Estimating heterogeneous treatment effects for spatio-temporal causal inference

Lingxiao Zhou
with Kosuke Imai, Jason Lyall, Georgia Papadogeorgou

Department of Statistics,
University of Florida

April 16, 2025

Motivation

- Causality in spatio-temporal settings
  - the causal effect of an intervention on the assignment of treatment
- Challenges
  - Point pattern treatment and outcome
  - Spatial spillover
  - temporal carryover effect
- Causal framework for definition, identification and estimation of average treatment effects (**P** *et al.*, 2022)
- Questions remain :
  - How do different communities respond to the intervention ?
    $\rightsquigarrow$ (effect heterogeneity)
- Defining, identifying and estimating conditional treatment effects in the spatio-temporal setting

Motivation

- Causality in spatio-temporal settings
  - the causal effect of an intervention on the assignment of treatment
- Challenges
  - Point pattern treatment and outcome
  - Spatial spillover
  - temporal carryover effect
- Causal framework for definition, identification and estimation of average treatment effects (**P** *et al.*, 2022)
- Questions remain :
  - How do different communities respond to the intervention ?

    $\rightsquigarrow$ (effect heterogeneity)
- Defining, identifying and estimating conditional treatment effects in the spatio-temporal setting

Motivation

- Causality in spatio-temporal settings
  - the causal effect of an intervention on the assignment of treatment
- Challenges
  - Point pattern treatment and outcome
  - Spatial spillover
  - temporal carryover effect
- Causal framework for definition, identification and estimation of average treatment effects (**P** *et al.*, 2022)
- Questions remain :
  - How do different communities respond to the intervention ?
    $\rightsquigarrow$ (effect heterogeneity)
- Defining, identifying and estimating conditional treatment effects in the spatio-temporal setting

Our motivating setting

- Data in Iraq for 2007-08
- "Treatment" (or exposure) : Airstrikes
  (date, location, weapons type, and aircraft used)
- Outcome : Insurgent attacks
  (exact time, location, attack type)
- More airstrikes lead to more insurgent violence (**P** *et al.*, 2022)
- Do communities respond different to the increase in airstrikes based on prior humanitarian aid ?
- Potential effect modifier : US Aid Spending
  (district-level aid spending, 104 districts, during month prior to airstrikes)
- Hearts-and-minds theory :
  - How prior aid modifies the effect of airstrikes on insurgent violence

Our motivating setting

- Data in Iraq for 2007-08
- "Treatment" (or exposure) : Airstrikes
  (date, location, weapons type, and aircraft used)
- Outcome : Insurgent attacks
  (exact time, location, attack type)
- More airstrikes lead to more insurgent violence (**P** *et al.*, 2022)
- Do communities respond different to the increase in airstrikes based on prior humanitarian aid ?
- Potential effect modifier : US Aid Spending
  (district-level aid spending, 104 districts, during month prior to airstrikes)
- Hearts-and-minds theory :
  - How prior aid modifies the effect of airstrikes on insurgent violence

Our motivating setting

- Data in Iraq for 2007-08
- "Treatment" (or exposure) : Airstrikes
  (date, location, weapons type, and aircraft used)
- Outcome : Insurgent attacks
  (exact time, location, attack type)
- More airstrikes lead to more insurgent violence (**P** *et al.*, 2022)
- Do communities respond different to the increase in airstrikes based on prior humanitarian aid ?
- Potential effect modifier : US Aid Spending
  (district-level aid spending, 104 districts, during month prior to airstrikes)
- Hearts-and-minds theory :
  - How prior aid modifies the effect of airstrikes on insurgent violence

