

Randomized Least Squares Regression: Combining Model- and Algorithm-Induced Uncertainties

Ilse C.F. Ipsen

Joint work with: Jocelyn T. Chi

North Carolina State University
Raleigh, NC, USA

Research supported in part by NSF DGE and DMS

Least Squares/Regression Models

Given: Design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with $\text{rank}(\mathbf{X}) = p$

1 Gaussian linear regression model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$$

Unknown parameter vector $\boldsymbol{\beta}_0$

2 Least squares problem

$$\min_{\boldsymbol{\beta}} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{y}\|_2$$

Unique maximum likelihood estimator $\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$

3 Randomized (row compression) algorithm

$$\min_{\boldsymbol{\beta}} \|\mathbf{S}(\mathbf{X}\boldsymbol{\beta} - \mathbf{y})\|_2$$

Minimal norm solution $\tilde{\boldsymbol{\beta}} = (\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{y})$

Objective

Determine **combined** mean and variance of $\tilde{\beta}$ w.r.t.

- Gaussian linear model $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$
- Randomized row compression \mathbf{S}

Inspiration

Ping Ma, Michael Mahoney, Bin Yu

A statistical perspective on algorithmic leveraging

J. Mach. Learn. Res., vol. 16, pp 861-911 (2015)

Overview

Existing work

Examples

Structural perturbation bounds

Model-induced uncertainty, conditioned on algorithm-induced uncertainty

Combined model-induced and algorithm-induced uncertainty

Summary

Existing Work

(Numerical) row-compression methods for least squares

Drineas, Mahoney, Muthukrishnan 2006

Zhou, Lafferty, Wasserman 2007

Boutsidis, Drineas 2009

Drineas, Mahoney, Muthukrishnan, Sarlós 2011

Meng, Saunders, Mahoney 2014

Bartels, Hennig 2016

Becker, Kawas, Petrick, Ramamurthy 2017

Statistical properties

Ma, Mahoney, Yu 2015

Raskutti, Mahoney 2016

Ahfock, Astle, Richardson 2017

Thanei, Heinze, Meinshausen 2017

Lopes, Wang, Mahoney 2018

Examples

Model-Induced Uncertainty

Gaussian linear model: $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}_n)$

Design matrix:

$$\mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

Left inverse: $\mathbf{X}^\dagger \mathbf{X} = \mathbf{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$

Maximum likelihood estimator: $\hat{\boldsymbol{\beta}} = \mathbf{X}^\dagger \mathbf{y}$

Unbiased estimator: $\mathbb{E}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \mathbf{X}^\dagger \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}^\dagger \mathbf{X} \boldsymbol{\beta}_0 = \boldsymbol{\beta}_0$

Variance: $\text{Var}_{\mathbf{y}}[\hat{\boldsymbol{\beta}}] = \mathbf{X}^\dagger \text{Var}_{\mathbf{y}}[\mathbf{y}] (\mathbf{X}^\dagger)^T = \sigma^2 \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix}$

Model-Induced Uncertainty Conditioned on Algorithm-Induced Uncertainty (1)

$\min_{\beta} \|\mathbf{S}\mathbf{X}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$, $\mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}\beta_0$

Sketching preserves rank: $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$

$$\mathbf{S}\mathbf{X} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = (\mathbf{S}\mathbf{X})^\dagger$$

Unbiased estimator:

$$\mathbb{E}_{\mathbf{y}} \left[\tilde{\beta} \mid \mathbf{S} \right] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{X}\beta_0 = \beta_0$$

Variance has increased:

$$\text{Var}_{\mathbf{y}} \left[\tilde{\beta} \mid \mathbf{S} \right] = \sigma^2 (\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \left((\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \right)^T = \sigma^2 \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \succcurlyeq \sigma^2 \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix}$$

Model-Induced Uncertainty Conditioned on Algorithm-Induced Uncertainty (2)

$\min_{\beta} \|\mathbf{S}(\mathbf{X}\beta - \mathbf{y})\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$, $\mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}\beta_0$

