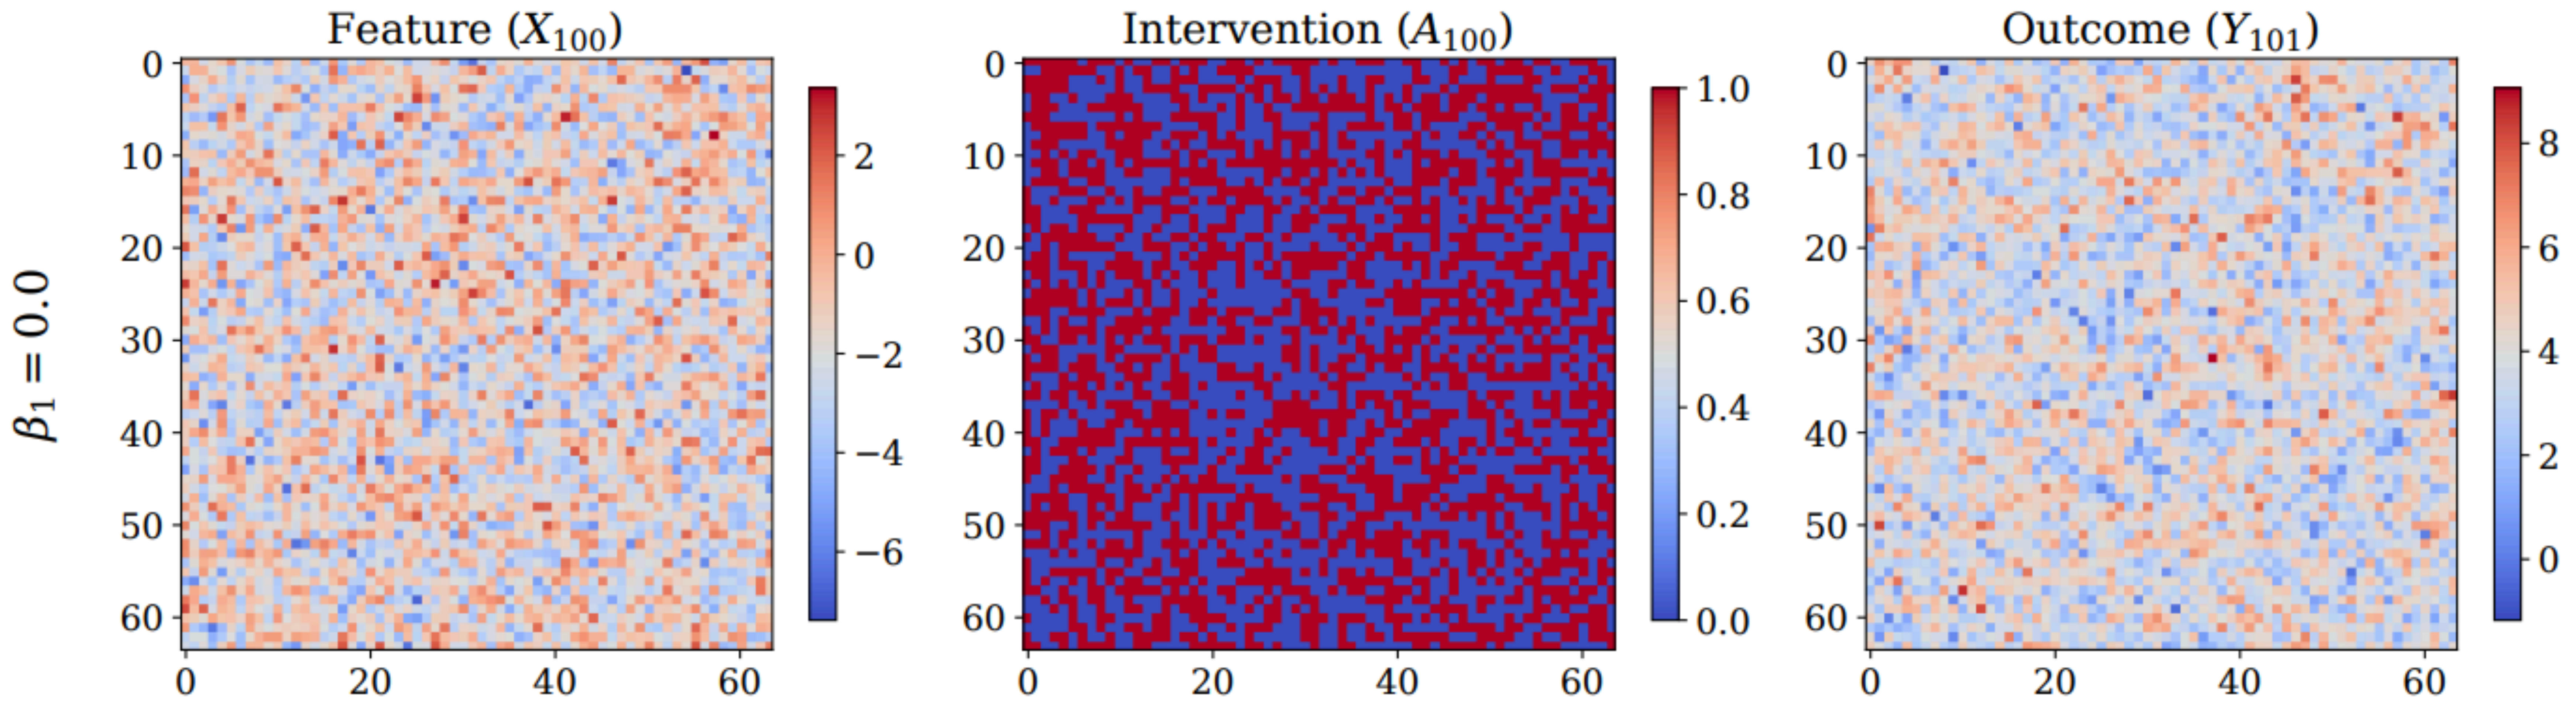

Time-varying confounders grow => GST-UNET becomes better than the baselines.

