Optimal Policy Learning Under Spatial Dependence

With Applications to Groundwater in Wisconsin

Xindi Lin

Joint work with Chan Park, Christopher Zahasky, Hyunseung Kang
University of Wisconsin-Madison

BLAST Seminar
September, 2025

Introduction

- The U.S. Environmental Protection Agency (EPA) sets health-based enforcement standards for many environmental contaminants (e.g. 9 μg/m^3 for PM 2.5, 0.014 μg/L for PFAS)

Introduction

- The U.S. Environmental Protection Agency (EPA) sets health-based enforcement standards for many environmental contaminants (e.g. 9 $\mu g/m^3$ for PM 2.5, 0.014 $\mu g/L$ for PFAS)
- The primary goal of mitigation policies is to reduce contaminant concentrations to levels below these enforcement standards.
- Using nitrate, a common groundwater contaminant, as an example, we introduce a method for optimal policy learning in environmental public health.

Outline

Spatial Policy Learning: Motivation

- **Nitrate in Groundwater**
  - Nitrate (NO_3^-) has been one of the top three contaminant in public drinking water supplies for over two decades (Pennino et al. 2020).
  - Nitrate is originally from nitrogen (N) from surfance land. As nitrogen converts to nitrate, it dissolves in water and seeps into groundwater by rain.

Spatial Policy Learning: Motivation

- **Nitrate in Groundwater**
  - Nitrate (NO_3^-) has been one of the top three contaminant in public drinking water supplies for over two decades (Pennino et al. 2020).
  - Nitrate is originally from nitrogen (N) from surfance land. As nitrogen converts to nitrate, it dissolves in water and seeps into groundwater by rain.
- **Negative Effect of Nitrate**
  - Low level of nitrate is naturally found in food and environment.
  - High level of nitrate poses serious health risks to birth defects in infants, and several cancers in adults.

Spatial Policy Learning: Motivation

- **Nitrate in Groundwater**
  - Nitrate (NO_3^-) has been one of the top three contaminant in public drinking water supplies for over two decades (Pennino et al. 2020).
  - Nitrate is originally from nitrogen (N) from surfance land. As nitrogen converts to nitrate, it dissolves in water and seeps into groundwater by rain.
- **Negative Effect of Nitrate**
  - Low level of nitrate is naturally found in food and environment.
  - High level of nitrate poses serious health risks to birth defects in infants, and several cancers in adults.
  - U.S. Environmental Protection Agency (EPA) set the enforcement standard of nitrate to be 10mg/L.

Spatial Policy Learning: Motivation

- **Nitrate Regulation in Wisconsin** (Wisconsin Groundwater Report, DNR 2024)
  - More than 240 public water supply system and more than 10 percent of private wells have nitrate level exceeding 10mg/L.
  - By state law, drinking water wells must be replaced if nitrate levels exceed the *enforcement standard* of 10mg/L.
  - The total estimated cost for replacing non-compliant wells is over *$466 million*.

Spatial Policy Learning: Motivation

- **Nitrate Regulation in Wisconsin** (Wisconsin Groundwater Report, DNR 2024)
  - More than 240 public water supply system and more than 10 percent of private wells have nitrate level exceeding 10mg/L.
  - By state law, drinking water wells must be replaced if nitrate levels exceed the *enforcement standard* of 10mg/L.
  - The total estimated cost for replacing non-compliant wells is over *$466 million*.
- **Our Goal: Optimal Policy Learning Under Spatial Dependence**
  - A key way to reduce nitrate levels in drinking water wells is by increasing well depth.
  - But, increasing well depth significantly raises costs.
  - Our goal is to determine the *minimum well depth* required to meet the *10mg/L enforcement standard*, reducing nitrate levels in drinking water in a cost-effective manner.

Spatial Policy Learning: Motivation

- **Dataset Compiling**
  - Groundwater nitrate measurements and the corresponding well depth from 2014 to 2024 in Wisconsin.

Spatial Policy Learning: Motivation

- **Dataset Compiling**
  - Groundwater nitrate measurements and the corresponding well depth from 2014 to 2024 in Wisconsin.

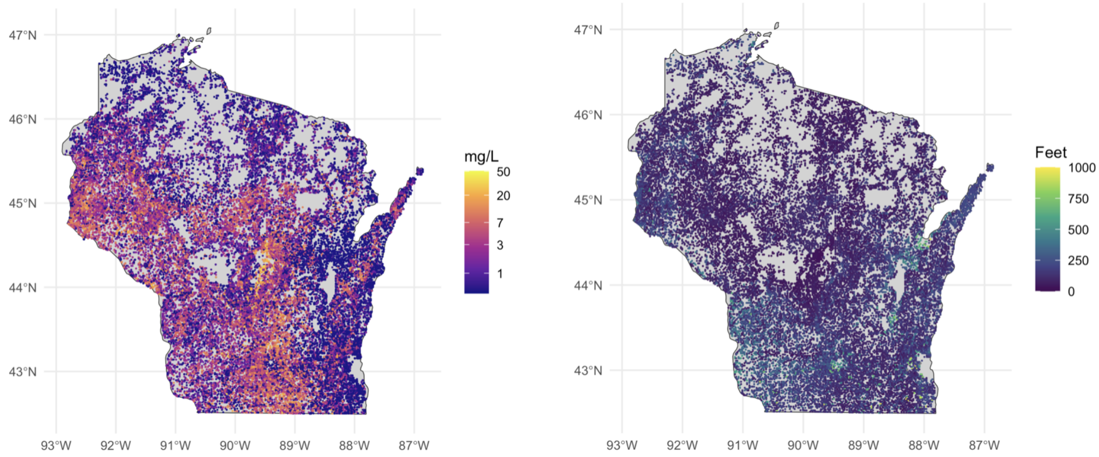

Figure 1. Left: nitrate measurement in Wisconsin. Right: well depth in Wisconsin.

Spatial Policy Learning: Dataset Visualization

- **Spatial predictive variables**
  - Nitrates are primarily from agricultural fertilizers and livestock manure in the surface land.
  - We include land use and livestocks numbers in farms as spatial predictive variables.

Spatial Policy Learning: Dataset Visualization

- **Spatial predictive variables**
  - Nitrates are primarily from agricultural fertilizers and livestock manure in the surface land.
  - We include land use and livestocks numbers in farms as spatial predictive variables.