Our data

Airstrike

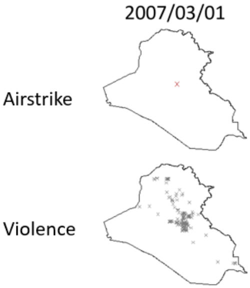
2007/03/01
Airstrike
Violence

Our data

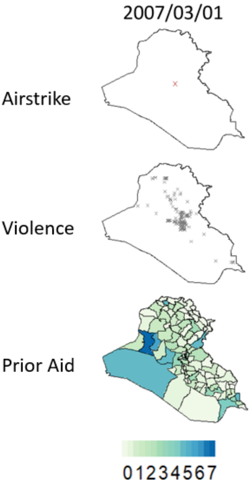

Our data

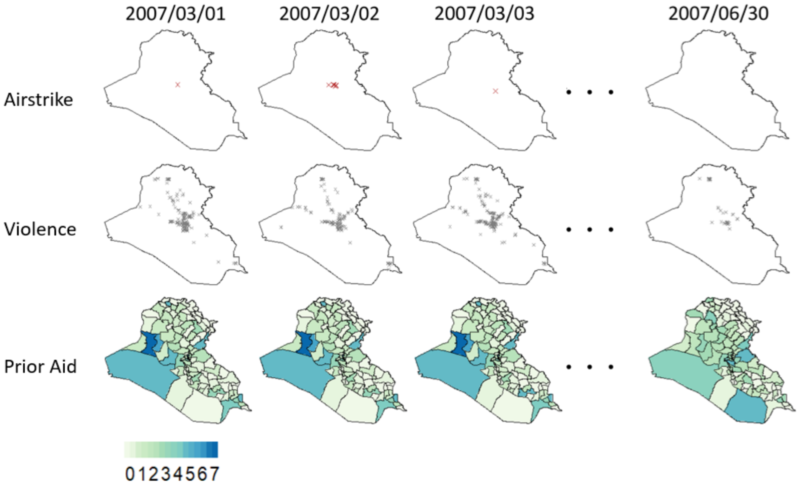

2007/03/01
2007/03/02
2007/03/03
2007/06/30
Airstrike
Violence
Prior Aid
. . .
01234567

Challenges

- Spatial spillover
- Temporal carryover effect
- Point pattern treatment and outcome
- Defining, identifying and estimating conditional treatment effects in this spatio-temporal setting

Challenges

- Spatial spillover
- Temporal carryover effect
- Point pattern treatment and outcome
- Defining, identifying and estimating conditional treatment effects in this spatio-temporal setting

Some notation

- Ω : geography of interest
- Treatment and outcome events are **spatial point patterns** measured at discrete time periods $t \in \mathcal{T} = \{1, 2, \dots, T\}$

- W_t : treatment point pattern at time t
  - realization w_t
  - $\overline{\boldsymbol{W}}_t = (W_1, W_2, \dots, W_t)$ treatment pattern until time t

- $Y_t(\overline{\boldsymbol{w}}_t)$: potential outcome at time t under treatment path $\overline{\boldsymbol{w}}_t$
  - No assumptions on the interference structure

- $\boldsymbol{X}_t$: time-varying variables occurring prior to W_t
- $\boldsymbol{R}_t \subset \boldsymbol{X}_t$: the potential effect modifier

Some notation

- Ω : geography of interest
- Treatment and outcome events are **spatial point patterns** measured at discrete time periods $t \in \mathcal{T} = \{1, 2, \ldots, T\}$

- W_t : treatment point pattern at time t
  - realization w_t
  - $\overline{\boldsymbol{W}}_t = (W_1, W_2, \ldots, W_t)$ treatment pattern until time t

- $Y_t(\overline{\boldsymbol{w}}_t)$: potential outcome at time t under treatment path $\overline{\boldsymbol{w}}_t$
  - No assumptions on the interference structure

- $\boldsymbol{X}_t$: time-varying variables occurring prior to W_t
- $\boldsymbol{R}_t \subset \boldsymbol{X}_t$: the potential effect modifier

Some notation

- Ω : geography of interest
- Treatment and outcome events are **spatial point patterns** measured at discrete time periods $t \in \mathcal{T} = \{1, 2, \ldots, T\}$

- W_t : treatment point pattern at time t
  - realization w_t
  - $\overline{\boldsymbol{W}}_t = (W_1, W_2, \ldots, W_t)$ treatment pattern until time t

- $Y_t(\overline{\boldsymbol{w}}_t)$: potential outcome at time t under treatment path $\overline{\boldsymbol{w}}_t$
  - No assumptions on the interference structure