Sketching causes loss of rank: $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$

$$\mathbf{S}\mathbf{X} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = (\mathbf{S}\mathbf{X})^\dagger$$

Biased estimator:

$$\mathbb{E}_{\mathbf{y}} [\tilde{\beta} \mid \mathbf{S}] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{X}\beta_0 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \beta_0 \neq \beta_0$$

Variance is singular:

$$\text{Var}_{\mathbf{y}} [\tilde{\beta} \mid \mathbf{S}] = \sigma^2 (\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \left((\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \right)^T = \sigma^2 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

Model-Induced Uncertainty Conditioned on Algorithm-Induced Uncertainty (3)

$\min_{\beta} \|\mathbf{S}(\mathbf{X}\beta - \mathbf{y})\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$, $\mathbb{E}_{\mathbf{y}}[\mathbf{y}] = \mathbf{X}\beta_0$

Summary: Model-induced uncertainty conditioned on \mathbf{S}

- Sketching **preserves rank**: $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$

Left inverse: $(\mathbf{S}\mathbf{X})^\dagger = (\mathbf{X}^T \mathbf{S}^T \mathbf{S}\mathbf{X})^{-1} (\mathbf{S}\mathbf{X})^T$

$\tilde{\beta}$ is an **unbiased** estimator: $\mathbb{E}_{\mathbf{y}}[\tilde{\beta} \mid \mathbf{S}] = \beta_0$

- Sketching causes **loss of rank**: $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$

No left inverse: $(\mathbf{X}^T \mathbf{S}^T \mathbf{S}\mathbf{X})^{-1}$ does not exist

$\tilde{\beta}$ is a **biased** estimator: $\mathbb{E}_{\mathbf{y}}[\tilde{\beta} \mid \mathbf{S}] \neq \beta_0$

Variance $\text{Var}_{\mathbf{y}}[\tilde{\beta} \mid \mathbf{S}]$ is **singular**

Combined Uncertainty: Uniform Sampling with Replacement (1)

Sampling 2 out of 4 rows: $\mathbf{s}_{ij} = \sqrt{2} \begin{pmatrix} \mathbf{e}_i^T \\ \mathbf{e}_j^T \end{pmatrix}$, $1 \leq i, j \leq 4$

$$\mathbf{s}_{11}\mathbf{X} = \sqrt{2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} = \sqrt{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

$$\mathbf{s}_{42}\mathbf{X} = \sqrt{2} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix} = \sqrt{2} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

Unbiased estimator of the identity:

$$\mathbb{E}_{\mathbf{s}}[\mathbf{S}^T \mathbf{S}] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} \mathbf{s}_{ij}^T \mathbf{s}_{ij} = \mathbf{I} = \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}$$

Combined Uncertainty: Uniform Sampling with Replacement (2)

$\min_{\beta} \|\mathbf{S}(\mathbf{X}\beta - \mathbf{y})\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$

Conditional mean: $\mathbb{E}_{\mathbf{y}}[\tilde{\beta} \mid \mathbf{S}] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{X}\beta_0$

Sequential conditioning:

$$\mathbb{E}[\tilde{\beta}] = \mathbb{E}_{\mathbf{s}} \left[\mathbb{E}_{\mathbf{y}} \left[\tilde{\beta} \mid \mathbf{S} \right] \right] = \mathbb{E}_{\mathbf{s}} \left[(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{X} \right] \beta_0$$

Biased estimator:

$$\mathbb{E}_{\mathbf{s}} \left[(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{X} \right] = \sum_{i=1}^4 \sum_{j=1}^4 \frac{1}{16} (\mathbf{S}_{ij}\mathbf{X})^\dagger (\mathbf{S}_{ij}\mathbf{X}) = \frac{1}{16} \begin{pmatrix} 12 & 0 \\ 0 & 7 \end{pmatrix} \neq \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

12 out of 16 sketched matrices $\mathbf{S}_{ij}\mathbf{X}$ are rank deficient

Structural Perturbation Bounds

Perturbed Solution

Exact problem $\min_{\beta} \|\mathbf{X}\beta - \mathbf{y}\|_2$

Hat matrix $\mathbf{P}_x = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T = \mathbf{X} \mathbf{X}^\dagger$