Real-Wold Case Study

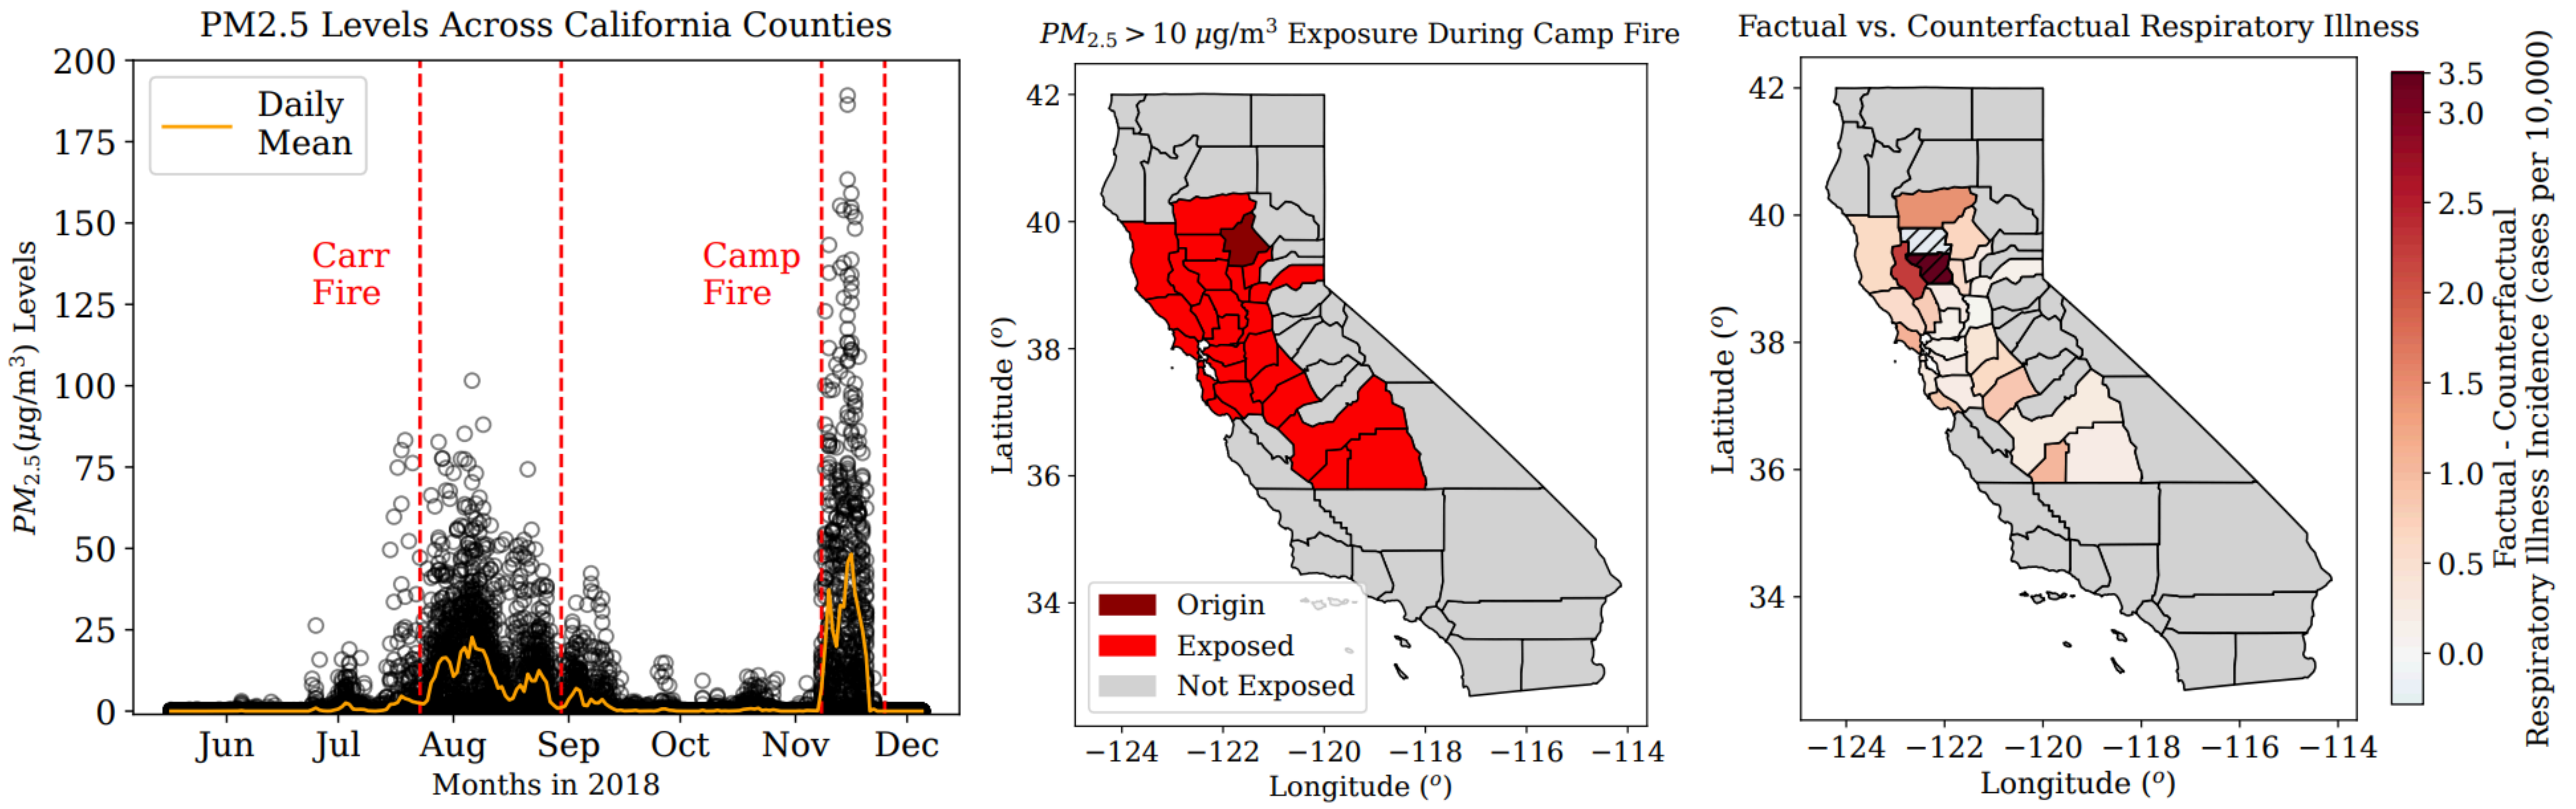

Figure 3. **(Left)** Daily $PM_{2.5}$ levels across California from May to December 2018, with red lines marking major wildfires. **(Center)** Counties exposed to average $PM_{2.5} > 10\,\mu g/m^3$ during the Camp Fire (red), origin county in dark red. **(Right)** Factual minus CAPO-predicted daily respiratory admissions during peak Camp Fire. Hashed areas indicate small-population counties (< 30,000).

