Free Discontinuity Regression

With an Application to the Economic Effects of Internet Shutdowns

David Van Dijcke
University of Michigan\Ipsos Public Affairs

Econometric Society IF AI+ML Conference
Cornell University

joint with Florian Gunsilius (Emory)

Introduction

- Discontinuous changes in regression surfaces can reveal fundamental insights into underlying problem structure
- Especially when both **location** and **size** of discontinuity are unknown
  - **Location**: treatment assignment mechanisms, structural breaks, change points...
  - **Size**: causal effects of treatment (under extra assumptions), magnitude of regime shift, etc.
- **This paper: statistical method for estimating regression surface together with discontinuities in any dimension**
  - Statistical reformulation of celebrated Mumford-Shah functional from computer vision
  - No pointwise smoothness assumptions on discontinuity set
  - *Segmentation:* balances function estimation, denoising, and thresholding
- Examples:
  - Confidential assignment algorithms: loan/financial aid disbursement (Argyle et al., 2020; Carneiro et al., 2019), school admission (Brunner et al., 2021), customer segmentation for marketing (Hartmann et al., 2011), bandit algorithms (Li et. al., 2010)
  - Behavioral discontinuities: tax brackets ("fuzzy" thresholds) (Saez, 2010; Alvero and Xiao, 2023), racial segregation (Card et al., 2008), tipping points Hospital Stress Responses (Kuntz et al., 2015)
  - Complex Systems/Tipping Points: Climate Change (Scheffer et al., 2009), Financial Markets (Hansen, 2017)
  - Geographic Discontinuities: Broadcast Signal Penetration (Sonin and Wright, 2022; Kern and Hainmueller, 2009; Olken, 2009), Epidemic Spread (Ambrus et al., 2020)

Existing approaches

- **Change point detection**: Page (1954), Killick et al. (2012), Porter and Yu (2015), Donoho and Johnstone (1994), Harchaoui et al. (2008).
  - Focus on estimating discontinuity locations, less emphasis on estimating jump sizes.
- **Structural breaks**: Andrews (1993), Bai and Perron (2003), Delgado and Hidalgo (2000).
  - Time as a primary input variable — one-dimensional.
- **Extensions to multivariate domains** (Park, 2022; Herlands et al., 2018, 2019; Zhu et al., 2014; Madrid Padilla et al, 2022)
  - No statistical results, focus on discontinuity locations
  - Multivariate discontinuities in economics: Cheng (2023), Narita and Yata (2021), Abdulkadiroglu et al. (2022), Cattaneo and Titiunik (2022).
- **Regression surfaces with jumps:** Qiu and Yandell (1997), Korostelev and Tsybakov (1993), O'Sullivan and Qian (1994), Muller and Song (1994), Donoho (1999), Li and Ghosal (2017), Qiu (1998).
  - Multidimensional jump estimation under strong assumptions on discontinuity set.
- **Fused Lasso**: Tibshirani et al. (2005), Rinaldo (2009), Harchaoui and Lévy-Leduc (2010).
  - Piecewise constant functions on multivariate domains.

Our contribution

- **Statistical convergence** results for convex estimator
- **Jointly estimate** discontinuity locations and size with regression surface
- **Identification** results: identify true jump locations and sizes at limit
- For **general discontinuity** set in any dimension
- **Application** to internet shutdowns, implemented in PyTorch library

Preview of Results

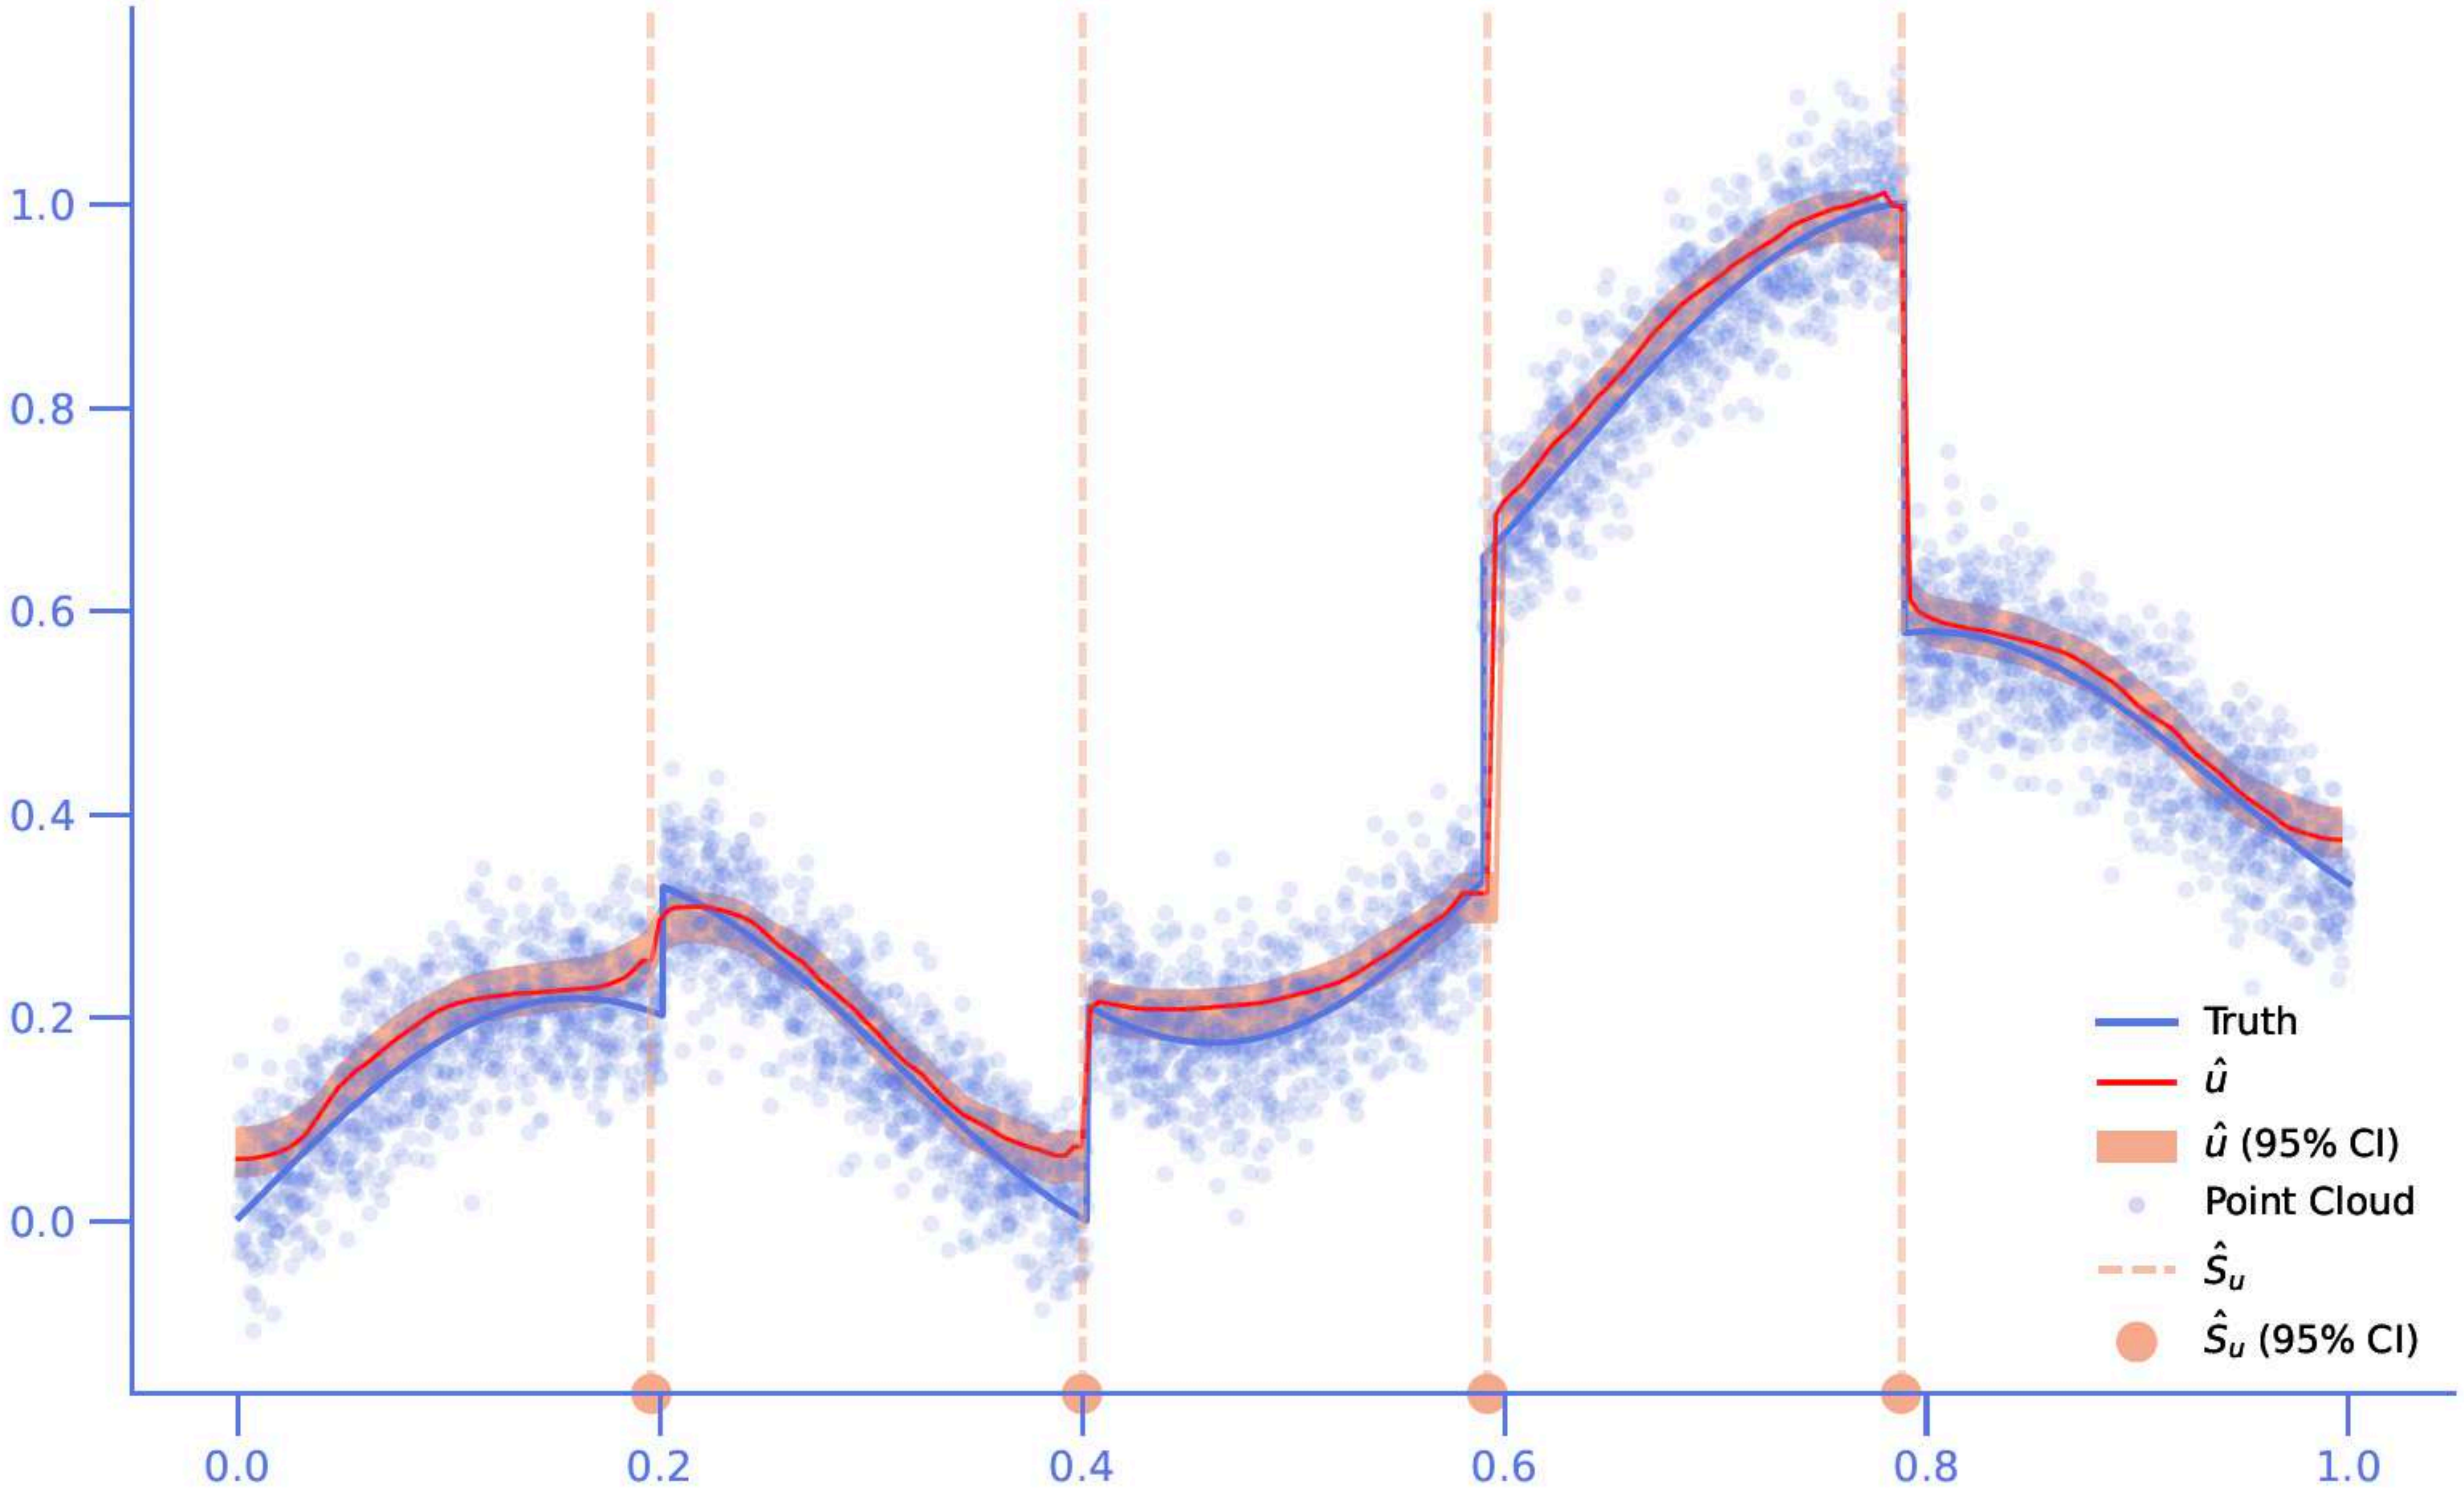

(a) 1D: Truth + Point Cloud + $\hat{u}$ + Boundary

Preview of Results

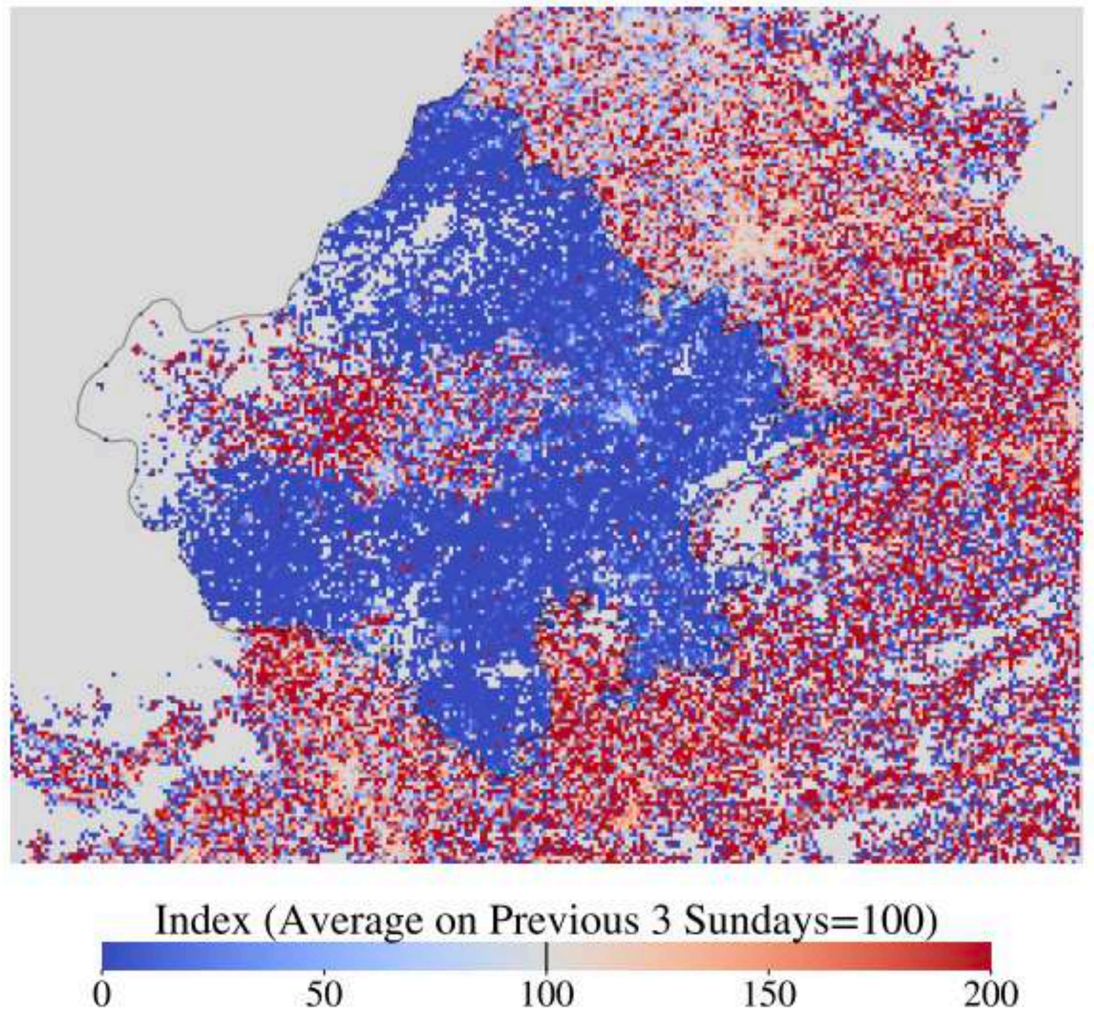

(a) Mobile: Raw Data

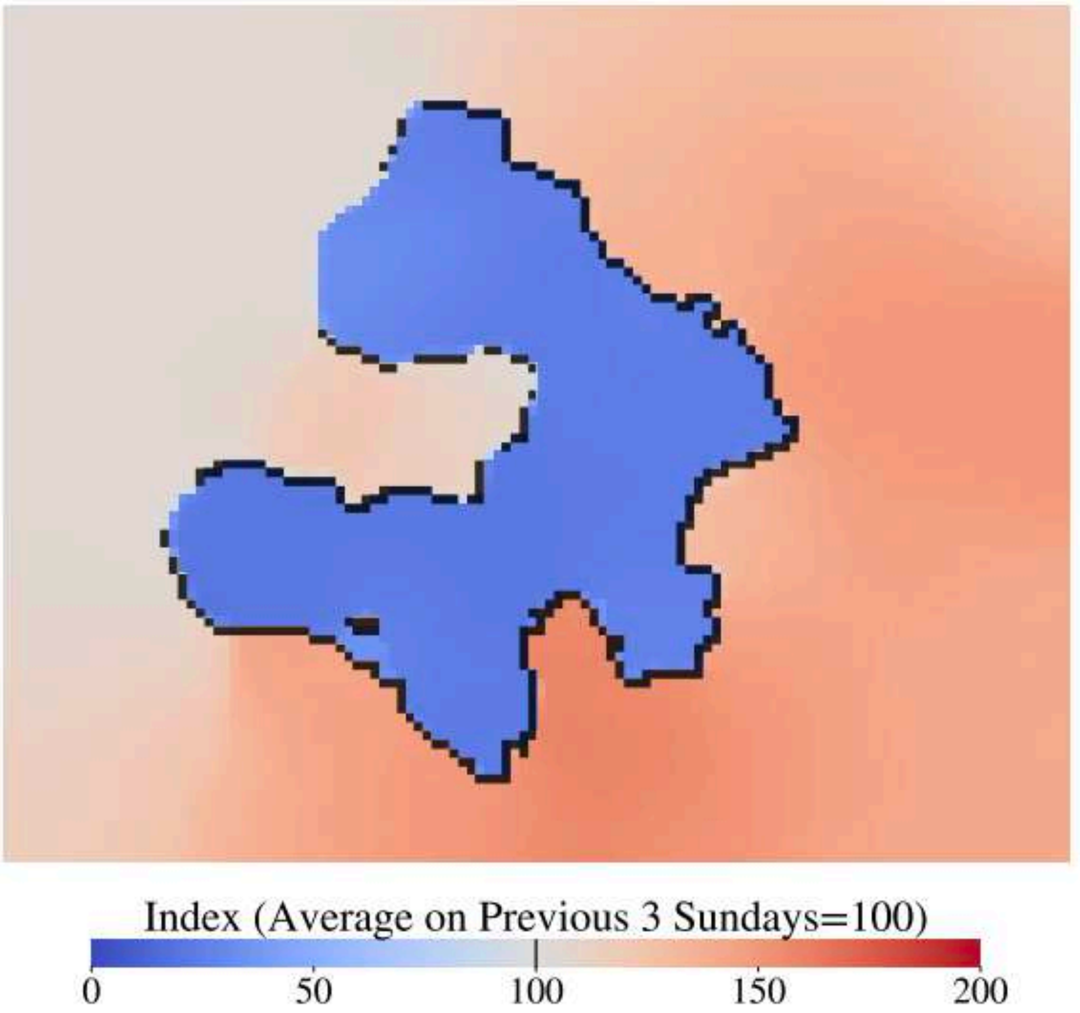

(b) Mobile: $\hat{u}$ & S_u

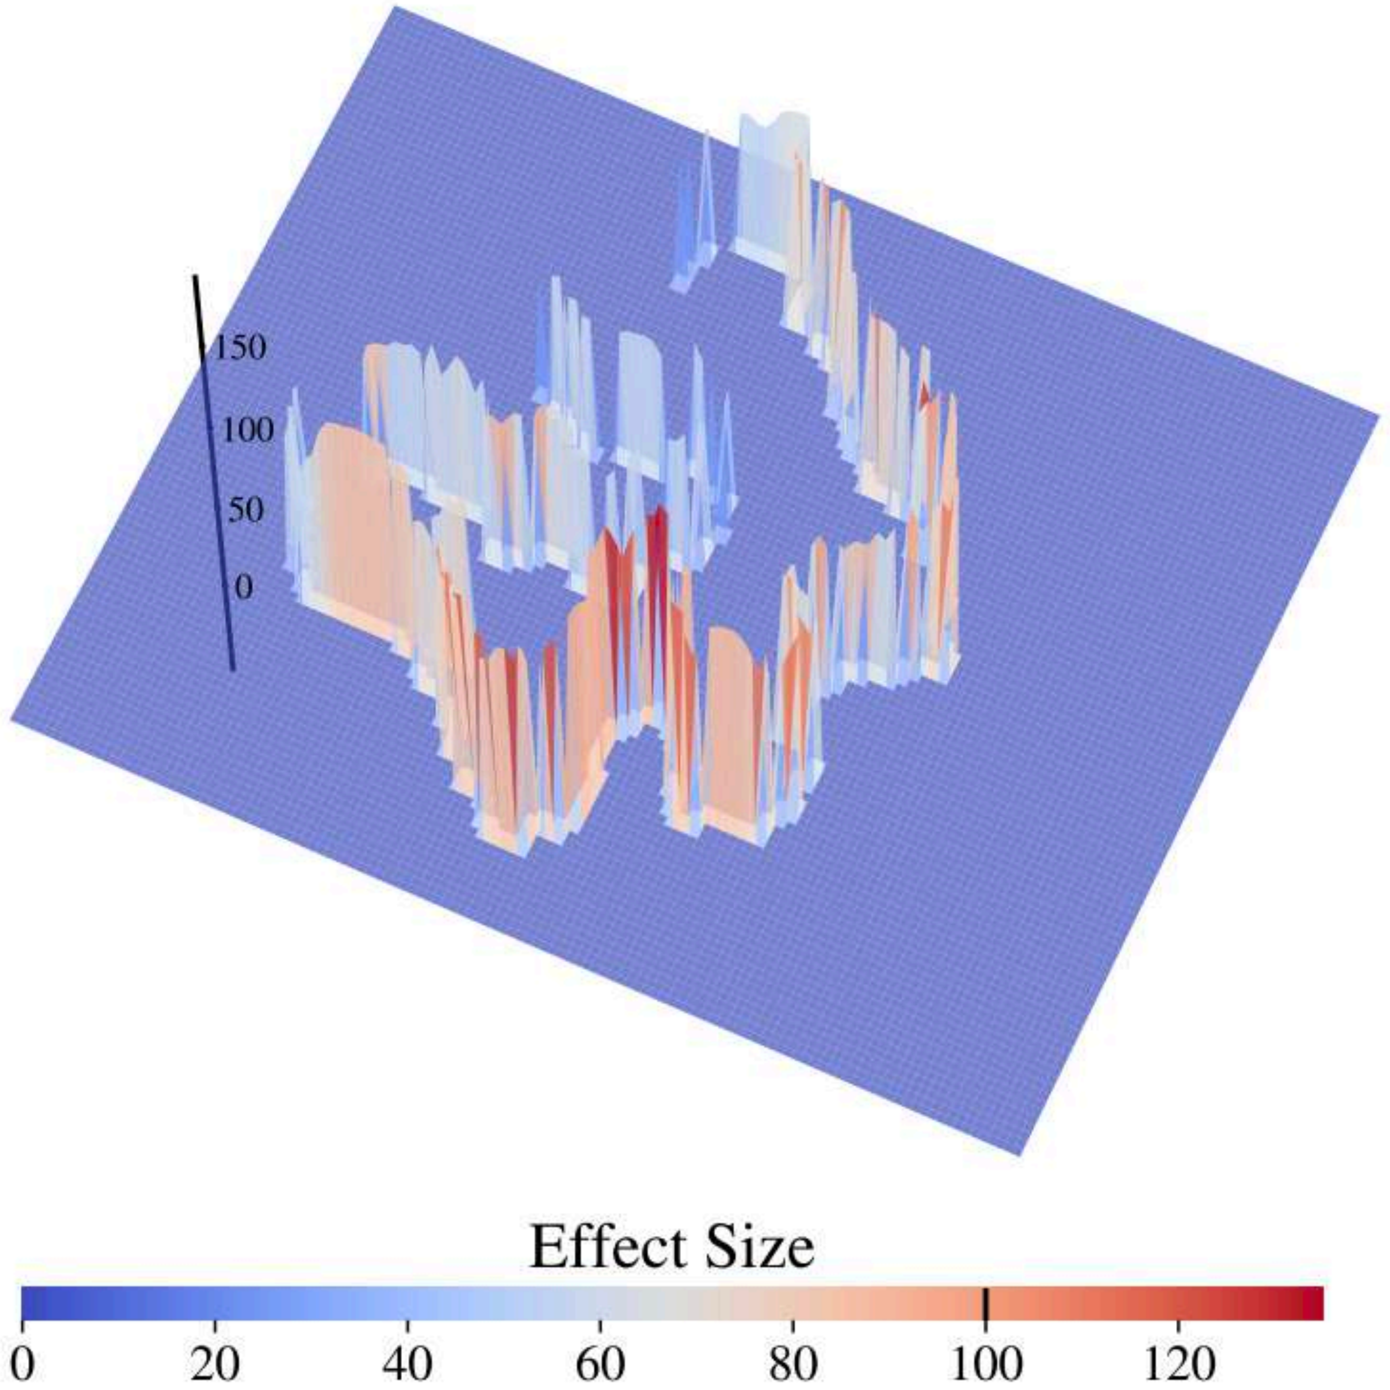

Regression framework

- $i = 1, \ldots, n$ units
- $X_i \in \mathbb{R}^d$ potentially multivariate regressor
- $Y_i \in \mathbb{R}$ outcome of interest (can be extended to $\mathbb{R}^d$)
- Regression model:

$$Y_i = f(X_i) + \varepsilon_i, \qquad \mathbb{E}[\varepsilon_i \mid X_i] = 0, \qquad \varepsilon_i \sim \text{i.i.d.}$$

