Quantile-based Iterative Methods for Corrupted Systems of Linear Equations

Liza Rebrova

UCLA, Department of Mathematics & LBL, Computational Research Division

Applied Mathematics & Statistics seminar
Johns Hopkins University
February 12, 2021

What do I do?

- Study the structure of large high-dimensional objects in the presence of randomness
- Use this understanding to develop (randomized) methods that work with complex data efficiently

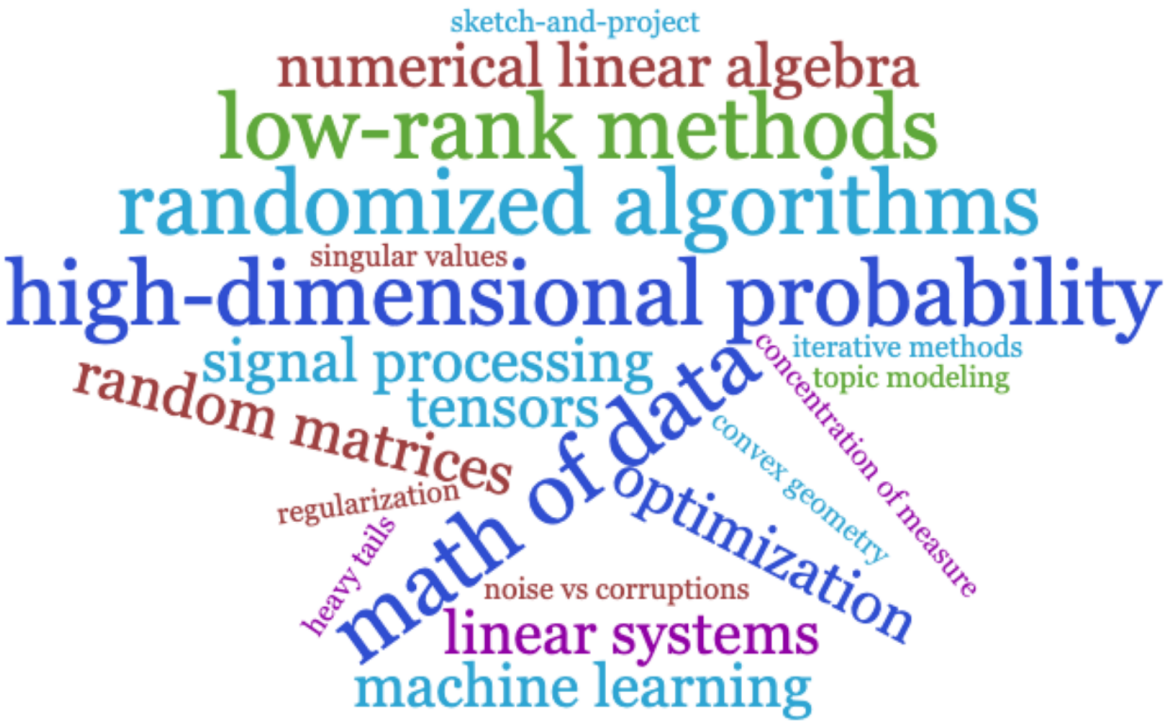

Large high-dimensional objects:

- Sets
- Matrices
- Tensors
- Graphs
- Systems of linear equations
- Neural nets
- ...

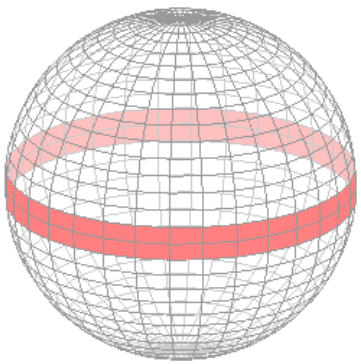

Concentration of measure phenomenon

High-dimensional probability helps revealing their structure:

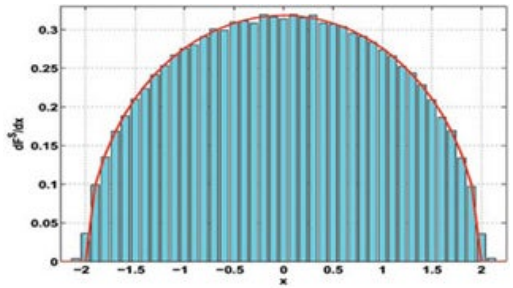

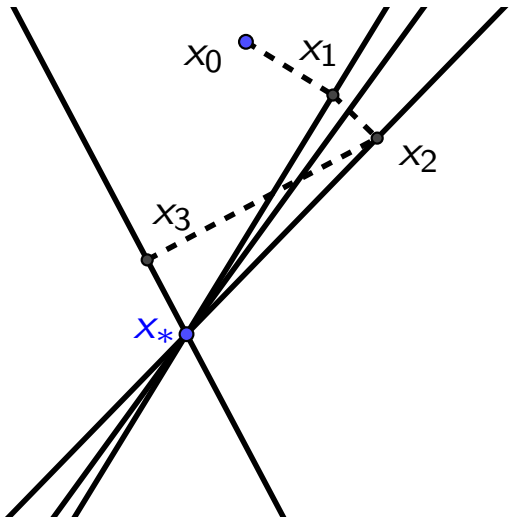

Solving linear systems

$$A\mathbf{x} = \mathbf{b}, \qquad \text{find } \mathbf{x} \in \mathbb{R}^n \text{ given } A \in \mathbb{R}^{m\times n}, \mathbf{b} \in \mathbb{R}^m.$$

Solving linear systems

$$A\mathbf{x} = \mathbf{b}, \qquad \text{find } \mathbf{x} \in \mathbb{R}^n \text{ given } A \in \mathbb{R}^{m\times n}, \mathbf{b} \in \mathbb{R}^m.$$

- Direct methods: $\mathbf{x} = A^\dagger \mathbf{b}$, LU, Cholesky, ...
- Methods exploiting the structure of A: H-, HSS solvers, ...
- Iterative methods: SGD, GMRES, projection based methods, ...
- Sketch-and-project (Gower and Richtarik, 2015)

Solving linear systems

$$A\mathbf{x} = \mathbf{b}, \qquad \text{find } \mathbf{x} \in \mathbb{R}^n \text{ given } A \in \mathbb{R}^{m\times n}, \mathbf{b} \in \mathbb{R}^m.$$

- Direct methods: $\mathbf{x} = A^\dagger \mathbf{b}$, LU, Cholesky, ...
- Methods exploiting the structure of A: H-, HSS solvers, ...
- Iterative methods: SGD, GMRES, projection based methods, ...
- Sketch-and-project (Gower and Richtarik, 2015)

Disclaimer: I will focus on overdetermined full-rank systems $m \geq n$ either we assume that solution exists (redundant equations) or not (then we search for a least square solution)