- $\boldsymbol{X}_t$: time-varying variables occurring prior to W_t
- $\boldsymbol{R}_t \subset \boldsymbol{X}_t$: the potential effect modifier

Some notation

- Ω : geography of interest
- Treatment and outcome events are **spatial point patterns** measured at discrete time periods $t \in \mathcal{T} = \{1, 2, \ldots, T\}$

- W_t : treatment point pattern at time t
  - realization w_t
  - $\overline{\boldsymbol{W}}_t = (W_1, W_2, \ldots, W_t)$ treatment pattern until time t

- $Y_t(\overline{\boldsymbol{w}}_t)$: potential outcome at time t under treatment path $\overline{\boldsymbol{w}}_t$
  - No assumptions on the interference structure

- $\boldsymbol{X}_t$: time-varying variables occurring prior to W_t
- $\boldsymbol{R}_t \subset \boldsymbol{X}_t$: the potential effect modifier

Estimands under stochastic interventions

- For point pattern treatments, we use stochastic interventions to represent useful treatment assignments
- F_h represents a hypothetical treatment assignment strategy (distribution over treatment point patterns)
- $h = c\phi_0$
  - What would happen if treatment intensity increased?
- Estimands represent contrasts of outcomes under different treatment assignments for locations with a specific effect modifier value

Estimand

- Expected number of outcome events in pixel S_i at time t under intervention F_h

$$N_{it}(F_h) \;=\; \int_{\mathcal{W}} N_{S_i}(Y_t(\overline{\boldsymbol{W}}_{t-1}, w_t))\, \mathrm{d}F_h(w_t).$$

- Treatment effect of $F_{h'}$ versus $F_{h''}$ for time period t and pixel S_i :

$$\tau_{it}(F_{h'}, F_{h''}) \;=\; N_{it}(F_{h''}) - N_{it}(F_{h'})$$

- The conditional average treatment effect (CATE) when the moderator takes a specific value $\boldsymbol{r} \in \mathcal{R}$ at time t :

$$\tau_{t,h',h''}(\boldsymbol{r}) \;=\; \frac{1}{\sum_{i=1}^{p} I(\boldsymbol{R}_{it} = \boldsymbol{r})} \sum_{i=1}^{p} \tau_{it}(F_{h'}, F_{h''}) I(\boldsymbol{R}_{it} = \boldsymbol{r})$$

Estimand

- Projection estimand : $\tau^{\text{Proj.}}_{t,h',h''}(r;\boldsymbol{\beta}^*_t)$ with

$$\boldsymbol{\beta}^*_t = \arg\min_{\boldsymbol{\beta}_t} \sum_{i=1}^{p} (\tau_{t,h',h''}(\boldsymbol{R}_{it}) - \tau^{\text{Proj.}}_{t,h',h''}(\boldsymbol{R}_{it};\boldsymbol{\beta}_t))^2,$$

- Overall CATE :

$$\tau^{\text{Proj.}}_{h',h''}(\boldsymbol{r};\boldsymbol{\beta}^*_M,\ldots,\boldsymbol{\beta}^*_T) = \frac{1}{T}\sum_{t=1}^{T} \tau^{\text{Proj.}}_{t,h'h''}(\boldsymbol{r};\boldsymbol{\beta}^*_t).$$

- Extend for interventions over multiple time periods

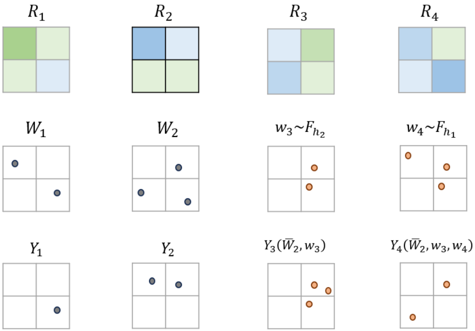

Figure – An illustration of the moderator (top row) over four pixels and four time periods, along with the treatment (middle row) and outcome (bottom row) events with a stochastic intervention $F_{\boldsymbol{h}} = F_{h_1} \times F_{h_2}$ over the last two time periods.