- Quantify extra respiratory hospitalizations caused by the 2018 Camp Fire smoke across California.
- Spatial grid: 10 km x 10 km.
- Counterfactual treatment: Set every cell to “no-smoke” for 8 – 17 Nov 2018 (10 days ahead).
- Outcome: model attributes ≈ 4 650 excess admissions to the fire-driven pollution.

Summary

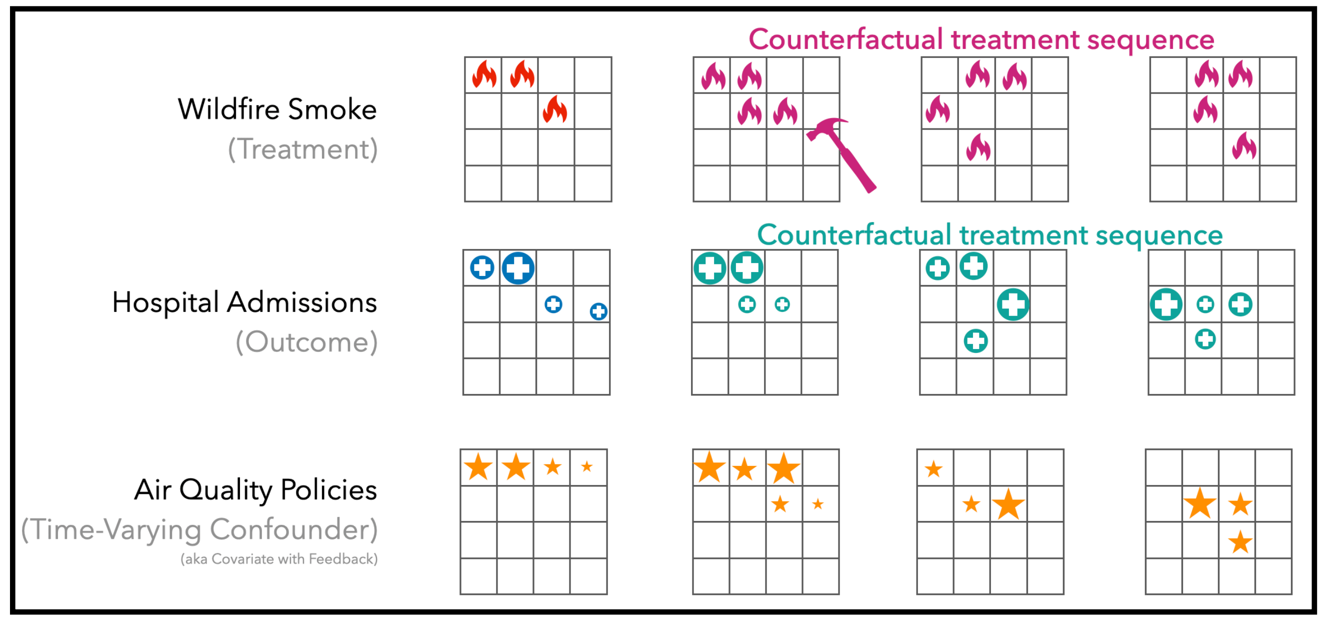