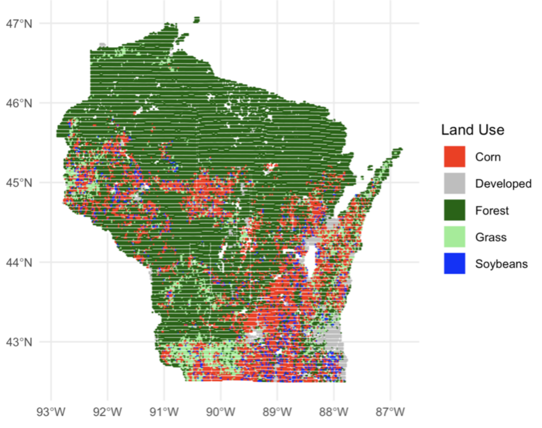

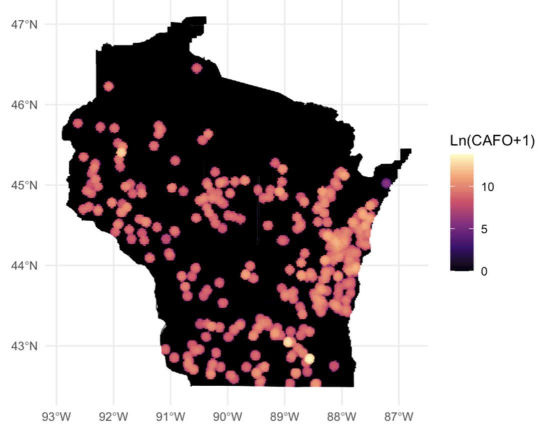

Figure 2. Left: land use in Wisconsin. Right: animal manure in Wisconsin, caculated by livestock numbers and scaled factor.

Spatial Policy Learning: Dataset Visualization

- **Spatial Predictive Variable**
  - Nitrate primarily gets into groundwater when water from rain moves down through soil, carrying nitrate with it .
  - We include soil drainage level and the average precipitation level as covariates.

Spatial Policy Learning: Dataset Visualization

- **Spatial Predictive Variable**
  - Nitrate primarily gets into groundwater when water from rain moves down through soil, carrying nitrate with it .
  - We include soil drainage level and the average precipitation level as covariates.

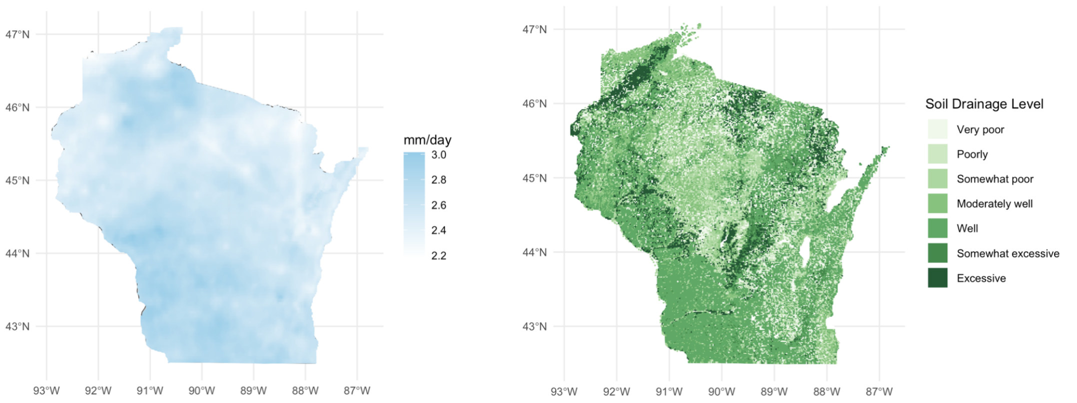

Figure 3. Left: precipitation in Wisconsin. Right: soil drainage level in Wisconsin.

Data Sources

- Nitrate measurements: Wisconsin Groundwater Retrieval Network (GRN).
- Well information: Wisconsin Groundwater Retrieval Network (GRN).
- Land use: Cropland Data Layer (CDL), National Agricultural Statistics Service (NASS).
- Precipitation: United States Geological Survey (USGS).
- Soil drainage: Soil Survey Geographic Database (SSURGO).
- Livestock operation: Wisconsin Pollutant Discharge Elimination System (WPDES).

Outline

Setup

- Let $\mathcal{S} = \{S_1, S_2, \dots, S_n\} \subset \mathcal{D} \subset \mathbb{R}^2$ be the observed well locations. For each $S_i \in \mathcal{S}$:
  - Y_i : Observed nitrate level.
  - $A_i \in \mathcal{A}$: Observed well depth, where $\mathcal{A}$ is the set of possible well depths.
  - $Y_i(a)$: Potential nitrate level if the well were to have depth $a \in \mathcal{A}$.
  - X_i: Measured spatial covariates at location S_i (e.g., land use, drainage level).
  - U_i: Unmeasured spatial covariate at location S_i.

Setup

- Let $\mathcal{S} = \{S_1, S_2, \ldots, S_n\} \subset \mathcal{D} \subset \mathbb{R}^2$ be the observed well locations. For each $S_i \in \mathcal{S}$:
  - Y_i : Observed nitrate level.
  - $A_i \in \mathcal{A}$: Observed well depth, where $\mathcal{A}$ is the set of possible well depths.
  - $Y_i(a)$: Potential nitrate level if the well were to have depth $a \in \mathcal{A}$.
  - X_i: Measured spatial covariates at location S_i (e.g., land use, drainage level).
  - U_i: Unmeasured spatial covariate at location S_i.
- **Definition**
  - *Target Threshold* $\mathcal{T}$: predefined maximum acceptable nitrate level (e.g., $\mathcal{T}$=10 mg/L).
  - *Optimal Policy* $\theta^*(S_0)$: for any location S_0, let Y_0, X_0, U_0 be corresponding variable, *the minimum well depth required for the nitrate level to meet the target threshold* $\mathcal{T}$, defined as

$$\theta^*(S_0) = \arg\min_{a \in \mathcal{A}} a, \text{ such that } E[Y_0(a) | X = X_0, U = U_0] \leq \mathcal{T}.$$

Identification

- **Key Assumptions**
  - (A1) Causal consistency: $Y_i = Y_i(A_i)$ almost surely.
  - (A2) Strong ignorability: $A_i \perp Y_i(a) \mid X_i, U_i$ and $p(a \mid X_i = x, U_i = u) > 0$ for all a, x, u.
  - (A3) Spatial unconfounded: U_i is fixed measurable function, and $(Y_i(a), A_i) \perp S_i \mid X_i, U_i$.
  - (A4) Monotonicity: $E[Y_i(a)|X_i, U_i]$ is monotonically decreasing in a for any X_i, U_i.

