Sparse Additive Interaction Learning

Sahir Bhatnagar

Department of Epidemiology, Biostatistics and Occupational Health
Department of Diagnostic Radiology

October 29, 2021

https://sahirbhatnagar.com/sail
https://sahirbhatnagar.com/papers

McGill UNIVERSITY

Outline

- Betting on Sparsity
  - A Thought Experiment
- Motivating Example: The Nurse Family Partnership
- `sail`: Strong Additive Interaction Learning
- Algorithm
- Theory
- Simulations
- `sail` R package
- Real Data Application
- Discussion
  - Current and Future Work
- Acknowledgements

Betting on Sparsity
A Thought Experiment

Motivating Example: The Nurse Family Partnership

sail: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

sail R package

Real Data Application

Discussion
Current and Future Work

Acknowledgements

Bet on Sparsity Principle

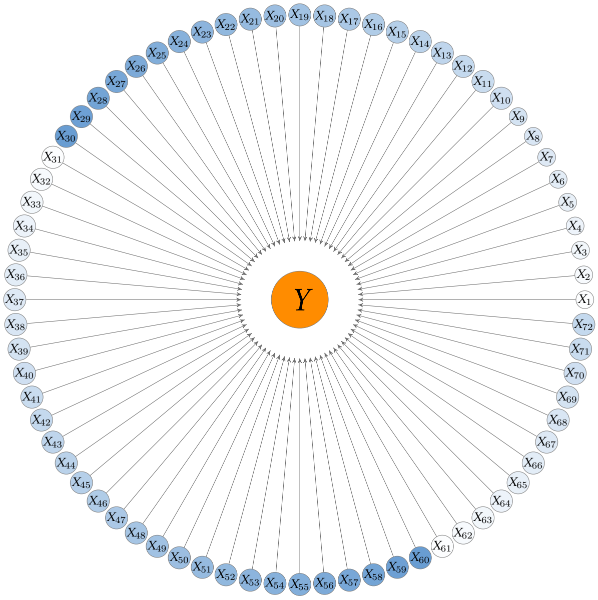

Bet on Sparsity Principle

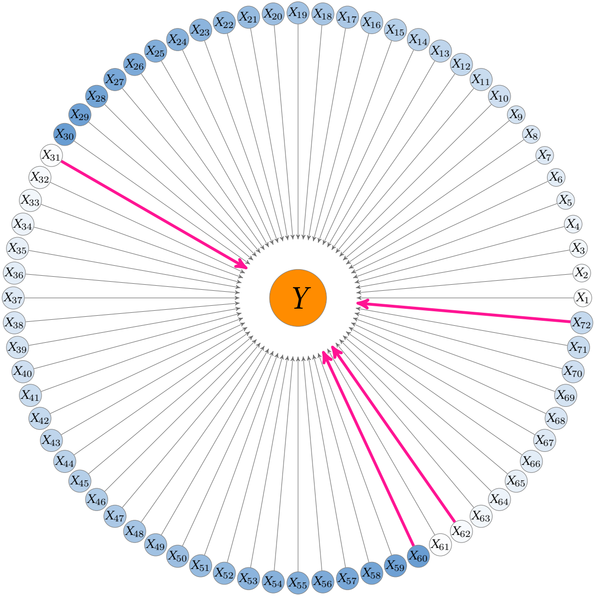

Bet on Sparsity Principle

Use a procedure that does well in sparse problems, since no procedure does well in dense problems.[1]

- We often don't have enough data to estimate so many parameters
- Even when we do, we might want to identify a **relatively small number of predictors** ($k < N$) that play an important role
- Faster computation, easier to understand, and stable predictions on new datasets.

[1] The elements of statistical learning. Springer series in statistics, 2001.

How would you schedule a meeting of 20 people?

How would you schedule a meeting of 20 people?

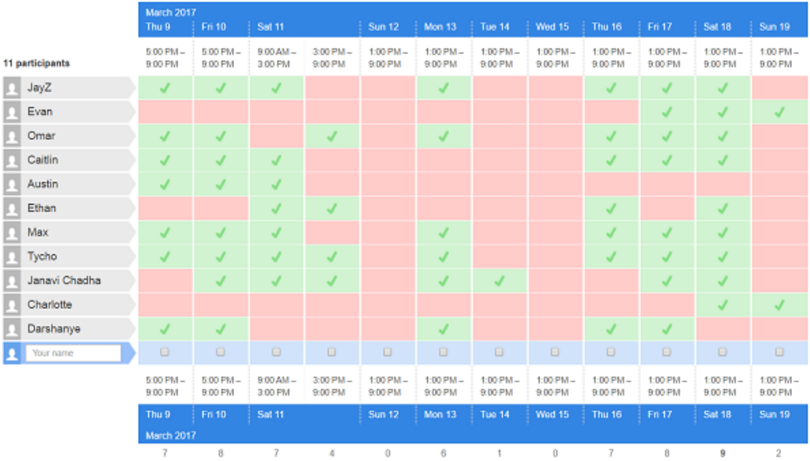

11 participants	Thu 9 5:00 PM – 9:00 PM	Fri 10 5:00 PM – 9:00 PM	Sat 11 9:00 AM – 3:00 PM	Sat 11 3:00 PM – 9:00 PM	Sun 12 1:00 PM – 9:00 PM	Mon 13 1:00 PM – 9:00 PM	Tue 14 1:00 PM – 9:00 PM	Wed 15 1:00 PM – 9:00 PM	Thu 16 1:00 PM – 9:00 PM	Fri 17 1:00 PM – 9:00 PM	Sat 18 1:00 PM – 9:00 PM	Sun 19 1:00 PM – 9:00 PM
JayZ	✓	✓	✓			✓			✓	✓	✓	
Evan										✓	✓	✓
Omar	✓	✓		✓		✓			✓	✓	✓	
Caitlin	✓	✓	✓						✓	✓	✓	
Austin	✓	✓	✓									
Ethan			✓	✓					✓		✓	
Max	✓	✓	✓			✓			✓	✓	✓	
Tycho	✓	✓	✓	✓		✓			✓	✓	✓	
Janavi Chadha		✓	✓	✓		✓	✓			✓	✓	
Charlotte											✓	✓
Darshanye	✓	✓				✓			✓	✓		
Your name												
March 2017	7	8	7	4	0	6	1	0	7	8	9	2

Doctors Bet on Sparsity Also

Doctors Bet on Sparsity Also

Betting on Sparsity
A Thought Experiment
Motivating Example: The Nurse Family Partnership
sail: Strong Additive Interaction Learning
Algorithm
Theory
Simulations
sail R package
Real Data Application
Discussion
Current and Future Work
Acknowledgements

Nurse-Family Partnership

Helping First-Time Parents Succeed®

Nurse-Family Partnership is an evidence-based, community health program with over 40 years of evidence showing significant improvements in the health and lives of first-time moms and their children living in poverty.

Human Brain Development

Synapse formation dependent on early experiences

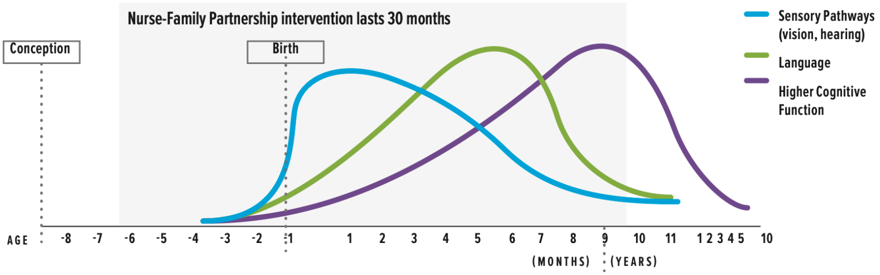

Source: Nelson, C.A., In Neurons to Neighborhoods (2000).