- Goal: estimate regression surface $f(x)$ including location and jump sizes of discontinuity set S_f

Multivariate discontinuities

Definition Let $u \in L^1(\mathcal{X})$, we say that u is a **function of bounded variation** in $\mathcal{X} \subset \mathbb{R}^d$ if the distributional derivative Du is representable by a finite Radon measure in $\mathcal{X}$, i.e. if

$$\int_{\mathcal{X}} u \operatorname{div} \varphi dx = -\sum_{i=1}^{d} \int_{\mathcal{X}} \varphi_i dD_i u \quad \forall \varphi_i \in C_c^1(\mathcal{X})$$

for some $\mathbb{R}^n$-valued Radon measure $Du = (D_1 u, \ldots, D_d u)$ in $\mathcal{X}$.

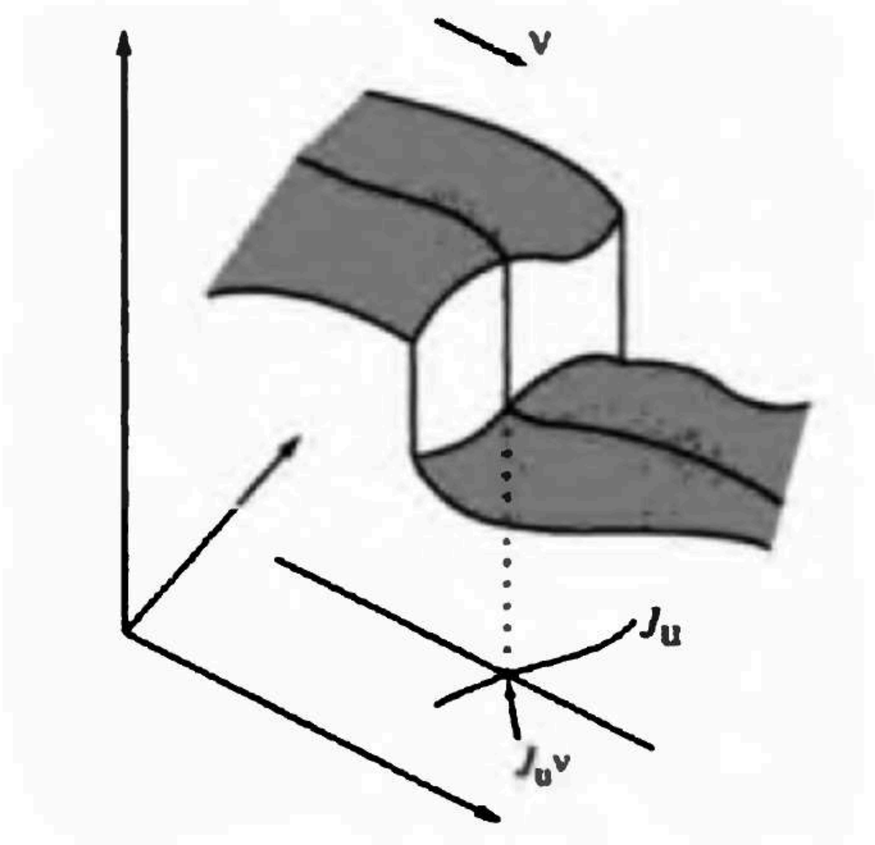

Ambrosio. Fusco. Pallara. 2000

Definition Let $u \in L^1_{loc}(\mathcal{X})$. Then $x \in \mathcal{X}$ is an **approximate jump point** of u if there exists $a \neq b \in \mathbb{R}, \nu \in \mathbf{S}^{d-1}$ such that

$$\lim_{\varepsilon \downarrow 0} ⨍_{B_\varepsilon^+(x,\nu)} |u(y) - a| dy = 0. \quad \lim_{\varepsilon \downarrow 0} ⨍_{B_\varepsilon^-(x,\nu)} |u(y) - b| dy = 0.$$

$$\begin{cases} B_\varepsilon^+(x,v) := \{y \in B_\varepsilon(x) : \langle y - x, v \rangle > 0\} \\ B_\varepsilon^-(x,v) := \{y \in B_\varepsilon(x) : \langle y - x, v \rangle < 0\} \end{cases}$$

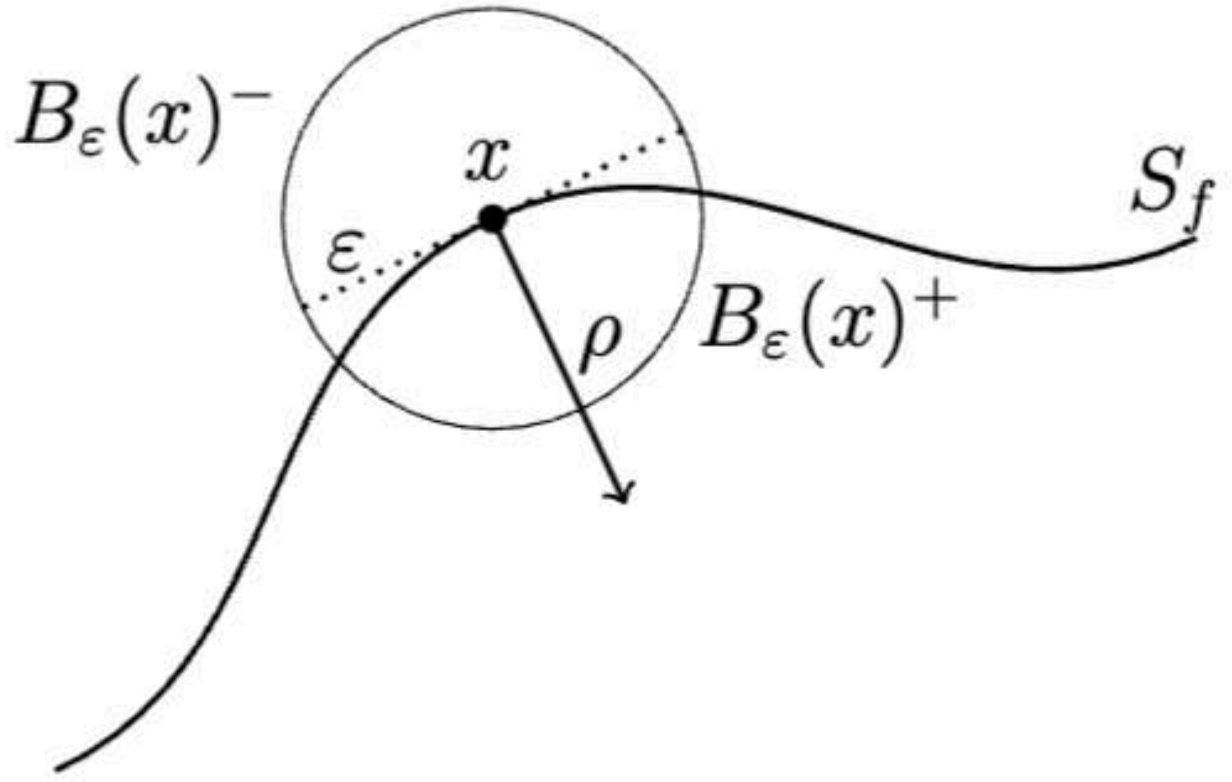

Then we denote the **set of jump points of function** u **as** S_u

Free Discontinuity Regression via the Mumford-Shah functional

Our statistical version of the Mumford-Shah functional is defined as:

$$E(u) = \lambda \underbrace{\int_{\mathcal{X}} (f-u)^2 f_X(x)\, dx}_{\text{Regression}} + \underbrace{\int_{\mathcal{X}\setminus S_u} |\nabla u|^2 f_X(x)\, dx}_{\text{Roughness penalty away from discontinuity}} + \underbrace{\nu \mathcal{H}^{d-1}(S_u)}_{\text{Boundary regularity penalty}}$$

where $\mathcal{X}$ is the support of $u(x)$, $f_X(x)$ is the density of X, S_u is its discontinuity set, and $\mathcal{H}^{d-1}$ is the Hausdorff measure in d-1 dimensions.

Natural improvement of naive edge detection approaches. Benefits:

1. **one-step denoising + thresholding**
2. **imposes boundary regularity using global function information**
3. **explicitly estimates the discontinuity set**

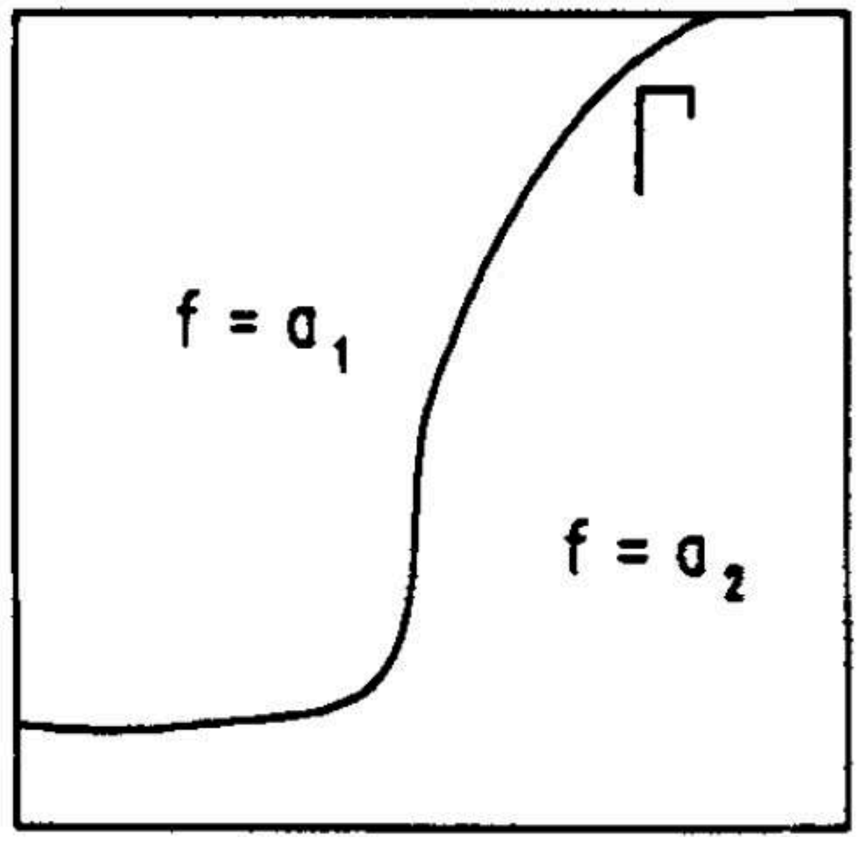

Mumford & Shah, 1989

Convexifying the problem

The MS functional is not convex $\Longrightarrow$ local solutions and bad artifacts, **issue in particular for statistical inference!**

Solution: Convexification via the **calibration method** (Alberti, Bouchitté, dal Maso 2003; Pock et al 2009)

$$E(u) = \sup_{p \in K} \int_{\mathcal{X} \times \mathbb{R}} p \cdot D1_u \qquad \text{with } 1_u := 1_{u(x)>t}(x,t)$$

$$K = \left\{ p \in C_0\left(\mathcal{X} \times \mathbb{R}, \mathbb{R}^{d+1}\right) : p^t(x,t) \geq \frac{\left|p^x(x,t)\right|^2}{4\, f_X(x)} - \lambda\, f_X(x)(t - f(x))^2 \right.$$

$$\left. \text{and} \quad \left| \int_{t_1}^{t_2} p^x(x,s)ds \right| \leq \nu \right\}$$

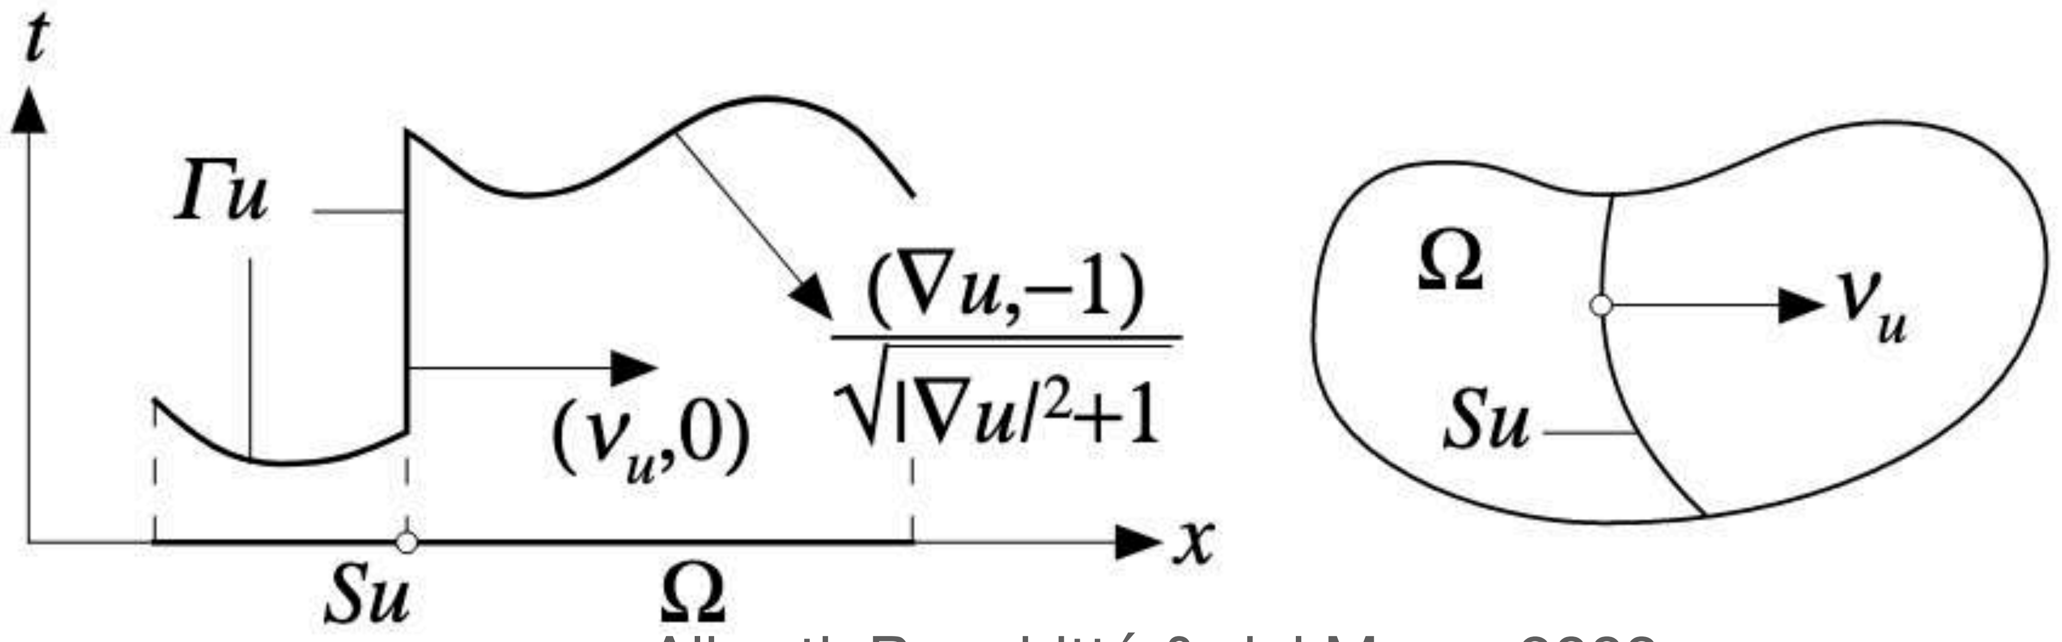

Alberti, Bouchitté & dal Maso, 2003

Additional convexification: replace 1_u with $v \in C$ where

$$C = \left\{ v \in SBV(\mathcal{X} \times \mathbb{R}) \middle| \; v : \mathcal{X} \times \mathbb{R} \to [0,1], \lim_{t \to -\infty} v(x,t) = 1, \lim_{t \to +\infty} v(x,t) = 0 \right\}$$

Convexifying the problem

Example of “functional lifting”

(a) Original image

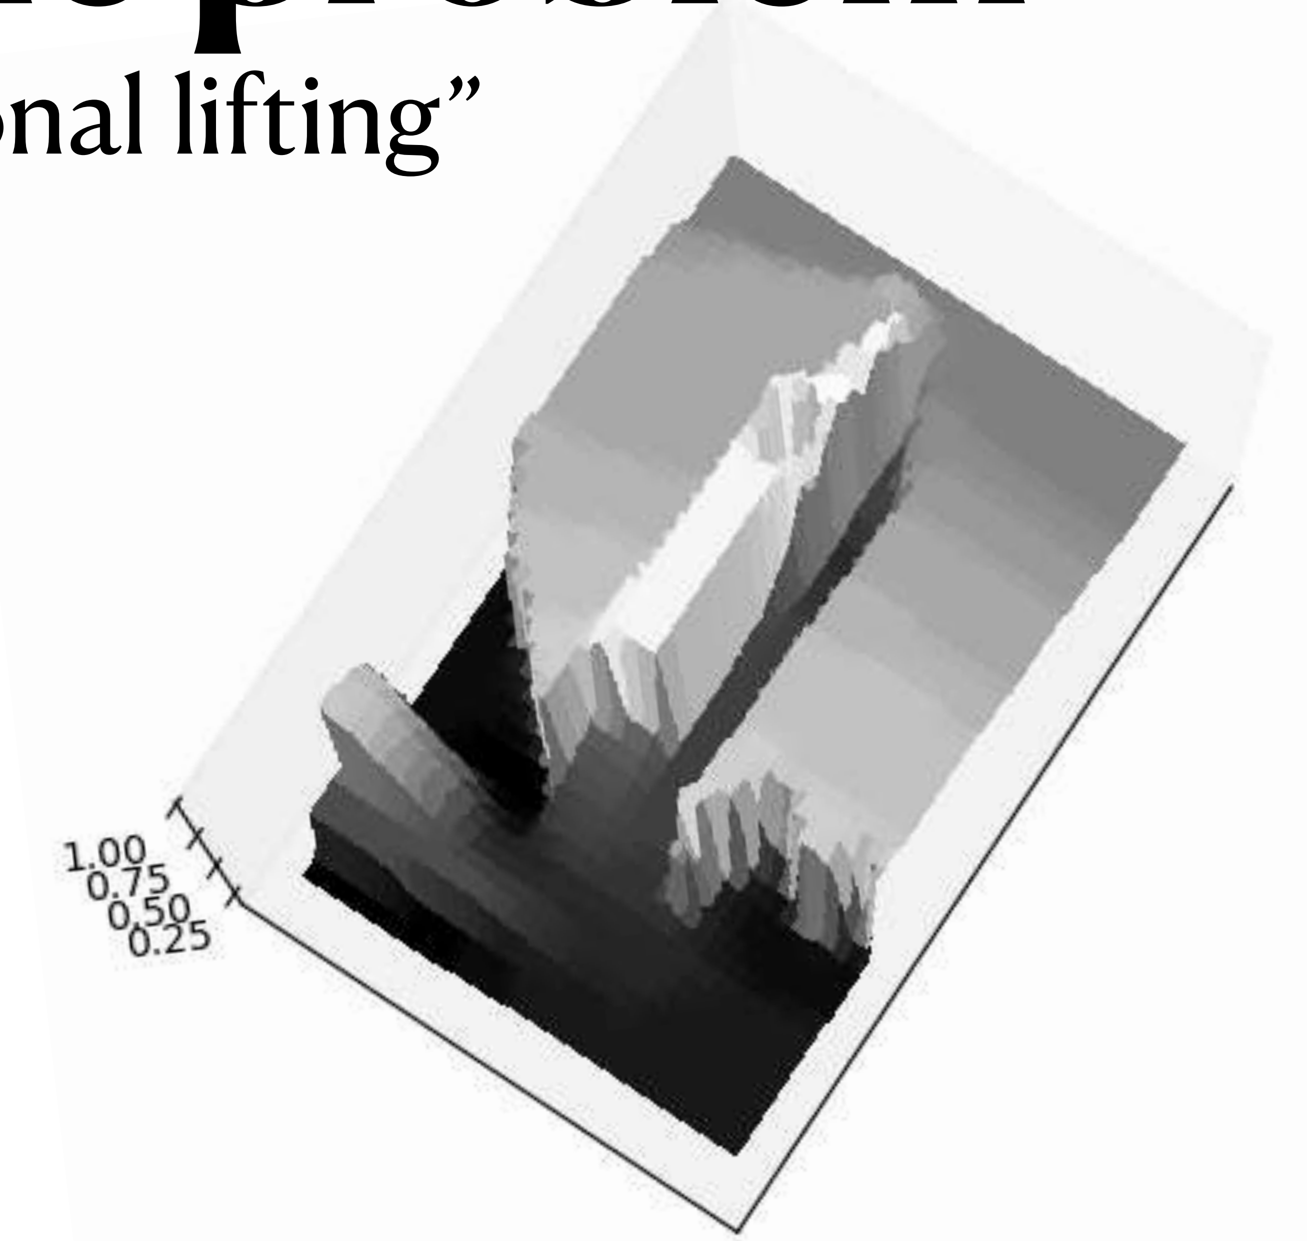

(b) Lifted estimate

(c) Estimated isosurface

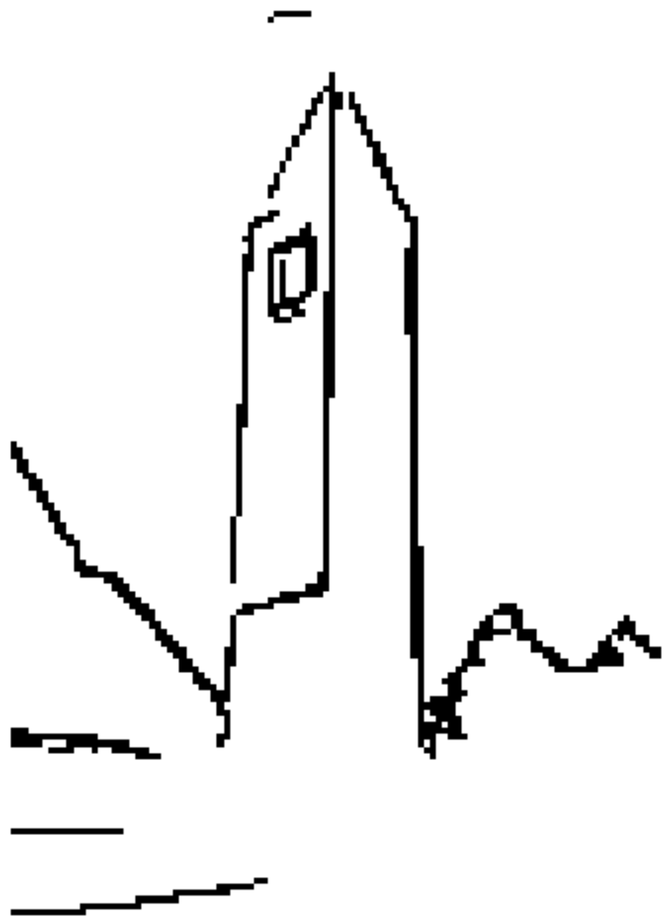

(d) Estimated boundary

Convexifying the problem

Example of “functional lifting”