Model: overdetermined linear system

$$A \cdot x_* = b, \quad m > n$$

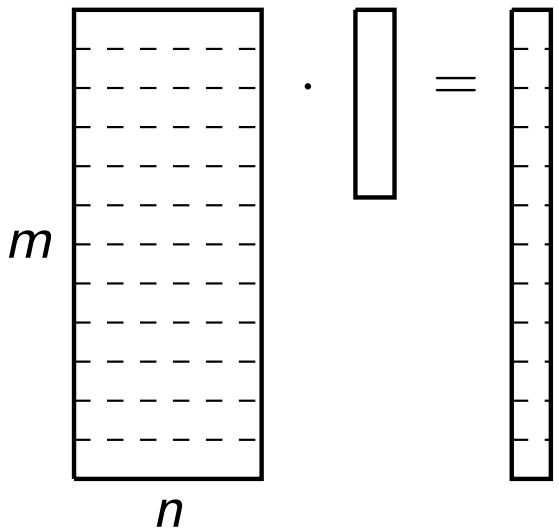

Sketch-and-project

instead of $A\mathbf{x} = b$, solve $S^T A\mathbf{x} = S^T b$

$S = m \times s$ sketch matrix, if $s \ll m$ (sketched system is easier)
Iteration:

$$\mathbf{x}_k = \mathbf{x}_{k-1} + (S^T A)^\dagger (S^T \mathbf{b} - S^T A\mathbf{x}_k)$$

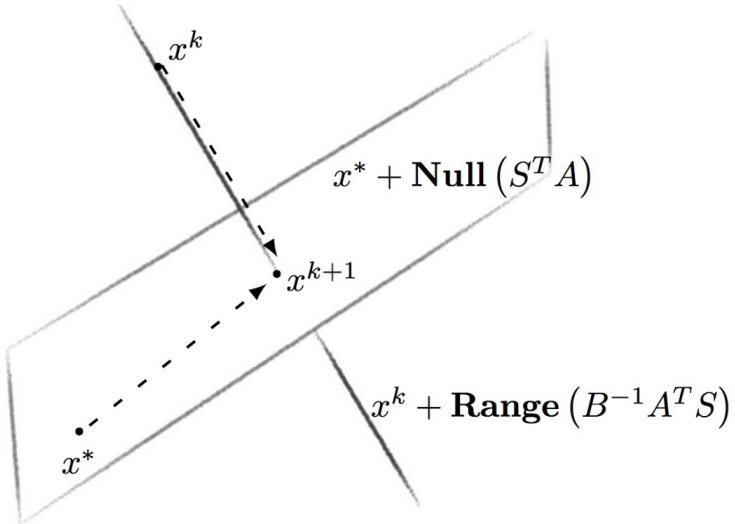

Discrete random sketches and Kaczmarz methods

$$A_i = (0, \ldots, 0, 1, 0, \ldots, 0) \cdot A$$

$$A_\tau = \left[\begin{array}{c|c|c} 0 & \mathrm{Id} & 0 \end{array}\right] \cdot A = S^T A; \quad \mathbf{b}_\tau = S^T b$$

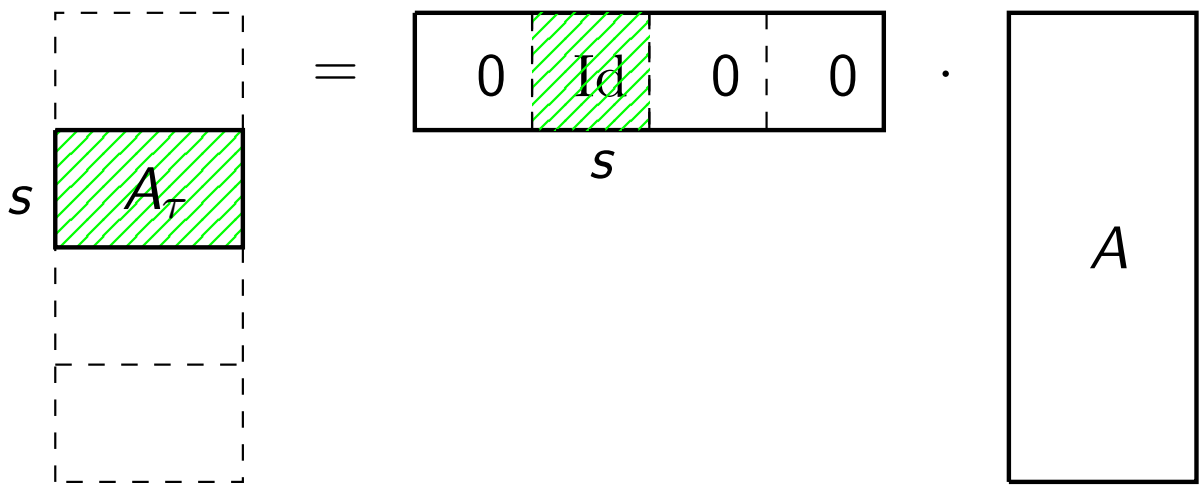

Sketch-and-project methods with $S =$ (randomly placed identity completed by zeroes) are randomized Kaczmarz methods

Randomized Kaczmarz (RK) method

Assume that for all the rows $\|\mathbf{A_i}\| = 1$.

1. initialize $\mathbf{x_0}$
2. project current iterate to $\mathbf{A_i}$:
   $\mathbf{x}_{k+1} = \mathbf{x}_k - (\langle \mathbf{A_i}, \mathbf{x}_k \rangle - \mathbf{b}_i)\mathbf{A_i}$,
   where $i \sim \mathit{Unif}\{1, \ldots, m\}$

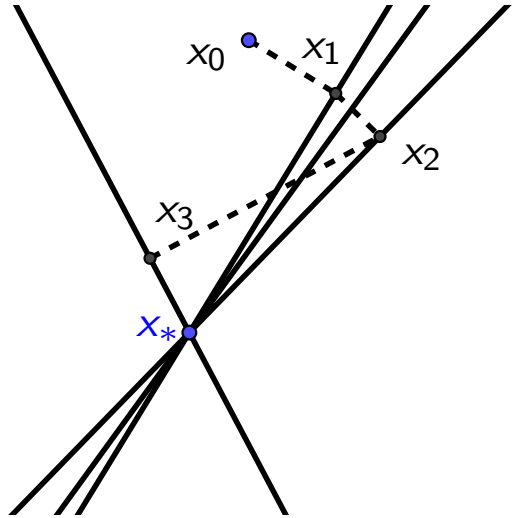

Geometrically, each index i corresponds to a hyperplane in $\mathbb{R}^n$. RK projects orthogonally onto a randomly chosen hyperplane

Convegence theorem (Strohmer - Vershynin 2009)

For a system $A\mathbf{x}_* = b$, RK converges to x_* linearly in expectation:

$$\mathbb{E}||\mathbf{x}_k - \mathbf{x}_*||_2^2 \leq \left(1 - \frac{\sigma_{min}^2(A)}{\|A\|_F^2}\right)^k ||\mathbf{x}_0 - \mathbf{x}_*||_2^2.$$

Block Kaczmarz Method

Starting at $\mathbf{x}_0 \in \mathbb{R}^n$:

1. Choose A_τ a block row subset at random, $\tau = \tau(k) \subset [m]$, $|\tau| = s$
2. Define $\mathbf{x}_k := \mathbf{x}_{k-1} + (A_\tau)^\dagger(\mathbf{b}_\tau - A_\tau \mathbf{x}_k)$

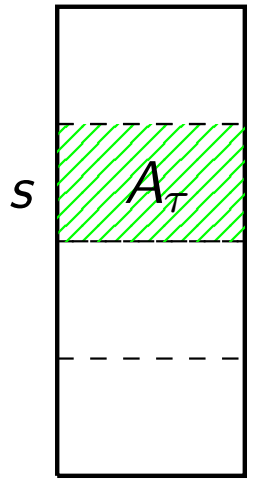

Convegence theorem (Needell - Tropp 2012)

The block Kaczmarz converges to $\mathbf{x}_*$ *in expectation with accelerated rate*

$$\mathbb{E}||\mathbf{x}_k - \mathbf{x}_*||_2^2 \leq \left(1 - c\frac{\sigma_{min}^2(A)}{||A||^2 \log m}\right)^k ||\mathbf{x}_0 - \mathbf{x}_*||_2^2,$$

if all blocks are well-conditioned: for some $\delta \in (0,1)$,
number of blocks $\cdot \max_\tau \|A_\tau\|_2^2 \lesssim \|A\|_2^2 \log(m)\frac{1}{\delta^2} \cdot (1+\delta)$.