Identification

- **Key Assumptions**
  - (A1) Causal consistency: $Y_i = Y_i(A_i)$ almost surely.
  - (A2) Strong ignorability: $A_i \perp Y_i(a) \mid X_i, U_i$ and $p(a \mid X_i = x, U_i = u) > 0$ for all a, x, u.
  - (A3) Spatial unconfounded: U_i is fixed measurable function, and $(Y_i(a), A_i) \perp S_i \mid X_i, U_i$.
  - (A4) Monotonicity: $E[Y_i(a)|X_i, U_i]$ is monotonically decreasing in a for any X_i, U_i.

- Assumption (A3) is a straight-forward extension of Gilbert et al. 2021, which enables the use of S_i as the proxy of U_i in Assumption (A2).

Identification

- **Key Assumptions**
  - (A1) Causal consistency: $Y_i = Y_i(A_i)$ almost surely.
  - (A2) Strong ignorability: $A_i \perp Y_i(a) \mid X_i, U_i$ and $p(a \mid X_i = x, U_i = u) > 0$ for all a, x, u.
  - (A3) Spatial unconfounded: U_i is fixed measurable function, and $(Y_i(a), A_i) \perp S_i \mid X_i, U_i$.
  - (A4) Monotonicity: $E[Y_i(a)|X_i, U_i]$ is monotonically decreasing in a for any X_i, U_i.
- Assumption (A3) is a straight-forward extension of Gilbert et al. 2021, which enables the use of S_i as the proxy of U_i in Assumption (A2).
- **Proposition**
  - Let $\mu^*(a, x, s) := E[Y|A = a, X = x, S = s]$. Under the above assumptions, $\theta^*(S_0)$ can be identified by

$$\theta^*(S_0) = \arg\min_{a \in \mathcal{A}} a, \text{ such that } \mu^*(a, X_0, S_0) \leq \mathcal{T}. \tag{1}$$

Identification

- **Key Assumptions**
  - (A1) Causal consistency: $Y_i = Y_i(A_i)$ almost surely.
  - (A2) Strong ignorability: $A_i \perp Y_i(a) \mid X_i, U_i$ and $p(a \mid X_i = x, U_i = u) > 0$ for all a, x, u.
  - (A3) Spatial unconfounded: U_i is fixed measurable function, and $(Y_i(a), A_i) \perp S_i \mid X_i, U_i$.
  - (A4) Monotonicity: $E[Y_i(a)|X_i, U_i]$ is monotonically decreasing in a for any X_i, U_i.

- Assumption (A3) is a straight-forward extension of Gilbert et al. 2021, which enables the use of S_i as the proxy of U_i in Assumption (A2).
- **Proposition**
  - Let $\mu^*(a, x, s) := E[Y|A = a, X = x, S = s]$. Under the above assumptions, $\theta^*(S_0)$ can be identified by

$$\theta^*(S_0) = \arg\min_{a \in \mathcal{A}} \ a, \ \text{such that } \mu^*(a, X_0, S_0) \leq \mathcal{T}. \quad (1)$$

- **Indirect Method**
  - An estimator $\hat{\theta}(S_0)$ can be obtained by replacing μ^* with an estimate $\hat{\mu}$ in equation (1) and doing grid search:

$$\hat{\theta}(S_0) = \inf_{a \in \mathcal{A}} \{a \mid \hat{\mu}(a, X_0, S_0) \leq \mathcal{T}\}$$

Indirect Method

- **Spatial Mixed Effect Model**

$$Y_i = f^*(A_i, X_i) + U_i + \epsilon_i.$$

  - f^* is a nonparametric function.
  - U_i is a zero-mean Gaussian process (GP) with a parametric covariance function $C(\cdot;\psi)$.

Indirect Method

- **Spatial Mixed Effect Model**

$$Y_i = f^*(A_i, X_i) + U_i + \epsilon_i.$$

  - f^* is a nonparametric function.
  - U_i is a zero-mean Gaussian process (GP) with a parametric covariance function $C(\cdot; \psi)$.

- **Estimation**
  - *(Initial regression)* Fit Y_i on A_i and X_i to obtain the estimate $\hat{f}(a, x)$.
  - *(GP regression)* Use the fitted residual $\hat{U}_i = Y_i - \hat{f}(A_i, X_i)$ to obtain the estimate $\hat{\psi}$ by maximum-likelihood estimation
  - *(Kriging)* For any location S_0, obtain the estimate $\hat{U}_0$ by kriging using the estimated covariance function $C(\cdot; \hat{\psi})$. The estimated outcome regression at location S_0 is

