Package 'spaci'

September 21, 2026

Title Causal Effect Estimation Under Spatial Confounding and Interference

Version 0.1.1

Description Implements the distance-adjusted propensity score with interference (iDAPS) and recoverU+ methods for estimating the average treatment effect on the treated (ATT) from spatial observational data in the presence of both spatial confounding and spatial interference. iDAPS matches units on a data-driven composite of propensity-score distance, spatial proximity and neighbourhood-exposure distance. recoverU+ is a doubly robust estimator that augments the propensity-score and control-outcome models with a partially recovered spatial confounder and a neighbourhood-exposure term. The package also provides the naive propensity score, DAPS and recoverU comparators, and a simulator for the spatial confounding/interference data-generating process.

License MIT + file LICENSE

Encoding UTF-8

Depends R (>= 4.1.0)

Imports stats

Suggests geoR, ggplot2, readxl, testthat (>= 3.0.0), knitr, rmarkdown

Config/testthat/edition 3

VignetteBuilder knitr

URL https://github.com/Ogunsolaia/spaci

BugReports https://github.com/Ogunsolaia/spaci/issues

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author Isqeel Ogunsola [aut, cre],
Olatunji Johnson [aut],
Thomas House [aut]

Maintainer Isqeel Ogunsola <isqeel.ogunsola@postgrad.manchester.ac.uk>

Repository CRAN

Date/Publication 2026-09-21 21:30:02 UTC

Contents

`daps` *Distance-adjusted propensity score (DAPS) matching ATT*

Description

Estimates the ATT by matching on a convex combination of propensity-score distance and spatial (Euclidean) distance, $\alpha D^{PS} + (1-\alpha) D^{Spatial}$, adjusting for spatial confounding but not interference. The weight α is chosen over a grid to minimise covariate imbalance in the matched sample (Papadogeorgou et al., 2019).

Usage

```
daps(
  Y,
  Z,
  X,
  coords,
  caliper = 0.25,
  alpha_grid_step = 0.1,
  seed = NULL,
  level = 0.95
)
```

Arguments

`Y`	Numeric outcome vector.
`Z`	Binary treatment vector (0/1).
`X`	Covariate matrix or data frame.
`coords`	Two-column matrix or data frame of spatial coordinates.
`caliper`	Maximum acceptable matching distance (default `0.25`).

`alpha_grid_step`
: Grid step for the α search over `[0, 1]`.

`seed`
: Optional integer seed for the (randomised) matching order.

`level`
: Confidence level for the reported interval.

Value

An `idaps_fit` object; the selected α is reported in `weights`.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
daps(sim$Y, sim$Z, sim$X, sim$coords)
```

`idaps` *Distance-adjusted propensity score with interference (iDAPS)*

Description

Estimates the ATT by matching on the composite distance of Equation (2.3), $\pi_1 D^{PS} + \pi_2 D^{Spatial} + \pi_3 D^{Interference}$, where the interference component is the absolute difference in neighbourhood exposure. The weights (π_1, π_2, π_3) sum to one and are chosen over a grid to minimise a composite balance score (covariate balance, spatial proximity and exposure balance), rather than tuned by hand. Setting $\pi_3 = 0$ recovers DAPS and $\pi_2 = \pi_3 = 0$ recovers naive PS.

Usage

```
idaps(
  Y,
  Z,
  X,
  coords,
  tau = 0.1,
  caliper = 0.25,
  pi_grid_step = 0.1,
  normalize = TRUE,
  seed = NULL,
  level = 0.95
)
```

Arguments

`Y`
: Numeric outcome vector.

`Z`
: Binary treatment vector (0/1).

`X`
: Covariate matrix or data frame.

`coords`
: Two-column matrix or data frame of spatial coordinates.

tau	Positive spatial decay parameter of the exposure kernel (default 0.1); see neighbourhood_exposure().
caliper	Maximum acceptable matching distance (default 0.25).
pi_grid_step	Grid step for the weight search over the simplex.
normalize	Logical; row-normalise the exposure kernel (default TRUE).
seed	Optional integer seed for the (randomised) matching order.
level	Confidence level for the reported interval.

Value

An idaps_fit object; the selected weights are reported in weights.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
idaps(sim$Y, sim$Z, sim$X, sim$coords, tau = 0.1)
```

---

naive_ps *Naive propensity-score matching ATT*

Description

Estimates the ATT by greedy 1:1 matching on the (normalised) difference in estimated propensity scores. This is the baseline that ignores both spatial confounding and spatial interference.

Usage

```
naive_ps(Y, Z, X, caliper = 0.25, seed = NULL, level = 0.95)
```

Arguments

Y	Numeric outcome vector.
Z	Binary treatment vector (0/1).
X	Covariate matrix or data frame.
caliper	Maximum acceptable matching distance (default 0.25).
seed	Optional integer seed for the (randomised) matching order.
level	Confidence level for the reported interval.