Gaussian sketching

$$A_{xi} := \xi^T \cdot A, \text{ where } \xi \sim N(0, \mathrm{Id})$$

$$A_S := S^T \cdot A, \text{ where } S \text{ is } m \times s \text{ gaussian random matrix}$$

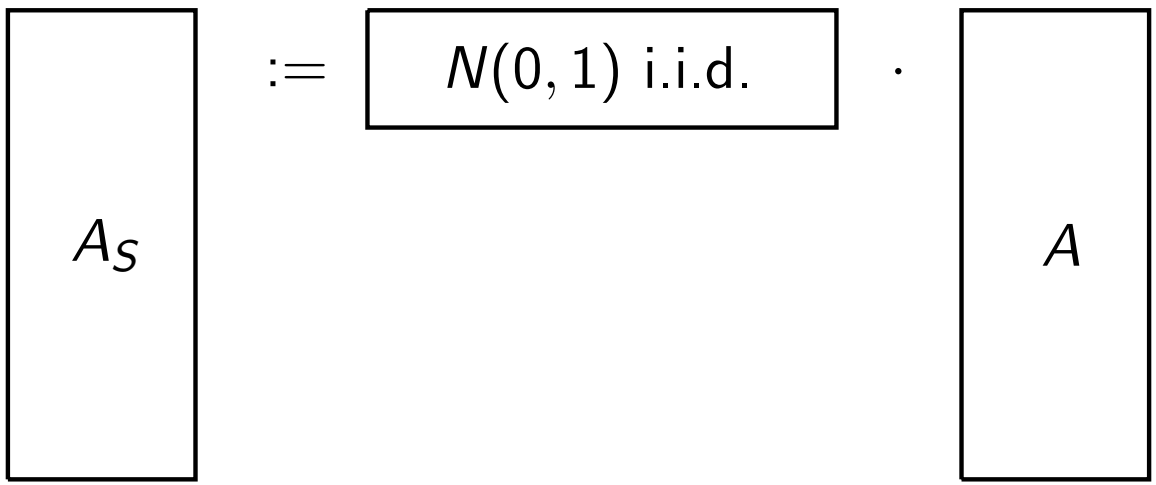

Gaussian sketch-and-project method takes gaussian random matrices S with i.i.d. entries as sketches.

BGK convergence rate

Convegence theorem (R. - Needell 2019)

The gaussian block Kaczmarz method converges to $\mathbf{x}_*$ *with the rate*

$$\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \left(1 - \frac{s\sigma_{min}^2(A)}{(9\sqrt{s}\|A\| + C\|A\|_F)^2}\right)^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2,$$

where $1 \leq s \leq m$ *is the dimension of the gaussian sketch* S*.*

- recovers "standard rate" $\sigma_{min}^2(A)/\|A\|_F^2$ for $s = 1$
- per iteration performance improves with increasing s
- actually, cputime performance also improves with increasing s

Better convergence for bigger sketch size

For $A = 50000 \times 500$ i.i.d. gaussian matrix:

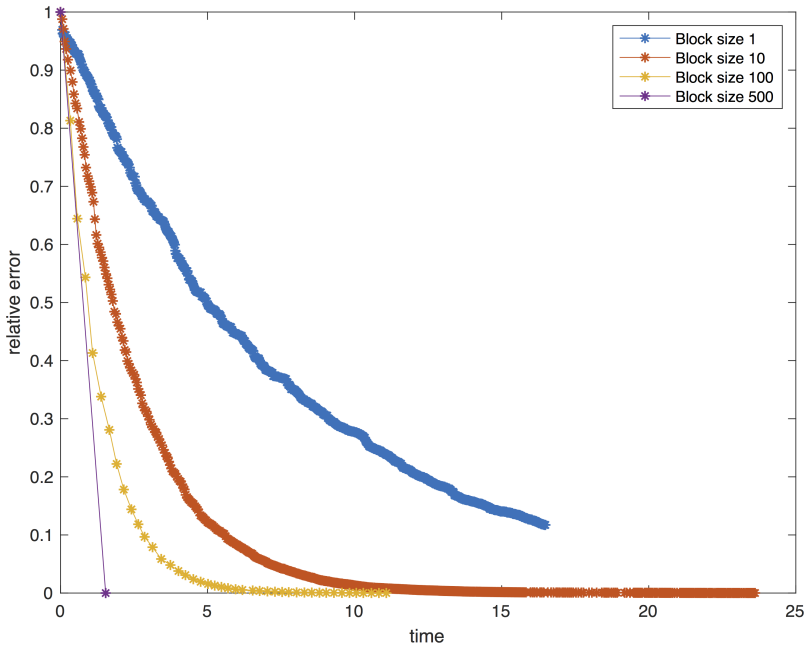

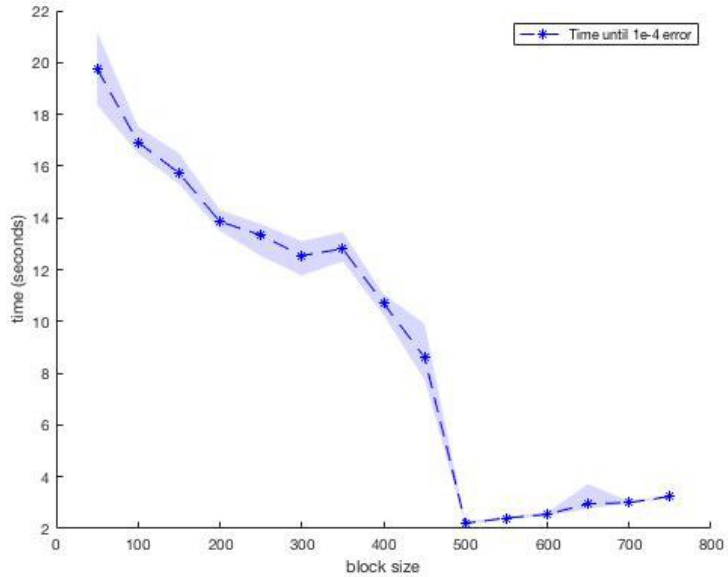

Left: time(s) vs relative error for the varying sketch size $s = 1, 10, 100, 500$; right: block size vs average time until relative error $1e-4$

Proof ideas: random matrices

$$\begin{aligned}\mathbf{x}_k &= \mathbf{x}_{k-1} + (S^T A)^\dagger (S^T \mathbf{b} - S^T A \mathbf{x}_k) \\ &= (\mathrm{Id} - (S^T A)^\dagger S^T A)\mathbf{x}_k + (S^T A)^\dagger S^T \mathbf{b}.\end{aligned}$$

1. We need to estimate $\mathbb{E}\|(S^T A)^\dagger \cdot S^T A \mathbf{x}\|_2^2$ from below - a product of two (dependent!) random matrices

Proof ideas: random matrices

$$\mathbf{x}_k = \mathbf{x}_{k-1} + (S^T A)^\dagger (S^T \mathbf{b} - S^T A \mathbf{x}_k)$$
$$= (\mathrm{Id} - (S^T A)^\dagger S^T A)\mathbf{x}_k + (S^T A)^\dagger S^T \mathbf{b}.$$

1. We need to estimate $\mathbb{E}\|(S^T A)^\dagger \cdot S^T A\mathbf{x}\|_2^2$ from below - a product of two (dependent!) random matrices

2. S is $m \times s$ standard normal i.i.d. matrix.

$$\mathbb{E}\|S^T A\mathbf{x}\|_2^2 = s\|A\mathbf{x}\|_2^2 \geq s\sigma_{min}^2(A)$$

But we need a high probability statement for any $s \geq 1$:

$$\mathbb{P}(\|S^T v\|_2^2 > \|v\|^2 s/10) \geq 0.5$$

for any $v \in \mathbb{R}^m$ and $s \geq 1$ - Cramér's concentration theorem.