$$\hat{\mu}(a, X_0, U_0) = \hat{f}(a, X_0) + \hat{U}_0$$

Indirect Method

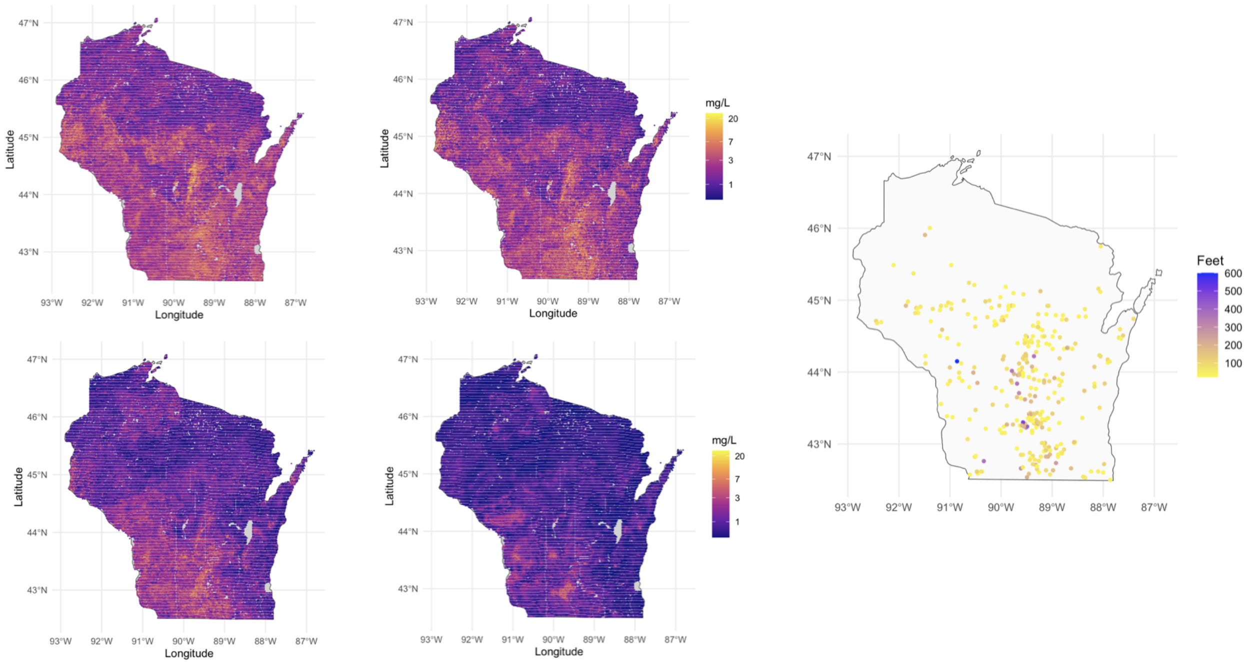

Figure 4. Left: nitrate prediction at depth 50, 100, 200, and 400 feet. Right: the required depth for well replacement.

Indirect Method

- **Drawback of Indirect Method**
  - The indirect method replies on the correct specification of the outcome regression model.
  - The indirect method suffers from potential over-fitting for small sample size.

Indirect Method

- **Drawback of Indirect Method**
  - The indirect method replies on the correct specification of the outcome regression model.
  - The indirect method suffers from potential over-fitting for small sample size.
- **Prediction Error of Nitrate Map**
  - Split the data into training and validation set. Verify the performance at validation set.
  - The prediction fails to capture the extreme value of nitrate contaminants in groundwater.

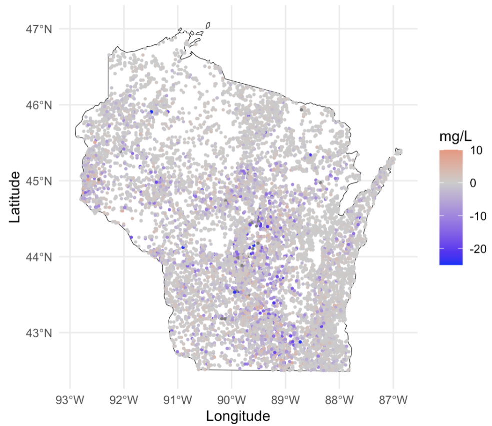

Figure 5. Prediction error of nitrate level on validation set.

Indirect Method

- **Drawback of Indirect Method**
  - The indirect method replies on the correct specification of the outcome regression model.
  - The indirect method suffers from potential over-fitting for small sample size.
- **Prediction Error of Nitrate Map**
  - Split the data into training and validation set. Verify the performance at validation set.
  - The prediction fails to capture the extreme value of nitrate contaminants in groundwater.
- **New Method**
  - Though the underlying distribution of groundwater nitrate is hard to model, we found it is easier to model the distribution of well depth.
  - This motivates us to propose a doubly robust method that utilizes the estimation of the distribution of the well depth.

Outline

Doubly Robust Risk Minimization

- **Loss Function**
  - Let $p^*(a|x, s)$ be the conditional density of $A \mid X = x, S = s$, i.e., generalized propensity score. Let $K_\delta(\cdot)$ is a surrogate kernel with bandwidth δ. Let $O_i := \{A_i, X_i, S_i, Y_i\}$. Define pseudo-outcome as

$$\psi^{(\delta)}(a, O_i) := \mu(a, X_i, S_i) + \frac{(Y_i - \mu(A_i, X_i, S_i))K_\delta(A_i - a)}{p(A_i|X_i, S_i)}$$

  - Define loss function as

$$L^{(\delta)}(\theta(S_i), O_i) = \int_{\mathcal{A}} (\mathcal{T} - \psi^{(\delta)}(a, O_i))1(a \leq \theta(S_i))\mathrm{d}a,$$

Doubly Robust Risk Minimization

- **Loss Function**
  - Let $p^*(a|x,s)$ be the conditional density of $A \mid X = x, S = s$, i.e., generalized propensity score. Let $K_\delta(\cdot)$ is a surrogate kernel with bandwidth δ. Let $O_i := \{A_i, X_i, S_i, Y_i\}$. Define pseudo-outcome as

$$\psi^{(\delta)}(a, O_i) := \mu(a, X_i, S_i) + \frac{(Y_i - \mu(A_i, X_i, S_i))K_\delta(A_i - a)}{p(A_i|X_i, S_i)}$$

  - Define loss function as

$$L^{(\delta)}(\theta(S_i), O_i) = \int_{\mathcal{A}} (\mathcal{T} - \psi^{(\delta)}(a, O_i))1(a \leq \theta(S_i))\mathrm{d}a,$$

- **Proposition** (Doubly Robust Risk Minimization)
  - Suppose $\mu^*(a, x, s)$ is uniformly continuous with respect to a. Define the risk function as:

$$R(\theta) := \lim_{\delta \to 0} E[L^{(\delta)}(\theta(S_i), O_i)]$$

  Then, θ^* is the minimizer of the risk function if either $\mu = \mu^*$ or $p = p^*$.