ALL CAUSE MORTALITY OVER 20 YEAR FOLLOW-UP

3x Mothers who did not receive nurse home visits were nearly **3 times more likely to die** from all causes of death than nurse visited mothers (3.7% versus 1.3%)[1]

8x Mothers that did not receive nurse home visits were **8 times more likely to die** from external causes – including unintentional injuries, suicide, drug overdose and homicide – than nurse visited mothers (1.7% versus 0.2%)[1]

PREVENTABLE CHILD MORTALITY OVER 20 YEAR FOLLOW-UP

- Among Nurse-Family Partnership participants, there were **lower rates of preventable child mortality** from birth until age 20.[1]
- 1.6% of the children not receiving nurse home visits died from preventable causes – including sudden infant death syndrome, unintentional injuries and homicide – while none of the nurse visited children died from these causes.[1]

Additional Maternal and Child Health Outcomes

Maternal Health Outcomes

35% fewer cases of pregnancy-induced hypertension[5]

18% fewer preterm births[6]

79% reduction in preterm delivery among women who smoke cigarettes[7]

31% reduction in very closely spaced (<6 months) subsequent pregnancies[8]

Child Health Outcomes

48% reduction in child abuse and neglect[9]

39% fewer health care encounters for injuries or ingestions in the first 2 years of life among children born to mothers with low psychological resources[10]

67% less behavioral and intellectual problems in children at age 6[11]

56% fewer emergency room visits for accidents and poisonings through age 2[12]

Interactions between Intervention and Genetics

Stanford–Binet Fifth Edition (SB5) classification[4]

IQ Range ("deviation IQ")	IQ Classification
145–160	Very gifted or highly advanced
130–144	Gifted or very advanced
120–129	Superior
110–119	High average
90–109	Average
80–89	Low average
70–79	Borderline impaired or delayed
55–69	Mildly impaired or delayed
40–54	Moderately impaired or delayed

$\sim$

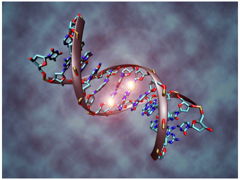

$\times$

Phenotype
IQ Score

Large Data
Genetic Markers

Environment
NFP Intervention

$$Y = \sum_{j=1}^{p} X_j \beta_j + \sum_{j=1}^{p} X_j X_E \tau_j + \varepsilon$$

$$Y = \sum_{j=1}^{p} X_j \beta_j + \sum_{j=1}^{p} X_j X_E \tau_j + \varepsilon$$

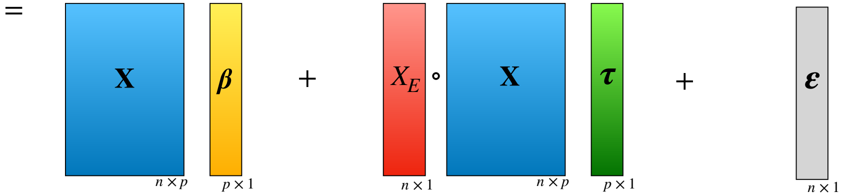

$$Y = \sum_{j=1}^{p} X_j \beta_j + \sum_{j=1}^{p} X_j X_E \tau_j + \varepsilon$$

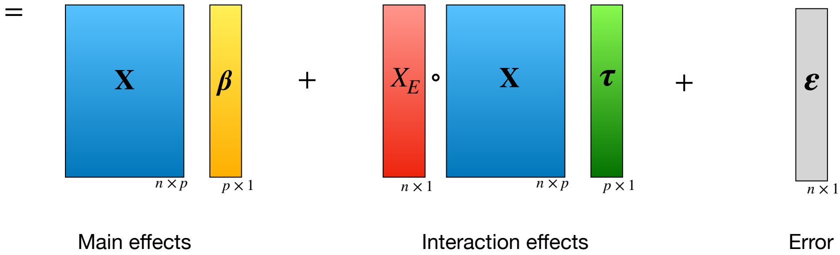

$$Y = \sum_{j=1}^{p} X_j \beta_j + \sum_{j=1}^{p} X_j X_E \tau_j + \varepsilon$$

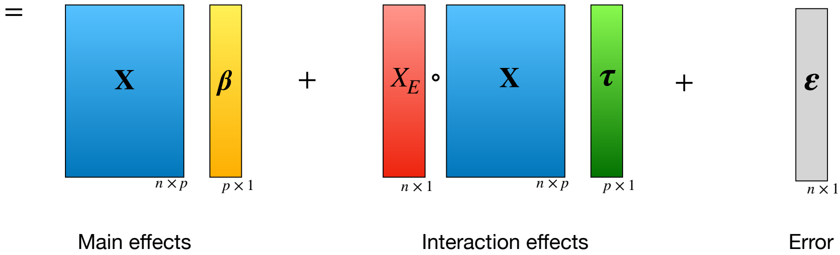

$$\text{Let} \quad Z_{jE} = X_E X_j$$

$$Y = \sum_{j=1}^{p} X_j \beta_j + \sum_{j=1}^{p} Z_{jE} \tau_j + \varepsilon$$

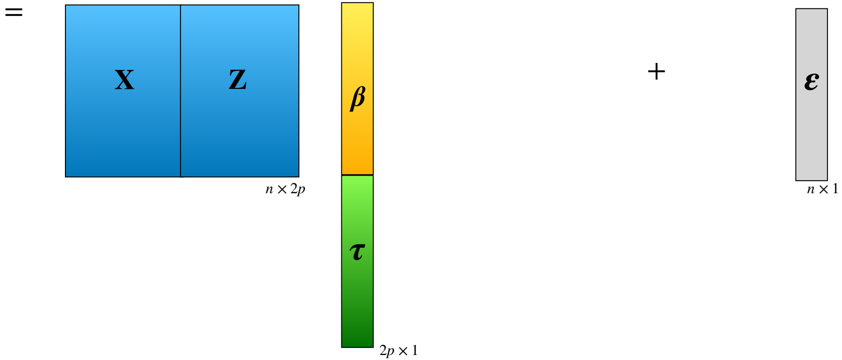

Betting on Sparsity
A Thought Experiment

Motivating Example: The Nurse Family Partnership

`sail`: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

`sail` R package

Real Data Application

Discussion
Current and Future Work

Acknowledgements

Motivation 1: Non-linear Interactions

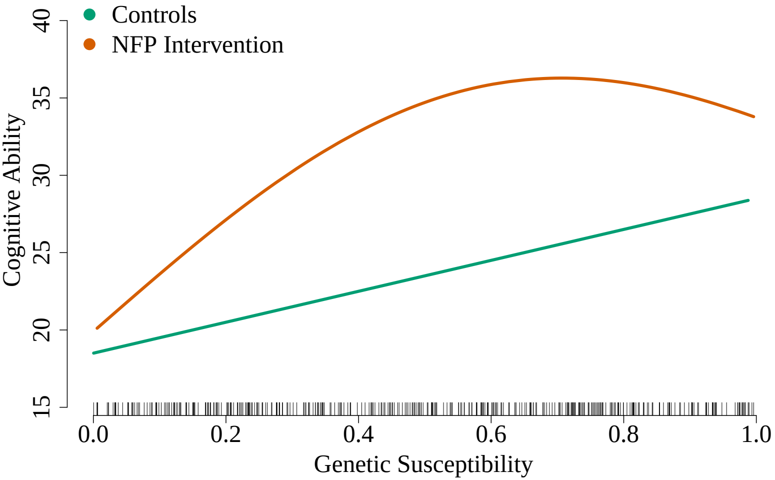

Motivation 2: Heredity Property

$$Y = \beta_0 \cdot \mathbf{1} + \underbrace{\sum_{j=1}^{p} \beta_j X_j + \beta_E X_E}_{\text{main effects}} + \underbrace{\sum_{j=1}^{p} \tau_j X_E X_j}_{\text{interactions}} + \varepsilon$$