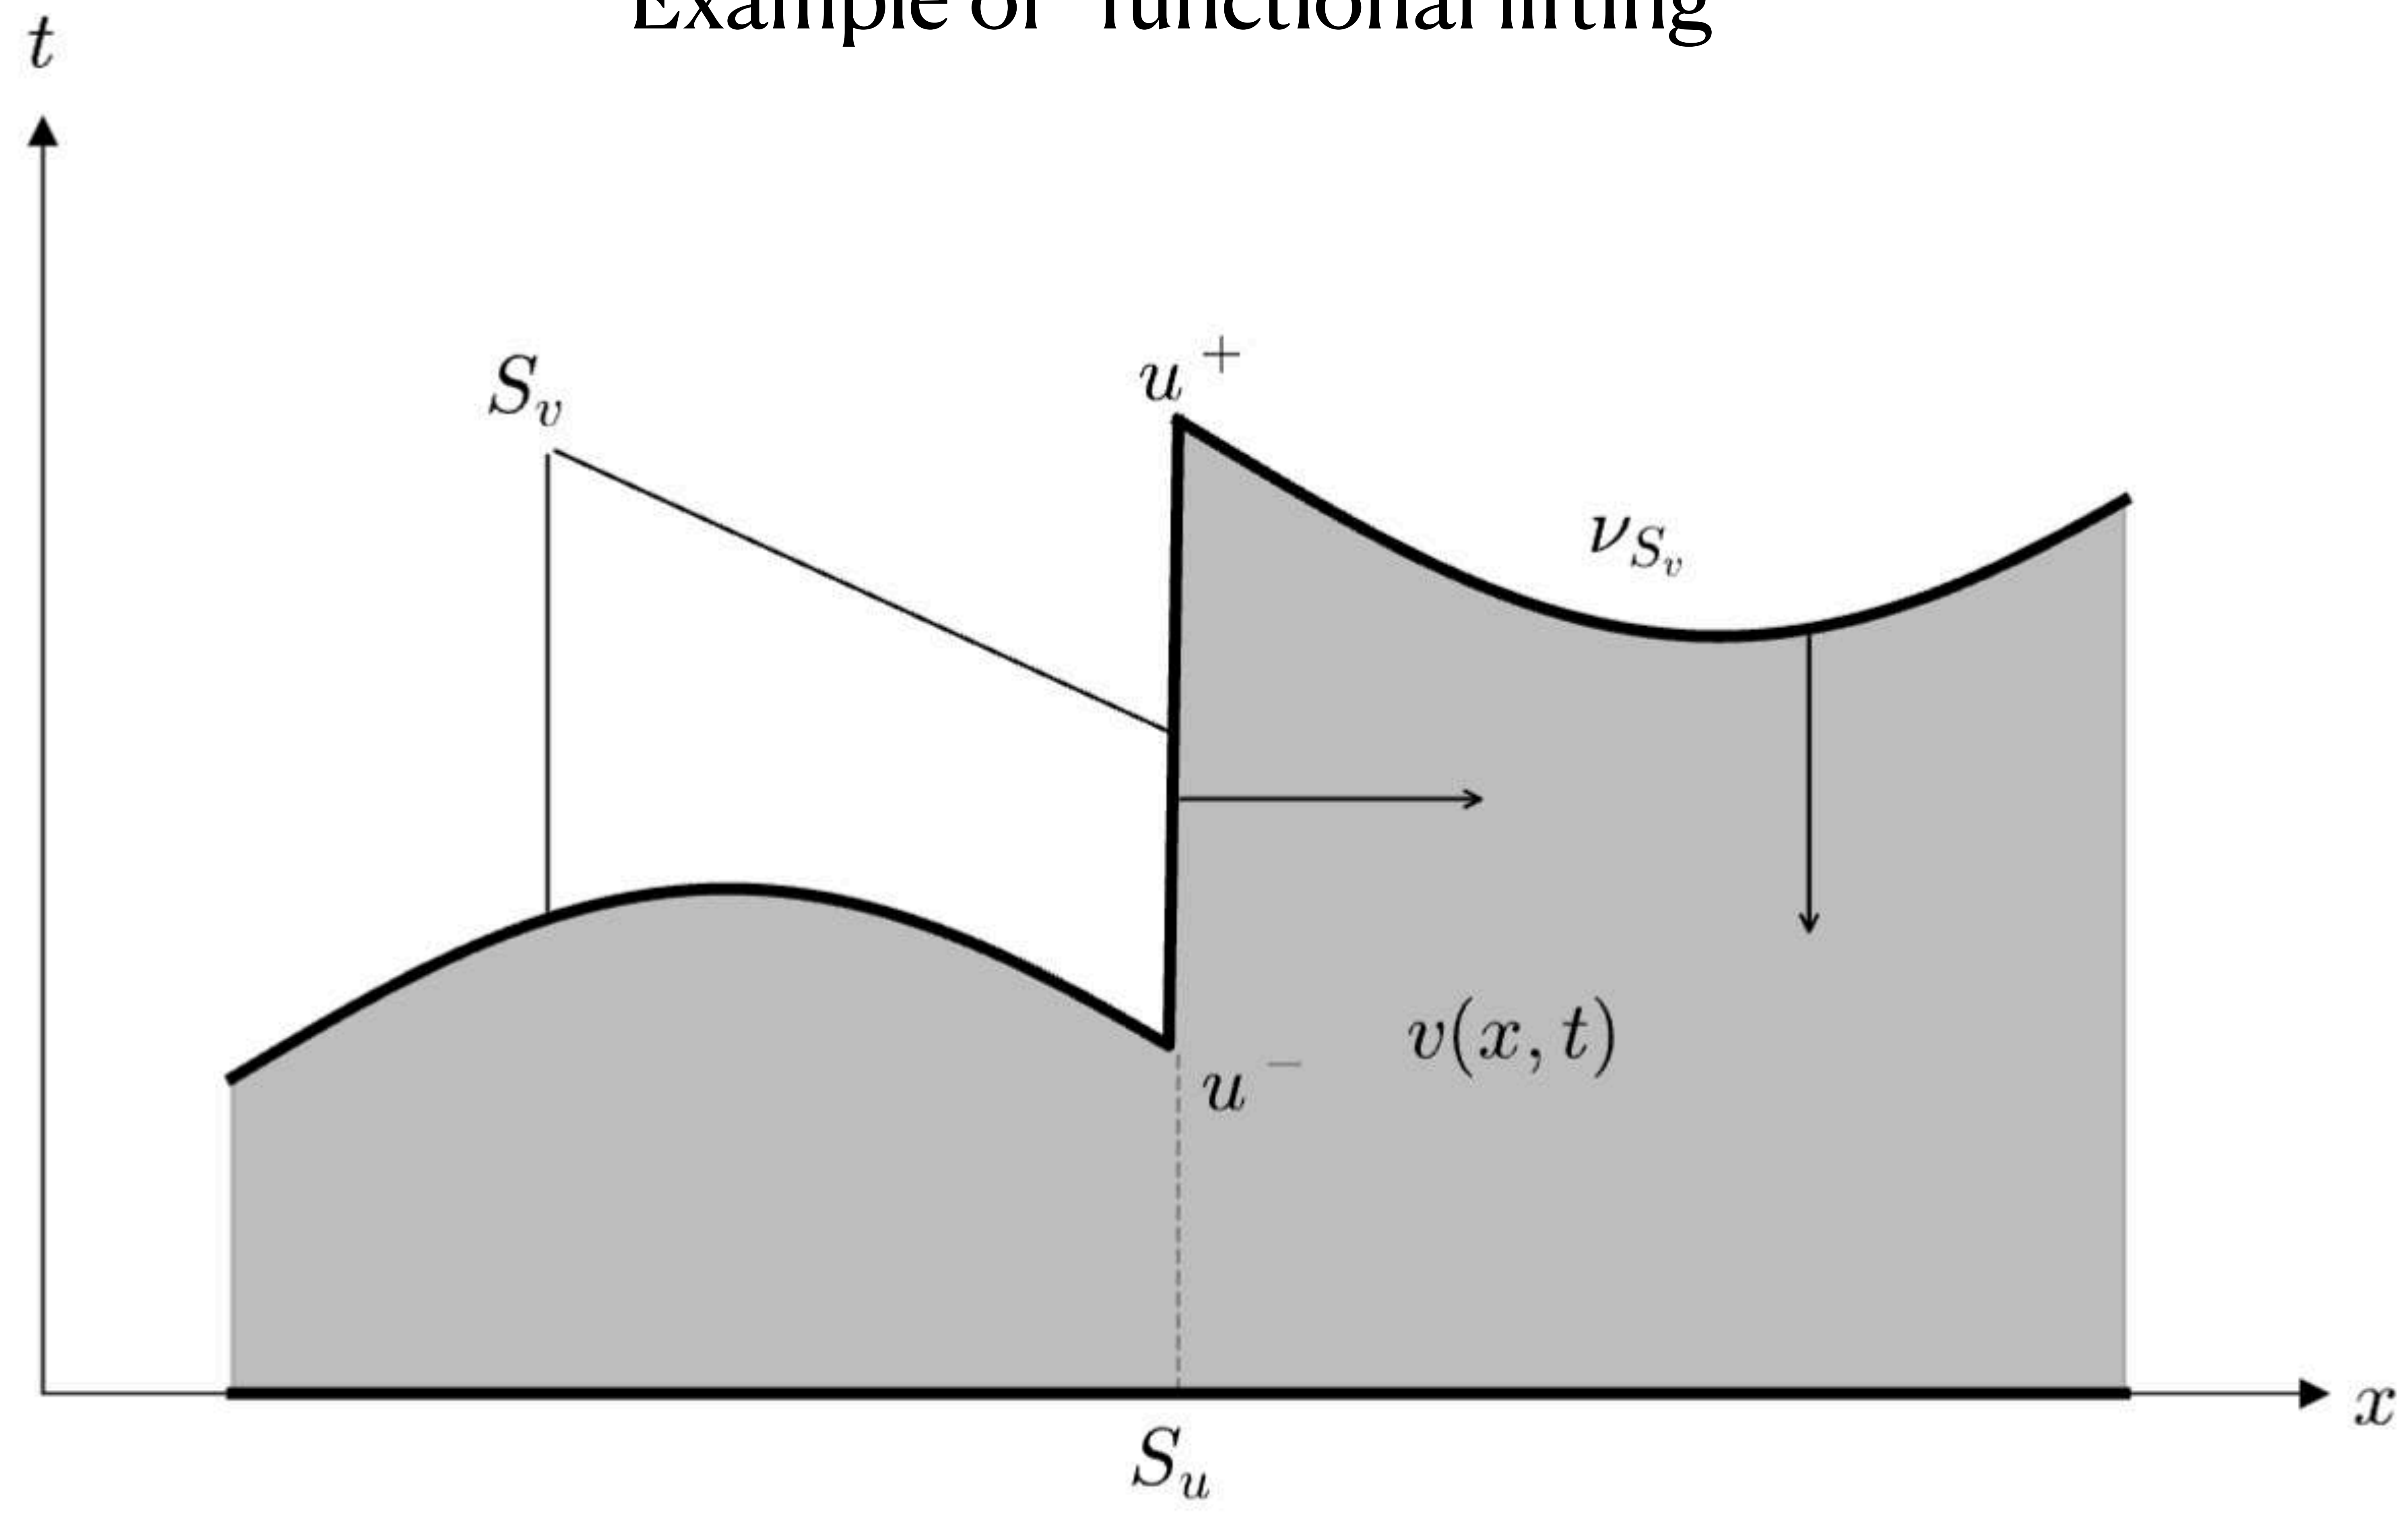

Figure A-1: Convex Relaxation Through Functional Lifting

Identification

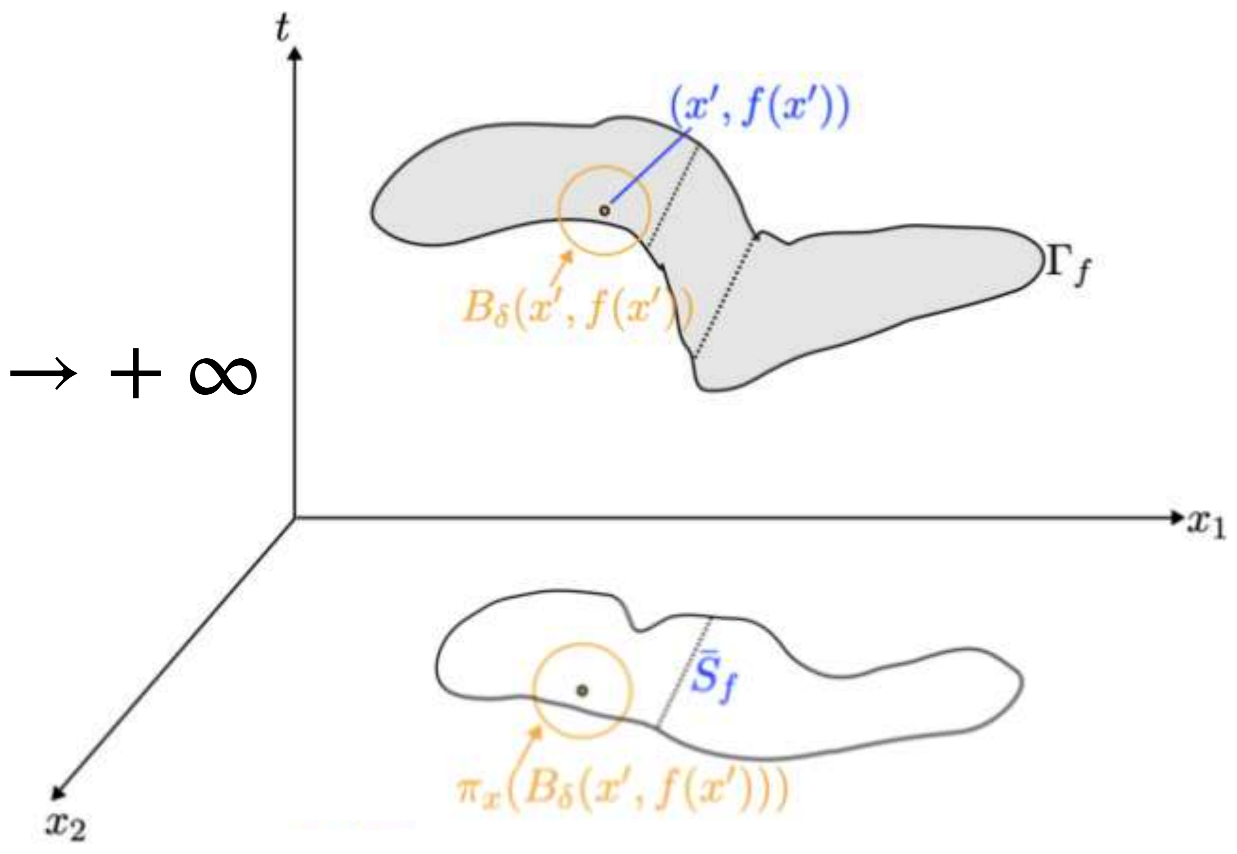

We prove that we identify **both the location of the discontinuity as well as the jump sizes in the limit as $\lambda \to +\infty$**

Assumption 2 (i) $f_X(x) \geq c > 0$ on its support $\mathscr{X}$; (ii) $f \in SBV(\mathscr{X})$ and bounded almost everywhere; (iii) $\int_{\mathscr{X}} |\nabla f|^2 dx + \mathscr{H}^{d-1}(S_f) < +\infty$.

Assumption 3 For any $x \in S_f$ it holds $\mathscr{H}^{d-1}\left(S_f \cap B_\rho(x)\right) > 0$ for all $\rho > 0$. Moreover, for any set $A \subset \mathscr{X}$ with $\mathrm{dist}(A, S_f) > 0$ there exists a constant $0 < L < +\infty$ such that $\left|f(x) - f(y)\right| \leq L|x - y|$ for all $x, y \in A$.

Theorem Let Assumptions 2 and 3 hold. Then for fixed $\nu > 0$ and in the limit as $\lambda \to +\infty$ every sequence of solutions $v^*(\lambda)$ to the optimization problem satisfies $\lim_{\lambda \to +\infty} \nabla v^*(\lambda) = 0$ $\mathscr{L}^{d+1}$-almost everywhere. Moreover, the jump set $J_{v^*}(\lambda)$ converges in Hausdorff distance to the graph Γ_f of f, i.e.

$$\lim_{\lambda \to +\infty} d_H(J_{v^*}, \Gamma_f) = 0.$$

This problem complements existing deep results for the classical MS functional (Morini 2001, Richardson 1992), but is more general, because we get convergence of **the entire graph in Hausdorff distance**. Existing results show identification of the discontinuity set S_f.

Empirical implementation

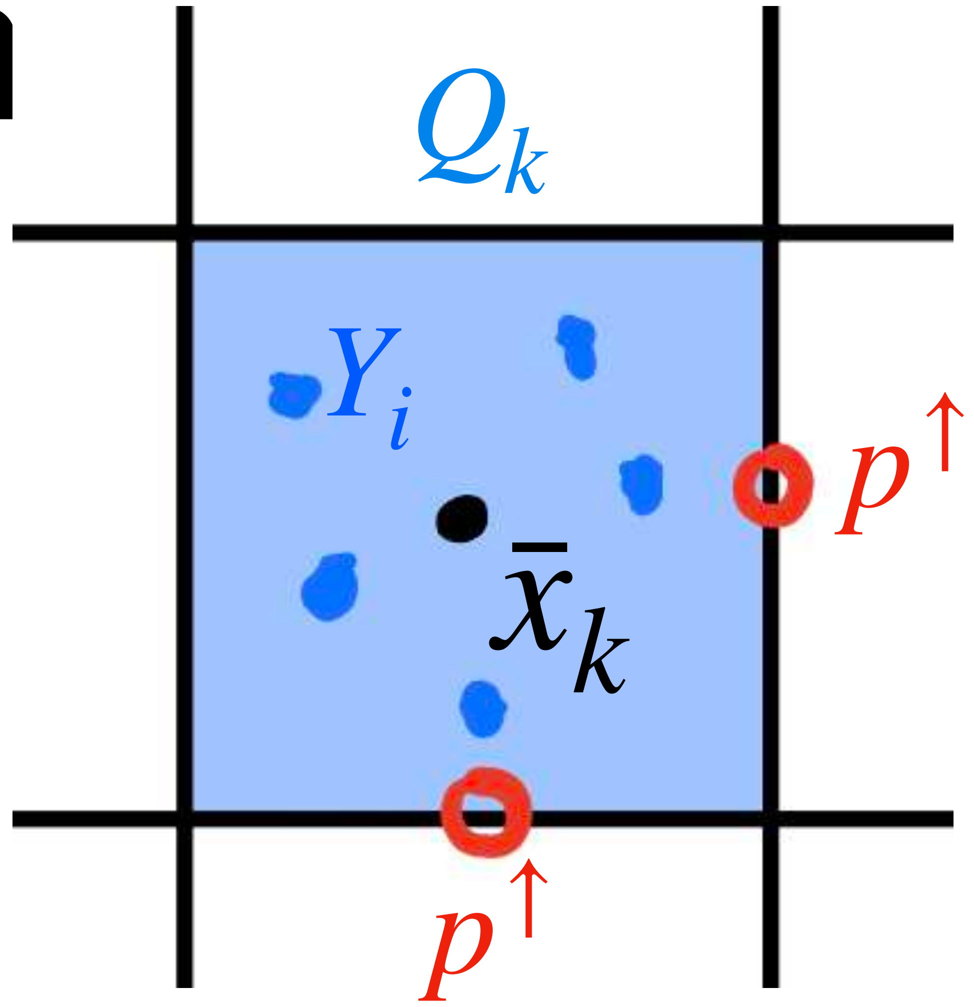

Random data points on grid.

$$\hat{f}_{X,Nn}(\bar{x}_k) = \frac{1}{nh}\sum_{i}^{n} K\left(\frac{\bar{x}_k - X_i}{h}\right) \qquad \hat{f}_{Nn}(\bar{x}_k) = \sum_{i:X_i\in Q_k} w_i Y_i$$

$$\min_{v\in\tilde{C}_N}\max_{p\in\hat{K}_{Nn}} \langle p, D_N v\rangle_N$$

$$\hat{K}_{Nn} = \left\{ p = \left(p^x, p^t\right)^T \in Y : p^t(i_1,\ldots,i_d,k) \geq \frac{\left|p^x(i_1,\ldots,i_d,k)\right|^2}{4\,\hat{f}_{X,Nn}(i_1,\ldots,i_d)} - \lambda\,\hat{f}_{X,Nn}(i_1,\ldots,i_d)\left(\frac{k}{S} - \hat{f}_{Nn}(i_1,\ldots,i_d)\right)^2, \left|\frac{1}{N}\sum_{k_1\leq k\leq k_2} p^x(i_1,\ldots,i_d,k)\right| \leq \nu \right\},$$

$$\tilde{C}_N = \left\{ u \in X : v(i_1,\ldots,i_d,k) \in [0,1], v(i_1,\ldots,i_d,1) = 1, v(i_1,\ldots,i_d,S) = 0 \right\}$$

Solve this problem using a **primal-dual gradient descent-ascent** algorithm (Chambolle and Pock, 2011).