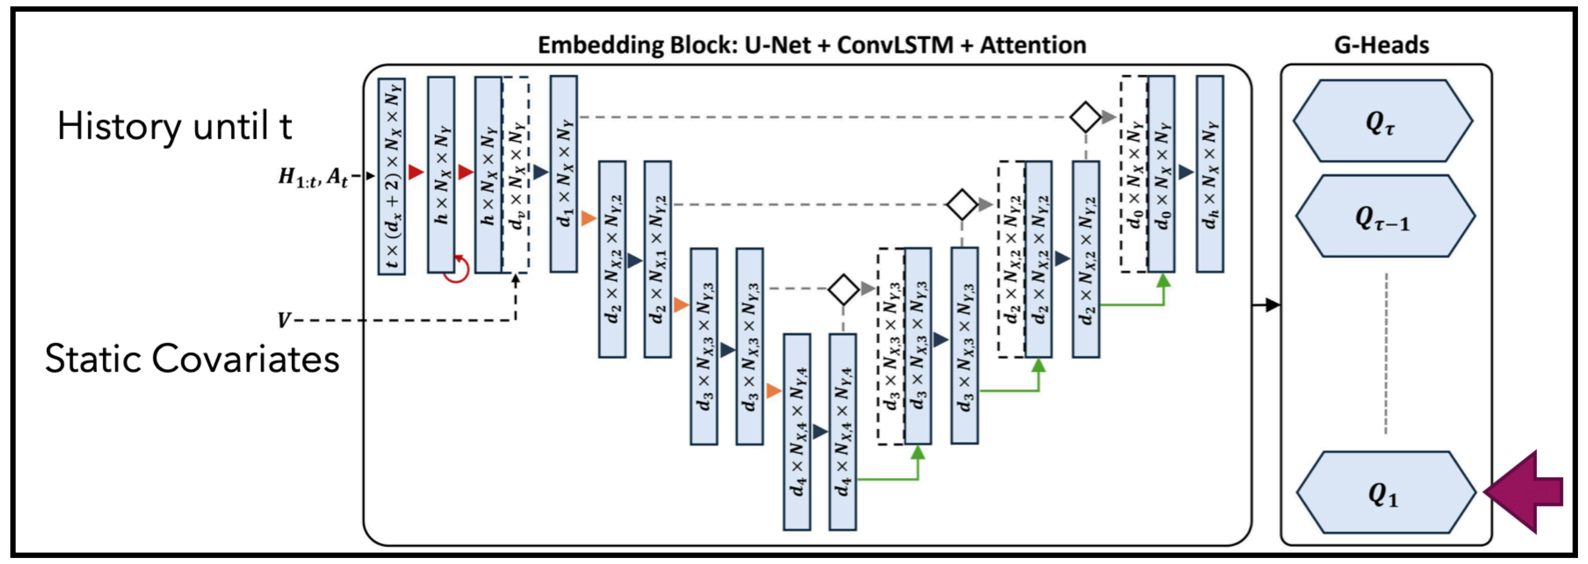

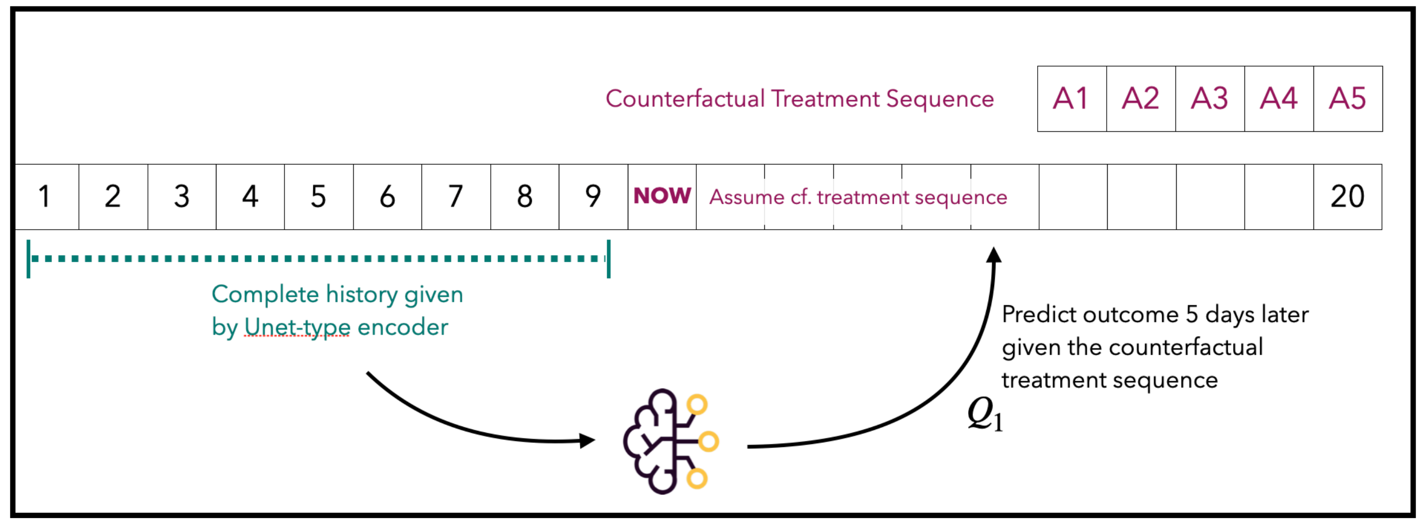

- How realistic are the assumptions?
- NN-architecture + training unprincipled?
- Is there a conceptual difference than just learning the transition operator?
- Cab you sample from the counterfactual distribution?
- Effects of grid size not discussed.