[1] Chipman. Canadian Journal of Statistics (1996)

[2] McCullagh and Nelder. Generalized Linear Models (1983)

[3] Cox. International Statistical Review (1984)

Motivation 2: Heredity Property

$$Y = \beta_0 \cdot \mathbf{1} + \underbrace{\sum_{j=1}^{p} \beta_j X_j + \beta_E X_E}_{\text{main effects}} + \underbrace{\sum_{j=1}^{p} \tau_j X_E X_j}_{\text{interactions}} + \varepsilon$$

Strong Heredity[1]

$$\hat{\tau}_j \neq 0 \qquad \Rightarrow \qquad \hat{\beta}_j \neq 0 \qquad \text{and} \qquad \hat{\beta}_E \neq 0$$

- Heredity property is desired for the purposes of **interpretability**[2]
- Large main effects are more likely to lead to appreciable interactions[3]

[1] Chipman. Canadian Journal of Statistics (1996)
[2] McCullagh and Nelder. Generalized Linear Models (1983)
[3] Cox. International Statistical Review (1984)

Lasso interaction model

- $Y \rightarrow$ response
- $X_E \rightarrow$ environment
- $X_j \rightarrow$ predictors, $j = 1, \ldots, p$

$$Y = \beta_0 \cdot \mathbf{1} + \sum_{j=1}^{p} \beta_j X_j + \beta_E X_E + \sum_{j=1}^{p} \tau_j X_E X_j + \varepsilon$$

$$\underset{\boldsymbol{\Theta} := (\beta_0, \boldsymbol{\beta}, \boldsymbol{\tau})}{\operatorname{argmin}} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(\|\boldsymbol{\beta}\|_1 + \|\boldsymbol{\tau}\|_1)$$

Strong Heredity Interactions: Current State of the Art

Type	Model	Software
Linear	`CAP` (Zhao et al. 2009, Ann. Stat)	✗
	`SHIM` (Choi et al. 2009, JASA)	✗
	`hiernet` (Bien et al. 2013, Ann. Stat)	`hierNet(x, y)`
	`GRESH` (She and Jiang 2014, JASA)	✗
	`FAMILY` (Haris et al. 2014, JCGS)	`FAMILY(x, z, y)`
	`glinternet` (Lim and Hastie 2015, JCGS)	`glinternet(x, y)`
	`RAMP` (Hao et al. 2016, JASA)	`RAMP(x, y)`
	`LassoBacktracking` (Shah 2018, JMLR)	`LassoBT(x, y)`
Non-linear	`VANISH` (Radchenko and James 2010, JASA)	✗
	`sail` (Bhatnagar et al. 2020+, in revision CSDA)	`sail(x, e, y, basis)`

Our Extension to Nonlinear Effects

Consider the basis expansion

$$f_j(X_j) = \sum_{\ell=1}^{m_j} \psi_{j\ell}(X_j)\beta_{j\ell}$$

$$f(X_1) = \underbrace{\begin{bmatrix} \psi_{11}(X_{11}) & \psi_{12}(X_{12}) & \cdots & \psi_{11}(X_{15}) \\ \vdots & \vdots & \cdots & \vdots \\ \vdots & \vdots & \cdots & \vdots \\ \psi_{11}(X_{i1}) & \psi_{12}(X_{i2}) & \cdots & \psi_{11}(X_{i5}) \\ \vdots & \vdots & \cdots & \vdots \\ \vdots & \vdots & \cdots & \vdots \\ \psi_{11}(X_{N1}) & \psi_{12}(X_{N2}) & \cdots & \psi_{11}(X_{N5}) \end{bmatrix}_{N \times 5}}_{\boldsymbol{\Psi}_1} \times \underbrace{\begin{bmatrix} \beta_{11} \\ \beta_{12} \\ \beta_{13} \\ \beta_{14} \\ \beta_{15} \end{bmatrix}_{5 \times 1}}_{\theta_1}$$

sail: Additive Interactions

- $\boldsymbol{\theta}_j = (\beta_{j1}, \ldots, \beta_{jm_j}) \in \mathbb{R}^{m_j}$
- $\boldsymbol{\tau}_j = (\tau_{j1}, \ldots, \tau_{jm_j}) \in \mathbb{R}^{m_j}$
- $\boldsymbol{\Psi}_j \to n \times m_j$ matrix of evaluations of the $\psi_{j\ell}$
- In our implementation, we use cubic `bsplines` with 5 degrees of freedom

Model

$$Y = \beta_0 \cdot \mathbf{1} + \sum_{j=1}^{p} \boldsymbol{\Psi}_j \boldsymbol{\theta}_j + \beta_E X_E + \sum_{j=1}^{p} (X_E \circ \boldsymbol{\Psi}_j) \boldsymbol{\tau}_j + \varepsilon$$

sail: Strong Heredity

Reparametrization[1]

$$\boldsymbol{\tau}_j = \gamma_j \beta_E \boldsymbol{\theta}_j$$

Model

$$Y = \beta_0 \cdot \mathbf{1} + \sum_{j=1}^{p} \boldsymbol{\Psi}_j \boldsymbol{\theta}_j + \beta_E X_E + \sum_{j=1}^{p} \gamma_j \beta_E (X_E \circ \boldsymbol{\Psi}_j) \boldsymbol{\theta}_j + \varepsilon$$

Objective Function

$$\underset{\boldsymbol{\Theta} := (\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma})}{\text{argmin}} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E |\beta_E| + \sum_{j=1}^{p} w_j \|\boldsymbol{\theta}_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE} |\gamma_j|$$

[1] Choi et al. JASA (2010)

sail: Weak Heredity

Reparametrization

$$\boldsymbol{\tau}_j = \gamma_j(\beta_E \cdot \mathbf{1}_{m_j} + \boldsymbol{\theta}_j)$$

Model

$$Y = \beta_0 \cdot \mathbf{1} + \sum_{j=1}^{p} \boldsymbol{\Psi}_j \theta_j + \beta_E X_E + \sum_{j=1}^{p} \gamma_j (X_E \circ \boldsymbol{\Psi}_j)(\beta_E \cdot \mathbf{1}_{m_j} + \boldsymbol{\theta}_j) + \varepsilon$$

Objective Function

$$\underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\arg\min} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

Toy example

- With a sample size of $n = 100$, we sample $p = 20$ covariates $X_1, \ldots X_p$ independently from a $N(0, 1)$ distribution truncated to the interval [0,1].
- Data were generated from a model which follows the strong heredity principle, but where only one covariate, X_2, is involved in an interaction with a binary exposure variable (E):

$$Y = f_1(X_1) + f_2(X_2) + 1.75E + 1.5E \cdot f_2(X_2) + \varepsilon.$$

- For illustration, function $f_1(\cdot)$ is assumed to be linear, whereas function $f_2(\cdot)$ is non-linear: $f_1(x) = -3x$, $f_2(x) = 2(2x - 1)^3$.