Range $\mathcal{R}(\mathbf{P}_x) = \mathcal{R}(\mathbf{X})$

Solution $\hat{\beta} = \mathbf{X}^\dagger \mathbf{y} = \mathbf{X}^\dagger \mathbf{P}_x \mathbf{y}$

Perturbed problem $\min_{\beta} \|\mathbf{S}(\mathbf{X}\beta - \mathbf{y})\|_2$

$(\mathbf{X}^T \mathbf{S}^T \mathbf{S} \mathbf{X})^{-1}$ does not work!

Comparison Hat matrix* $\mathbf{P} = \mathbf{X}(\mathbf{S} \mathbf{X})^\dagger \mathbf{S}$

Range $\mathcal{R}(\mathbf{P}) \subset \mathcal{R}(\mathbf{X}) = \mathcal{R}(\mathbf{P}_x)$

If $\text{rank}(\mathbf{S} \mathbf{X}) = \text{rank}(\mathbf{X})$ then $\mathcal{R}(\mathbf{P}) = \mathcal{R}(\mathbf{X})$

Solution $\tilde{\beta} = (\mathbf{S} \mathbf{X})^\dagger \mathbf{S} \mathbf{y} = \mathbf{X}^\dagger \mathbf{P} \mathbf{y}$

*More general than [Raskutti, Mahoney 2016]

Example: Hat Matrix, and Comparison Hat Matrix

Design matrix, and Hat matrix:

$$\mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 1 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Sketching matrix with $\text{rank}(\mathbf{SX}) = \text{rank}(\mathbf{X})$, and Comparison Hat matrix:

$$\mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{P} = \mathbf{X}(\mathbf{SX})^\dagger \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathcal{R}(\mathbf{P}) = \mathcal{R}(\mathbf{X})$$

Sketching matrix with $\text{rank}(\mathbf{SX}) < \text{rank}(\mathbf{X})$, and Comparison Hat matrix:

$$\mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{P} = \mathbf{X}(\mathbf{SX})^\dagger \mathbf{S} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathcal{R}(\mathbf{P}) \subset \mathcal{R}(\mathbf{X})$$

Perturbed Solution

Exact problem $\min_{\beta} \|\mathbf{X}\beta - \mathbf{y}\|_2$

Hat matrix $\mathbf{P}_x = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T = \mathbf{X} \mathbf{X}^\dagger$

Solution $\hat{\beta} = \mathbf{X}^\dagger \mathbf{P}_x \mathbf{y}$

Perturbed problem $\min_{\beta} \|\mathbf{S}(\mathbf{X}\beta - \mathbf{y})\|_2$

Comparison Hat matrix $\mathbf{P} = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}$

Solution $\tilde{\beta} = \mathbf{X}^\dagger \mathbf{P} \mathbf{y}$

Difference between perturbed and exact solution

$$\tilde{\beta} = \hat{\beta} + \mathbf{X}^\dagger (\mathbf{P} - \mathbf{P}_x) \mathbf{y}$$

proportional to difference between Hat and Comparison Hat matrix

Multiplicative Perturbation Bounds

Ingredients:

Condition number: $\kappa_2(\mathbf{X}) = \|\mathbf{X}\|_2 \|\mathbf{X}^\dagger\|_2$

Angle between \mathbf{y} and $\text{range}(\mathbf{X})$: $0 < \theta < \pi/2$

Relative error in perturbed solution:

$$\frac{\|\tilde{\hat{\beta}} - \hat{\beta}\|_2}{\|\hat{\beta}\|_2} \leq \underbrace{\frac{\kappa_2(\mathbf{X})}{\cos \theta}}_{\text{Amplifier}} \underbrace{\|\mathbf{P} - \mathbf{P}_x\|_2}_{\text{Perturbation}}$$

Least squares solution **insensitive** to multiplicative perturbations, if

- 1 Matrix \mathbf{X} **well-conditioned** with respect to (left) inversion
- 2 Righthand side \mathbf{y} **close** to $\text{range}(\mathbf{X})$