3.

$$\mathbb{E} \sup_{\mathbf{x}\in S^{n-1}} \|S^T A\mathbf{x}\|_2 \leq \sqrt{m}\|A\|_2$$

Can we get a better estimate? Yes!

$$\mathbb{E} \sup_{\mathbf{x}\in S^{n-1}} \|S^T A\mathbf{x}\|_2 = \mathbb{E} \sup_{w\in AS^{n-1}} \|S^T w\|_2 \leq \sqrt{s}\|A\| + C\|A\|_F.$$

3.

$$\mathbb{E} \sup_{\mathbf{x}\in S^{n-1}} \|S^T A\mathbf{x}\|_2 \le \sqrt{m}\|A\|_2$$

Can we get a better estimate? Yes!

$$\mathbb{E} \sup_{\mathbf{x}\in S^{n-1}} \|S^T A\mathbf{x}\|_2 = \mathbb{E} \sup_{w\in AS^{n-1}} \|S^T w\|_2 \le \sqrt{s}\|A\| + C\|A\|_F.$$

To show 3.: apply matrix deviation inequality:

$$\mathbb{E} \sup_{w\in U} \|S^T w\|_2 \le \sqrt{s} \sup_{w\in U} \|w\|_2 + C\gamma(U),$$

to the ellipse $U := AS^{n-1}$. Here, $\gamma(U)$ is gaussian complexity of the set U:

$$\gamma(U) := \mathbb{E} \sup_{w\in U} |\langle \xi, w\rangle|, \text{ where } \xi \sim N(0, I_n)$$

Sampling sketches from finite collection

We could select sketches from the pre-sampled collection of gaussian random matrices

Theorem (R. - Needell)

Let $\mathcal{S} = \{S_1, \dots, S_N\}$ be a set of $m \times s$ random matrices with i.i.d. standard normal entries, $m^{2.5} \leq N \leq e^{m/3}$. Then, with probability at least $1 - 3/m$, for any initial estimate $\mathbf{x}_0$, finite block gaussian Kaczmarz method converges with the rate

$$\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq \left(1 - \frac{s}{36m\kappa^2(A)}\right)^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2.$$

In practice, the collection $\mathcal{S}$ can be much smaller, about $|\mathcal{S}| \sim m/s$

Solving noisy systems

If the system is inconsistent, we can search for least-squares problem solution with gaussian block Kaczmarz method:

$$\mathbf{x}_* = \arg\min_{\mathbf{x}} \|A\mathbf{x} - b\|_2^2$$

and the error (noise) $e := A\mathbf{x}_* - b$.

Theorem (R - Needell)

The gaussian block Kaczmarz method converges to $\mathbf{x}_*$ *with the rate:*

$$\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|_2^2 \leq r^k \|\mathbf{x}_0 - \mathbf{x}_*\|_2^2 + \frac{\|e\|_2^2}{s_{min}^4(A)} \cdot \left[\frac{(9\sqrt{s}\|A\| + C\|A\|_F)^2}{(\sqrt{n} - \sqrt{s})^2}\right]$$

Structurally differs from the noiseless case: diverges when $s \sim n$

Dependence on the block size in the noisy case

$A = 50000 \times 500$ i.i.d. gaussian matrix,
$e =$ random gaussian noise, normalized: $\|e\|_2 = 0.05 * \|b\|_2$

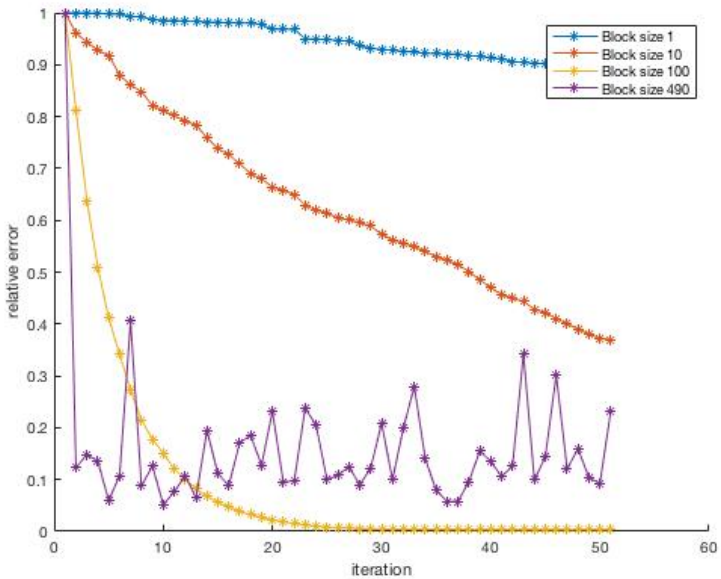

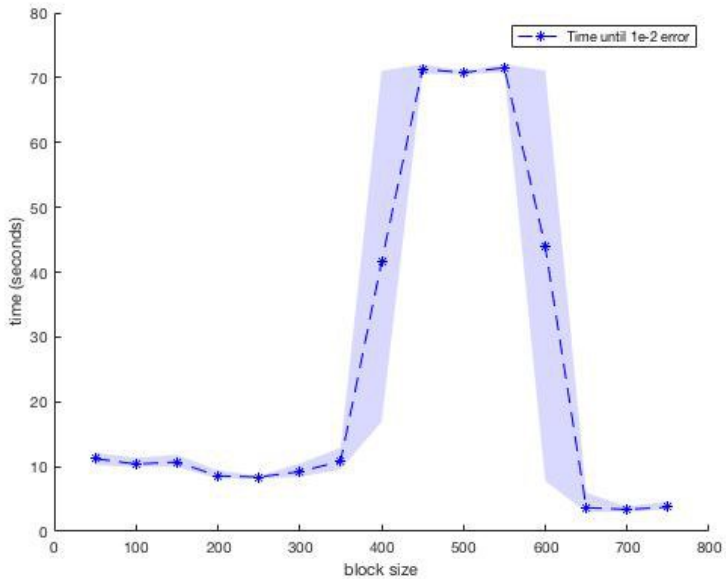

Left: iteration vs relative error for the sketch size $s = 1, 10, 100, 490$; right: block size vs average time until relative error $1e-2$; 70 sec is max allowed time

Is gaussian sketching practical?