Toy example - Solution path

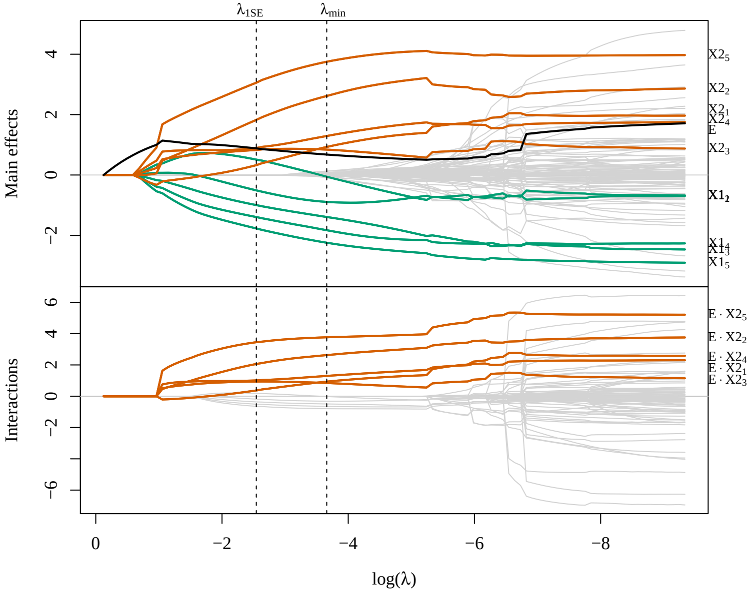

Toy example - Estimated effects

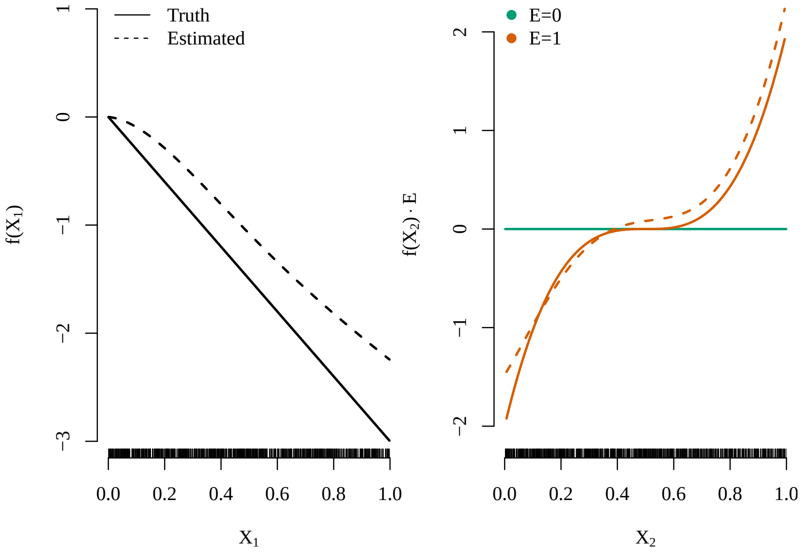

Betting on Sparsity
A Thought Experiment
Motivating Example: The Nurse Family Partnership
`sail`: Strong Additive Interaction Learning
Algorithm
Theory
Simulations
`sail` R package
Real Data Application
Discussion
Current and Future Work
Acknowledgements

Block Relaxation (De Leeuw, 1994)

Algorithm 1: Block Relaxation Algorithm

Set the iteration counter $k \leftarrow 0$ and fix $\alpha \in (0, 1)$;

for each λ **do**

repeat

$$\boldsymbol{\gamma}^{(k+1)} \leftarrow \underset{\boldsymbol{\gamma}}{\text{argmin}} \quad Q_\lambda \left(\boldsymbol{\gamma}, \beta_E^{(k)}, \boldsymbol{\theta}^{(k)} \right)$$

$$\boldsymbol{\theta}^{(k+1)} \leftarrow \underset{\boldsymbol{\theta}}{\text{argmin}} \quad Q_\lambda \left(\boldsymbol{\theta}, \beta_E^{(k)}, \boldsymbol{\gamma}^{(k+1)} \right)$$

$$\beta_E^{(k+1)} \leftarrow \underset{\boldsymbol{\beta_E}}{\text{argmin}} \quad Q_\lambda \left(\boldsymbol{\theta}^{(k+1)}, \beta_E, \boldsymbol{\gamma}^{(k+1)} \right)$$

$k \leftarrow k + 1$

until convergence criterion is satisfied;

end

Implementation

Objective Function

$$\underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\arg\min} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

[1] https://cran.r-project.org/package=sail

Implementation

Objective Function

$$\underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\text{argmin}} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j||\theta_j||_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

Lasso problem

$$\underset{\boldsymbol{\gamma}}{\text{argmin}} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j||\theta_j||_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

[1] https://cran.r-project.org/package=sail

Implementation

Objective Function

$$\underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\arg\min} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

[1] https://cran.r-project.org/package=sail

Implementation

Objective Function

$$\underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\text{argmin}} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

Group Lasso problem

$$\underset{\beta_E, \boldsymbol{\theta}}{\text{argmin}} \quad \mathcal{L}(Y; \boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

[1] https://cran.r-project.org/package=sail

Betting on Sparsity
A Thought Experiment

Motivating Example: The Nurse Family Partnership

`sail`: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

`sail` R package

Real Data Application

Discussion
Current and Future Work

Acknowledgements

Sparsity

Theorem 1

$$\widehat{\boldsymbol{\Theta}}_n = \underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\text{argmin}} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

$$\mathcal{A}_1 = \{j : \theta_j \neq 0, \beta_j \neq 0\}$$
$$\mathcal{A}_2 = \{k : \gamma_k \neq 0\}, \qquad \mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$$

Under certain regularity conditions and the existence of a local minimizer $\widehat{\boldsymbol{\Theta}}_n$ that is $\sqrt{n}$-consistent

$$P\left(\widehat{\boldsymbol{\Theta}}_{\mathcal{A}^c} = 0\right) \to 1$$

Sparsity

Theorem 1

$$\widehat{\mathbf{\Theta}}_n = \underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\arg\min} \quad \mathcal{L}(\mathbf{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

$$\mathcal{A}_1 = \{j : \theta_j \neq 0, \beta_j \neq 0\}$$
$$\mathcal{A}_2 = \{k : \gamma_k \neq 0\}, \qquad \mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$$

Under certain regularity conditions and the existence of a local minimizer $\widehat{\mathbf{\Theta}}_n$ that is $\sqrt{n}$-consistent

$$P\left(\widehat{\mathbf{\Theta}}_{\mathcal{A}^c} = 0\right) \to 1$$

Theorem 1 shows that when the tuning parameters for the nonzero coefficients converge to 0 faster than $n^{-1/2}$ `sail` can consistently remove the noise terms with probability tending to 1.

Asymptotic normality

Theorem 2

$$\widehat{\boldsymbol{\Theta}}_n = \underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\arg\min} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

Under certain regularity conditions, the component $\widehat{\boldsymbol{\Theta}}_{\mathcal{A}}$ of the local minimizer $\widehat{\boldsymbol{\Theta}}_n$ satisfies

$$\sqrt{n}\left(\widehat{\boldsymbol{\Theta}}_{\mathcal{A}} - \boldsymbol{\Theta}_{\mathcal{A}}\right) \to_d \mathcal{N}\left(0, \mathbf{I}^{-1}\left(\boldsymbol{\Theta}_{\mathcal{A}}\right)\right)$$

Theorem 2 shows that the `sail` estimates for nonzero coefficients in the true model have the same asymptotic distribution as they would have if the zero coefficients were known in advance.