Assumption (Unconfoundedness)

$f(W_t \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\mathcal{Y}}_T, \overline{\mathcal{X}}_T) = f(W_t \mid \overline{H}_{t-1})$

Assumption (Bounded relative overlap)

There exists $\delta_W > 0$ *such that* $e_t(w) > \delta_W f_h(w)$ *for all* $w \in \mathcal{W}$ *where* $e_t(w) = f(W_t = w \mid \overline{H}_{t-1})$

where $\overline{H}_t$ is the observed history up to and including time t

Estimation procedure

Step 1 : Create the pseudo outcomes

- Pseudo outcome (IPW weights) for intervention $F_{\boldsymbol{h}}$:

$$\widetilde{Y}_{it}^{I}(F_{\boldsymbol{h}};\hat{\gamma}) = \prod_{j=t-M+1}^{t} \frac{f_{\boldsymbol{h}}(W_j)}{e_j(W_j;\hat{\gamma})} \; N_{S_i}(Y_t)$$

$\rightsquigarrow$ $\widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}};\hat{\gamma})$: use stabilized IPW weights
$\rightsquigarrow$ for $F_{\boldsymbol{h}'}$ versus $F_{\boldsymbol{h}''}$: $\widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}''};\hat{\gamma}) - \widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}'};\hat{\gamma})$

Step 2 : Regress pseudo outcomes on the effect modifier

- Fit the regression model for each time period and then average the results over all time periods.
- Consider linear models with spline basis functions of the effect modifier.

Estimation procedure

Step 1 : Create the pseudo outcomes

- Pseudo outcome (IPW weights) for intervention $F_{\boldsymbol{h}}$:

$$\widetilde{Y}_{it}^{I}(F_{\boldsymbol{h}};\hat{\gamma}) \;=\; \prod_{j=t-M+1}^{t} \frac{f_{\boldsymbol{h}}(W_j)}{e_j(W_j;\hat{\gamma})} \;\; N_{S_i}(Y_t)$$

$\rightsquigarrow$ $\widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}};\hat{\gamma})$: use stabilized IPW weights
$\rightsquigarrow$ for $F_{\boldsymbol{h}'}$ versus $F_{\boldsymbol{h}''}$: $\widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}''};\hat{\gamma}) - \widetilde{Y}_{it}^{H}(F_{\boldsymbol{h}'};\hat{\gamma})$

Step 2 : Regress pseudo outcomes on the effect modifier

- Fit the regression model for each time period and then average the results over all time periods.
- Consider linear models with spline basis functions of the effect modifier.

Asymptotic distributions of the estimator

Theorem (Asymptotic Normality of the Hájek Estimator Using the Estimated Propensity Score)

Suppose that unconfoundedness and overlap assumptions hold, along with regularity conditions. Then as $T \to \infty$, we have that

$$\frac{1}{\sqrt{T-M+1}} \sum_{t=M}^{T} (\hat{\boldsymbol{\beta}}_t^H - \boldsymbol{\beta}_t^*) \xrightarrow{d} N\big(\mathbf{0}, V\big).$$

- Combining theory on inverse weighting, estimated propensity scores, martingale difference series
- Cannot consistently estimate the asymptotic variance due to the fact that we only observe a single time series
- Asymptotic efficiency under the estimated propensity score

Theorem (Asymptotic Normality of the Hájek Estimator Using the Estimated Propensity Score)

Suppose that unconfoundedness and overlap assumptions hold, along with regularity conditions. Then as $T \to \infty$, we have that

$$\frac{1}{\sqrt{T-M+1}} \sum_{t=M}^{T} (\hat{\boldsymbol{\beta}}_t^H - \boldsymbol{\beta}_t^*) \xrightarrow{d} N\big(\mathbf{0}, V\big).$$

- Combining theory on inverse weighting, estimated propensity scores, martingale difference series
- Cannot consistently estimate the asymptotic variance due to the fact that we only observe a single time series
- Asymptotic efficiency under the estimated propensity score

Asymptotic distributions of the estimator

Theorem (Asymptotic Normality of the Hájek Estimator Using the Estimated Propensity Score)