Doubly Robust Risk Minimization

- **Loss Function**
  - Let $p^*(a|x,s)$ be the conditional density of $A \mid X = x, S = s$, i.e., generalized propensity score. Let $K_\delta(\cdot)$ is a surrogate kernel with bandwidth δ. Let $O_i := \{A_i, X_i, S_i, Y_i\}$. Define pseudo-outcome as

$$\psi^{(\delta)}(a, O_i) := \mu(a, X_i, S_i) + \frac{(Y_i - \mu(A_i, X_i, S_i))K_\delta(A_i - a)}{p(A_i|X_i, S_i)}$$

  - Define loss function as

$$L^{(\delta)}(\theta(S_i), O_i) = \int_{\mathcal{A}} (\mathcal{T} - \psi^{(\delta)}(a, O_i))1(a \le \theta(S_i))\mathrm{d}a,$$

- **Proposition** (Doubly Robust Risk Minimization)
  - Suppose $\mu^*(a, x, s)$ is uniformly continuous with respect to a. Define the risk function as:

$$R(\theta) := \lim_{\delta \to 0} E[L^{(\delta)}(\theta(S_i), O_i)]$$

    Then, θ^* is the minimizer of the risk function if either $\mu = \mu^*$ or $p = p^*$.
  - As long as $E[\psi^{(\delta)}(a, O_i)] = \mu^*(a, X_i, S_i)$. θ^* is the minimizer of the risk function.

Doubly Robust Risk Minimization

- **Loss Function**
  - Let $p^*(a|x,s)$ be the conditional density of $A \mid X = x, S = s$, i.e., generalized propensity score. Let $K_\delta(\cdot)$ is a surrogate kernel with bandwidth δ. Let $O_i := \{A_i, X_i, S_i, Y_i\}$. Define pseudo-outcome as

$$\psi^{(\delta)}(a, O_i) := \mu(a, X_i, S_i) + \frac{(Y_i - \mu(A_i, X_i, S_i))K_\delta(A_i - a)}{p(A_i|X_i, S_i)}$$

  - Define loss function as

$$L^{(\delta)}(\theta(S_i), O_i) = \int_{\mathcal{A}} (\mathcal{T} - \psi^{(\delta)}(a, O_i))1(a \leq \theta(S_i))\mathrm{d}a,$$

- **Proposition** (Doubly Robust Risk Minimization)
  - Suppose $\mu^*(a, x, s)$ is uniformly continuous with respect to a. Define the risk function as:

$$R(\theta) := \lim_{\delta \to 0} E[L^{(\delta)}(\theta(S_i), O_i)]$$

    Then, θ^* is the minimizer of the risk function if either $\mu = \mu^*$ or $p = p^*$.
  - As long as $E[\psi^{(\delta)}(a, O_i)] = \mu^*(a, X_i, S_i)$. θ^* is the minimizer of the risk function.
  - Thus, we can let $\psi^{(\delta)}(a, O_i)$ to be $\mu(a, X_i, S_i)$ or $Y_i K_\delta(A_i - a)/p(A_i|X_i, S_i)$.

An Overview of the Estimation Procedure

- **Step 1: Data Splitting**
  - Randomly partition the full dataset into two disjoint sets, $\mathcal{D}_1$ and $\mathcal{D}_2$.
- **Step 2: Estimating Nuisance Parameters**
  - Use $\mathcal{D}_1$ to obtain estimate $\hat{\mu}(a, x, s)$ and $\hat{p}(a|x, s)$ assuming a spatial mixed effect model.
- **Step 3: SVM Estimator of Optimal Policy**
  - Plug the nuisance estimates from $\mathcal{D}_1$ to construct loss $\hat{L}^{(\delta_n)}$.
  - Let $\mathcal{H}_{\gamma_n}$ be the corresponding reproducing kernel Hilbert space (RKHS),

$$\hat{\theta}_n = \min_{\theta \in \mathcal{H}_{\gamma_n}} \left\{ \frac{1}{|\mathcal{D}_2|} \sum_{i \in \mathcal{D}_2} \hat{L}^{(\delta_n)}(\theta(X_i, S_i), O_i) + \lambda_n \|\theta\|_{\mathcal{H}_{\gamma_n}} \right\}$$

  - The bandwidth γ_n and penalty λ_n are selected via cross-validation on $\mathcal{D}_2$.
  - The optimization can be solved with difference of convex functions (DC) algorithm.
  - If intepretability is preferred, we can also use parametric method to estimate the optimal policy.
- **Repeat steps 1-3 and do median aggregation**

DC Algorithm of Optimization

- **Computational Challenge**
  - The doubly robust loss function is not convex and might have local maximum points.

DC Algorithm of Optimization

- **Computational Challenge**
  - The doubly robust loss function is not convex and might have local maximum points.
- **Key Idea of DC Algorithm**
  - The loss function $L^{(\delta)}(t, O_i)$ can be decomposed into the difference of two convex function $L_{+}^{(\delta)}(t, O_i)$ and $L_{-}^{(\delta)}(t, O_i)$.
  - The optimization can be solved by **repeatly solving a series of convex optimization**:

$$\hat{\theta}_t = \arg\min \left\{ \sum_{i=1}^{n} L_{+}^{(\delta)}(\theta(X_i, S_i), O_i) - \left(\sum_{i=1}^{n} \nabla L_{-}^{(\delta)}(\theta(X_i, S_i), O_i) \right)^{\top} (\theta - \theta_{t-1}) \right\}$$

Bound on Excess Risk

- **Key Assumptions**
  - Suppose θ^* belongs to a Besov space $\mathcal{B}^{\beta}_{1,\infty}(\mathbf{R}^d)$. d is the dimension of (X, U).
  - Suppose $\hat{\mu} - \mu^* = r_{\mu,n}$ and $\hat{p} - p^* = r_{p,n}$ with probability $1 - \Delta_n$.
- **Theorem (Execess Risk of SVM Estimator)**
  - Let σ_n be the largest eigenvalue of the covariance matrix of U_i, $i = 1, \ldots, n$. Then, with appropriately chosen γ_n and λ_n, the SVM estimator satisfies

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

Bound on Excess Risk

- **Key Assumptions**
  - Suppose θ^* belongs to a Besov space $\mathcal{B}^{\beta}_{1,\infty}(\mathbf{R}^d)$. d is the dimension of (X, U).
  - Suppose $\hat{\mu} - \mu^* = r_{\mu,n}$ and $\hat{p} - p^* = r_{p,n}$ with probability $1 - \Delta_n$.
- **Theorem (Execess Risk of SVM Estimator)**
  - Let σ_n be the largest eigenvalue of the covariance matrix of U_i, $i = 1, \ldots, n$. Then, with appropriately chosen γ_n and λ_n, the SVM estimator satisfies

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

  - σ_n is upper bounded for common GP covariance if locations are separated by a minimum distance.