Asymptotic normality

Theorem 2

$$\widehat{\boldsymbol{\Theta}}_n = \underset{\beta_E, \boldsymbol{\theta}, \boldsymbol{\gamma}}{\operatorname{argmin}} \quad \mathcal{L}(\boldsymbol{\Theta}) + \lambda(1-\alpha)\left(w_E|\beta_E| + \sum_{j=1}^{p} w_j \|\theta_j\|_2 \right) + \lambda\alpha \sum_{j=1}^{p} w_{jE}|\gamma_j|$$

Under certain regularity conditions, the component $\widehat{\boldsymbol{\Theta}}_{\mathcal{A}}$ of the local minimizer $\widehat{\boldsymbol{\Theta}}_n$ satisfies

$$\sqrt{n}\left(\widehat{\boldsymbol{\Theta}}_{\mathcal{A}} - \boldsymbol{\Theta}_{\mathcal{A}}\right) \to_d \mathcal{N}\left(0, \mathbf{I}^{-1}\left(\boldsymbol{\Theta}_{\mathcal{A}}\right)\right)$$

Theorem 2 shows that the `sail` estimates for nonzero coefficients in the true model have the same asymptotic distribution as they would have if the zero coefficients were known in advance.
Theorem 1 + 2 –> Oracle property (Fan and Li, 2001)

Betting on Sparsity
    A Thought Experiment

Motivating Example: The Nurse Family Partnership

`sail`: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

`sail` R package

Real Data Application

Discussion
    Current and Future Work

Acknowledgements

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

2. **Truth obeys weak hierarchy**

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

2. **Truth obeys weak hierarchy**
3. **Truth only has interactions**

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

2. **Truth obeys weak hierarchy**
3. **Truth only has interactions**
4. **Truth is linear**

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

2. **Truth obeys weak hierarchy**
3. **Truth only has interactions**
4. **Truth is linear**
5. **Truth only has main effects**

Simulation Scenarios

1. **Truth obeys strong hierarchy (right in our wheel house):**

$$Y = \sum_{j=1}^{4} f_j(X_j) + \beta_E \cdot X_E + X_E \times (f_3(X_3) + f_4(X_4)) + \varepsilon$$

2. **Truth obeys weak hierarchy**
3. **Truth only has interactions**
4. **Truth is linear**
5. **Truth only has main effects**

- $n_{train} = n_{tuning} = 200$, $n_{test} = 800$, $p = 1000$, $\beta_E = 1$, $SNR = 2$
- $X_j \sim$ `truncnorm(0,1)`, $j = 1, \ldots, 1000$, $E \sim$ `truncnorm(-1,1)`
- `sail` needs to estimate $1000 \times 5 \times 2 = 10$k parameters