Suppose that unconfoundedness and overlap assumptions hold, along with regularity conditions. Then as $T \to \infty$*, we have that*

$$\frac{1}{\sqrt{T-M+1}} \sum_{t=M}^{T} (\hat{\boldsymbol{\beta}}_t^H - \boldsymbol{\beta}_t^*) \xrightarrow{d} N\big(\mathbf{0}, V\big).$$

- Combining theory on inverse weighting, estimated propensity scores, martingale difference series
- Cannot consistently estimate the asymptotic variance due to the fact that we only observe a single time series
- Asymptotic efficiency under the estimated propensity score

Empirical Analysis

- Treatment effect heterogeneity for the effect of airstrikes on insurgent violence in Iraq, during the time period February 23, 2007 to July 05, 2008
- Outcome events : Improvised Explosive Device (IED) and Small Arms Fire (SAF) insurgent attacks
- Effect modifier : Aid per capita during the previous month
- Conditional average treatment effect of increasing the expected number of airstrikes from one per day to six per day
- $F_{\boldsymbol{h}'}$ v.s. $F_{\boldsymbol{h}''}$, where $h' = 1\phi_0$ and $h'' = 6\phi_0$.
  $\rightsquigarrow$ ϕ_0 is estimated based on historic data
- Working model :

$$\tau_{t,\boldsymbol{h'},\boldsymbol{h''}}(r;\boldsymbol{\beta}_t) = \beta_{t,0} + \sum_{l=1}^{4} \beta_{t,l} z_l(r) + \beta_{t,5} I\{r = 0\},$$

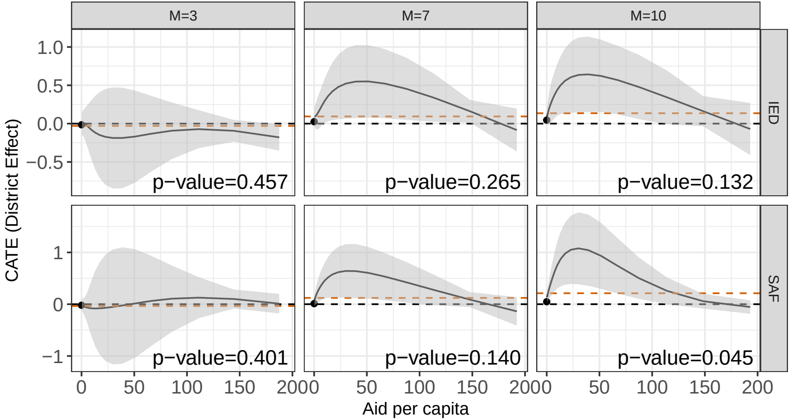

Figure – The estimated CATE for given value of aid per capita in the previous month with the shaded region indicating the 95% confidence intervals.

Summary

- We propose a method for studying the treatment heterogeneity in the spatio-temporal setting
  - Point pattern treatment and outcome
  - Spatial or spatio-temporal moderator
- Two step estimation procedure
  - No assumption on temporal carryover and spatial spillover effects
  - Utilize stablized IPW weights
- Empirical study
  - Examining how prior humanitarian aid influences the impact of airstrikes on insurgent violence.
- Publicly available software & data, and corresponding software manuscript (Mukaigawara *et al.*, 2024b,a)

Thank You !

Questions ? Comments ?

References

Mukaigawara, M., Zhou, L., **P**, **G**., Lyall, J., and Imai, K. (2024a). geocausal : An r package for spatio-temporal causal inference. Tech. rep., Center for Open Science.

Mukaigawara, M., Zhou, L., **P**, **G**., Lyall, J., and Imai, K. (2024b). geocausal : Causal inference with spatio-temporal data .

P, **G**., Imai, K., Lyall, J., and Li, F. (2022). Causal inference with spatio-temporal data : estimating the effects of airstrikes on insurgent violence in iraq. *Journal of the Royal Statistical Society Series B : Statistical Methodology* **84**, 5, 1969–1999.

Zhou, L., Imai, K., Lyall, J., and Papadogeorgou, G. (2024). Estimating heterogeneous treatment effects for spatio-temporal causal inference : How economic assistance moderates the effects of airstrikes on insurgent violence. *arXiv preprint arXiv :2412.15128* .