Value

An idaps_fit object.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
naive_ps(sim$Y, sim$Z, sim$X)
```

`neighbourhood_exposure` *Neighbourhood exposure via a spatial kernel*

Description

Constructs the neighbourhood-exposure mapping $E_i = \sum_{j \neq i} G_{ij} A_j$ that summarises spatial interference, where G_{ij} is an exponential kernel of the distance between units i and j, $G_{ij} = \exp(-d_{ij}/\tau)$ with $G_{ii} = 0$. This is the exposure used by `idaps()` and the recoverU family.

Usage

```
neighbourhood_exposure(coords, Z, tau = 0.1, normalize = TRUE)
```

Arguments

`coords`	A two-column matrix or data frame of spatial coordinates.
`Z`	Binary treatment vector (0/1), one entry per row of `coords`.
`tau`	Positive spatial decay (bandwidth) parameter of the exponential kernel. Smaller values concentrate exposure on nearer neighbours.
`normalize`	Logical; if `TRUE` (default) each row of the kernel is normalised to sum to one, so that exposure is a weighted average of neighbours' treatment. If `FALSE`, the raw kernel weights are used, matching the unnormalised definition in Equation (2.4) of the report.

Value

A list with components

E Numeric vector of neighbourhood exposures, one per unit.

G The (possibly row-normalised) kernel matrix with zero diagonal.

Examples

```
set.seed(1)
coords <- cbind(runif(20), runif(20))
Z <- rbinom(20, 1, 0.5)
ex <- neighbourhood_exposure(coords, Z, tau = 0.1)
head(ex$E)
```

plot.idaps_fit *Plot a single estimate*

Description

Plot a single estimate

Usage

```
## S3 method for class 'idaps_fit'
plot(x, ...)
```

Arguments

x	An idaps_fit object.
...	Passed to plot_ate().

Value

The plotted data frame, invisibly.

plot_ate *Forest plot of ATT estimates across methods*

Description

Draws a forest (caterpillar) plot of the estimated ATT and its confidence interval for each method, with a reference line at zero and, optionally, at the true effect. This reproduces the style of Figure 2.5 of the report and is the recommended way to compare methods visually.

Usage

```
plot_ate(
  x,
  true_att = NULL,
  null_line = 0,
  col = "#4C72B0",
  pt_col = "#1A1A1A",
  xlab = "Average treatment effect on the treated (ATT)",
  main = "Estimated ATT and confidence intervals",
  ...
)
```

Arguments

x	Either the data frame returned by spatial_ate() (columns Method, ATT, Lower, Upper) or a named list of idaps_fit objects.
true_att	Optional true effect; drawn as a dashed vertical reference line when supplied.
null_line	Position of the "no effect" reference line (default 0); set to NA to omit.
col, pt_col	Colours for the confidence intervals and point estimates.
xlab, main	Axis label and title.
...	Passed to graphics::plot().

Value

The plotted data frame, invisibly.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
res <- spatial_ate(sim$Y, sim$Z, sim$X, sim$coords, seed = 1)
plot_ate(res, true_att = sim$true_att)
```

---

print.idaps_fit *Print method for idaps fit objects*

Description

Print method for idaps fit objects

Usage

```
## S3 method for class 'idaps_fit'
print(x, digits = 4, ...)
```

Arguments

x	An idaps_fit object.
digits	Number of significant digits.
...	Ignored.

Value

x, invisibly.

recoverU *recoverU: doubly robust ATT with a recovered spatial confounder*

Description

Doubly robust ATT estimator that adjusts for spatial confounding through a partially recovered spatial confounder (Pokal et al., 2023), but does not adjust the propensity-score / control-outcome models for spatial interference. Provided as a comparator to `recoverUplus()`.

Usage

```
recoverU(
  Y,
  Z,
  X,
  coords,
  tau = 0.1,
  normalize = TRUE,
  matern_method = c("mle", "geoR"),
  level = 0.95
)
```

Arguments

`Y`	Numeric outcome vector.
`Z`	Binary treatment vector (0/1).
`X`	Covariate matrix or data frame.
`coords`	Two-column matrix or data frame of coordinates.
`tau`	Positive spatial decay parameter for the exposure kernel used in the initial outcome model (default `0.1`).
`normalize`	Logical; row-normalise the exposure kernel (default `TRUE`).
`matern_method`	Matern estimation engine: `"mle"` (default, no external dependency) or `"geoR"` (reproduces the reference analysis).
`level`	Confidence level for the reported interval.