Scenario 1: Main Effects for 500 Simulations

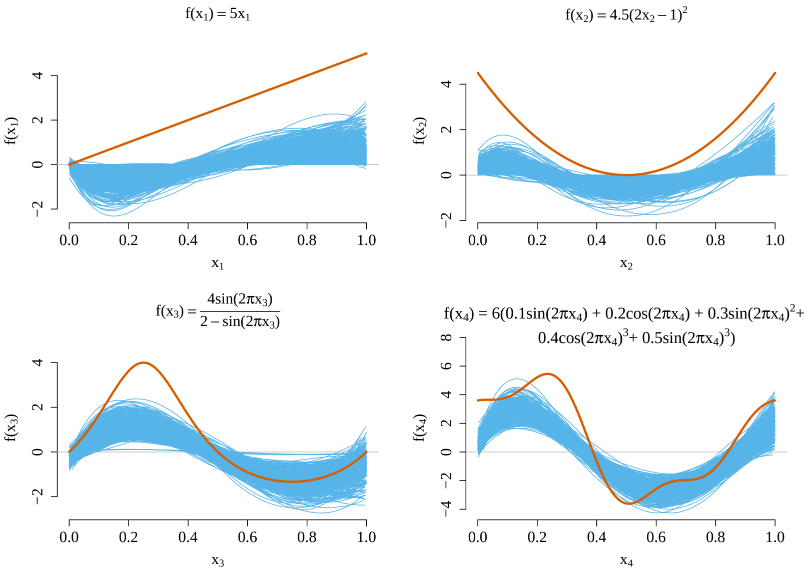

Scenario 1: Estimated Interaction Effects for $E \cdot f(X_3)$

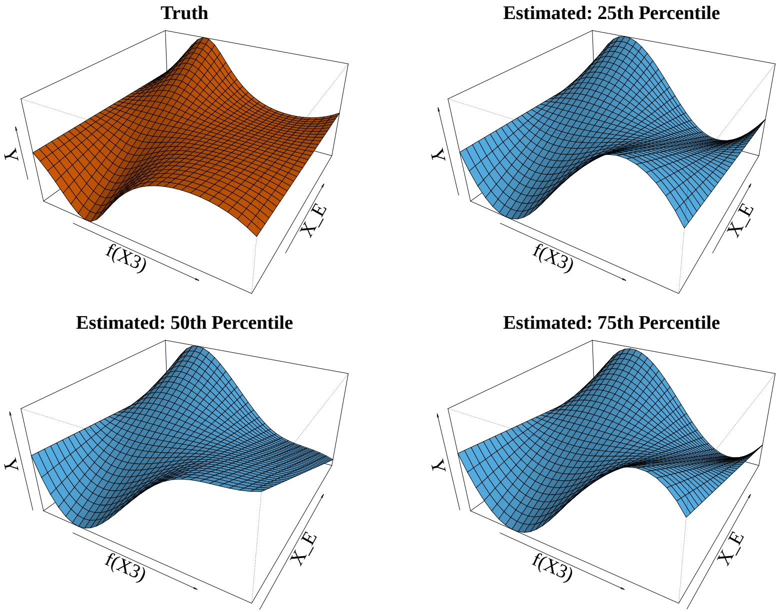

Scenario 1: Estimated Interaction Effects for $E \cdot f(X_4)$

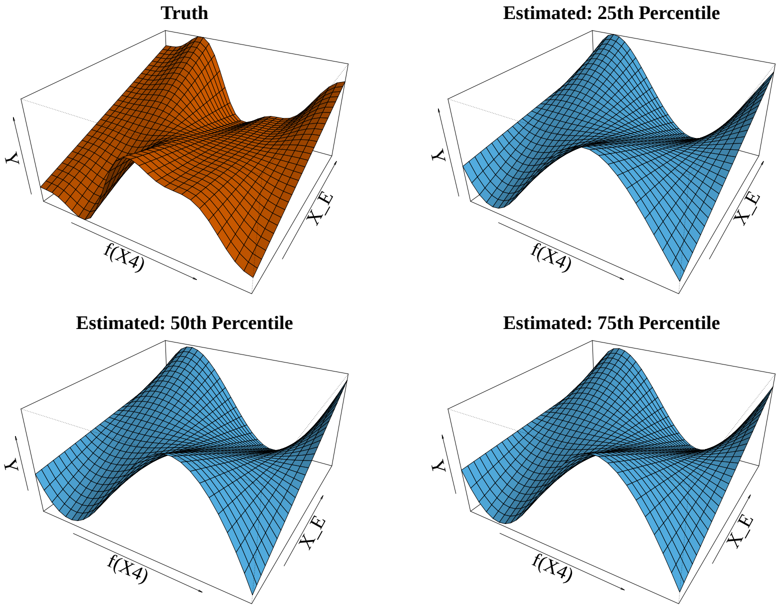

Right in Our Wheel House Simulation Results

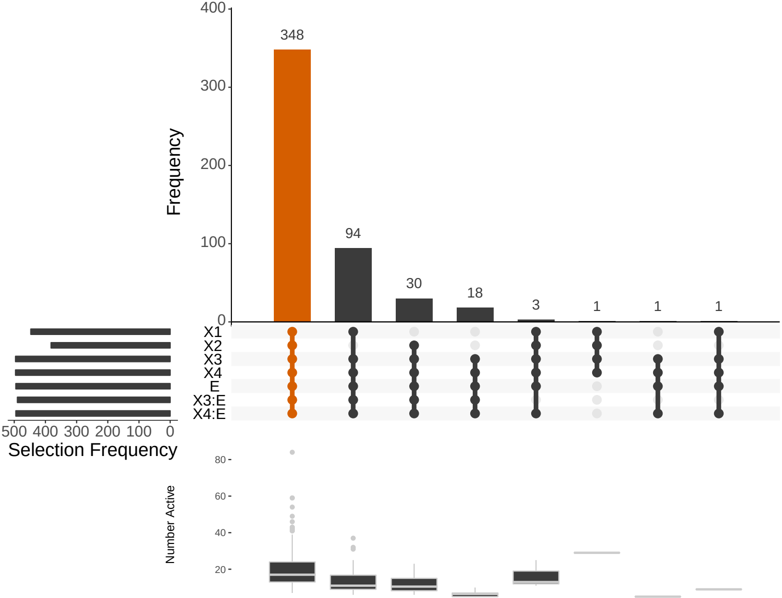

Strong Heredity

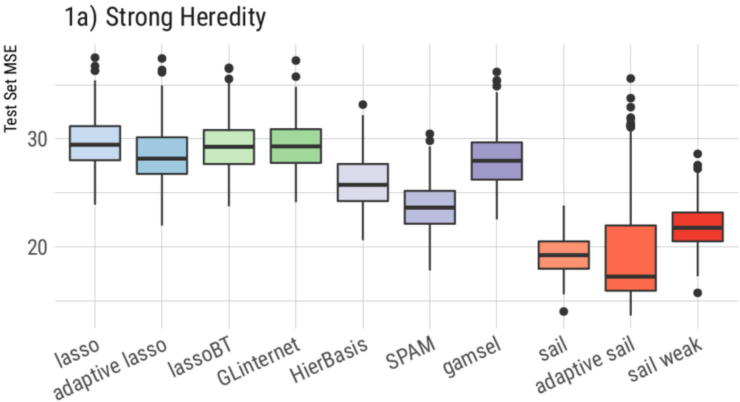

Main Effects Only

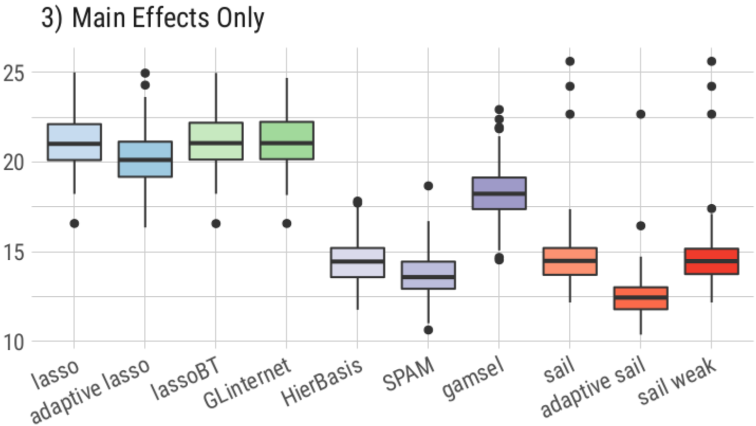

Betting on Sparsity
A Thought Experiment

Motivating Example: The Nurse Family Partnership

`sail`: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

`sail` R package

Real Data Application

Discussion
Current and Future Work

Acknowledgements

sail R package: Solution Path results

```
f.basis <- function(x) splines::bs(x, degree = 5)
fit <- sail(x, y, e, basis = f.basis)
plot(fit)
```

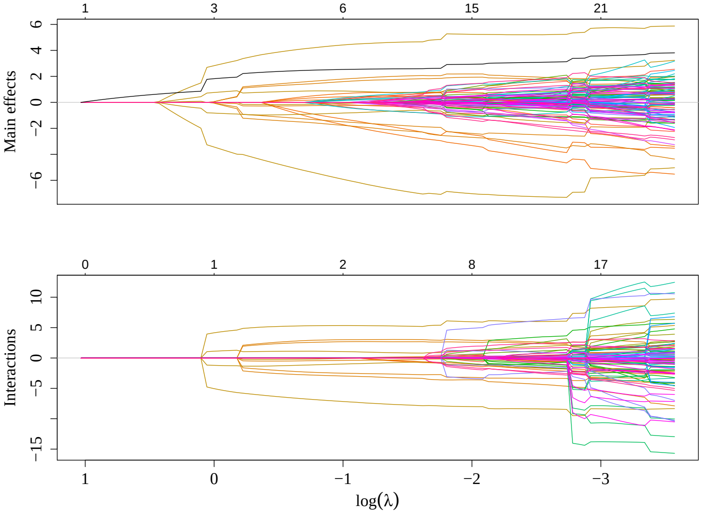

sail R package: Cross-validation results

```r
sail::plot(cvfit)
```

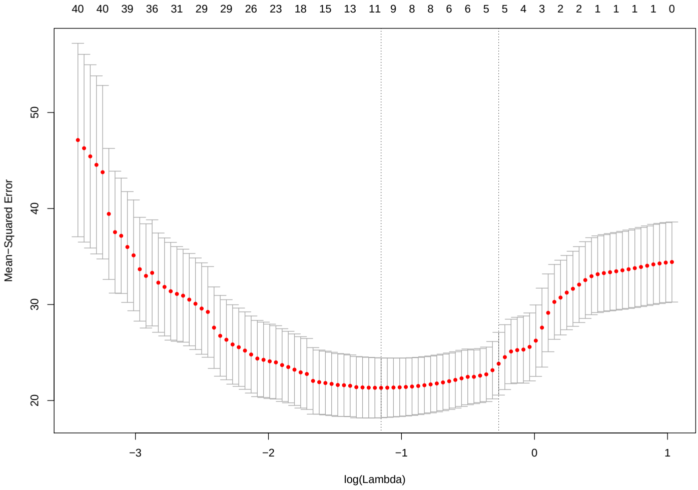

Betting on Sparsity
A Thought Experiment
Motivating Example: The Nurse Family Partnership
`sail`: Strong Additive Interaction Learning
Algorithm
Theory
Simulations
`sail` R package
Real Data Application
Discussion
Current and Future Work
Acknowledgements

Nurse Family Partnership Program

- Early intervention in young children has been shown to positively impact intellectual abilities.
- Genome-wide association studies (GWAS) suggest that 20% of the variance in educational attainment (years of education) may be accounted for by common genetic variation.
- An interesting query that arises is how the environment interacts with these genetics variants to predict measures of cognitive function.

Nurse Family Partnership Program

- Early intervention in young children has been shown to positively impact intellectual abilities.
- Genome-wide association studies (GWAS) suggest that 20% of the variance in educational attainment (years of education) may be accounted for by common genetic variation.
- An interesting query that arises is how the environment interacts with these genetics variants to predict measures of cognitive function.

Nurse Family Partnership Program

- The Stanford Binet IQ scores at 4 years of age were collected for 189 subjects born to women randomly assigned to control (n = 100) or nurse-visited intervention groups (n = 89).

Nurse Family Partnership Program

- The Stanford Binet IQ scores at 4 years of age were collected for 189 subjects born to women randomly assigned to control ($n = 100$) or nurse-visited intervention groups ($n = 89$).
- For each subject, we calculated a polygenic risk score (PRS) for educational attainment at different p-value thresholds using weights from a previous GWAS.