Tighter than [Drineas, Mahoney, Muthukrishnan, Sarlós 2011]

Model-Induced Uncertainty, Conditioned on Algorithm-Induced uncertainty

Conditioning on \mathbf{S} : Mean

$\min_{\beta} \|\mathbf{S}\mathbf{X}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$

Conditional mean

$$\mathbb{E}_{\mathbf{y}} [\tilde{\beta} \mid \mathbf{S}] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \mathbb{E}_{\mathbf{y}}[\mathbf{y}] = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S} \mathbf{X} \beta_0$$

- If $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$ then $\tilde{\beta}$ is unbiased estimator

$$\mathbb{E}_{\mathbf{y}} [\tilde{\beta} \mid \mathbf{S}] = \underbrace{(\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X})}_I \beta_0 = \beta_0$$

- If $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$ then $\tilde{\beta}$ is biased estimator

$$\mathbb{E}_{\mathbf{y}} [\tilde{\beta} \mid \mathbf{S}] = \beta_0 + \underbrace{\left(I - (\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) \right)}_{\text{Rank deficiency of } \mathbf{S}\mathbf{X}} \beta_0$$

Bias of $\tilde{\beta}$ increases with rank deficiency of $\mathbf{S}\mathbf{X}$

Conditioning on \mathbf{S} : Variance

$\min_{\beta} \|\mathbf{S}\mathbf{X}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = \mathbf{X}^\dagger \mathbf{P} \mathbf{y}$

Ingredients:

Hat matrix: $\mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger$

Comparison Hat matrix: $\mathbf{P} = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}$

Model variance: $\text{Var}_{\mathbf{y}}[\hat{\beta}] = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1} = \sigma^2 \mathbf{X}^\dagger (\mathbf{X}^\dagger)^T$

Conditional variance:

$$\begin{aligned} \text{Var}_{\mathbf{y}} \left[\tilde{\beta} \mid \mathbf{S} \right] &= \sigma^2 \mathbf{X}^\dagger \mathbf{P} \mathbf{P}^T (\mathbf{X}^\dagger)^T \\ &= \text{Var}_{\mathbf{y}}[\hat{\beta}] + \underbrace{\sigma^2 \mathbf{X}^\dagger (\mathbf{P} \mathbf{P}^T - \mathbf{P}_x) (\mathbf{X}^\dagger)^T}_{\text{Deviation of } \mathbf{P} \text{ from orthogonal projector}} \end{aligned}$$

\mathbf{P}_x is orthogonal projector onto $\mathcal{R}(\mathbf{X})$ with $\mathbf{P}_x^T = \mathbf{P}_x = \mathbf{P}_x^2$

Conditioning on \mathbf{S} : Summary

$\min_{\beta} \|\mathbf{S}\mathbf{X}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = \mathbf{X}^\dagger \mathbf{P} \mathbf{y}$

Comparison Hat matrix $\mathbf{P} = \mathbf{X}(\mathbf{S}\mathbf{X})^\dagger \mathbf{S}$

Model-induced uncertainty of $\tilde{\beta}$ conditioned on \mathbf{S}
governed by $\text{rank}(\mathbf{S}\mathbf{X})$

- Bias increases with deviation of $\mathbf{S}\mathbf{X}$ from full column-rank
- If $\text{rank}(\mathbf{S}\mathbf{X}) = \text{rank}(\mathbf{X})$ then $\tilde{\beta}$ is unbiased
- Conditional variance close to model variance, if
 - \mathbf{P} close to being an orthogonal projector onto $\mathcal{R}(\mathbf{X})$
 - \mathbf{X} well-conditioned with respect to inversion
- If $\text{rank}(\mathbf{S}\mathbf{X}) < \text{rank}(\mathbf{X})$ then $\text{Var}_{\mathbf{y}}[\tilde{\beta} \mid \mathbf{S}]$ is singular