Obtaining the boundary

Existing methods back out the **discontinuity set** by **thresholding the gradient** of the optimal u^*:

$$J_u = \left\{ (i,j) : |\nabla u^*(i,j)| \geq \sqrt{\nu} \right\}.$$

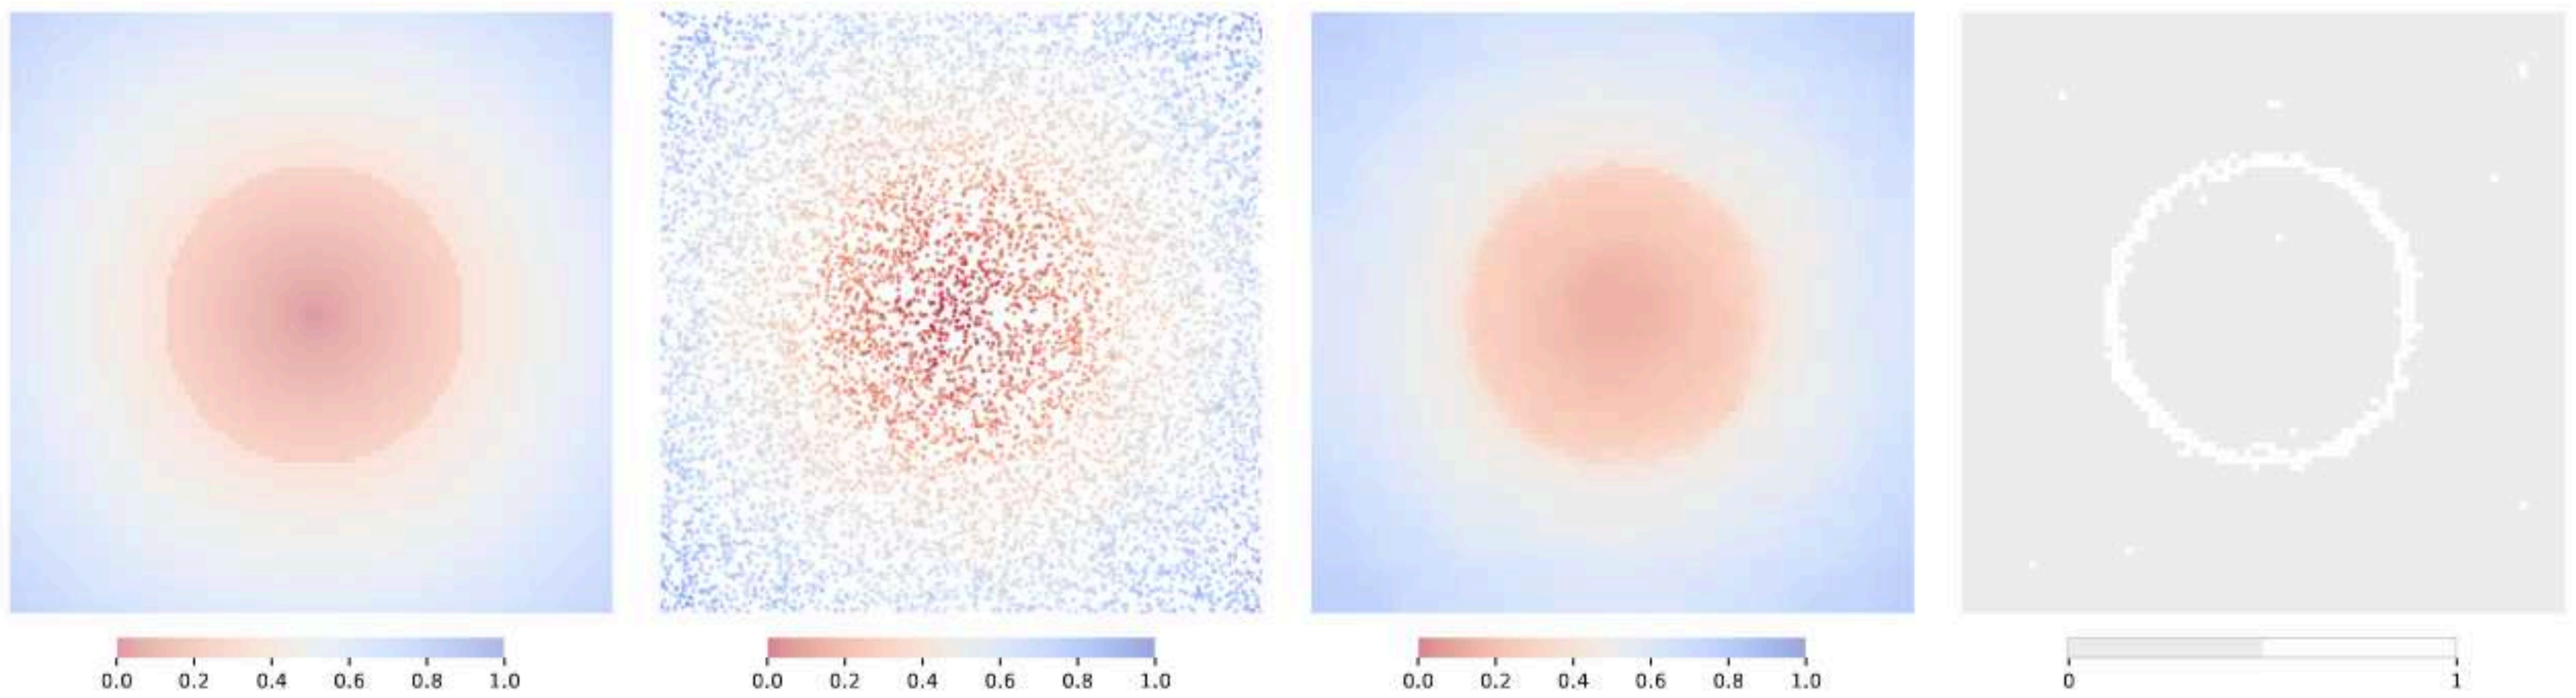

(b) 2D: Truth (c) 2D: Point Cloud (d) 2D: $\hat{u}$ (e) 2D: Boundary

How to pick a good ν in practice? **Hyperparameter selection via a version of Stein's Unbiased Risk Estimate (SURE)**, see also Lucas et. al. (2023) for classical MS.

Statistical consistency

We prove consistency via Γ-convergence (dal Maso, 2012).

Assumption 5: *The density f_X is bounded away from zero on its compact support $\mathcal{X}$. The function $f(x) \in BV\left([0,1]^{d+1}\right)$ is bounded.*

Theorem: *Let Assumption 1 hold and let $N(n) \to 0$ with $Nn \to +\infty$ as $n \to +\infty$. Then*

$$\hat{E}_{Nn}(v) = \sup_{p \in \hat{K}_{Nn}} \langle p, D_N v \rangle_N$$

Γ-converges in the weak-topology in probability to*

$$E(v) = \begin{cases} \sup_{p \in K} \int_{[0,1]^{d+1}} p \cdot Dv & \text{if } v \in C \\ +\infty & \text{else} \end{cases}.$$

This is the first convergence result for convexified MS and complements recent convergence results in this area:

Chambolle & Pock (2021), Caroccia et. al. (2020), García-Trillos & Slepcev (2016), Chambolle et.al. (2017)

Together with **compactness criterion** (simple) this proves **consistency of the minimizer**.

Hyperparameter tuning

- **Data-driven** choice for $\theta = (\lambda, \nu)$: SURE (Stein, 1981)
- Data-generating process:

$$\tilde{f} = f + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2) \in \mathbb{R}^n$$

- Given an estimator $\hat{u}_\theta(\tilde{f})$ for f, the SURE is,

$$\eta(\hat{u}_\theta(\tilde{f})) = \frac{1}{N}\|\tilde{f} - \hat{u}_\theta(\tilde{f})\|^2 - \sigma^2 + 2\sigma^2 \operatorname{div}_f \hat{u}_\theta(\tilde{f}) .$$

- Stein's lemma proves that $\hat{u}_\theta(\tilde{f})$ is an **unbiased estimator of the mean squared error**,

$$MSE(\hat{u}_\theta(f)) := \frac{1}{N}\|\hat{u}_\theta(f) - y\|^2.$$

- In practice, we compute $\operatorname{div}_f \hat{u}_\theta(\tilde{f})$ using a **Monte-Carlo perturbation approach**

$$\operatorname{div}_{\mathbf{y}} \{\mathbf{f}_\lambda(\mathbf{y})\} = \lim_{\varepsilon \to 0} E_{\mathbf{b}'} \left\{ \mathbf{b}^{\mathrm{T}} \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \right\},$$

which gives an asymptotically unbiased estimator of the divergence.

- Implemented in parallel in the Python *Ray* library, currently takes 4-12 hours RAY

Practical implementation

Details

- Implemented in **PyTorch** with support for Nvidia (CUDA), Intel, and Apple Silicon GPU acceleration
  - Extensions to R, STATA in progress
  - Benchmark:
    - Nvidia Tesla A100 Tensor Core GPU
    - n=90,000; $N_{grid} = 12{,}750$ (25%); S=32
    - Execution time: **69.37 seconds**
    - Can be improved by using approximations of the cost function ("sublabel") (Mollenhoff and Cremers, 2017)

Simulations

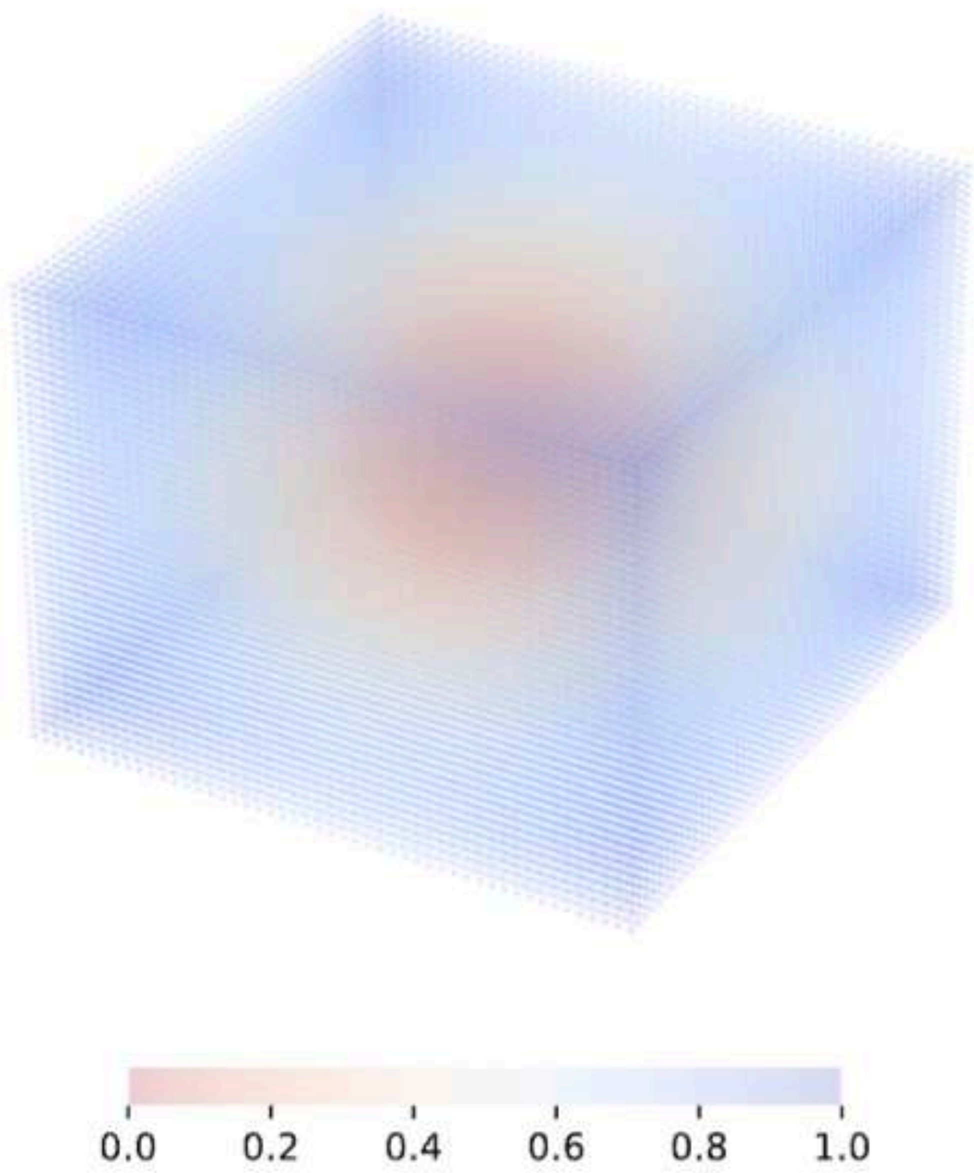

(f) 3D: Truth

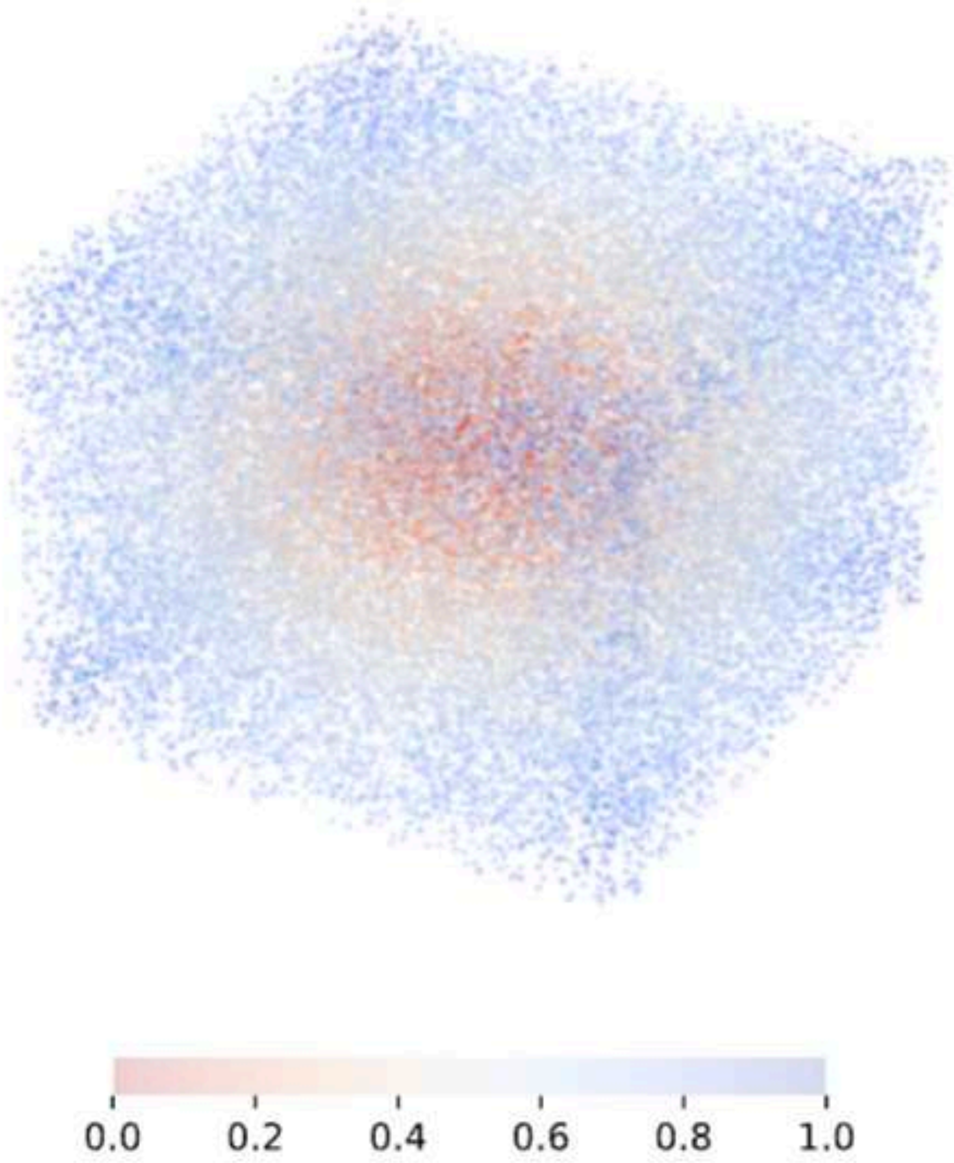

(g) 3D: Point Cloud

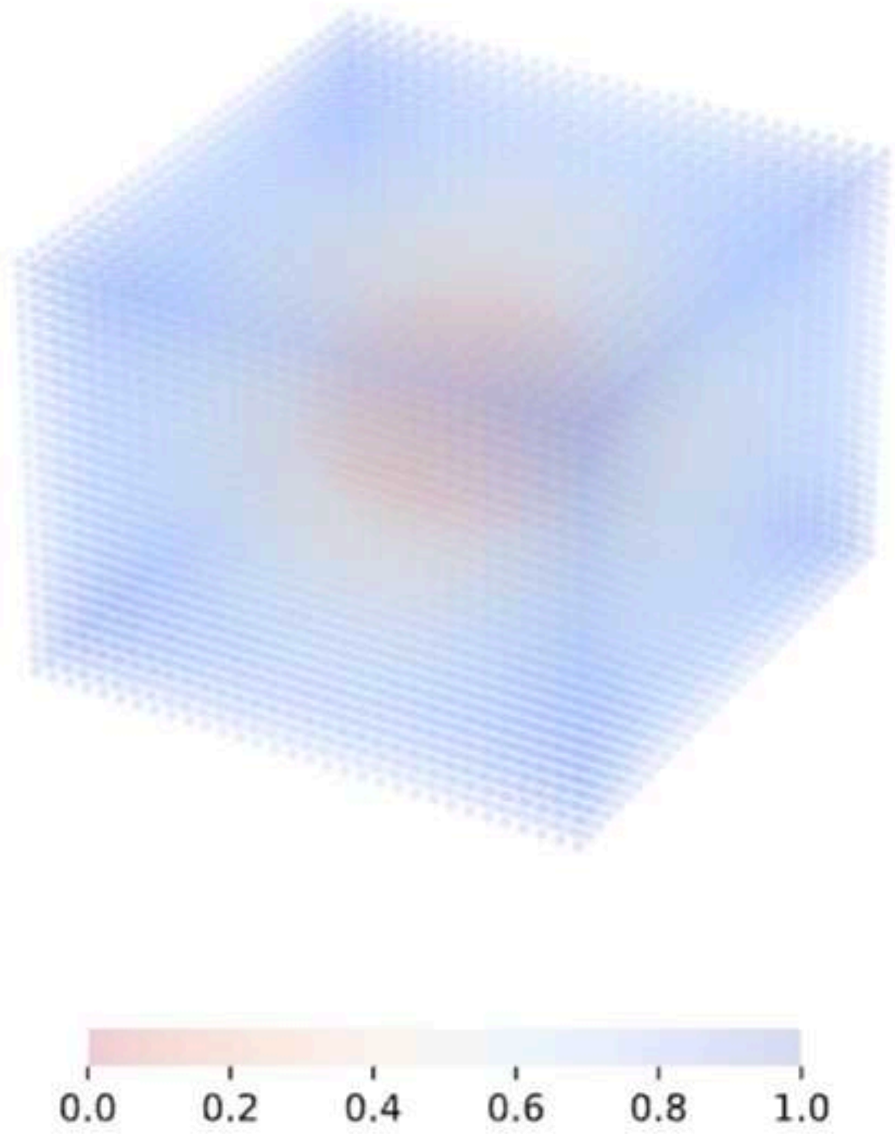

(h) 3D: $\hat{u}$

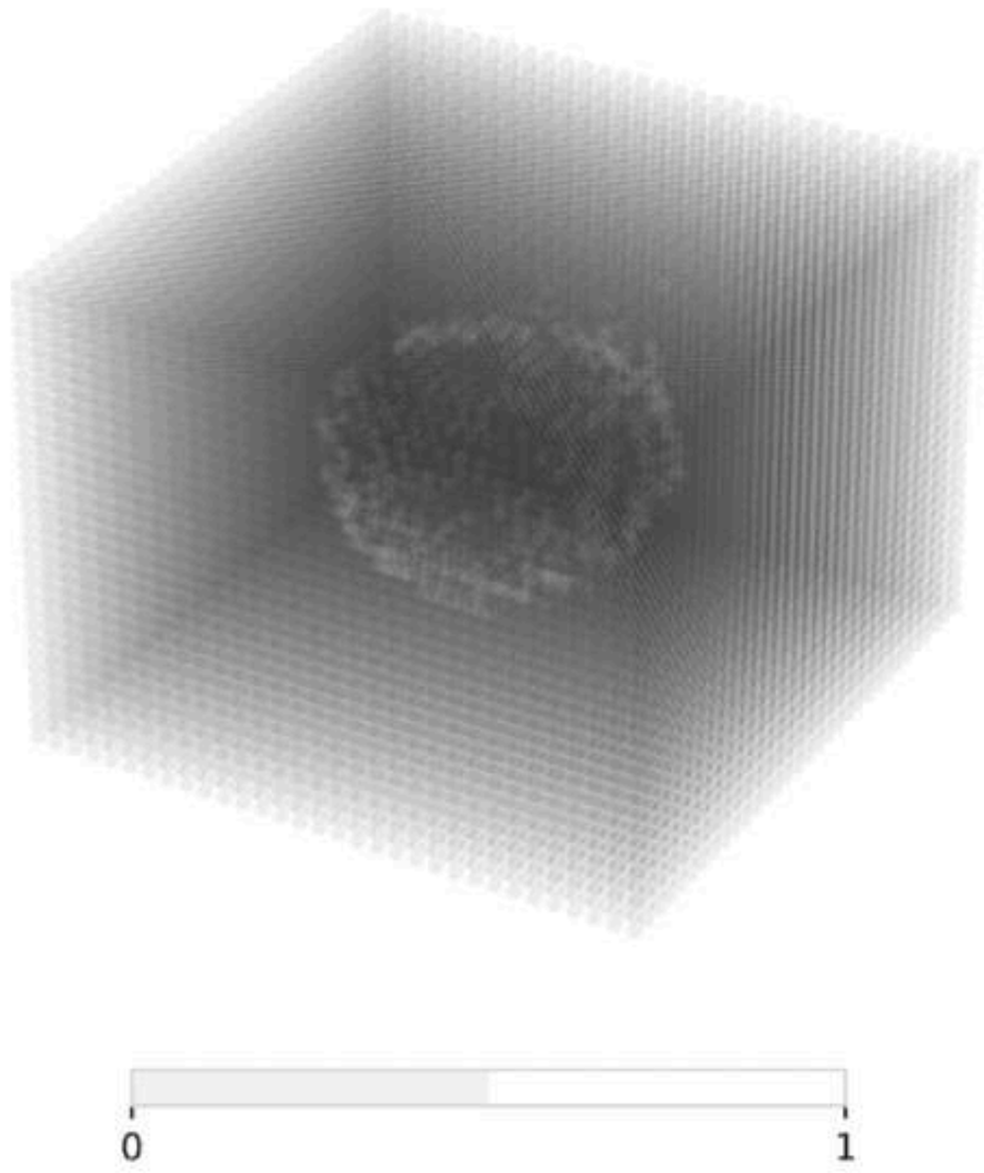

(i) 3D: Boundary

Table 1: Monte Carlo Simulations

(a) 1D

n	MSE	MSE τ_{FD}	Bias τ_{FD}	FNR	FPR
500	0.0789	0.0358	-0.084	-0.0437	0.1933
1000	0.0461	0.0305	-0.0659	0	0.1099
5000	0.0012	0.0099	-0.053	0	0.018

SURE: $\lambda = 98.6712$, $\nu = 0.0001$

(b) 2D

d = 0.25

n	α	$\hat{\alpha}$	MSE	MSE τ_{FD}	Bias τ_{FD}	FNR	FPR
1000	0.0375	0.0698	0.0022	0.0027	0.0323	0.0098	0.6757
5000	0.0375	0.0351	0.001	0.0002	-0.0024	0.0735	0.2319
10000	0.0375	0.0297	0.0009	0.0001	-0.0078	0.269	0.0391

SURE: $\lambda = 65.6470$, $\nu = 0.0010$

d = 0.50

n	α	$\hat{\alpha}$	MSE	MSE τ_{FD}	Bias τ_{FD}	FNR	FPR
1000	0.075	0.0859	0.0025	0.0021	0.0108	0.0159	0.561
5000	0.075	0.056	0.0011	0.001	-0.0191	0.0848	0.0837
10000	0.075	0.0545	0.001	0.0007	-0.0206	0.1476	0.0072

SURE: $\lambda = 107.6274$, $\nu = 0.0019$

d = 0.75

n	α	$\hat{\alpha}$	MSE	MSE τ_{FD}	Bias τ_{FD}	FNR	FPR
1000	0.1126	0.0953	0.003	0.0027	-0.0173	0.0093	0.5902
5000	0.1126	0.0732	0.0016	0.0036	-0.0394	0.0588	0.1676
10000	0.1126	0.078	0.0012	0.0027	-0.0346	0.0901	0.0181

SURE: $\lambda = 209.4585$, $\nu = 0.0020$

Application: Internet Shutdowns

veraset

- Various district governments of Rajasthan state, India shut down mobile internet on September 26, 2021 from 6am to 6pm to prevent cheating on the Rajasthan Eligibility Exam for Teachers (REET).
- Mobile network coverage in India is much better than WiFi/wired: major disruption of economic activity
  - WiFi 0.08% of wireless connectivity, wired connection only 3.74% of total internet subscriptions
  - Estimated 80,000 shops closed in Jaipur alone (The Economist, 2021)
  - Mobile payments ubiquitous (UPI), large digital economy, disruption of business operations, rural businesses no wired connection
- Estimate impact on mobile device signal using mobile device data from Veraset, a mobility data provider
  - 126 million pings in area on day of shutdown
  - 3.8 million unique devices (10% of mobile population)
- Perfect setting for our estimator:
  - **Unknown discontinuity**: diffuse area, not all districts that shut down issued official mandate, disconnected at point of contact
  - **Multivariate assignment:** latitude-longitude
- Effect on connectivity can serve as first-stage for downstream outcomes, we focus on economic activity

Application: Internet Shutdowns

- Main outcomes:

$$\overline{Pings_i} := \frac{Pings_{it_0}}{\frac{1}{3}\sum_{t=t_0-3}^{t_0-1} Pings_{it}},$$

- Index of total pings inside 5x5 km grid cells, relative to prior month average in same time window
- Measure economic impact: share of pings inside 40x40 km grid cells that fall in "economic areas"
  - "Economic areas": commercial or public OSM POI + SafeGraph shops — 108K polygons
- Idea: use remaining "cached" pings from satellite that get uploaded once internet connection is restored
- Assumption: disruption of numerator (economic pings) and denominator (total pings) is same
  - 1) economic activity of mobile users with continuous background location is representative: ~67% of Android users have it
  - 2) set of apps that cache background location is not skewed: by construction (continuous) + sourced from 1,000+ mobile apps (SDK)