Nurse Family Partnership Program

- The Stanford Binet IQ scores at 4 years of age were collected for 189 subjects born to women randomly assigned to control ($n = 100$) or nurse-visited intervention groups ($n = 89$).
- For each subject, we calculated a polygenic risk score (PRS) for educational attainment at different p-value thresholds using weights from a previous GWAS.
- In this context, individuals with a higher PRS have a propensity for higher educational attainment.

Nurse Family Partnership Program

- The Stanford Binet IQ scores at 4 years of age were collected for 189 subjects born to women randomly assigned to control ($n = 100$) or nurse-visited intervention groups ($n = 89$).
- For each subject, we calculated a polygenic risk score (PRS) for educational attainment at different p-value thresholds using weights from a previous GWAS.
- In this context, individuals with a higher PRS have a propensity for higher educational attainment.
- The goal of this analysis was to determine if there was an interaction between genetic predisposition to educational attainment (X) and maternal participation in the NFP program (E) on child IQ at 4 years of age (Y).

Application of `sail` to NFP data

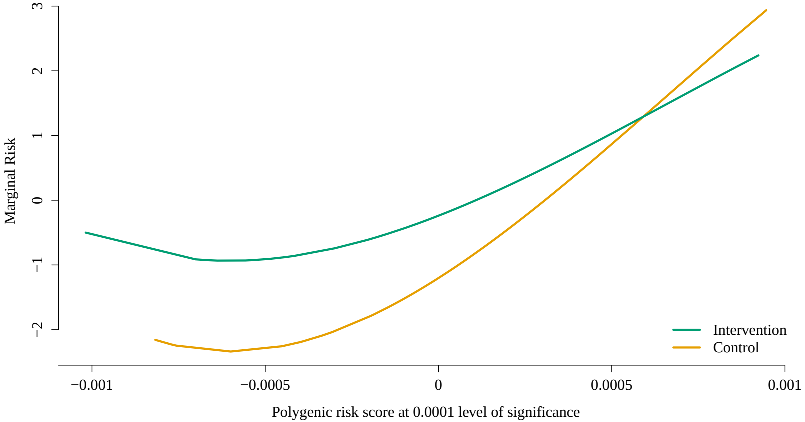

Fig.: The selected model, chosen via 10-fold cross-validation, contained three variables: the main effects for the intervention and the PRS for educational attainment using genetic variants significant at the 0.0001 level, as well as their interaction.

Betting on Sparsity
A Thought Experiment

Motivating Example: The Nurse Family Partnership

`sail`: Strong Additive Interaction Learning

Algorithm

Theory

Simulations

`sail` R package

Real Data Application

Discussion
Current and Future Work

Acknowledgements

Strengths and Limitations

Strengths

- Non-linear environment interactions with strong heredity property in $p >> N$
- `sail` allows for flexible modeling of input variables

Strengths and Limitations

Strengths

- Non-linear environment interactions with strong heredity property in $p >> N$
- `sail` allows for flexible modeling of input variables

Limitations

- `sail` can currently only handle $E \cdot f(X)$ or $f(E) \cdot X$
- Does not allow for $f(X_1, E)$ or $f(X_1, X_2)$
- Memory footprint is an issue

Dynamic Treatment Regimes (DTRs)

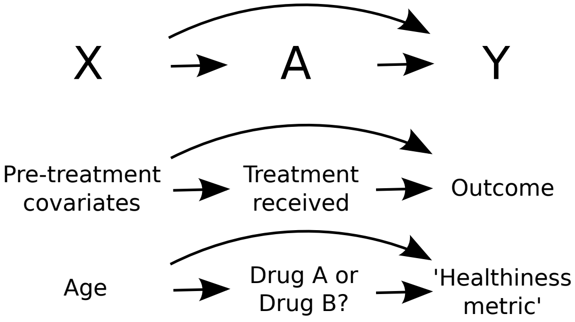

Dynamic Treatment Regimes (DTRs)

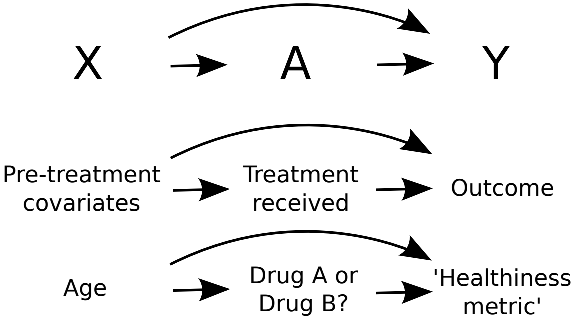

$$\mathbb{E}[Y \mid \mathbf{X}, A; \boldsymbol{\psi}, \boldsymbol{\beta}] = \underbrace{\mathbf{X}\boldsymbol{\beta}}_{\substack{\text{Impact of patient history} \\ \text{in the absence of treatment}}} + \underbrace{\psi_0 A + \boldsymbol{\psi} A \mathbf{X}}_{\substack{\text{Impact of treatment} \\ \text{on outcome}}}$$

Extension of `sail` to DTRs

Cornell University

arXiv.org > stat > arXiv:2101.07359

Statistics > Methodology

[Submitted on 18 Jan 2021]

Variable Selection in Regression-based Estimation of Dynamic Treatment Regimes

Zeyu Bian, Erica EM Moodie, Susan M Shortreed, Sahir Bhatnagar

Dynamic treatment regimes (DTRs) consist of a sequence of decision rules, one per stage of intervention, that finds effective treatments for individual pat
between treatment and a small number of covariates which are often chosen a priori. However, with increasingly large and complex data being collected,
driven approach of selecting these covariates might improve the estimated decision rules and simplify models to make them easier to interpret. We propo
method has the strong heredity property, that is, an interaction term can be included in the model only if the corresponding main terms have also been se
property, and the newly proposed methods compare favorably with other variable selection approaches.

Subjects: **Methodology (stat.ME)**; Computation (stat.CO)
Cite as: arXiv:2101.07359 **[stat.ME]**
(or arXiv:2101.07359v1 **[stat.ME]** for this version)

[1] *In revision at Biometrics.* https://arxiv.org/abs/2101.07359

Hierarchical Penalty Structure

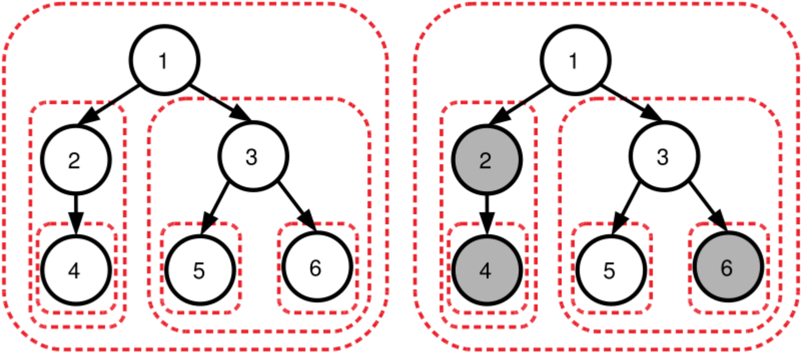

[1] Bach, Jenatton, Mairal and Obozinski (2011). Optimization with Sparsity-Inducing Penalties.