Bound on Excess Risk

- **Key Assumptions**
  - Suppose θ^* belongs to a Besov space $\mathcal{B}^{\beta}_{1,\infty}(\mathbf{R}^d)$. d is the dimension of (X, U).
  - Suppose $\hat{\mu} - \mu^* = r_{\mu,n}$ and $\hat{p} - p^* = r_{p,n}$ with probability $1 - \Delta_n$.
- **Theorem (Execess Risk of SVM Estimator)**
  - Let σ_n be the largest eigenvalue of the covariance matrix of U_i, $i = 1, \ldots, n$. Then, with appropriately chosen γ_n and λ_n, the SVM estimator satisfies

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

  - σ_n is upper bounded for common GP covariance if locations are separated by a minimum distance.
  - When θ^* is sufficiently smooth, the SVM optimal rate is close to $n^{-1/2}$.

Bound on Excess Risk

- **Key Assumptions**
  - Suppose θ^* belongs to a Besov space $\mathcal{B}^{\beta}_{1,\infty}(\mathbf{R}^d)$. d is the dimension of (X, U).
  - Suppose $\hat{\mu} - \mu^* = r_{\mu,n}$ and $\hat{p} - p^* = r_{p,n}$ with probability $1 - \Delta_n$.
- **Theorem (Execess Risk of SVM Estimator)**
  - Let σ_n be the largest eigenvalue of the covariance matrix of U_i, $i = 1, \ldots, n$. Then, with appropriately chosen γ_n and λ_n, the SVM estimator satisfies

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

  - σ_n is upper bounded for common GP covariance if locations are separated by a minimum distance.
  - When θ^* is sufficiently smooth, the SVM optimal rate is close to $n^{-1/2}$.
  - SVM on irregularly spatially dependent data has not been explored before. The result might be of independent interest.

Outline

Simulation

- **General setting**
  - $[0, 1]^2$ region with 50×50 uniform grid. Split data into training and validation set.
  - Use spatial mixed effect model as data generating model of observed outcome.
- **Scenario I: Nonparametric policy (RKHS)**
  - Nonlinear outcome regression model estimated using GAM (misspecification). Beta distributed treatment estimated using beta regression (correct model).

Simulation

- **General setting**
  - $[0, 1]^2$ region with 50×50 uniform grid. Split data into training and validation set.
  - Use spatial mixed effect model as data generating model of observed outcome.
- **Scenario I: Nonparametric policy (RKHS)**
  - Nonlinear outcome regression model estimated using GAM (misspecification). Beta distributed treatment estimated using beta regression (correct model).
  - Our method can mitigate the model misspecification and captures spatial dependence.

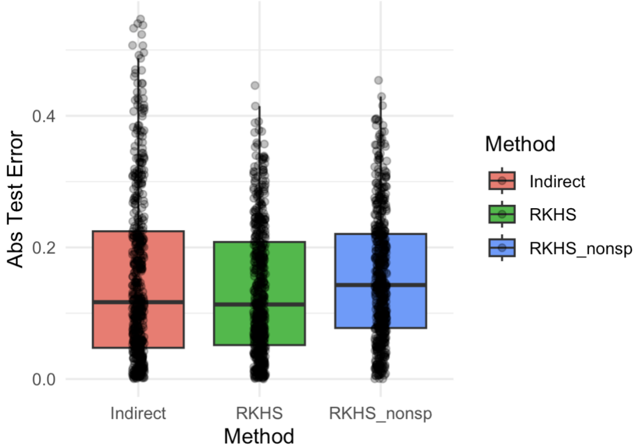

Figure 5. Boxplot of Absolute Error

Method	rMSE	Bias	Var
RKHS	0.1674	-0.004	0.0280
RKHS-nonsp	0.1833	0.0041	0.0337
Indirect	0.2006	0.009	0.0402

Table 1. Test errors summary. RHKS_nonsp: spatial dependence/confounding is not accounted for.

Simulation

- **Scenario II: Parametric policy**
  - Linear outcome regression model estimated using random forest (overfitting). Normally distributed treatment with linear expectation, estimated using linear regression (correct model).

- **Scenario II: Parametric policy**
  - Linear outcome regression model estimated using random forest (overfitting). Normally distributed treatment with linear expectation, estimated using linear regression (correct model).
  - The direct method can mitigate the overfitting. The doubly robust loss mitigate the model misspecification.

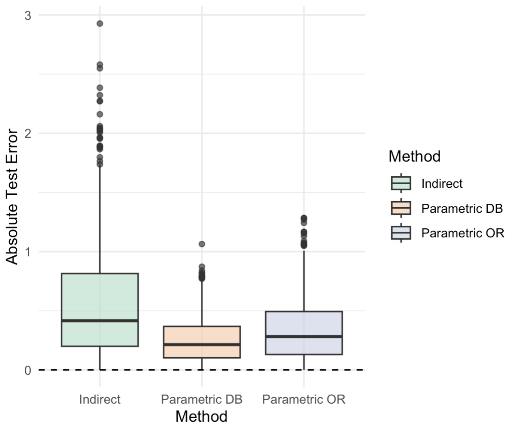

Figure 6. Boxplot of Absolute Error

Method	rMSE	Bias	Var
Parametric DB	0.3136	0.0317	0.0975
Parametric OR	0.4221	0.1029	0.1678
Indirect	0.7416	0.1420	0.5306

Table 2. Test errors summary. Parametric DB/OR: direct method uses pseudo- outcome/outcome regression in loss function.

Simulation

- **Simulation Conclusion**
  - Our method mitigates outcome regression misspecification and overfitting from indirect methods, and effectively captures spatial dependence/confounding.