$A = 50000 \times 500$ i.i.d. matrix:
$N(0,1)$ model (thin) and $Unif[0.8,1]$ model (bold lines)

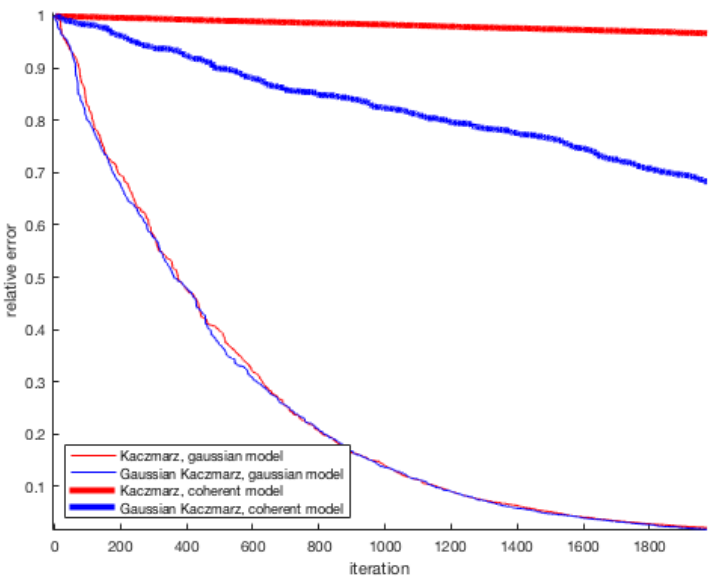

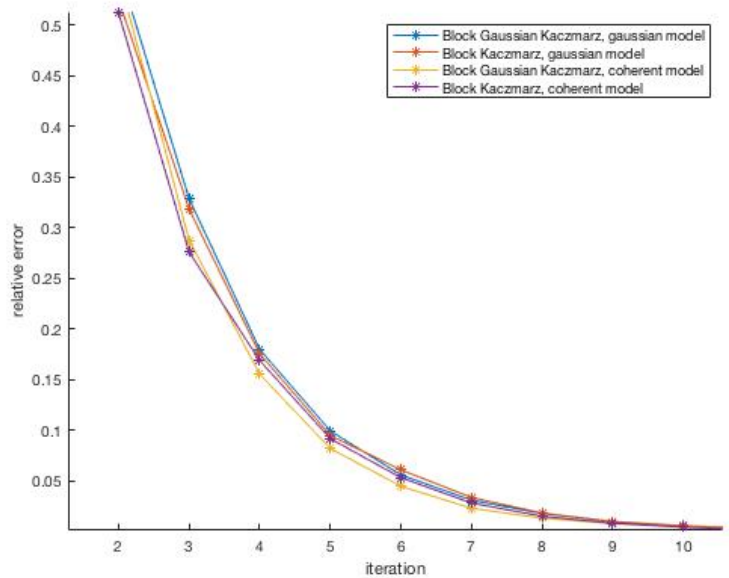

Left: $s = 1$, right: $s = 223$; blue = with gaussian sketching, red = without it

Gaussian sketching improves regular Kaczmarz for highly coherent systems when $s = 1$, but loses the advantage on bigger block sizes

Gaussian sketching reduces variance

$A = 50000 \times 500$ i.i.d. matrix,
$e =$ spiky noise, 10 random spikes of size 50.

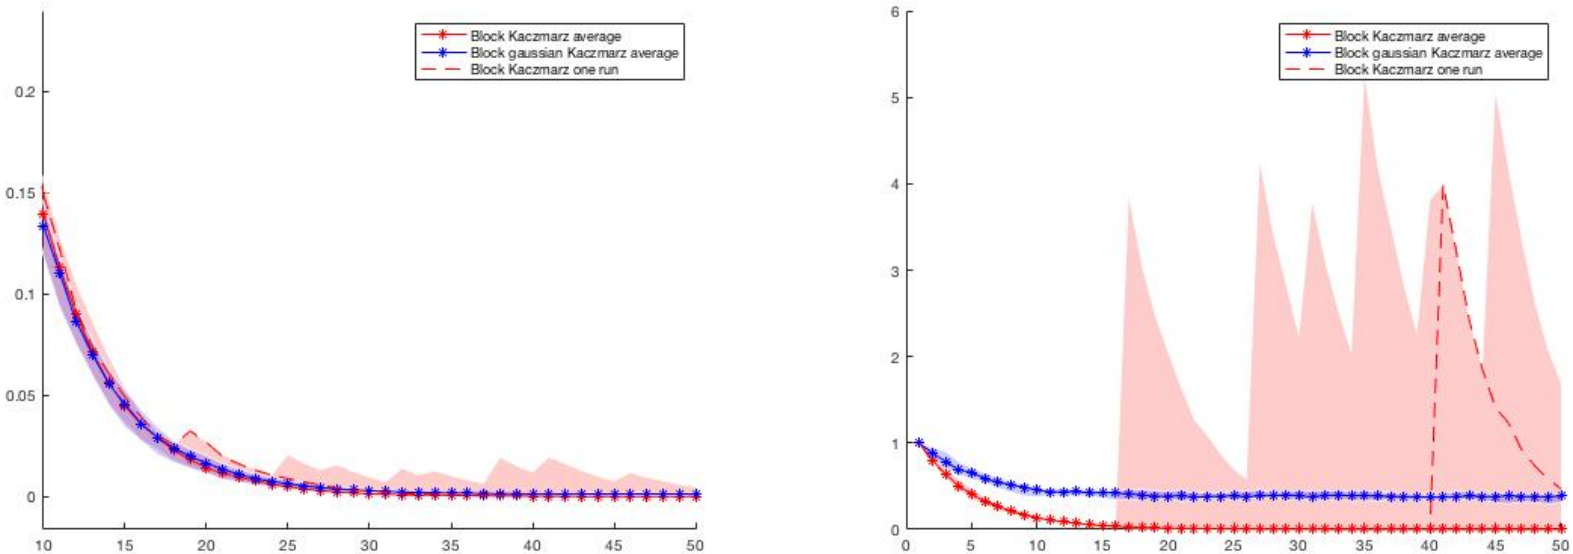

Iteration vs relative error (median and range over 10 runs). Left: gaussian model, right: coherent model; blue = with gaussian sketching, red = without it.

Noise vs corruptions

Noise: $A\mathbf{x} = \tilde{\mathbf{b}}$ and $\left\|\tilde{\mathbf{b}} - \mathbf{b}\right\| \leq \mathbf{e}$

- E.g., for RK: [Needell '10], [Needell, Tropp '12], [R., Needell '19] The result contains convergence horizon depending on the size of noise, $\|\mathbf{e}\|$ or $\|\mathbf{e}\|_\infty$.

Noise vs corruptions

Noise: $A\mathbf{x} = \tilde{\mathbf{b}}$ and $\left\|\tilde{\mathbf{b}} - \mathbf{b}\right\| \leq \mathbf{e}$

- E.g., for RK: [Needell '10], [Needell, Tropp '12], [R., Needell '19] The result contains convergence horizon depending on the size of noise, $\|\mathbf{e}\|$ or $\|\mathbf{e}\|_\infty$.

Corruptions: $\mathbf{b}$ is obtained by large changes on some of the entries of $\tilde{\mathbf{b}}$

- Convergence horizon can be huge... Idea: look at the residual (distances from an iterate to the equation hyperplanes) to detect suspicious candidates.
  [Haddock, Needell '18] – run RK for several iterations to approach $\mathbf{x}_*$, then examine the residual and remove the equations with the largest. This works with few corruptions only (up to $5 - 10\%$).