Bi-level selection

- Bi-level selection:

$$f(X_1) = \underbrace{\begin{bmatrix} X_{11} & \psi_{11}(X_{11}) & \psi_{12}(X_{12}) & \cdots & \psi_{11}(X_{15}) \\ & \vdots & \vdots & \cdots & \vdots \\ & \vdots & \vdots & \cdots & \vdots \\ X_{i1} & \psi_{11}(X_{i1}) & \psi_{12}(X_{i2}) & \cdots & \psi_{11}(X_{i5}) \\ & \vdots & \vdots & \cdots & \vdots \\ & \vdots & \vdots & \cdots & \vdots \\ X_{N1} & \psi_{11}(X_{N1}) & \psi_{12}(X_{N2}) & \cdots & \psi_{11}(X_{N5}) \end{bmatrix}_{N\times 5}}_{\mathbf{\Psi_1}} \times \underbrace{\begin{bmatrix} \beta_{\text{linear}} \\ \beta_{11} \\ \beta_{12} \\ \beta_{13} \\ \beta_{14} \\ \beta_{15} \end{bmatrix}_{6\times 1}}_{\theta_1}$$

Betting on Sparsity
A Thought Experiment
Motivating Example: The Nurse Family Partnership
`sail`: Strong Additive Interaction Learning
Algorithm
Theory
Simulations
`sail` R package
Real Data Application
Discussion
Current and Future Work
Acknowledgements

Acknowledgements


Zeyu Bian, PhD (c)

CANSSI INCASS

NSERC CRSNG

CIHR IRSC

MiCM McGill initiative in Computational Medicine

Mitacs

Acknowledgements

- Tianyuan Lu (McGill)
- Yi Yang (McGill)
- Celia Greenwood (Lady Davis Institute)
- Erica Moodie (McGill)
- Kieran O'Donnell (Yale)

Calcul Québec

compute canada | calcul canada

References

1. Bhatnagar, SR, Lu, T, Lovato, A, Olds, DL, Kobor, MS, Meaney, MJ, O'Donnell, K, Yang, Y, and Greenwood, CMT (2021+). A Sparse Additive Model for High-Dimensional Interactions with an Exposure Variable. bioRxiv. DOI 10.1101/445304. *In revision at Computational Statistics and Data Analysis.*
2. **Bian Z**, Moodie EEM, Shortreed S, Bhatnagar SR (2021+). Variable Selection in Regression-based Estimation of Dynamic Treatment Regimes. `https://arxiv.org/abs/2101.07359`. *In revision at Biometrics.*
3. De Leeuw, J. (1994). Block-relaxation algorithms in statistics. In Information systems and data analysis (pp. 308-324). Springer Berlin Heidelberg.
4. Choi, N. H., Li, W., & Zhu, J. (2010). Variable selection with the strong heredity constraint and its oracle property. Journal of the American Statistical Association, 105(489), 354-364.
5. Chipman, H. (1996). Bayesian variable selection with related predictors. Canadian Journal of Statistics, 24(1), 17-36.

`sahirbhatnagar.com`

Session Info

```
R version 4.1.1 (2021-08-10)
Platform: x86_64-pc-linux-gnu (64-bit)
Running under: Pop!_OS 21.04

Matrix products: default
BLAS:   /usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3
LAPACK: /usr/lib/x86_64-linux-gnu/openblas-pthread/libopenblasp-r0.3.13.so

attached base packages:
[1] stats     graphics  grDevices utils     datasets  methods   base

other attached packages:
[1] xtable_1.8-4      rpart.plot_3.1.0  rpart_4.1-15      data.table_1.14.2
[5] ISLR_1.2          ggplot2_3.3.5     knitr_1.36

loaded via a namespace (and not attached):
 [1] pillar_1.6.4       compiler_4.1.1     highr_0.9          tools_4.1.1
 [5] digest_0.6.28      evaluate_0.14      lifecycle_1.0.1    tibble_3.1.5
 [9] gtable_0.3.0       pkgconfig_2.0.3    rlang_0.4.12       DBI_1.1.1
[13] xfun_0.26          withr_2.4.2        dplyr_1.0.7        stringr_1.4.0
[17] generics_0.1.0     vctrs_0.3.8        grid_4.1.1         tidyselect_1.1.1
[21] glue_1.4.2         R6_2.5.1           fansi_0.5.0        pacman_0.5.1
[25] purrr_0.3.4        RSkittleBrewer_1.1 blob_1.2.1         magrittr_2.0.1
[29] scales_1.1.1       ellipsis_0.3.2     assertthat_0.2.1   colorspace_2.0-2
[33] utf8_1.2.2         stringi_1.7.5      munsell_0.5.0      crayon_1.4.1
```

B-Spline Expansion

```
x <- truncnorm::rtruncnorm(1000, a = 0, b = 1)
B <- splines::bs(x, df = 5, degree=3, intercept = FALSE)
```

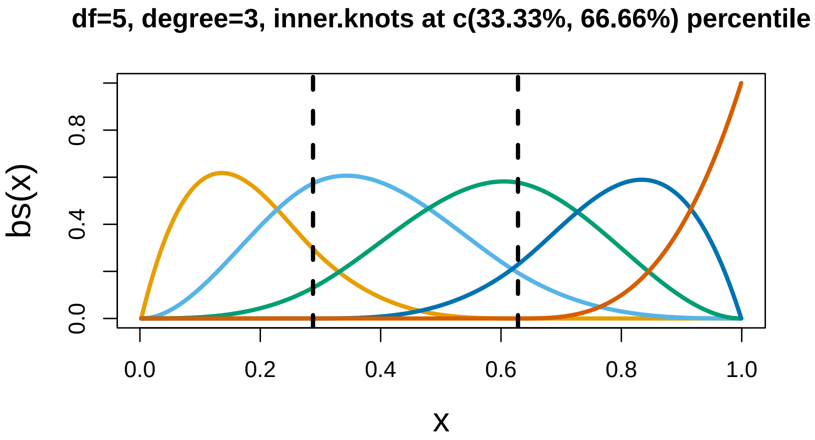

sail A Note on the Second Tuning Parameter results

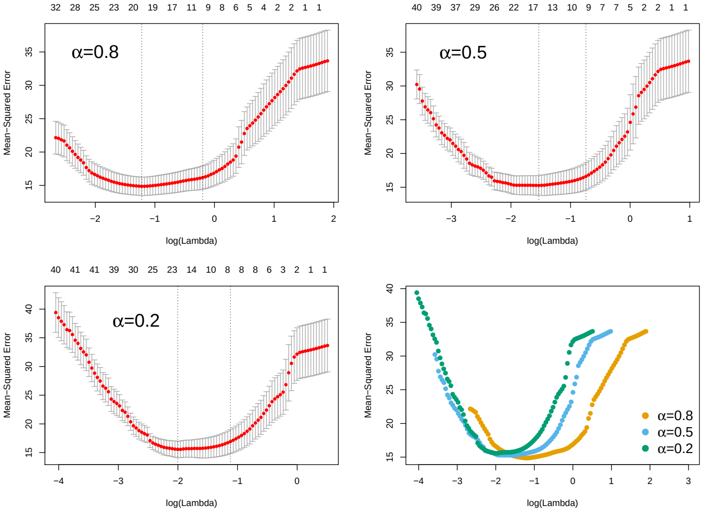

Why the L1 norm ?

- For a fixed real number $q \geq 0$ consider the criterion

$$\tilde{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} x_{ij}\beta_j \right)^2 + \lambda \sum_{j=1}^{p} |\beta_j|^q \right\}$$

- Why do we use the ℓ_1 norm? Why not use the $q = 2$ (Ridge) or any ℓ_q norm?

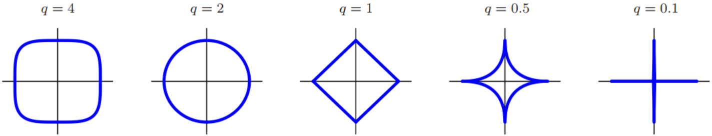

- $q = 1$ is the smallest value that yields a sparse solution and yields a **convex** problem $\rightarrow$ scalable to high-dimensional data
- For $q < 1$ the constrained region is **nonconvex**

Linear Effects Simulation - Comparison