- **Simulation Conclusion**
  - Our method mitigates outcome regression misspecification and overfitting from indirect methods, and effectively captures spatial dependence/confounding.
  - The improvement from the doubly robust strategy is dependent on the treatment assignment.
  - If the observed treatment does not overlap with the optimal policy, the improvement is small.
  - Revist risk bound

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

▶ Simulation Conclusion

- Our method mitigates outcome regression misspecification and overfitting from indirect methods, and effectively captures spatial dependence/confounding.
- The improvement from the doubly robust strategy is dependent on the treatment assignment.
- If the observed treatment does not overlap with the optimal policy, the improvement is small.
- Revist risk bound

$$R(\hat{\theta}) - R(\theta^*) = \underbrace{O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)}}_{\text{optimal rate for SVM}} + \underbrace{O_p(r_{p,n} r_{\mu,n})}_{\text{doubly robust plug-in error}} + \underbrace{O_p(\delta + \delta r_{\mu,n})}_{\text{Surrogate Rate}}.$$

▶ Real Data Analysis

- Statistical modelling of nitrate is very challenging.
- Parametric policy is insufficient.
- Nonparametric policy is computationally expensive.

Conclusions

- **Problem:** Finding the **minimum required well depth** to ensure nitrate levels in Wisconsin groundwater are below the 10mg/L enforcement.

Conclusions

- **Problem:** Finding the **minimum required well depth** to ensure nitrate levels in Wisconsin groundwater are below the 10mg/L enforcement.
- **Our Approach:** Statistical model for optimal policy learning that explicitly models **spatial dependence/counfounding**.
  - Spatial mixed-effect model and a **doubly robust risk minimization** to identify the optimal policy.
  - The policy is estimated using a Support Vector Machine (SVM) approach with cross-fitting to ensure robustness.
- **Key Results & Takeaways:**
  - **Theoretically**, we established an excess risk bound for our estimator which demonstrates an optimal statistical rate and includes a doubly robust error term.
  - **Empirically**, our outperforms non-spatial and indirect method in simulation studies.
  - **Main takeaway:** Accounting for spatial dependence/counfounding is **crucial** for creating accurate and effective environmental policies. Our framework provides a robust tool for this purpose.

Acknowledgements

- Thank you for listening. Comments are highly appreciated!
  - xlin268@wisc.edu
- Acknowledgment
  - We are grateful to Matthias Katzfuss, Brian Reich, Elizabeth Ogburn, Abhi Datta, and Xinran Miao for their insightful feedback and suggestions.

Indirect Method

- **Spatial Mixed Effect Model**

$$Y_i = f^*(A_i, X_i) + U_i + \epsilon_i.$$

  - f^* is a nonparametric function.
  - U_i is a zero-mean Gaussian process (GP) with a parametric covariance function $C(\cdot;\psi)$.

- **Estimation**
  - (Initial regression) Fit Y_i on A_i and X_i to obtain the estimate $\hat{f}(a,x)$.
  - (GP regression) Denote $\hat{U}_s = Y_s - \hat{f}(A_s, X_s)$, $\hat{\mathbf{U}} = (\hat{U}_1, \cdots, \hat{U}_n)$. $\mathbf{C}(\psi) = \{C(\|s_i - s_j\|;\psi)\}_{i,j=1}^n$. Obtain $\hat{\psi}$ by maximum-likelihood estimation

$$\hat{\psi} = \arg\max_{\psi} \left\{ -\frac{n}{2}\log(|\mathbf{C}(\psi)|) - \frac{1}{2}\sum_{i=1}^{n} \hat{\mathbf{U}}^\top \mathbf{C}(\psi)^{-1}\hat{\mathbf{U}} \right\}$$

  - (Leave-one-out Kriging) Let $\mathbf{C}_{-i}(\psi)$ be $\mathbf{C}(\psi)$ removing i-th row and column, $\mathbf{C}_i(\psi)$ be i-th row of $\mathbf{C}(\psi)$ removing i-th entry, $\hat{\mathbf{U}}_{-i}$ be $\hat{\mathbf{U}}$ removing i-th entry. We define

$$\hat{E}[U_s|\mathbf{U}_{-i}] = \mathbf{C}_i(\hat{\psi})\mathbf{C}_{-i}^{-1}(\hat{\psi})\hat{\mathbf{U}}_{-i}, \quad i = 1, \ldots, n$$

  - The estimated outcome regression is $\hat{\mu}(a, X_s, U_s) = \hat{f}(a, X_s) + \hat{E}[U_s|\mathbf{U}_{-i}]$

DC Algorithm of Optimization

- The loss function is not convex and might have local maximum points.
- Suppose the range of well depth $\mathcal{A} = [a_L, a_U]$. Let $(a)_+ = \max(a, 0), (a)_- = -\min(a, 0)$. The regression of surrogate kernel can be written as the subtraction of two positive convex function
$$\int_{a_L}^{t} K_\delta(A_i - a)da = \Psi_+^{(\delta)}(t, A_i) - \Psi_-^{(\delta)}(t, A_i).$$
Denote
$$r_i := \frac{(Y_i - \mu^*(A_i, X_i, U_i))}{p^*(A_i|X_i, U_i)},$$
- Then, we can decompose the **loss function $L^{(\delta)}(t, O_i)$ into the difference of two convex function $L_+^{(\delta)}(t, O_i)$** and $L_-^{(\delta)}(t, O_i)$

$$L_+^{(\delta)}(t, O_i) = \mathcal{T}(t - a_L) - \int_{a_L}^{t} \mu^*(a, X_i, U_i)\mathrm{d}a + (r_i)_-\Psi_+^{(\delta)}(t, A_i) + (r_i)_+\Psi_-^{(\delta)}(t, A_i) + \lambda_n\|\theta\|_{\mathcal{H}_{\gamma_n}},$$
$$L_-^{(\delta)}(t, O_i) = (r_i)_+\Psi_+^{(\delta)}(t, A_i) + (r_i)_-\Psi_-^{(\delta)}(t, A_i).$$

- The optimization can be solved by **repeatly solving a series of convex optimization**:
$$\hat{\theta}_t = \arg\min\left\{\sum_{i=1}^{n} L_+^{(\delta)}(\theta(X_i, S_i), O_i) - \left(\sum_{i=1}^{n} \nabla L_-^{(\delta)}(\theta(X_i, S_i), O_i)\right)^\top (\theta - \theta_{t-1})\right\}$$