Model: overdetermined linear system

$$A \cdot \mathbf{x}_* = b, \quad m > n$$

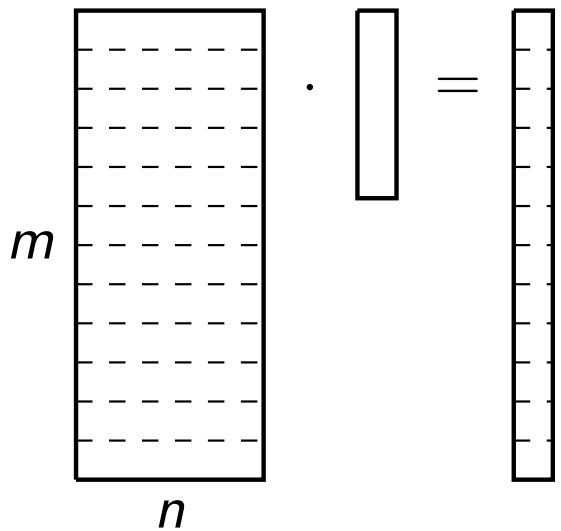

- $b_C \in \mathbb{R}^n$ has at most βm nonzero entries (β is the fraction of corrupted entries)
- Given knowledge of A and the corrupted measurements $\tilde{b} := Ax_{\mathbf{x}} * + b_C$, we would like an algorithm to recover $\mathbf{x}_{\mathbf{x}} *$.

RK on corrupted systems

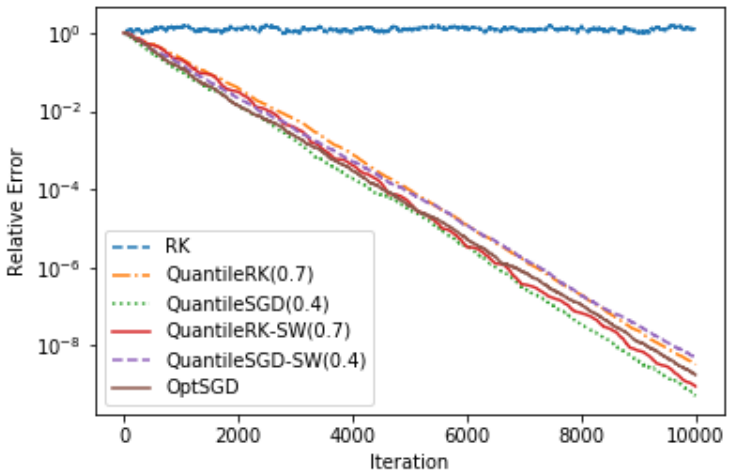

(a) Gaussian model

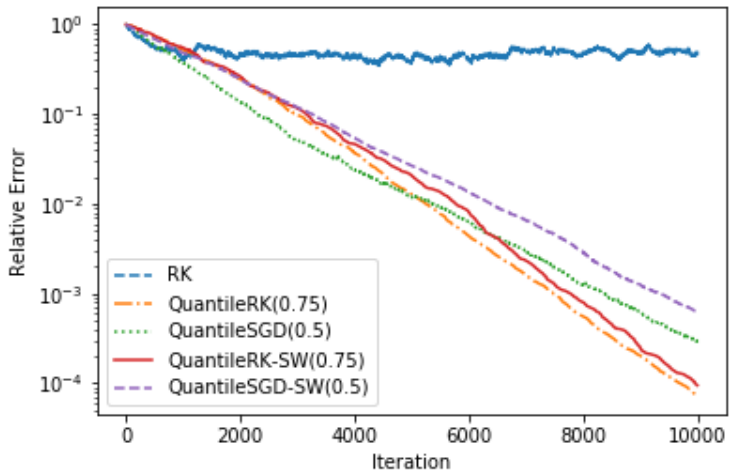

(b) Coherent model

Relative error as a function of iteration count plotted for a 50000×100 Gaussian and coherent model with a 0.2 corruption rate. The coherent system was generated by sampling entries uniformly in $[0, 1)$ and then normalizing the rows of the resulting matrix.

QuantileRK algorithm

Idea: If a sampled hyperplane looks corrupted, don't project!

1: sample several hyperplanes: $i_1, \dots i_T \sim \text{Uniform}(1, \dots, m)$
2: select a random next direction: $i \sim \text{Uniform}(1, \dots, m)$
3: **if** distance to the i-th hyperplane is at most
4: $|\langle \mathbf{A}_i, \mathbf{x}_k - b_i \rangle| \leq$ q-Quantile $(|\langle \mathbf{A}_{i_\ell}, \mathbf{x}_k \rangle - b_{i_\ell}| : \ell \in [T])$ **then**
5: project: $\mathbf{x}_{k+1} = \mathbf{x}_k - (\langle \mathbf{A}_\mathbf{i}, \mathbf{x}_k \rangle - b_i)\mathbf{A}_\mathbf{i}$,
6: **else**
7: hold: $\mathbf{x}_{k+1} = \mathbf{x}_k$

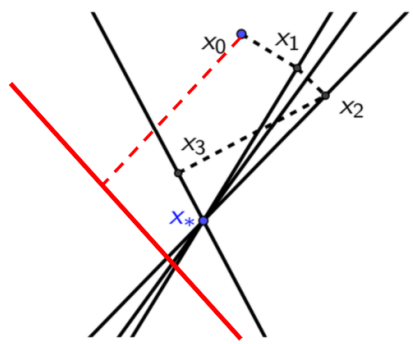

QuantileSGD algorithm

Idea: Adjust the step size of the "projection" according to the quantile.

1: sample several hyperplanes: $i_1, \dots i_T \sim \text{Uniform}(1, \dots, m)$
2: select a random next direction: $i \sim \text{Uniform}(1, \dots, m)$
3: define step size by $\gamma := \text{q-Quantile}\left(|\langle \mathbf{A}_{i_\ell}, \mathbf{x}_k \rangle - b_{i_\ell}| : \ell \in [T]\right)$
4: project cautiously:
5: $\quad \mathbf{x}_{k+1} = \mathbf{x}_k - \gamma \,\text{sign}(\langle \mathbf{A_i}, \mathbf{x}_k \rangle - b_i)\mathbf{A_i}$

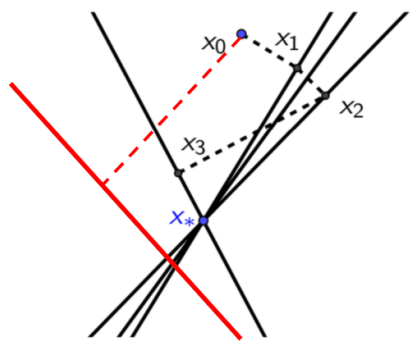

L_1-SGD

- Another approach is to minimize

$$\mathbf{x}_* = \arg\min_{\mathbf{x}} \|A\mathbf{x} - b\|_1^2.$$

This is a standard relaxation of the minimization of L_0-norm ([Candes, Tao, '05], [Candes, Rudelson, Tao, Vershynin, '05])

L_1-SGD

- Another approach is to minimize

$$\mathbf{x}_* = \arg\min_{\mathbf{x}} \|A\mathbf{x} - b\|_1^2.$$

This is a standard relaxation of the minimization of L_0-norm ([Candes, Tao, '05], [Candes, Rudelson, Tao, Vershynin, '05])