Combined Model-induced and Algorithm-Induced Uncertainty

Combined Uncertainty: Mean

$\min_{\beta} \|\mathbf{SX}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = (\mathbf{SX})^\dagger \mathbf{S}\mathbf{y}$

Total mean

$$\mathbb{E}[\tilde{\beta}] = \beta_0 + \mathbb{E}_s \left[(\mathbf{SX})^\dagger (\mathbf{SX}) - \mathbf{I} \right] \beta_0$$

- Deviation of combined uncertainties from model-induced uncertainties governed by
expected deviation of sketched matrix from rank deficiency
- Total bias proportional to
expected deviation of \mathbf{SX} from having full column-rank

Combined Uncertainty: Variance

$\min_{\beta} \|\mathbf{S}\mathbf{X}\beta - \mathbf{S}\mathbf{y}\|_2$ has solution $\tilde{\beta} = (\mathbf{S}\mathbf{X})^\dagger \mathbf{S}\mathbf{y}$

Deviation of total variance from model variance

$$\begin{aligned} \text{Var}[\tilde{\beta}] = \text{Var}_{\mathbf{y}}[\hat{\beta}] &+ \sigma^2 \mathbf{X}^\dagger \mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}] (\mathbf{X}^\dagger)^T \\ &+ \text{Var}_{\mathbf{s}} \left[\left((\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) - \mathbf{I} \right) \beta_0 \right] \end{aligned}$$

- $\mathbf{X}^\dagger \mathbb{E}_{\mathbf{s}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_{\mathbf{x}}] (\mathbf{X}^\dagger)^T$
Expected deviation of \mathbf{P} from orthogonal projector onto $\mathcal{R}(\mathbf{X})$, amplified by conditioning of \mathbf{X}
- $\text{Var}_{\mathbf{s}} \left[\left((\mathbf{S}\mathbf{X})^\dagger (\mathbf{S}\mathbf{X}) - \mathbf{I} \right) \beta_0 \right]$
Expected deviation of $\mathbf{S}\mathbf{X}$ from having full column-rank

Example: Best Case for Uniform Sampling

Columns of Hadamard matrix: Best Coherence

$$\mathbf{X} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

\mathcal{S} samples 2 rows uniformly and with replacement

- Expected deviation of \mathbf{SX} from full column-rank

$$\mathbb{E}_{\mathcal{S}} [(\mathbf{SX})^\dagger(\mathbf{SX}) - \mathbf{I}] = -\frac{4}{16} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

- Expected deviation of \mathbf{P} from orthogonal projector onto $\mathcal{R}(\mathbf{X})$

$$\mathbb{E}_{\mathcal{S}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x] = \frac{3}{16} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix}$$

Example: Worst Case for Uniform Sampling

Columns of identity matrix: Worst Coherence

$$\mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{P}_x = \mathbf{X}\mathbf{X}^\dagger = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

\mathbf{S} samples 2 rows uniformly and with replacement

- Expected deviation of $\mathbf{S}\mathbf{X}$ from full column-rank

$$\mathbb{E}_{\mathbf{S}} [(\mathbf{S}\mathbf{X})^\dagger(\mathbf{S}\mathbf{X}) - \mathbf{I}] = -\frac{9}{16} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

- Expected deviation of \mathbf{P} from orthogonal projector onto $\mathcal{R}(\mathbf{X})$

$$\mathbb{E}_{\mathbf{S}}[\mathbf{P}\mathbf{P}^T - \mathbf{P}_x] = \frac{9}{16} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Summary

- Randomized (row) sketching for full column-rank regression
- **Exact** expressions for uncertainties, induced by model and algorithm, under **very general** assumptions
- Introduced **Comparison Hat** matrix, to allow comparison between problems of different dimensions
- Tighter multiplicative perturbation bounds
- Total mean and variance governed by **expected deviation of**
 Sketched matrix from full column-rank
 Comparison Hat matrix from orthogonal projector
- Examples illustrate applicability
- J.T. Chi and I.C.F. Ipsen,
Randomized Least Squares Regression: Combining Model- and Algorithm-Induced Uncertainties, arXiv:1808.0594