Application: Internet Shutdowns

Mobile Signal

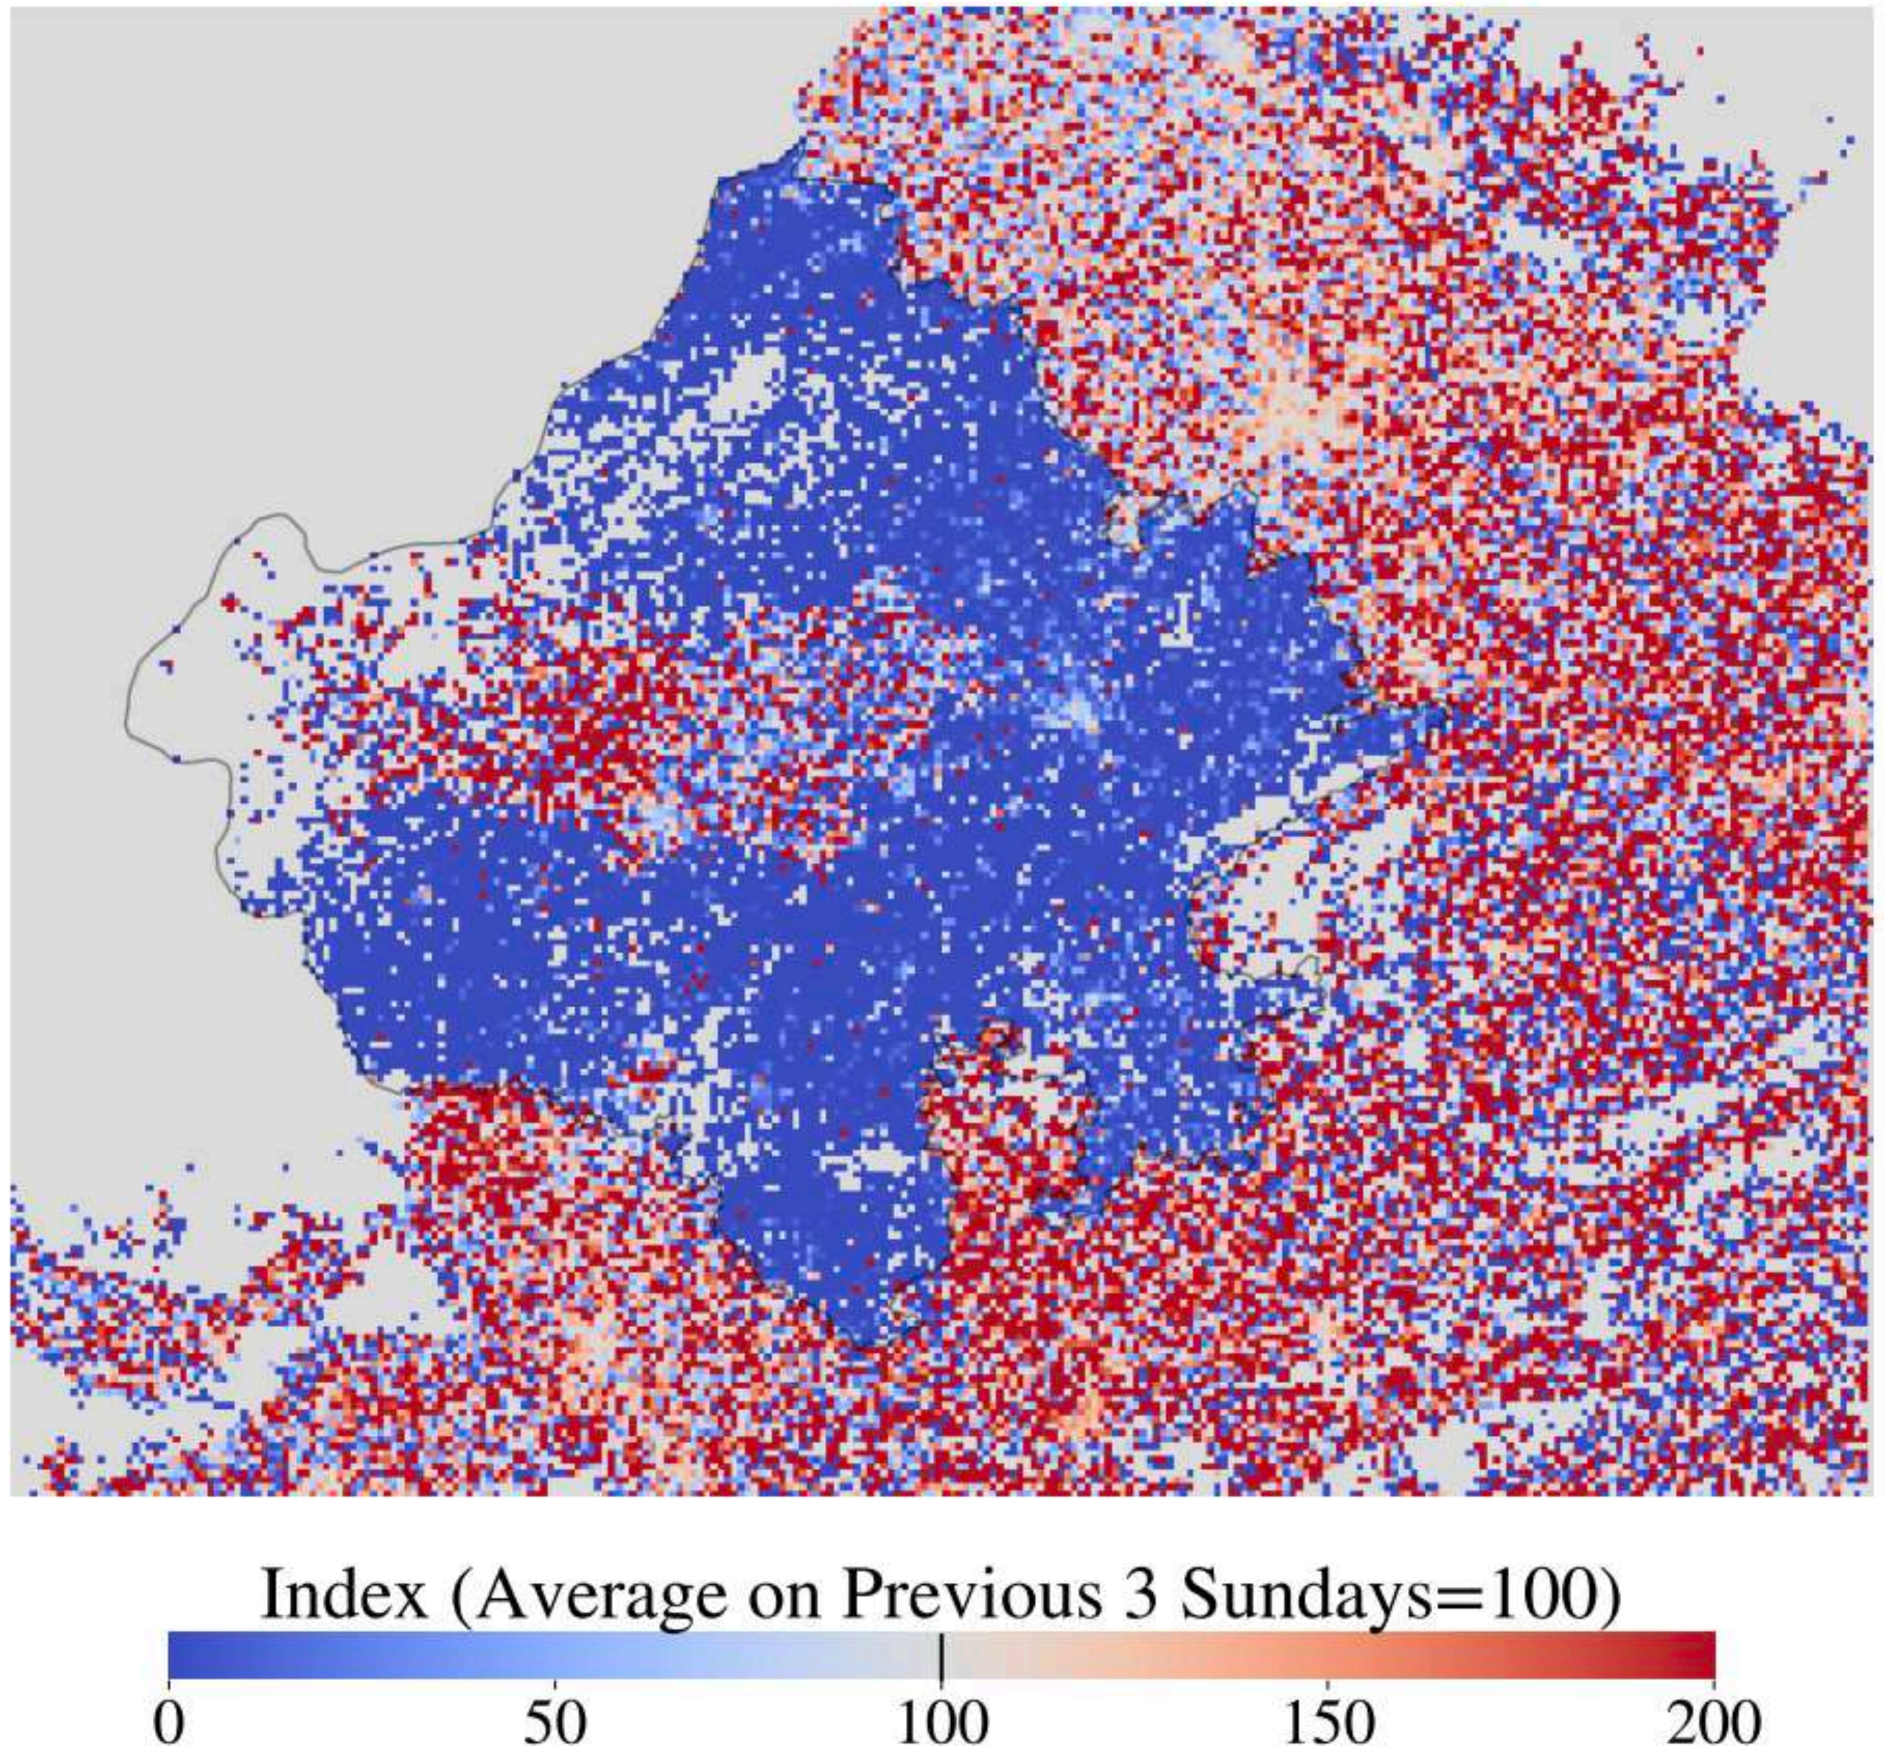

(a) Mobile: Raw Data

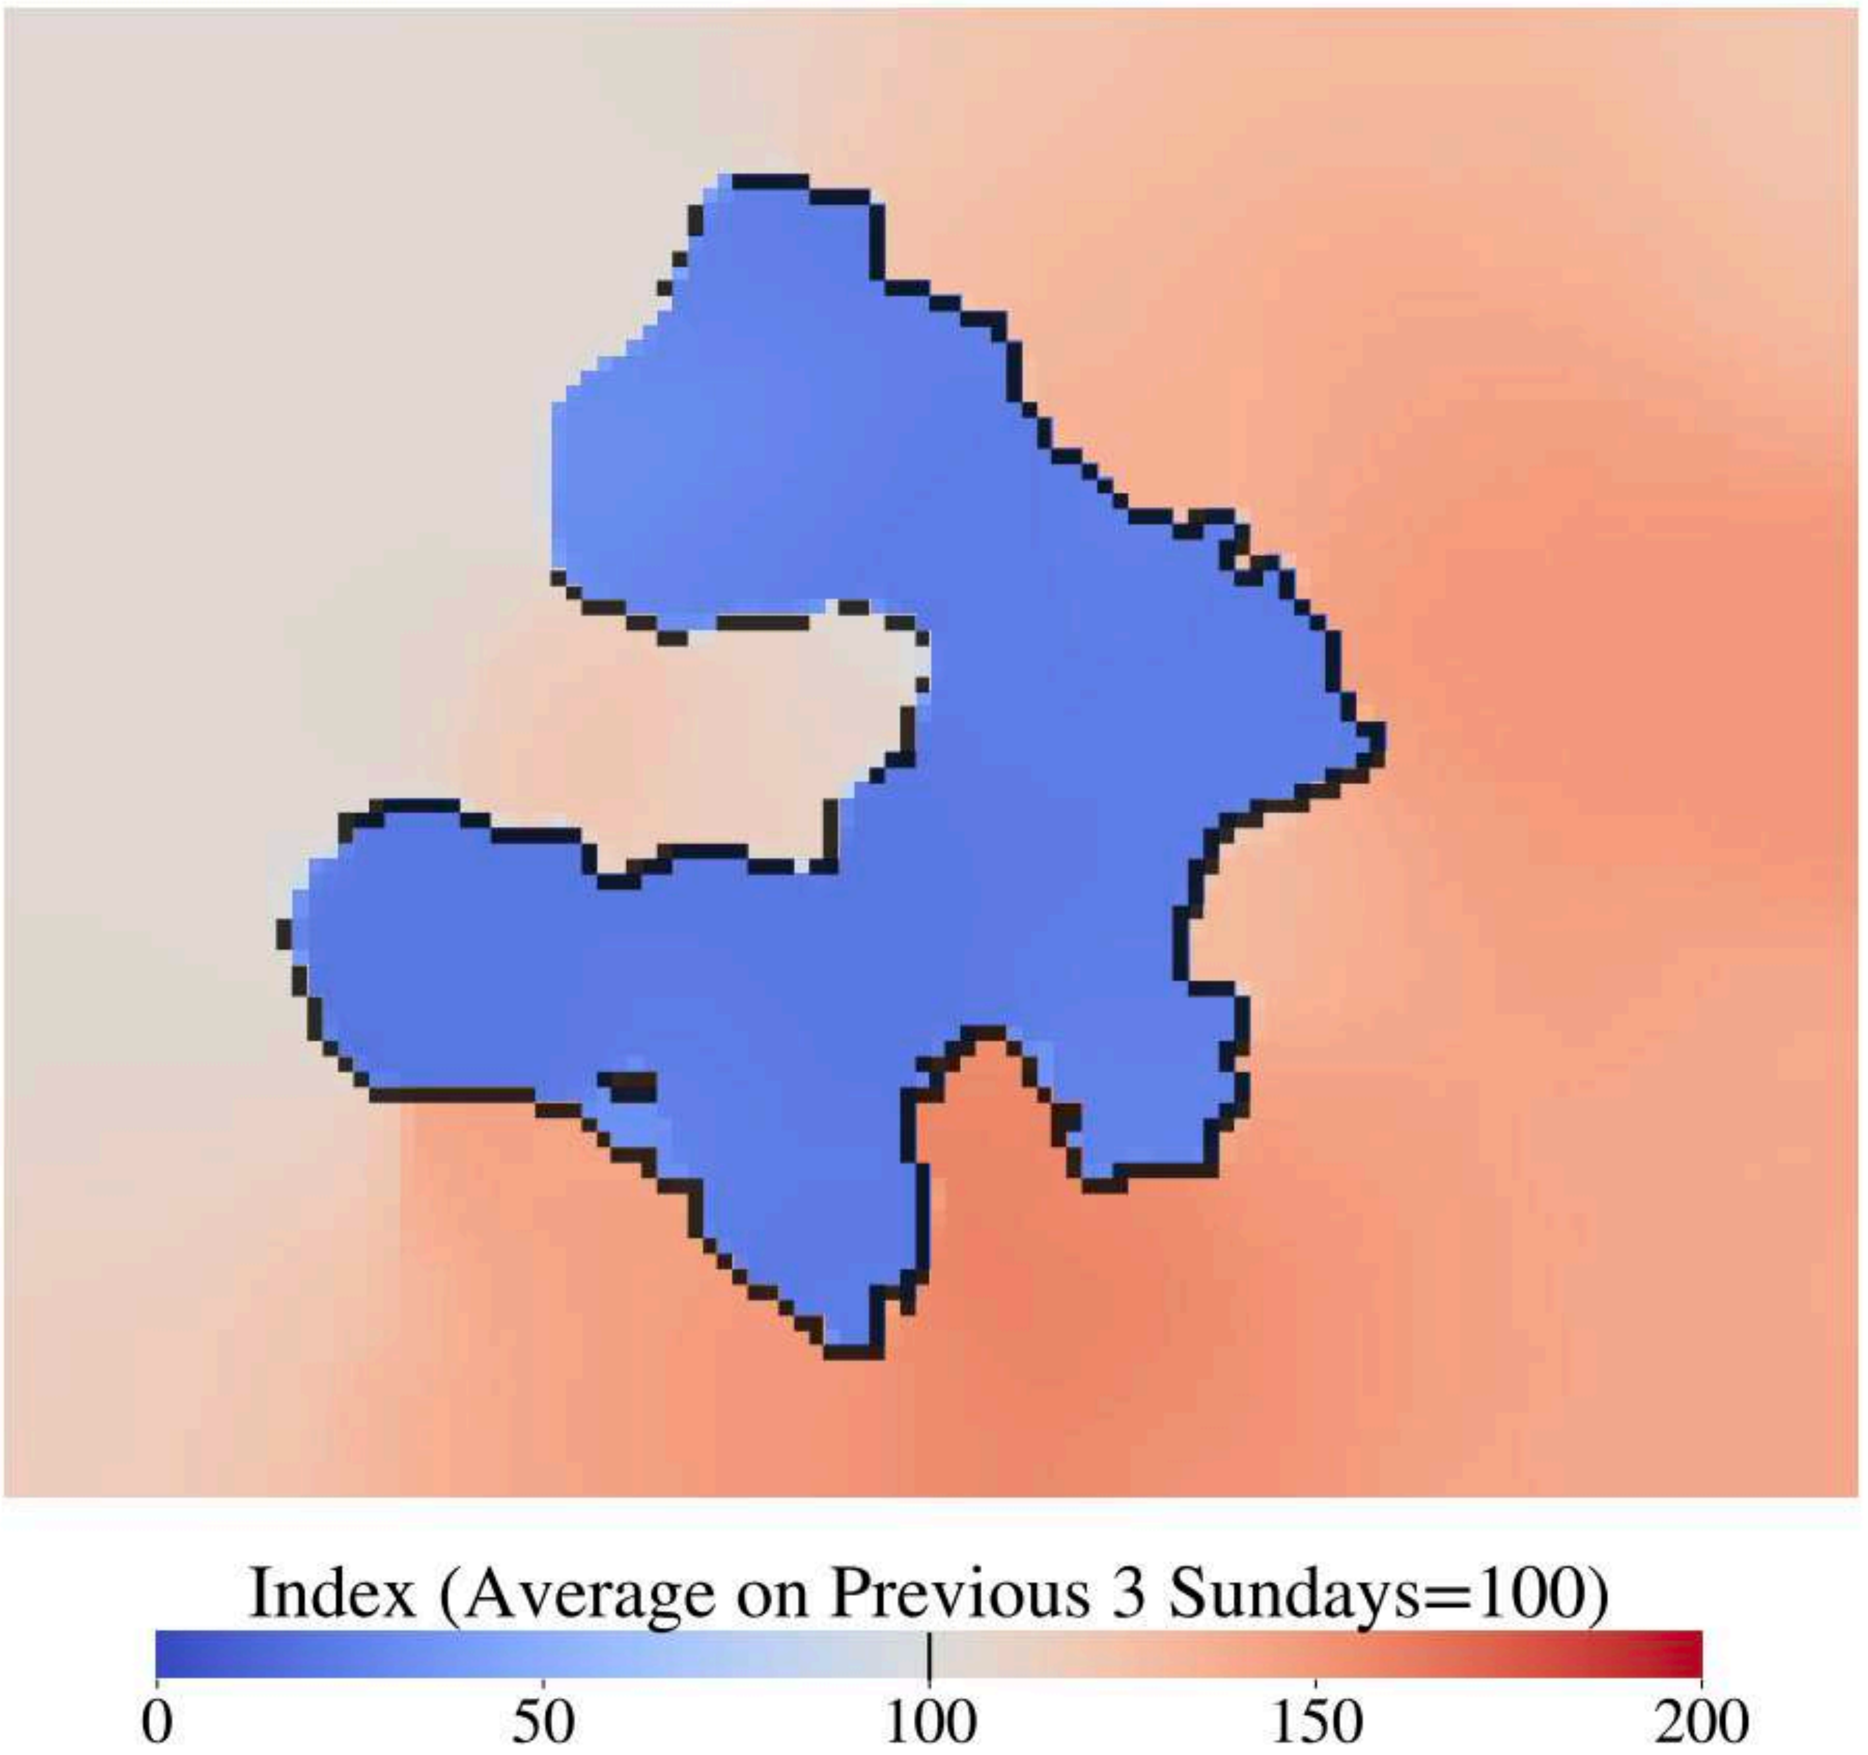

(b) Mobile: $\hat{u}$ & S_u

Application: Internet Shutdowns

Mobile signal: treatment Effect Curve

- Average drop of 100% relative to month average
- Increase in activity in connected areas: COVID rebound
- Overall: shutdown highly effective at disrupting mobile signal
- 25% of signal inside shutdown area remains due to satellite caching