- Stochastic Gradient Descent with L_1 loss has iteration

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \eta_k \operatorname{sign}(\langle \mathbf{A_i}, \mathbf{x}_k \rangle - b_i)\mathbf{A_i},$$

with randomly sampled $i \in [m]$ and some step size schedule η_k

L_1-SGD

- Another approach is to minimize

$$\mathbf{x}_* = \arg\min_{\mathbf{x}} \|A\mathbf{x} - b\|_1^2.$$

This is a standard relaxation of the minimization of L_0-norm ([Candes, Tao, '05], [Candes, Rudelson, Tao, Vershynin, '05])

- Stochastic Gradient Descent with L_1 loss has iteration

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \eta_k \operatorname{sign}(\langle \mathbf{A_i}, \mathbf{x}_k \rangle - b_i)\mathbf{A_i},$$

with randomly sampled $i \in [m]$ and some step size schedule η_k

- Note that RK can be viewed as SGD with L_2 loss and correct step size:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - (\langle \mathbf{A_i}, \mathbf{x}_k \rangle - \mathbf{b}_i)\mathbf{A}_i,$$

Optimal step size

- An optimal choice for η_k is such that

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) = \mathbb{E}(\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2) \quad \text{minimized}$$

Optimal step size

- An optimal choice for η_k is such that

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) = \mathbb{E}(\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2) \quad \text{minimized}$$

- Can be computed analytically:

$$\eta_k^* = \frac{1}{m}\sum_{j=1}^{m} \text{sign}(\langle \mathbf{A_i}, \mathbf{x}_k\rangle - \mathbf{b}_i)\langle \mathbf{e}_k, \mathbf{a}_i\rangle$$

and also

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) = \left(1 - \left(\frac{\eta_k^*}{\|\mathbf{e}_k\|}\right)^2\right)\|\mathbf{e}_k\|^2$$

Optimal step size

- An optimal choice for η_k is such that

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) = \mathbb{E}(\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2^2) \quad \text{minimized}$$

- Can be computed analytically:

$$\eta_k^* = \frac{1}{m}\sum_{j=1}^{m} \text{sign}(\langle \mathbf{A_i}, \mathbf{x}_k\rangle - \mathbf{b}_i)\langle \mathbf{e}_k, \mathbf{a}_i\rangle$$

and also

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) = \left(1 - \left(\frac{\eta_k^*}{\|\mathbf{e}_k\|}\right)^2\right)\|\mathbf{e}_k\|^2$$

- The last quantity $\left(\frac{\eta_k^*}{\|\mathbf{e}_k\|}\right)^2 \geq \frac{1}{n}$ for subgaussian matrices (which gives standard "Kaczmarz" linear convergence rate)

Clearly, $\mathbf{e}_k = \mathbf{x}_k - \mathbf{x}^*$ is not known.

Approximating optimal step size

Two key facts:

- Step sizes of similar order lead to similar convergence:

L_1-SGD with the stepsize η_k, satisfying $0 < c_1 \leq \frac{\eta_k}{\eta_k^*} \leq c_2 < 2$ at each iteration, converges with

$$\mathbb{E}(\|\mathbf{e}_{k+1}\|_2^2) \leq \left(1 - \frac{C}{n}\right) \|\mathbf{e}_k\|^2$$

-

Residual quantiles: q-quantile $(|\langle \mathbf{A}_{i_\ell}, \mathbf{x}_k \rangle - \mathbf{b}_{i_\ell}| : \ell \in [T]) \approx$

Uncorrupted quantiles: q-quantile $\{|\langle \mathbf{A}_{i_\ell}, \mathbf{x}_k - \mathbf{x}^* \rangle| : \ell \in [T]\} \approx$

Empirical mean : $\frac{1}{m} \sum_{i=1}^{m} |\langle \mathbf{x_k} - \mathbf{x}^*, \mathbf{A}_i \rangle| \approx$ Optimal step size η^*.

Convergence Theorem

Theorem (Haddock, Needell, R., Swartworth, 2020)

Let an $m \times n$ matrix A have subgaussian isotropic rows, and its entries have centered and bounded density functions. Then with probability $1 - ce^{-c_q m}$, **QuantileRK**(q) *and* **QuantileSGD**(q) *on the full residual converge with the standard convergence rate*

$$\mathbb{E}||x_k - x_*||_2^2 \leq \left(1 - C_q \frac{\sigma_{min}^2(A)}{\|A\|_F^2}\right)^k ||x_0 - x_*||_2^2$$

if β is smaller than some positive constant, $q \leq 0.5 - \beta$, and n and m/n are larger than fixed constants.

- The same order as standard RK rate for uncorrupted systems.
- Result essentially holds with subsampling as well.
- The corrupted entries and values may be chosen adversarially.

Model

Assumption

All the rows $\mathbf{A}_i$ of the matrix A have unit norm and are independent. Additionally, for all $i \in [m]$, $\sqrt{n}\mathbf{A}_i$ is mean zero isotropic and has uniformly bounded subgaussian norm $\left\|\sqrt{n}\mathbf{A}_i\right\|_{\psi_2} \leq K$.

Assumption

Each entry A_{ij} of A has probability density function ϕ_{ij} which satisfies $\phi_{ij}(t) \leq D\sqrt{n}$ for all $t \in \mathbb{R}$. (The quantity D is a constant which we will use throughout when referring to this assumption.)

Proof sketch for QuantileRK

1. Concentration of the residual quantiles $(|\langle \mathbf{A_i}, \mathbf{x_k} \rangle - \mathbf{b}_i| : i \in [m])$:

For n large enough, with probability at least $1 - e^{-cn}$, for every $\mathbf{x} \in \mathbb{R}^n$ for all but $\beta' m$ indices i the following holds $|\langle \mathbf{A_i}, \mathbf{x} \rangle| \leq C/(\beta' \sqrt{n}) \|\mathbf{x}\|$.

Proof sketch for QuantileRK

1. Concentration of the residual quantiles $(|\langle \mathbf{A_i}, \mathbf{x_k} \rangle - \mathbf{b}_i| : i \in [m])$:

For n large enough, with probability at least $1 - e^{-cn}$, for every $\mathbf{x} \in \mathbb{R}^n$ for all but $\beta' m$ indices i the following holds $|\langle \mathbf{A_i}, \mathbf{x} \rangle| \leq C/(\beta' \sqrt{n}) \|\mathbf{x}\|$.

2. Condition on choosing a row from the uncorrupted sub-matrix A'. Known RK rate holds here, since the condition number of A' is of the same order as of A by a non-trivial high probability uniform lower bound for the singular values of submatrices:

$$\Pr\left(\inf_{\substack{T \subseteq [m]: \\ |T| \geq \alpha m}} \sigma_{\min}(\mathbf{A}_T) \gtrsim \frac{\alpha^{3/2}}{24D} \sigma_{\min}(\mathbf{A}) \right) \geq 1 - 3e^{-c_2 \alpha m}$$

Proof sketch for QuantileRK

1. Concentration of the residual quantiles $(|\langle \mathbf{A_i}, \mathbf{x_k} \rangle - \mathbf{b}_i| : i \in [m])$:

For n large enough, with probability at least $1 - e^{-cn}$, for every $\mathbf{x} \in \mathbb{R}^n$ for all but $\beta' m$ indices i the following holds $|\langle \mathbf{A_i}, \mathbf{x} \rangle| \leq C/(\beta' \sqrt{n}) \|\mathbf{x}\|$.