Value

An `idaps_fit` object.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
recoverU(sim$Y, sim$Z, sim$X, sim$coords)
```

`recoverUplus` *recoverU+: doubly robust ATT under spatial confounding and interference*

Description

The recoverU+ estimator augments the doubly robust ATT with both a partially recovered spatial confounder and a neighbourhood-exposure term, so that the propensity-score and control-outcome models adjust for spatial confounding and spatial interference simultaneously (Equations 2.10-2.12). Compared with `recoverU()`, the neighbourhood exposure `G` is included in the PS and control-outcome models.

Usage

```
recoverUplus(
  Y,
  Z,
  X,
  coords,
  tau = 0.1,
  normalize = TRUE,
  matern_method = c("mle", "geoR"),
  level = 0.95
)
```

Arguments

`Y`	Numeric outcome vector.
`Z`	Binary treatment vector (0/1).
`X`	Covariate matrix or data frame.
`coords`	Two-column matrix or data frame of coordinates.
`tau`	Positive spatial decay parameter for the exposure kernel used in the initial outcome model (default `0.1`).
`normalize`	Logical; row-normalise the exposure kernel (default `TRUE`).
`matern_method`	Matern estimation engine: `"mle"` (default, no external dependency) or `"geoR"` (reproduces the reference analysis).
`level`	Confidence level for the reported interval.

Value

An `idaps_fit` object.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
recoverUplus(sim$Y, sim$Z, sim$X, sim$coords)
```

`simulate_spatial_causal` *Simulate spatial data with confounding and interference*

Description

Draws one data set from the data-generating process of the report (Equation 2.13 and Section 2.3): an unmeasured spatial confounder $U(s)$ drawn from an exponential Gaussian random field drives both treatment and outcome, and the outcome further depends on a neighbourhood-exposure term so that spatial interference is present. The true ATT is `true_att`.

Usage

```
simulate_spatial_causal(
  n = 150,
  true_att = 2,
  beta0 = 2.5,
  beta1 = 1,
  beta2 = 0.5,
  theta_spatial = 0.4,
  gamma_interference = 1.5,
  sigma_eps = 1,
  delta_u = 2,
  u_phi = 0.2,
  tau_exp = 0.1,
  normalize = TRUE,
  seed = NULL
)
```

Arguments

`n`	Number of units.
`true_att`	True average treatment effect on the treated.
`beta0`, `beta1`, `beta2`	Outcome-model intercept and covariate coefficients.
`theta_spatial`	Coefficient of the spatial confounder in the outcome.
`gamma_interference`	Coefficient of neighbourhood exposure in the outcome (interference strength).
`sigma_eps`	Outcome noise standard deviation.
`delta_u`	Strength of the spatial confounder in the treatment model (controls the degree of spatial confounding).
`u_phi`	Range parameter of the confounder's exponential covariance.
`tau_exp`	Spatial decay parameter of the exposure kernel.
`normalize`	Logical; row-normalise the exposure kernel (default `TRUE`).
`seed`	Optional integer seed.

Details

The Gaussian random field is generated directly from its covariance matrix (exponential covariance, i.e. Matern with smoothness 1/2), so the simulator has no external dependency.

Value

A list with `Y`, `Z`, `X` (matrix with columns `X1`, `X2`), `coords`, the exposure `E`, the latent confounder `U`, and `true_att`.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 42)
str(sim)
```

---

spatial_ate *Estimate the ATT with all available methods*

Description

Convenience wrapper that runs the naive propensity score, DAPS, iDAPS, recoverU and recoverU+ estimators on the same data and collects the point estimates, standard errors and confidence intervals in a single data frame.

Usage

```
spatial_ate(
  Y,
  Z,
  X,
  coords,
  tau = 0.1,
  caliper = 0.25,
  matern_method = c("mle", "geoR"),
  level = 0.95,
  seed = NULL,
  methods = c("Naive PS", "DAPS", "iDAPS", "recoverU", "recoverU+")
)
```

Arguments

`Y`	Numeric outcome vector.
`Z`	Binary treatment vector (0/1).
`X`	Covariate matrix or data frame.
`coords`	Two-column matrix or data frame of spatial coordinates.
`tau`	Positive spatial decay parameter of the exposure kernel (default `0.1`); see `neighbourhood_exposure()`.
`caliper`	Maximum acceptable matching distance (default `0.25`).

matern_method	Matern estimation engine for the recoverU family, "mle" (default) or "geoR".
level	Confidence level for the reported interval.
seed	Optional integer seed for the (randomised) matching order.
methods	Character vector selecting which estimators to run. Defaults to all five.

Value

A data frame with one row per method (Method, ATT, SE, Lower, Upper), with the fitted objects attached as the "fits" attribute.

Examples

```
sim <- simulate_spatial_causal(n = 50, seed = 1)
spatial_ate(sim$Y, sim$Z, sim$X, sim$coords, seed = 1)
```

Index