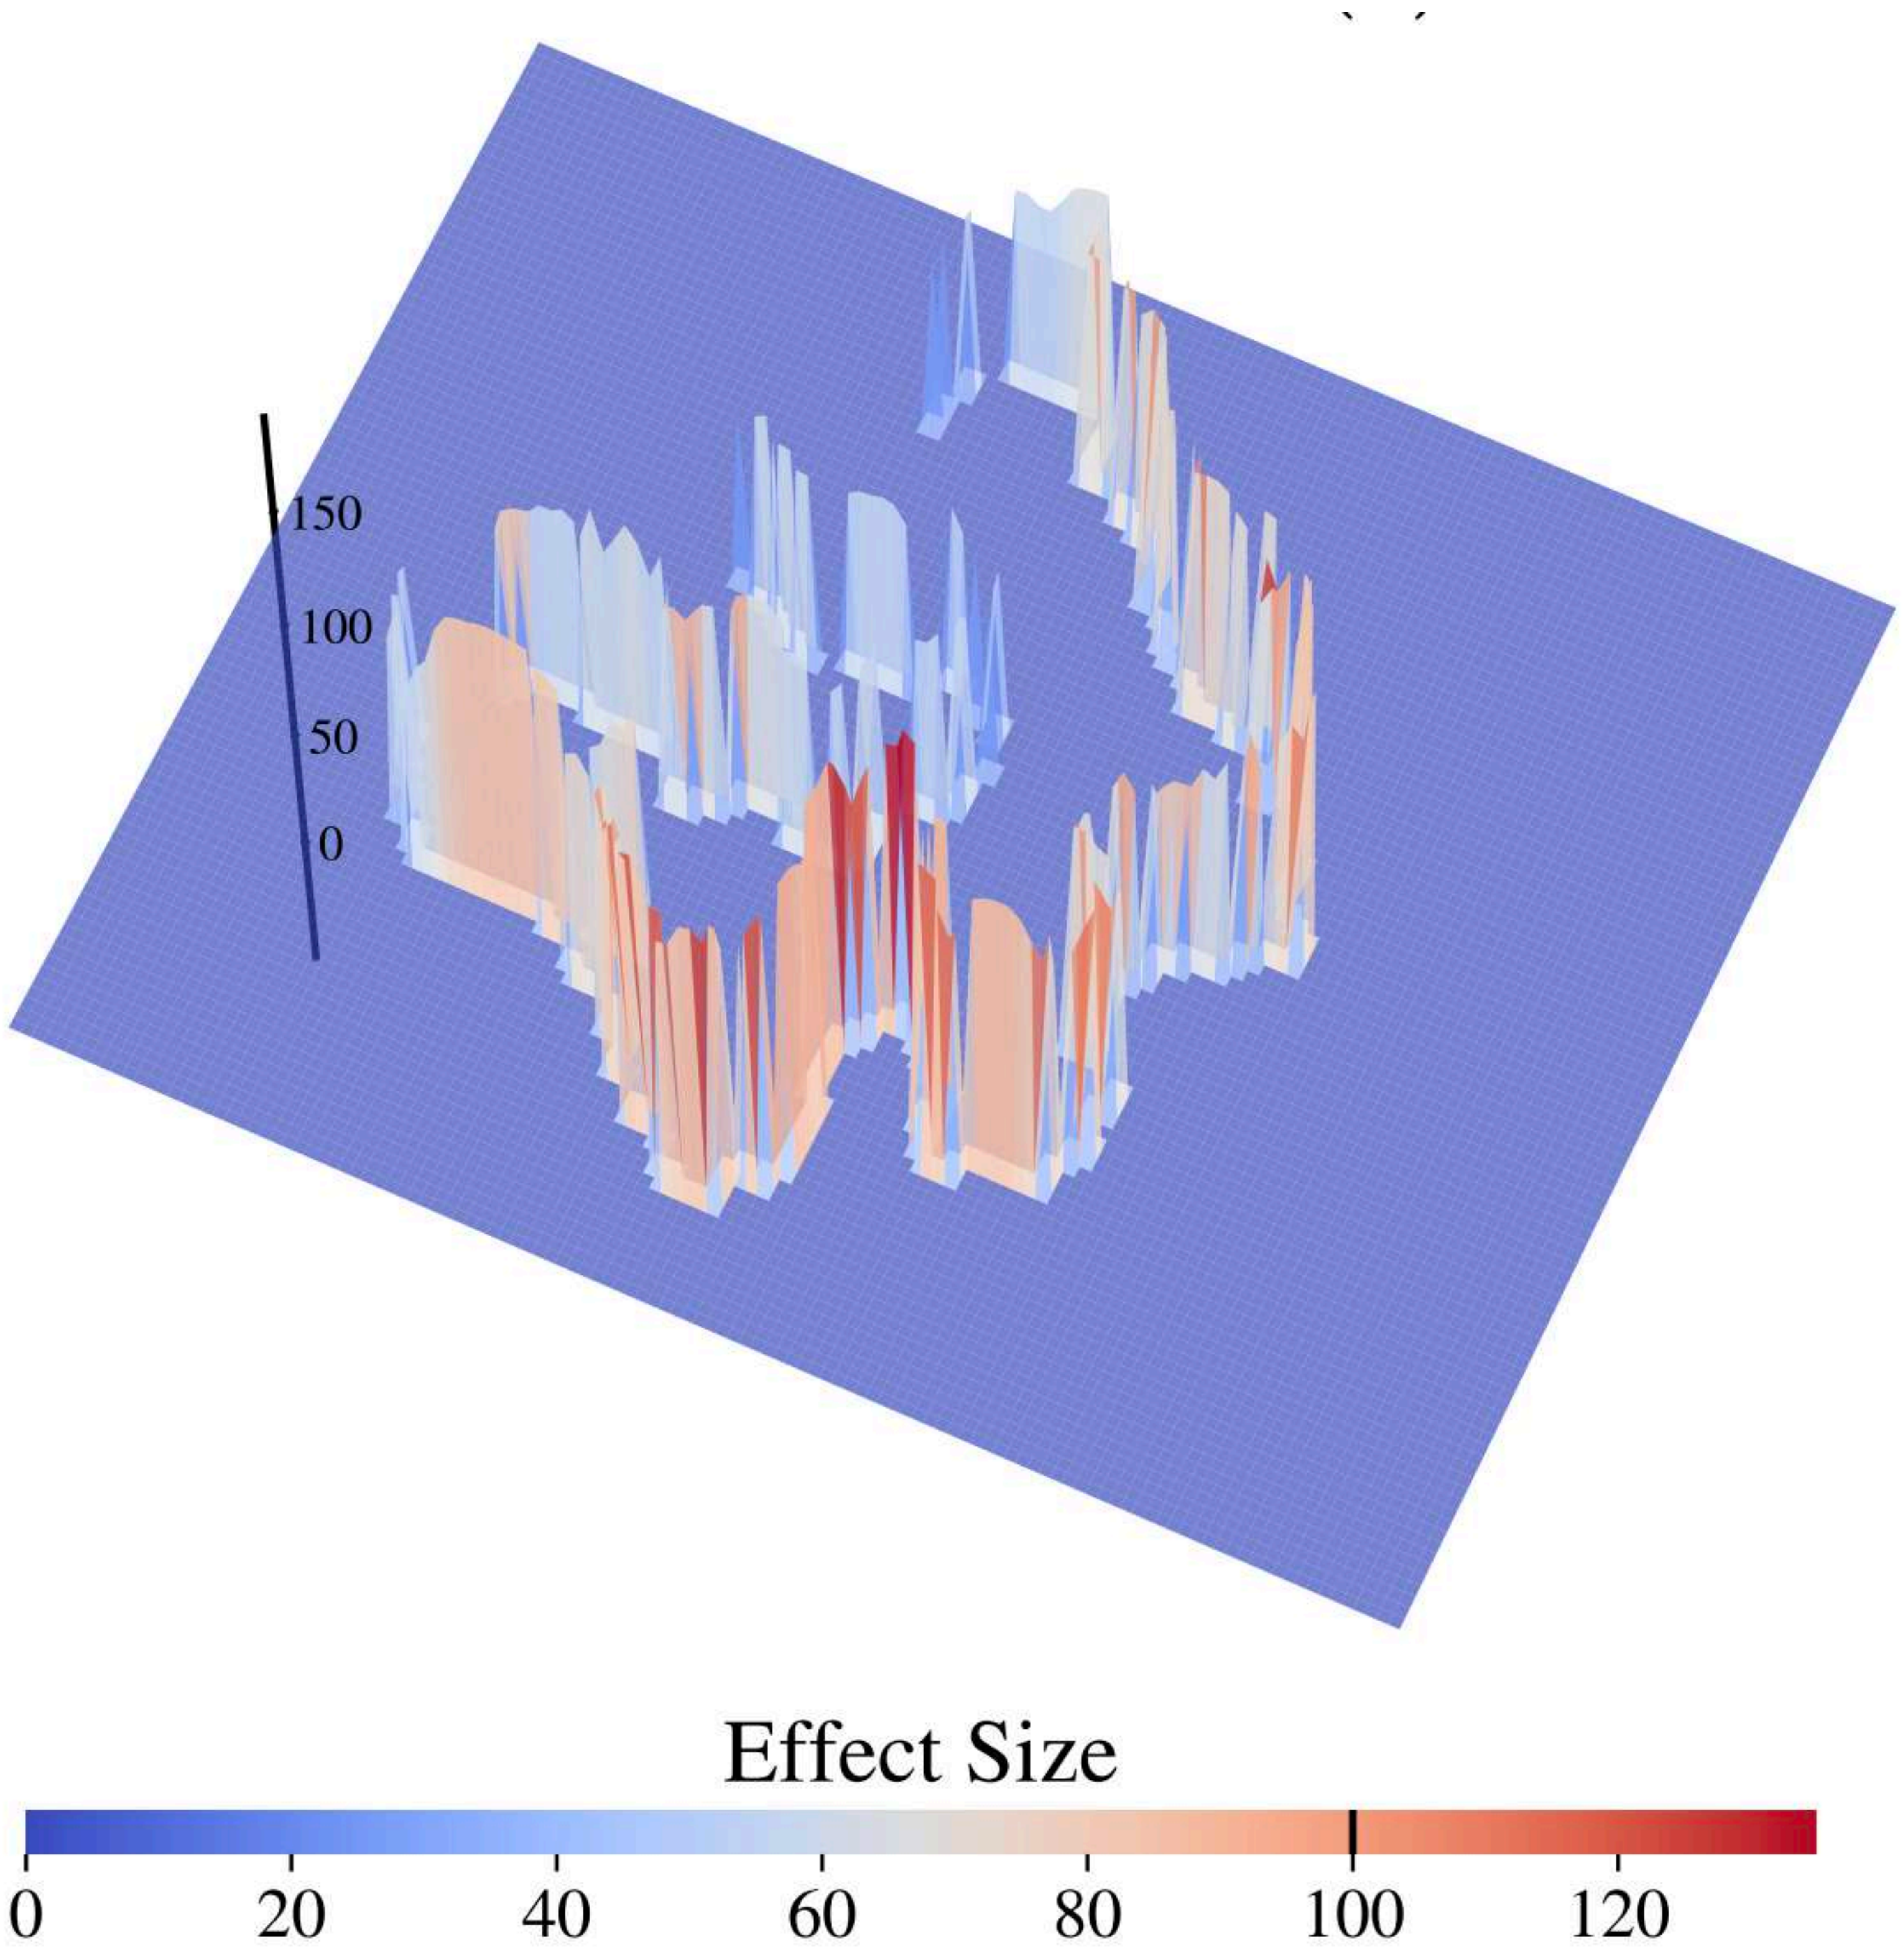

Application: Internet Shutdowns

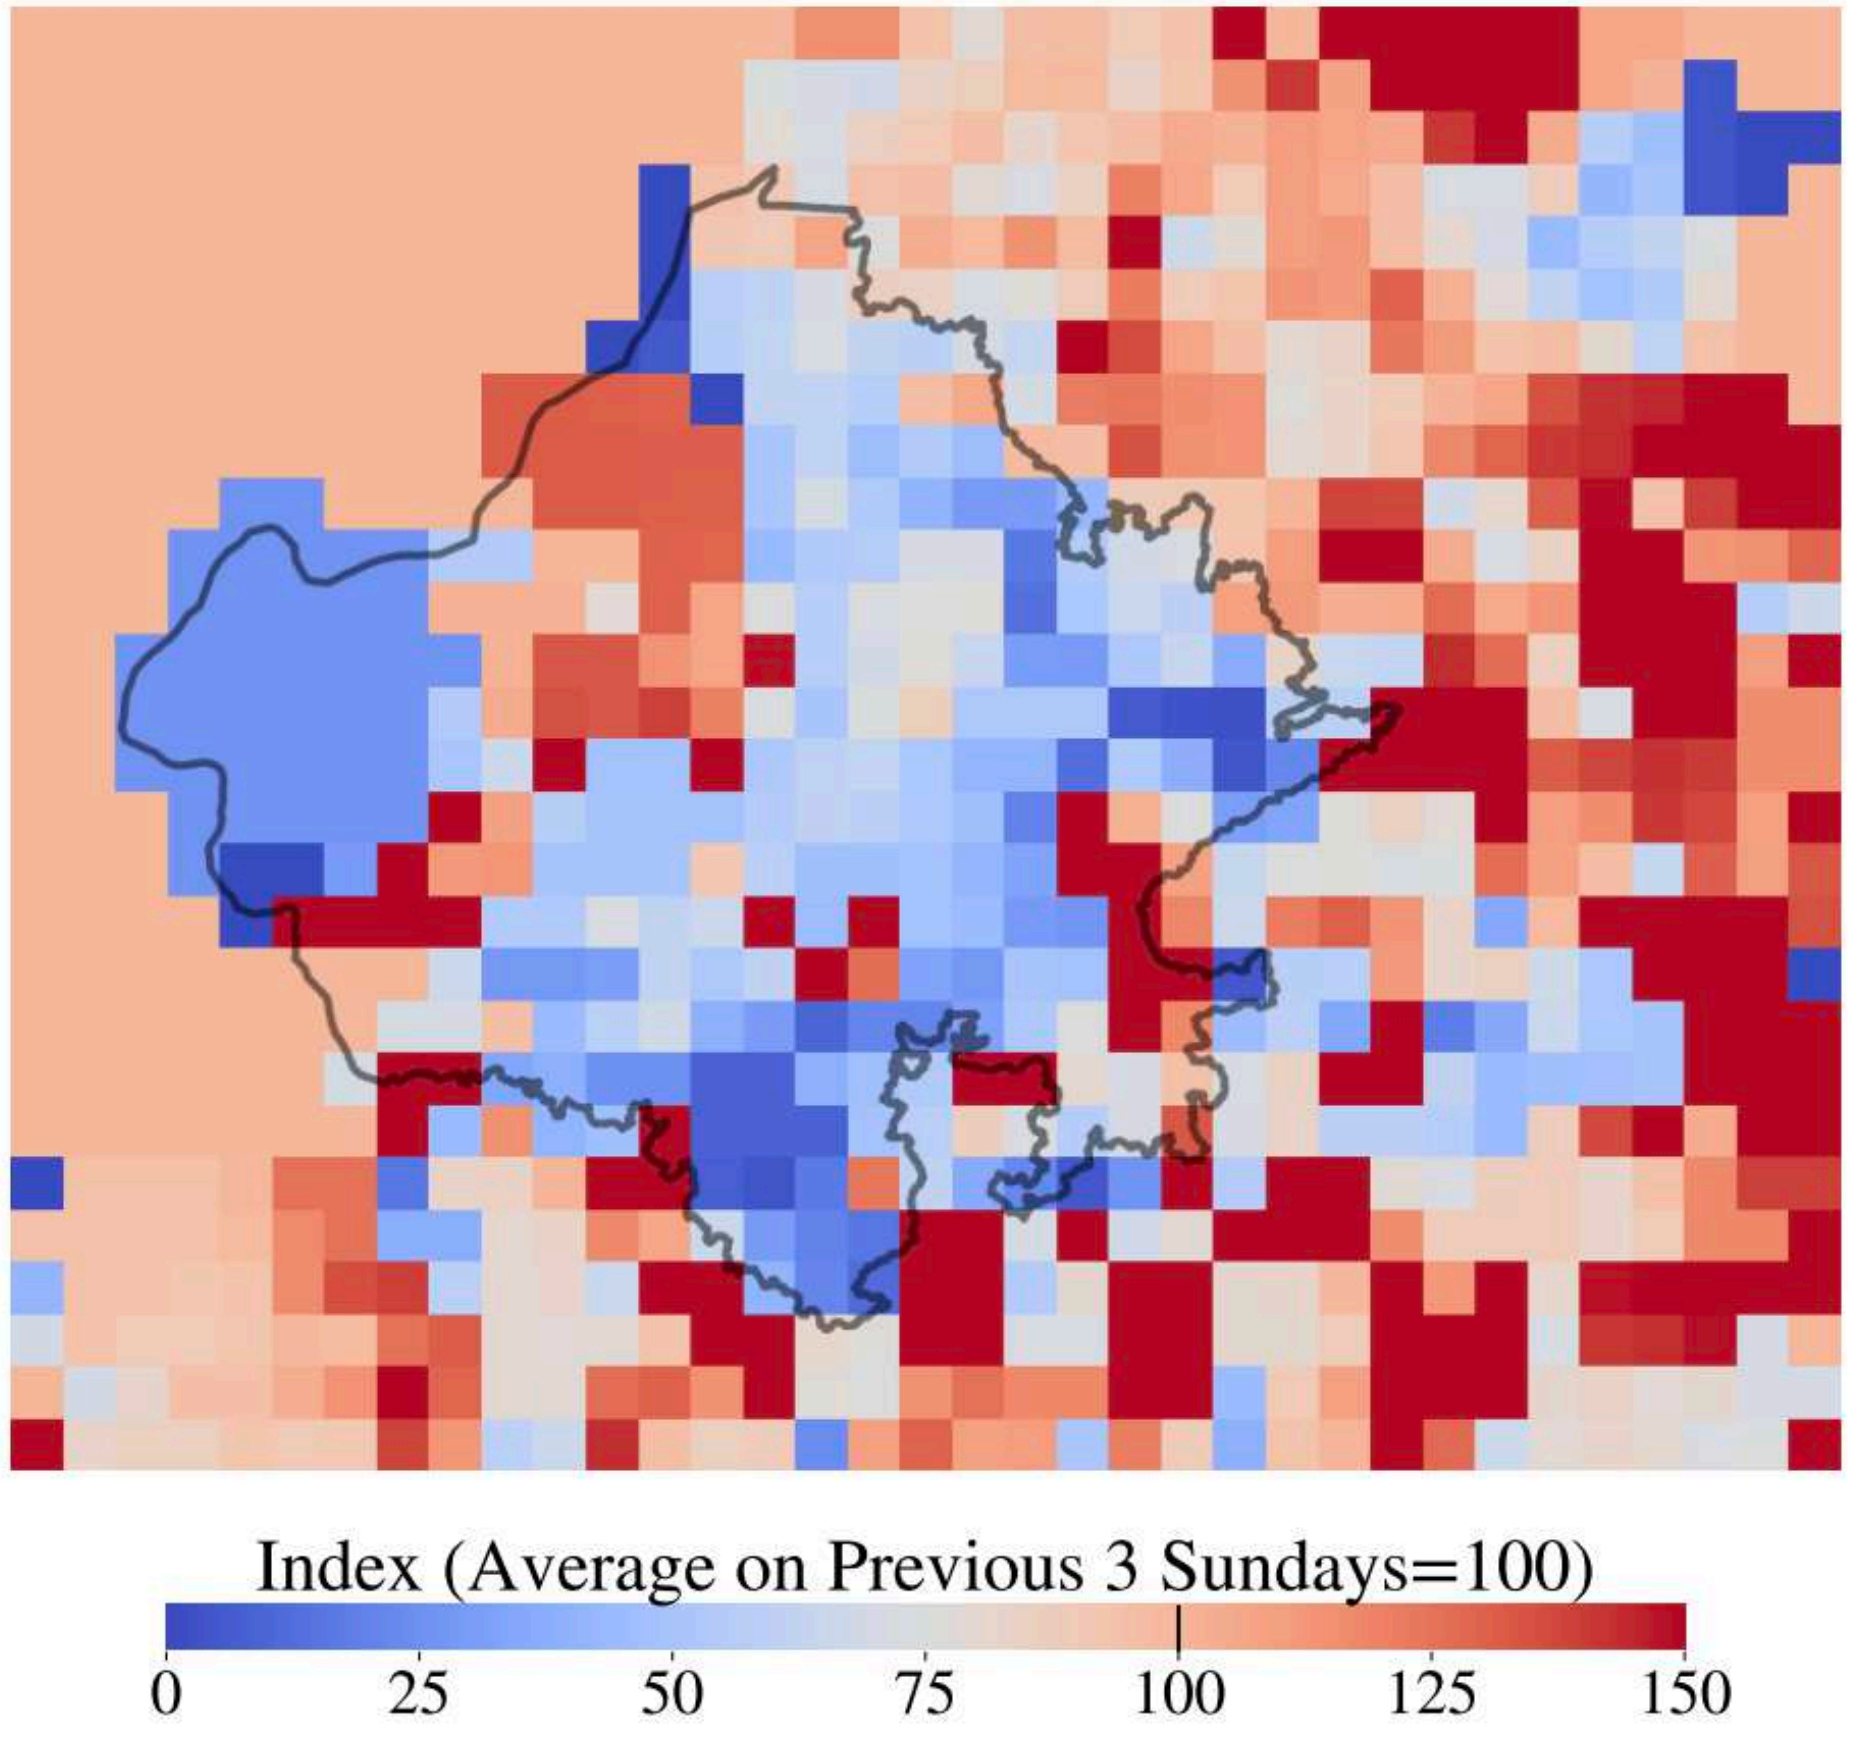

(a) Economic: Raw Data

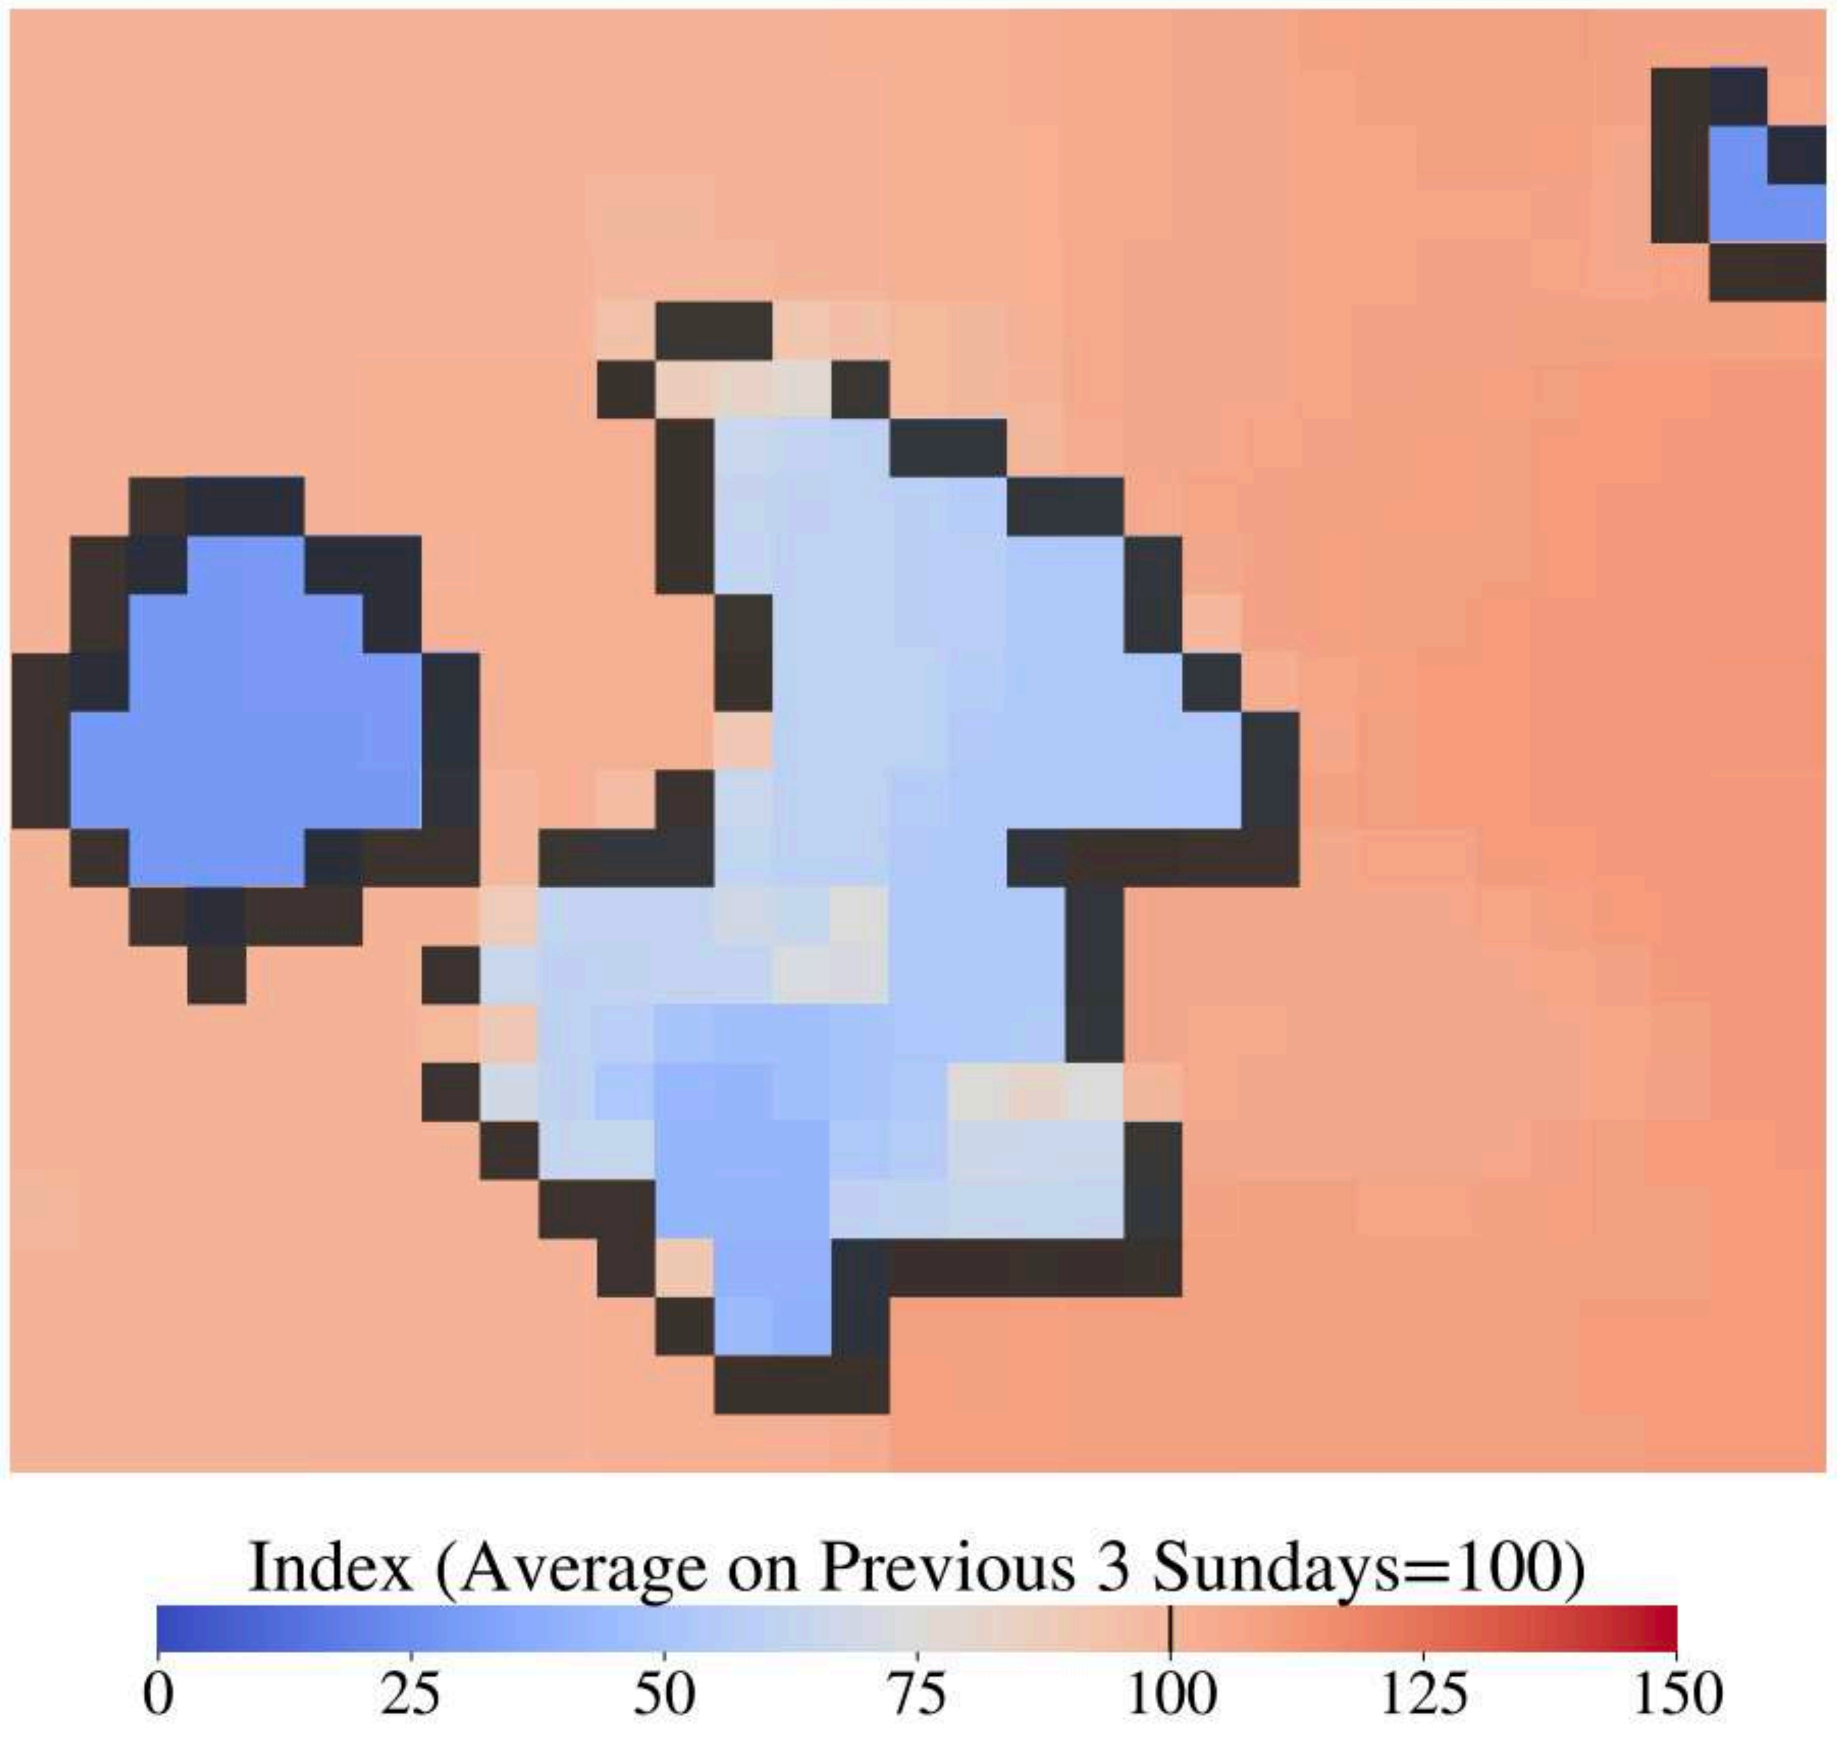

(b) Economic: $\hat{u}$ & S_u

Application: Internet Shutdowns

Economic Activity: Treatment Effect Curve

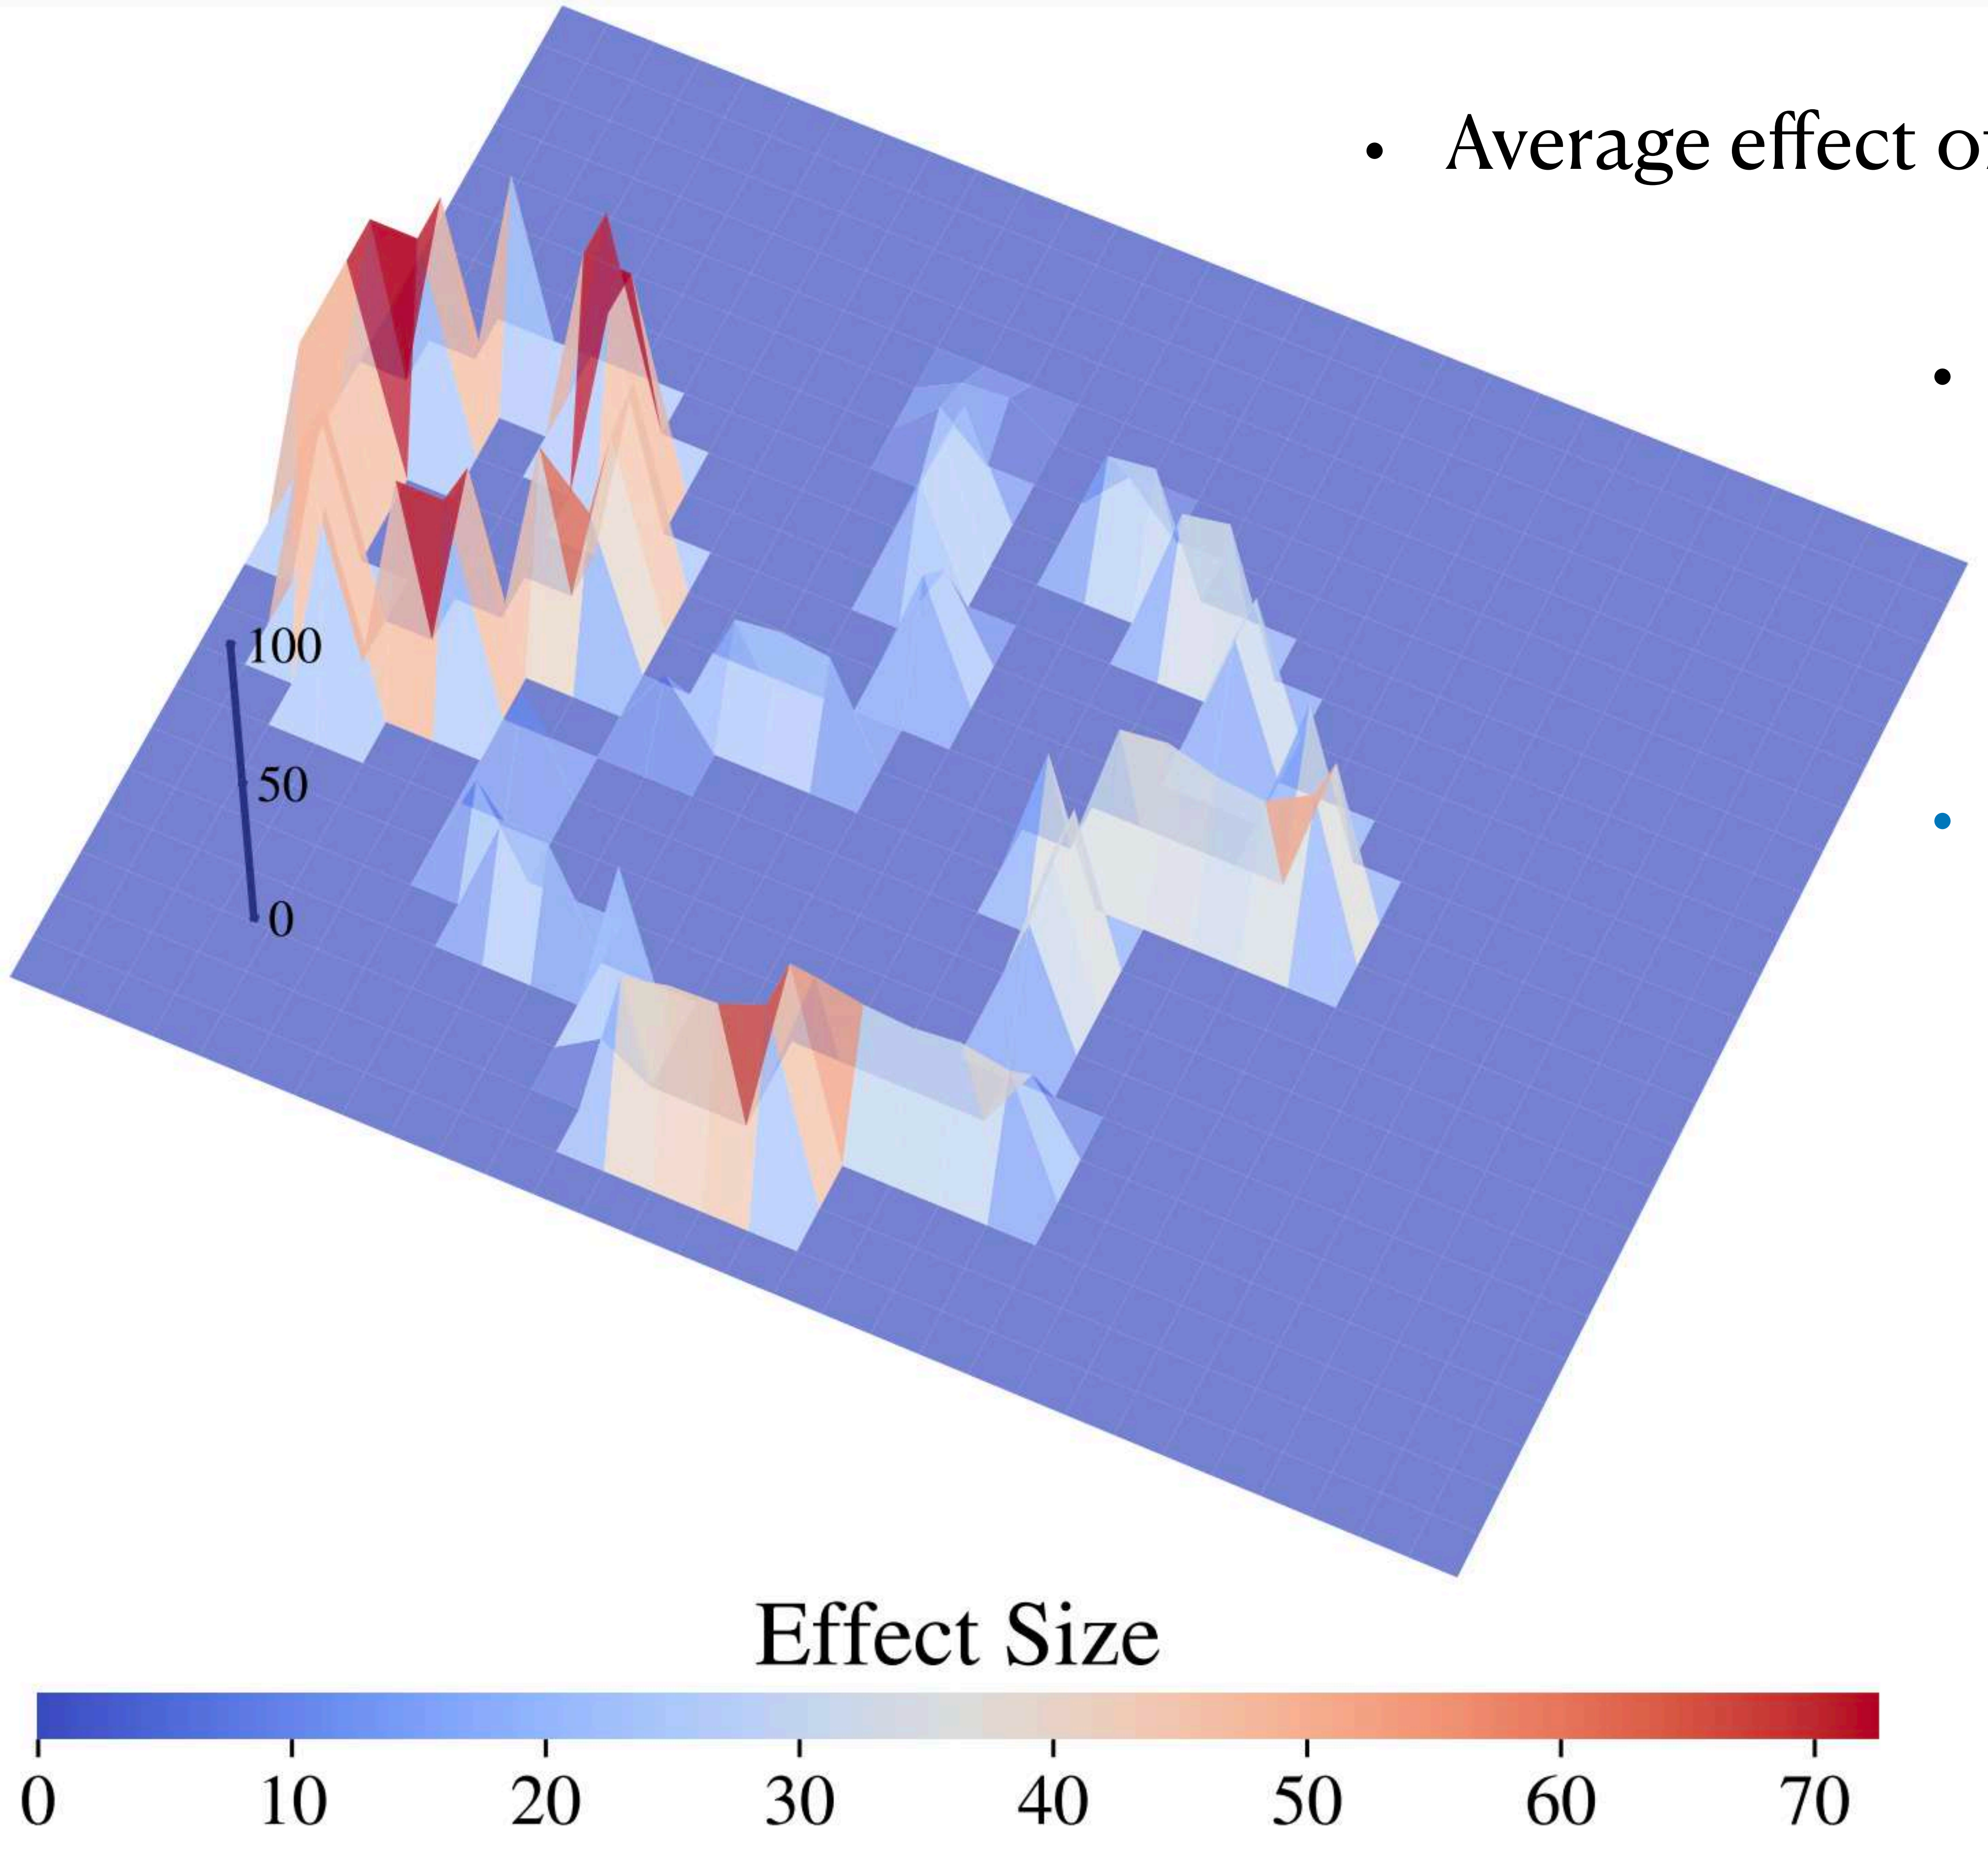

(c) Economic: $\tau_{\text{FD}}(x)$

- Average effect of 35% of monthly average (0-5% baseline increase)
- Correlation of at least 0.7 between mobility data and economic activity (Dong et. Al. 2017):

25%+ drop in economic activity

- 50% higher than highest prevalent estimate!
  - Based on old, single estimate of "digital multiplier"
  - Likely increased drastically since then
  - Digital disruptions can affect non-digitized sections of economy as well through its "enabling role" (Cybersecurity and Infrastructure Security Agency, n.d.)

Conclusion

- **Free discontinuity problem** is a penalized regression estimation problem **with unknown discontinuity in multiple dimensions**
- Connection to **Mumford-Shah functional, convexified via standard lifting approach (calibrations)**
- First mathematical and statistical **identification and consistency** results
- **Hyperparameter choice** via SURE
- Uncertainty quantification via **conformal inference on grid**
- **Application to economic effects of internet shutdowns** finds much stronger **short-term** effects than existing research
- Lots of applications in sciences (several applied papers in progress...)

References

Abdulkadiro˘glu, A., Angrist, J. D., Narita, Y., and Pathak, P. (2022). "Breaking ties: Regression discontinuity design meets market design." *Econometrica*, 90(1), 117–151.
Alvero, A., and Xiao, K. (2023). "Fuzzy bunching."
Ambrus, A., Field, E., and Gonzalez, R. (2020). "Loss in the time of cholera: Long-run impact of a disease epidemic on the urban landscape." *American Economic Review*, 110(2), 475–525.
Andrews, D. W. (1993). "Tests for parameter instability and structural change with unknown change point." *Econometrica: Journal of the Econometric Society*, pp. 821–856.
Argyle, B. S., Nadauld, T. D., and Palmer, C. J. (2020). "Monthly payment shrinkage."
Bai, J., and Perron, P. (2003). "Computation and analysis of multiple structural change models." *Journal of Applied Econometrics*, 18(1), 1–22.
Brunner, E. J., Dougherty, S. M., and Ross, S. L. (2021). "The effects of career and technical education: Evidence from the Connecticut technical high school system." *Review of Economics and Statistics*, pp. 1–46.
Cattaneo, M. D., and Titiunik, R. (2022). "Regression discontinuity designs." *Annual Review of Economics*, 14, 821–851.
Card, D., Mas, A., and Rothstein, J. (2008). "Tipping and the dynamics of segregation." *The Quarterly Journal of Economics*, 123(1), 177–218.
Carneiro, P., Galasso, E., and Ginja, R. (2019). "Tackling social exclusion: Evidence from Chile." *The Economic Journal*, 129(617), 172–208.
Cheng, A. (2023). "Regression discontinuity designs with multiple running variables." MIT working paper.
Delgado, M. A., and Hidalgo, J. (2000). "Nonparametric inference on structural breaks." *Journal of Econometrics*, 96(1), 113–144.
Donoho, D. L. (1999). "Wedgelets: Nearly minimax estimation of edges." *The Annals of Statistics*, 27(3), 859–897.
Donoho, D. L., and Johnstone, I. M. (1994). "Ideal spatial adaptation by wavelet shrinkage." *Biometrika*, 81(3), 425–455.
Hansen, B. E. (2017). "Regression kink with an unknown threshold." *Journal of Business & Economic Statistics*, 35(2), 228–240.
Harchaoui, Z., and Lévy-Leduc, C. (2008). "Kernel change-point analysis." *Advances in Neural Information Processing Systems*, 21.
Harchaoui, Z., and Lévy-Leduc, C. (2010). "Multiple change-point estimation with a total variation penalty." *Journal of the American Statistical Association*, 105(492), 1480–1493.