Bound on Excess Risk

- **Key Assumptions**
  - Suppose θ^* belongs to a Besov space $\mathcal{B}^{\beta}_{1,\infty}(\mathbf{R}^d)$. d is the dimension of (X, U).
  - Let σ_n be the largest eigenvalue of the covariance matrix of $U_s + \epsilon_i$.
  - Suppose $\hat{\mu} - \mu^* = r_{\mu,n}$ and $\hat{p} - p^* = r_{p,n}$ with probability $1 - \Delta_n$.
- **Theorem (Execess Risk of SVM Estimator)**
  - For any $\tau > 0$, $p \in (0, 1]$, with probability no less than $1 - \Delta_n - 8e^{-\tau}$, the excess risk of SVM estimator satisfies

$$R(\hat{\theta}) - R(\theta^*) \leq \underbrace{c_1 \lambda_n \gamma_n^{-d} + c_2 \gamma_n^{\beta}}_{\text{approximation error}} + \underbrace{c_3 (\gamma_n^d \lambda_n^p \sigma_n^{-1} n)^{-\frac{1}{2+2p}} + c_4 (\tau^{-1} \sigma_n^{-1} n)^{-\frac{1}{2}}}_{\text{estimation error}}$$

$$+ \underbrace{\delta + \delta r_{\mu,n}}_{\text{error from kernel surrogate function}} + \underbrace{r_{p,n} r_{\mu,n}}_{\text{plug-in error from nuisance estimators}}$$

  - Suppose $\gamma_n \propto (\sigma_n^{-1} n)^{-\frac{1}{2\beta+d}}$, $\lambda_n \propto (\sigma_n^{-1} n)^{-\frac{d+\beta}{2\beta+d}}$, and $\hat{\mu}, \mu^*, \hat{p}, p^*$ are very smooth. The bound reduces to

$$R(\hat{\theta}) - R(\theta^*) = O_P(\sigma_n^{-1} n)^{-\beta/(2\beta+d)} + O_p(r_{p,n} r_{\mu,n}).$$

Details of Nuisance Estimation

- **Spatial Mixed Effect Model**

$$\log(Y_i) = f^*(A_i, X_i) + U_i + \epsilon_i,$$
$$\log(A_i) = g^*(X_i) + V_i + \delta_i$$

- U_i and V_i are Gaussian processes. δ_i is i.i.d. gaussian error.
- f^* are estimated using gradient boosting with monotonic constraints.
- The estimation $\hat{f}$ is smoothed using cubic splines.

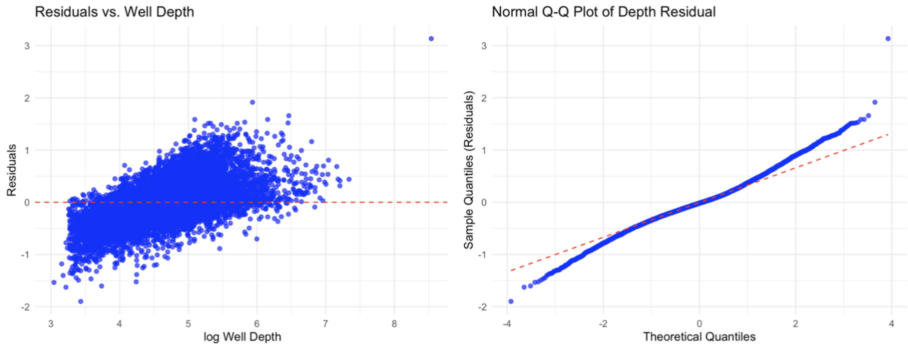

Figure 7. Diagnosis of fitted residual of log transformation of well depth.

Regularity Assumptions

- **Assumption (Doubly Robust Loss)**:
  - (B1) Monotonicity: for any x, u, $\mu^*(a, x, u)$ is monotonically decreasing with respect to a
  - (B2) Existence: For any x, u, there is a unique $a \in \mathcal{A}$ such that $\mu^*(a, x, u) = \mathcal{T}$
  - (B3) Continuity: for any x, u, $\mu^*(a, x, u)$ is continuous with respect to a.
- **Assumption (SVM Execess Risk)**
  - (C1) Smoothness: $\hat{\mu}, \mu^*, \hat{p}, p^*$ are s-order uniformly continuously differentiable w.r.t. a.
  - (C2) Boundedness: $\hat{\mu}, \mu^*, \hat{p}, p^*$ are uniformly bounded above. $\hat{p}, p^*$ are uniformly bounded away from zero.
  - (C3) Nuisance Error: Let $r_{p,n}$, $r_{\mu,n}$ and Δ_n be sequences of real numbers, which are $o(1)$ as $n \to \infty$. Then, we have $\hat{p} - p^* \leq r_{p,n}$, $\hat{\mu} - \mu^* \leq r_{\mu,n}$ with probability greater than $1 - \Delta_n$.

- **SVM Estimator**
  - Consider a support vector machine (SVM) estimator of θ^* using a Gaussian Kernel $\mathcal{K}\{(x, u), (x^{'}, u^{'})\}$ with bandwidth γ_n.
  - Let $\mathcal{H}_{\gamma_n}$ be the corresponding reproducing kernel Hilbert space (RKHS). The estimated policy can be written as $\hat{\theta}(s_0) = \sum_{j=1}^{n} \eta_j \mathcal{K}\left\{(X_{s_0}, U_{s_0}), (X_i, U_i)\right\} + b_j$ that solves

$$\hat{\theta} = \min_{\theta \in \mathcal{H}_\gamma} \left\{ \frac{1}{n} \sum_{i=1}^{n} \hat{L}_k^{(\delta_n)}(\theta(X_s, U_s), O_s) + \lambda_n \|\theta\|_{\mathcal{H}_{\gamma_n}} \right\}$$
$$= \min_{\eta, b} \left\{ \frac{1}{n} \sum_{i=1}^{n} \hat{L}_k^{(\delta_n)}(k_i^\top \eta + b, O_s) + \lambda_n \eta^\top K \eta \right\}$$

  - Here, $k_i = [\mathcal{K}\left\{(X_s, U_s), (X_1, U_1)\right\}, \ldots, \mathcal{K}\left\{(X_s, U_s), (X_n, U_n)\right\}]$, λ_n is the penalty parameter, $\eta = (\eta_1, \ldots, \eta_n)^\top$,