Hartmann, W. R., Nair, H. S., and Narayanan, S. (2011). "Identifying causal marketing mix effects using a regression discontinuity design." *Marketing Science*, 30(6), 1079–1097.
Herlands, W., McFowland III, E., Wilson, A. G., and Neill, D. B. (2018). "Automated local regression discontinuity design discovery." In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1512–1520.
Herlands, W., Neill, D. B., Nickisch, H., and Wilson, A. G. (2019). "Change surfaces for expressive multidimensional changepoints and counterfactual prediction." *J. Mach. Learn. Res.*, 20, 99–1.
Killick, R., Fearnhead, P., and Eckley, I. A. (2012). "Optimal detection of changepoints with a linear computational cost." *Journal of the American Statistical Association*, 107(500), 1590–1598.
Kern, H. L., and Hainmueller, J. (2009). "Opium for the masses: How foreign media can stabilize authoritarian regimes." *Political Analysis*, 17(4), 377–399.
Kuntz, L., Møller, A., and Østrup, R. (2015). "Mobile phone call data as a regional socio-economic proxy indicator." *PLoS ONE*, 10(4), e0124160.
Li, M., and Ghosal, S. (2017). "Bayesian detection of image boundaries." *The Annals of Statistics*, 45(5), 2190–2217.
Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). "A contextual-bandit approach to personalized news article recommendation." In *Proceedings of the 19th International Conference on World Wide Web*, pp. 661–670.
Madrid Padilla, C. M., et al. (2022). "Change-point detection for sparse and dense functional data in general dimensions." *Advances in Neural Information Processing Systems*, 35, 37121–37133.
Muller, H.-G., and Song, K.-S. (1994). "Maximin estimation of multidimensional boundaries." *Journal of Multivariate Analysis*, 50(2), 265–281.
Narita, Y., and Yata, K. (2021). "Algorithm is experiment: Machine learning market design and policy eligibility rules." arXiv:2104.12909.
O'Sullivan, F., and Qian, M. (1994). "A regularized contrast statistic for object boundary estimation-implementation and statistical evaluation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6), 561–570.
Olken, B. A. (2009). "Do television and radio destroy social capital? Evidence from Indonesian villages." *American Economic Journal: Applied Economics*, 1(4), 1–33.
Page, E. S. (1954). "Continuous inspection schemes." *Biometrika*, 41(1/2), 100–115.
Park, C. (2022). "Jump gaussian process model for estimating piecewise continuous regression functions." *The Journal of Machine Learning Research*, 23(1), 12737–12773.
Porter, J., and Yu, P. (2015). "Regression discontinuity designs with unknown discontinuity points: Testing and estimation." *Journal of Econometrics*, 189(1), 132–147.
Qiu, P. (1997). "Jump detection in regression surfaces." *Journal of Computational and Graphical Statistics*, 6(3), 332–354.
Qiu, P. (1998). "Discontinuous regression surfaces fitting." *The Annals of Statistics*, 26(6), 2218–2245.
Rinaldo, A. (2009). "Properties and refinements of the fused lasso." *The Annals of Statistics*, 37(5B), 2922–2952.
Saez, E. (2010). "Do taxpayers bunch at kink points?" *American Economic Journal: Economic Policy*, 2(3), 180–212.
Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., Van Nes, E. H., Rietkerk, M., and Sugihara, G. (2009). "Early-warning signals for critical transitions." *Nature*, 461(7260), 53–59.
Sonin, K., and Wright, A. L. (2022). "Information operations increase civilian security cooperation." *The Economic Journal*, 132(643), 1179–1199.
Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). "Sparsity and smoothness via the fused lasso." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(1), 91–108.
Zhu, H., Fan, J., and Kong, L. (2014). "Spatially varying coefficient model for neuroimaging data with jump discontinuities." *Journal of the American Statistical Association*, 109(507), 1084–1098.

Appendix

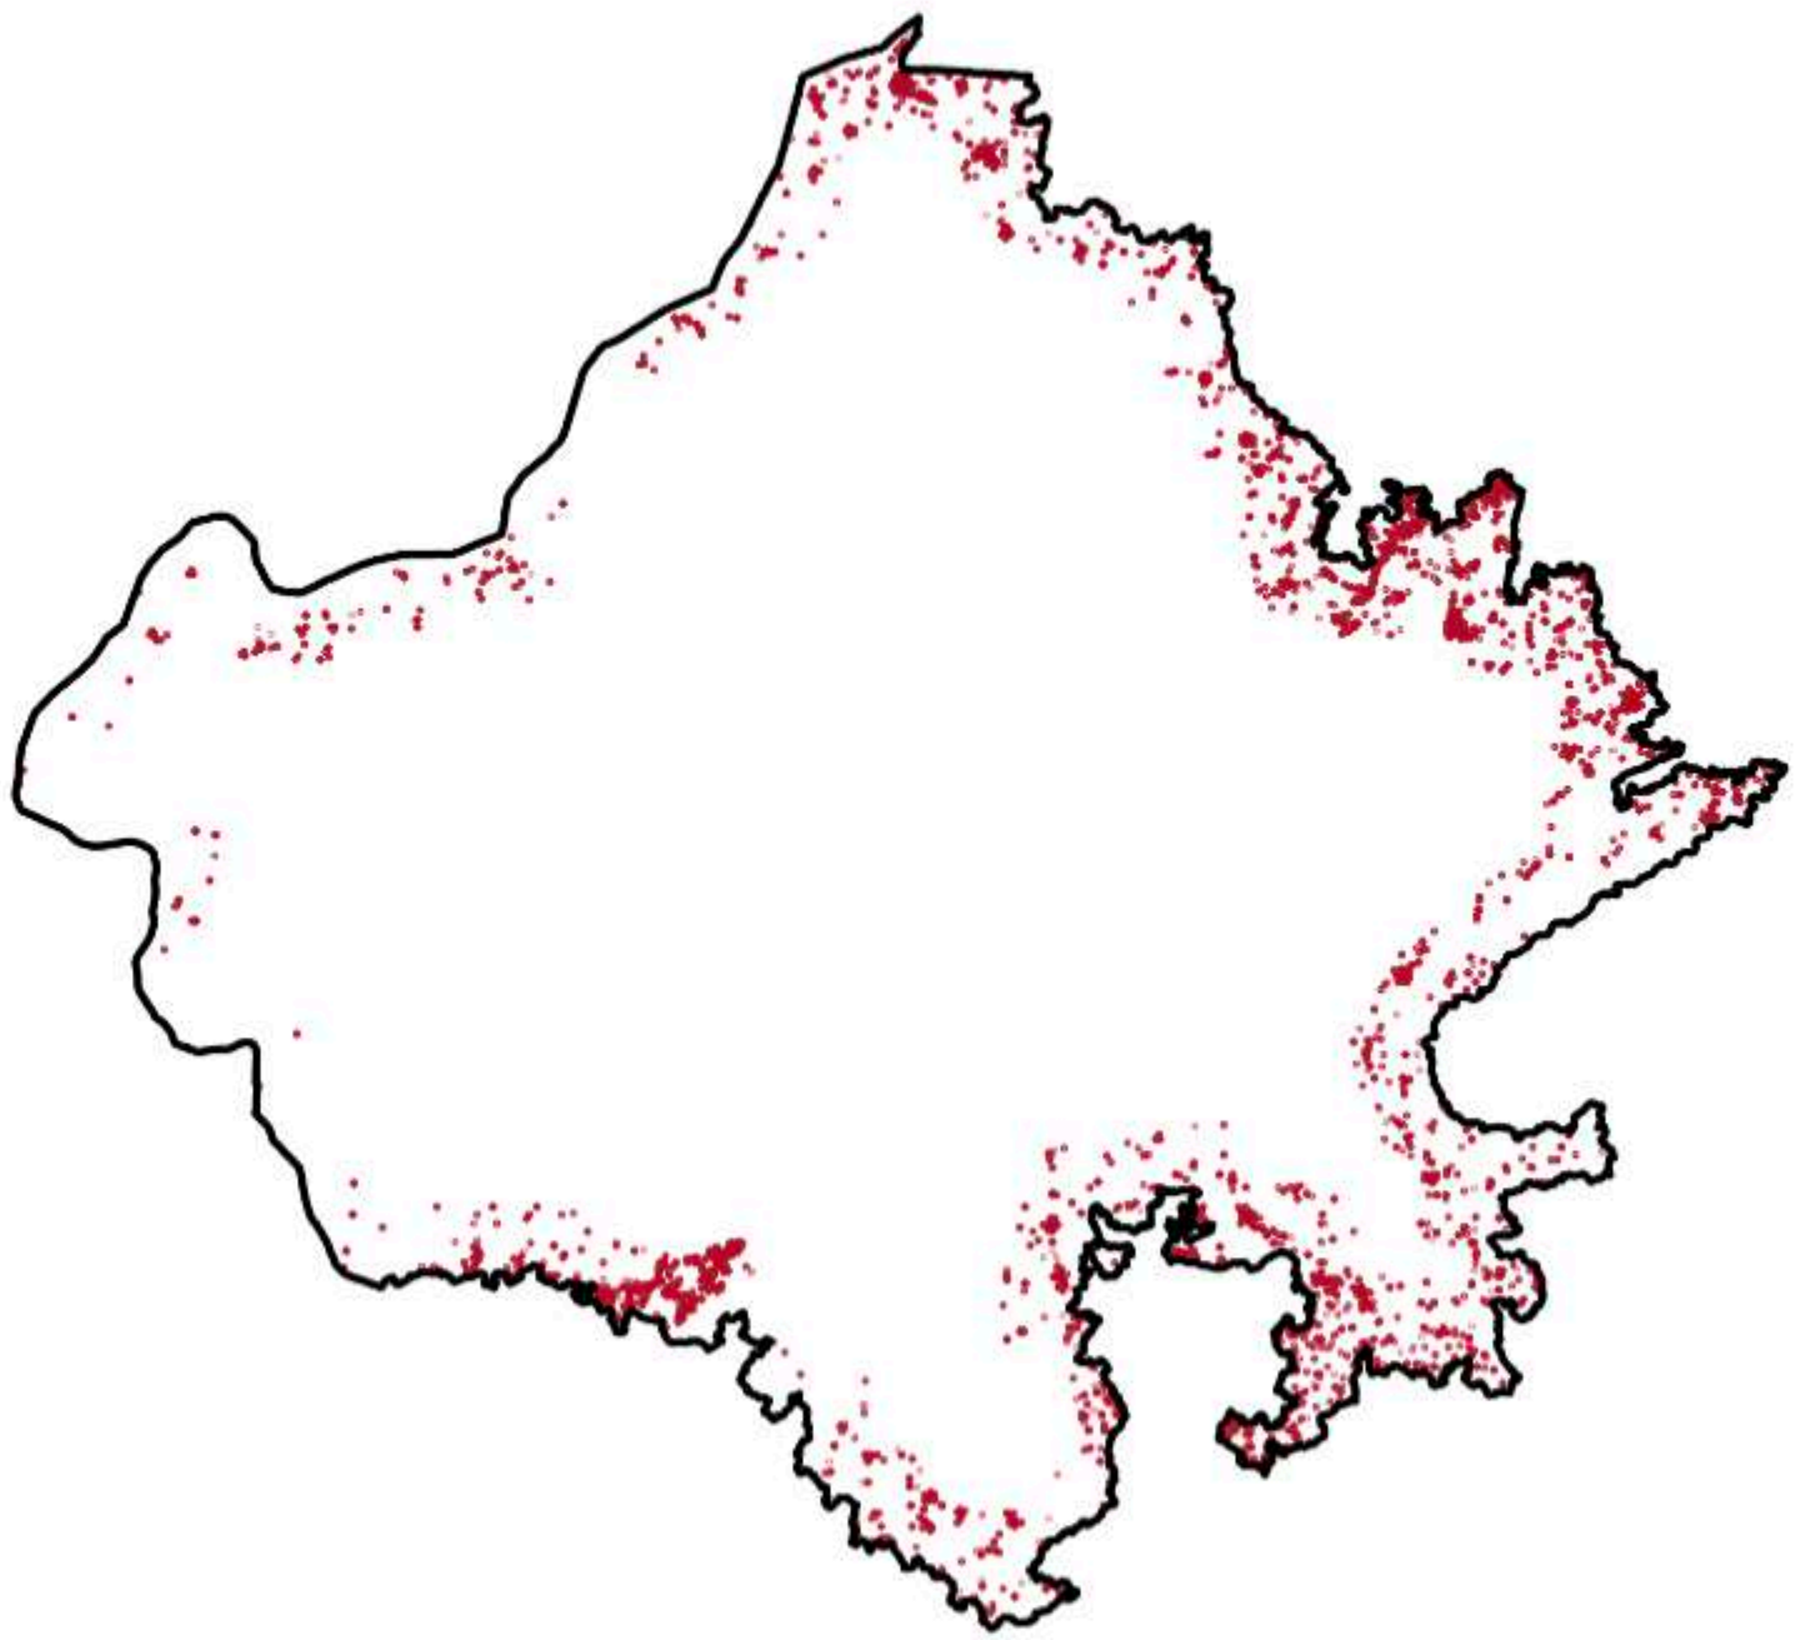

Figure A-1: Catchment Area for Estimating Degree of Selection

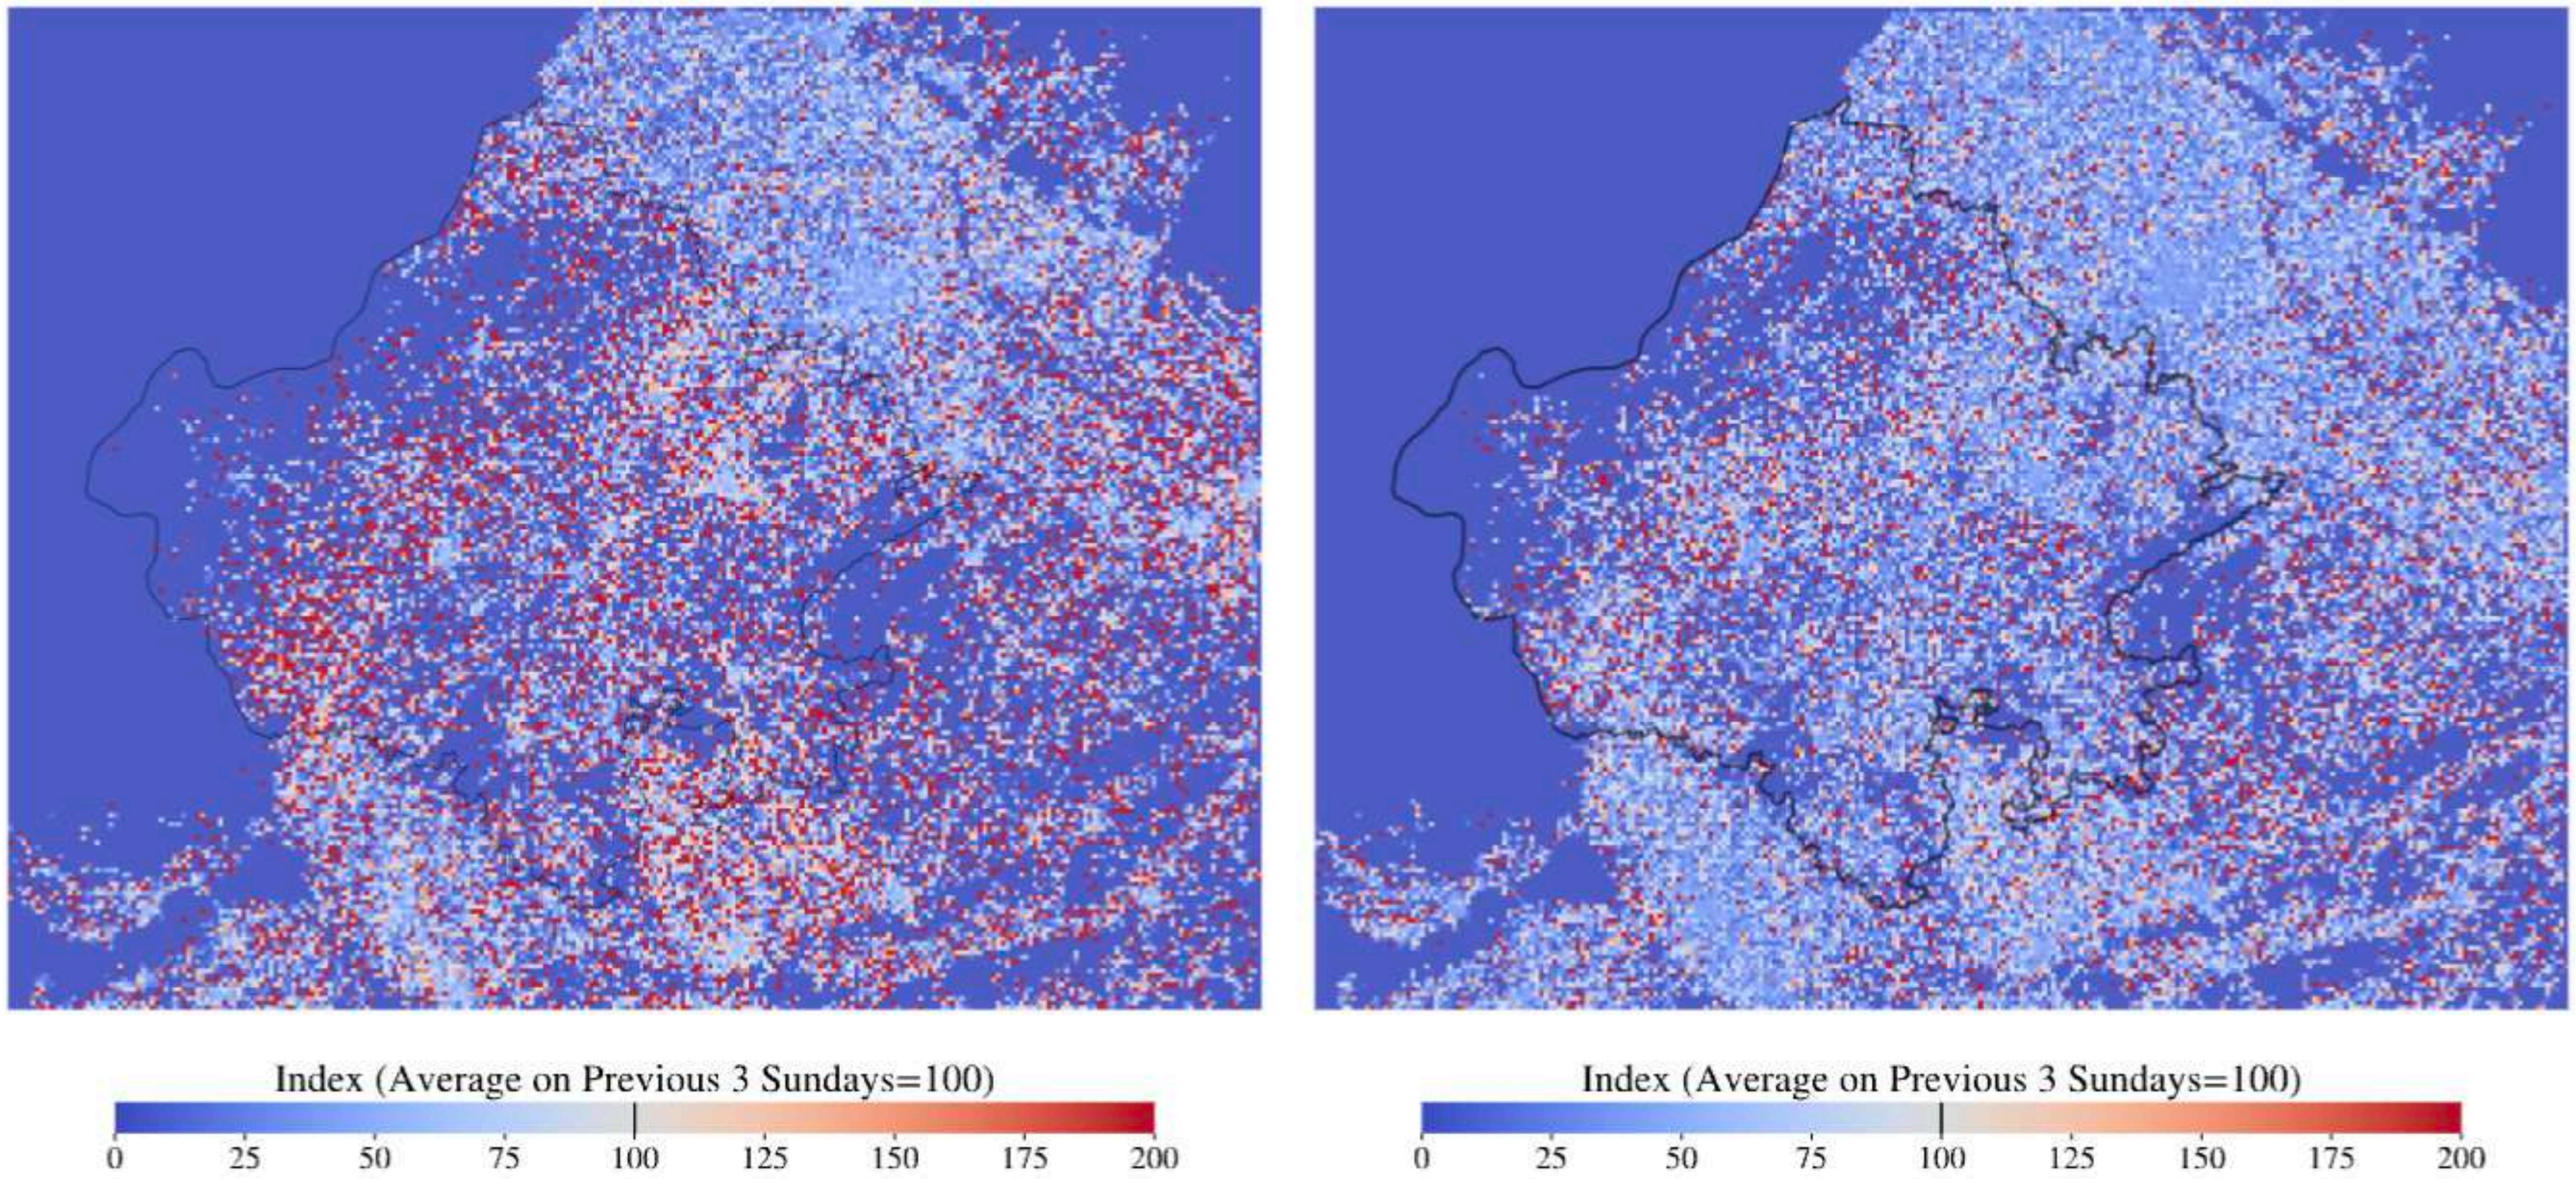

(a) Mobile: Raw Data, 6 pm – midnight

(b) Mobile: Raw Data, Day After, 6 am – 6 pm

Figure A-3: Post-Shutdown Activity

Note: Plot shows the raw mobile device data on a 5x5km grid with the fill color of each cell indicating the value of $\overline{Pings}$, for the hours between 6 pm and midnight on the day of the shutdown in (a) and for the time spanning the shutdown window the day after the shutdown in (b). The outline of Rajasthan state is indicated by black lines.

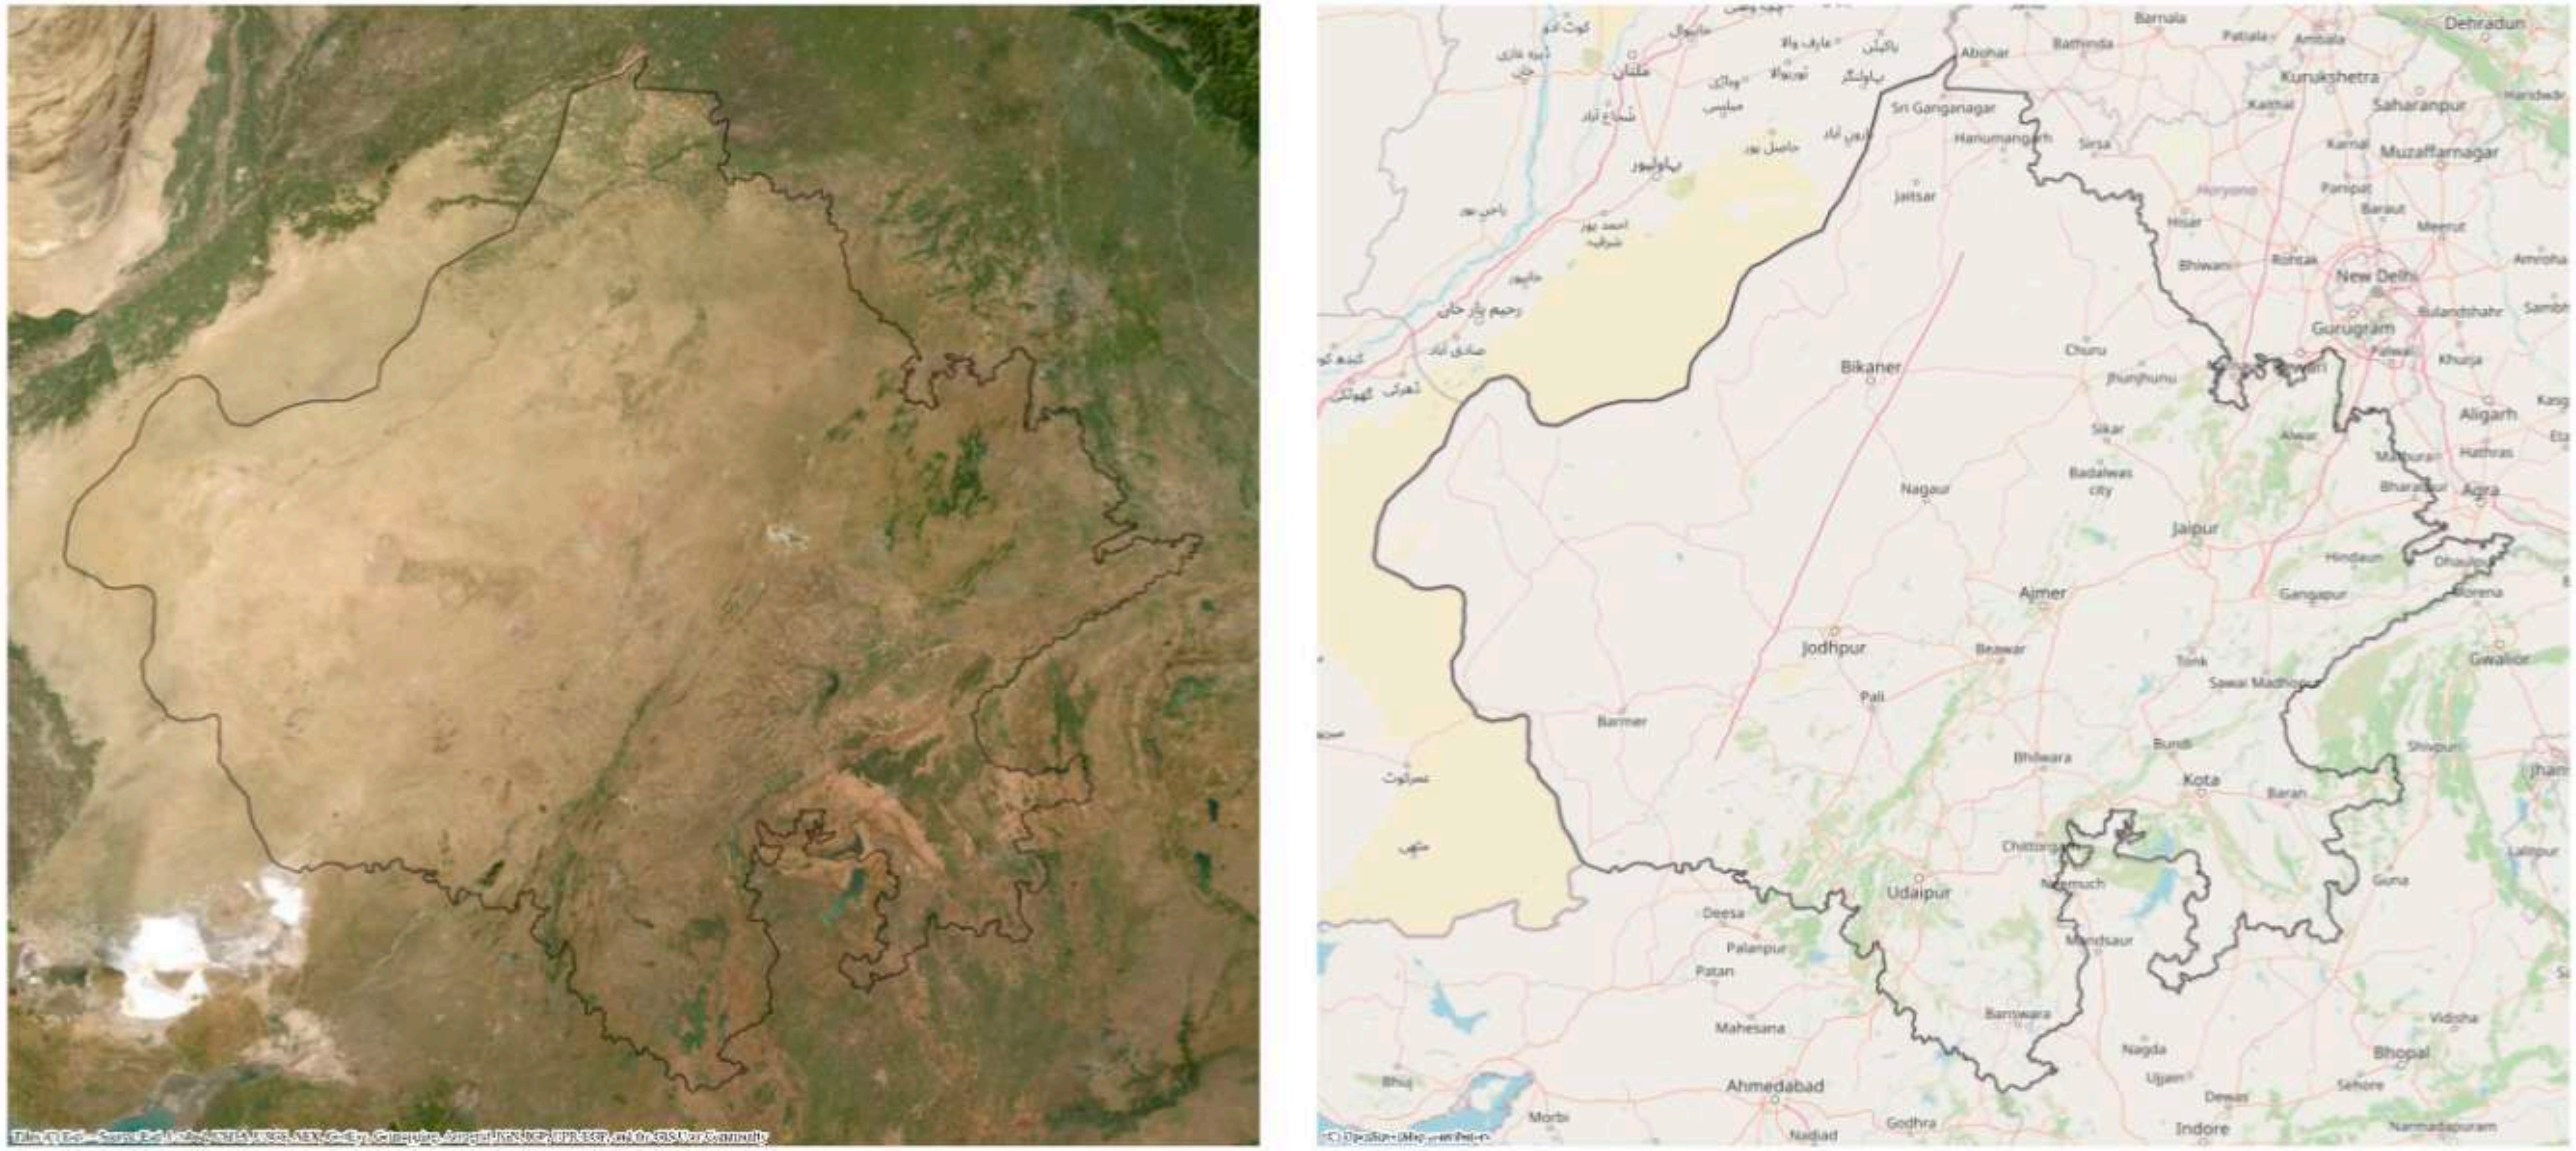

(a) Esri Satellite Imagery

(b) OpenStreetMap

Figure A-4: Rajasthan: Terrain View

Note: A-4a shows the satellite view of Rajasthan, obtained from Esri; A-4b shows the street map of Rajasthan obtained from OpenStreetMap. The outline of Rajasthan is depicted in black.