2. Condition on choosing a row from the uncorrupted sub-matrix A'. Known RK rate holds here, since the condition number of A' is of the same order as of A by a non-trivial high probability uniform lower bound for the singular values of submatrices:

$$\Pr\left(\inf_{\substack{T \subseteq [m]: \\ |T| \geq \alpha m}} \sigma_{\min}(\mathbf{A}_T) \gtrsim \frac{\alpha^{3/2}}{24D} \sigma_{\min}(\mathbf{A}) \right) \geq 1 - 3e^{-c_2 \alpha m}$$

3. Condition on choosing a corrupted row. This does not hurt too much by 1., since the step size is bounded by the quantile.

Concentration of the quantiles

Theorem

For n large enough, with probability at least $1 - e^{-cn}$, for every $\mathbf{x} \in \mathbb{R}^n$ for all but $\beta' m$ indices i the following holds $|\langle \mathbf{A_i}, \mathbf{x} \rangle| \leq C/(\beta' \sqrt{n}) \|\mathbf{x}\|$.

- Enough to show that with probability $1 - 2e^{-t^2 m}$ for every unit x

$$\frac{1}{m} \sum_{i=1}^{m} |\langle \mathbf{x}, \mathbf{A}_i \rangle| \leq \frac{1}{\sqrt{n}} + K \left(\frac{c_1}{\sqrt{m}} + \frac{c_2 t}{\sqrt{n}} \right)$$

- This follows from uniform concentration of $M(\mathbf{x}) := \frac{1}{m} \sum_{i=1}^{m} |\langle \mathbf{x}, \mathbf{A}_i \rangle|$ on the sphere (chaining + Dudley's inequality)
- And noticing that $(\mathbb{E} M(\mathbf{x}))^2 \leq \mathbb{E}(M(\mathbf{x}))^2 \sim \frac{1}{n}$, so there exists $\mathbf{x}$ at least that small.

Various quantiles

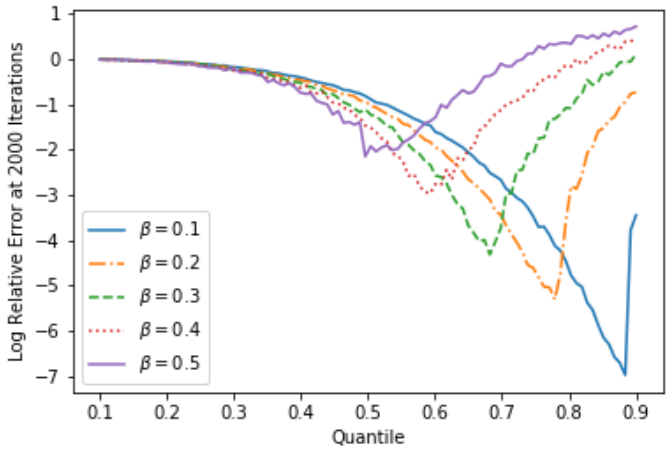

(a) QuantileRK

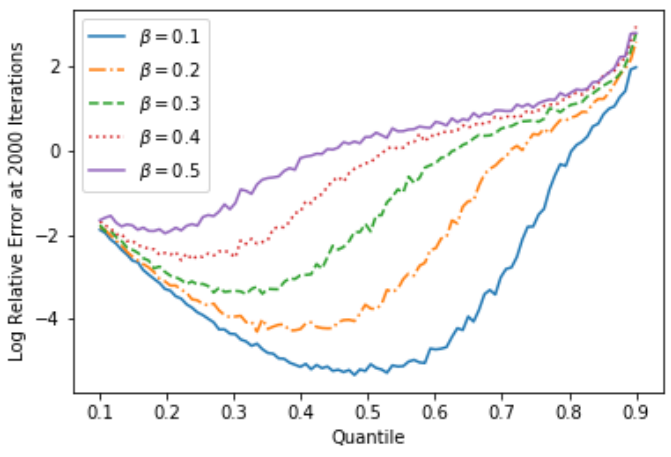

(b) QuantileSGD

$\log(\|\mathbf{x}_{2000} - \mathbf{x}^*\| / \|\mathbf{x}_0 - \mathbf{x}^*\|)$ for (a) QuantileRK and (b) QuantileSGD run on 50000×100 Gaussian system, with various corruption rates β and quantile choices.

Arbitrarily large corruptions

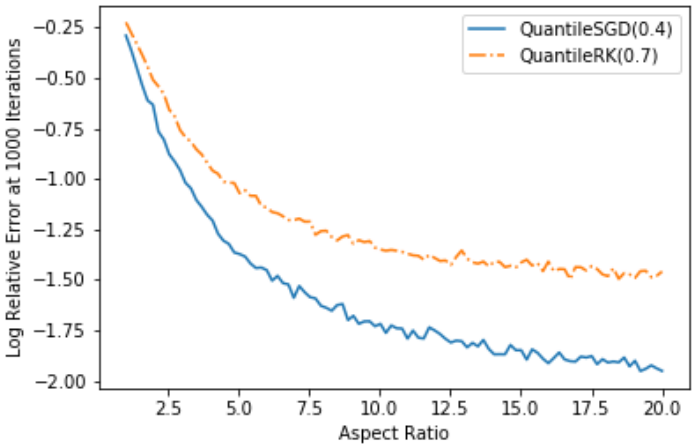

(a) Effect of aspect ratio

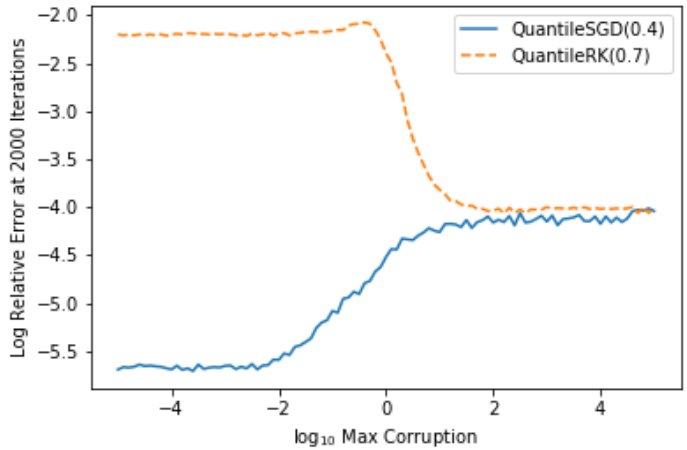

(b) Effect of corruption size

(a) Log relative error for QuantileSGD and QuantileRK after 1000 iterations on a $100a \times 100$ Gaussian system with a 0.2 corruption rate, where $a = m/n$ is the aspect ratio of the matrix.
(b) Log relative error for QuantileSGD and QuantileRK after 2000 iterations, as a function of corruption size. We use a 50000×100 Gaussian system and corrupt our system by adding a uniform value in $[-10^x, 10^x]$.

Future directions

- **More general corruption models**: byzantine corruptions, errors on the left hand side, mix of noise and corruptions
- **More general methods**: block methods for corrupted systems (with Ben Jarman, Deanna Needell)
- **More general problems**: randomized iterative methods for optimization problems beyond solving linear systems (with Will Swartworth, Han Lyu, Deanna Needell)

Literature

Quantile-based Iterative Methods for Corrupted Systems of Linear Equations
... J. Haddock, D. Needell, E. Rebrova, W. Swartworth
… submitted, arXiv:2009.08089

On block Gaussian sketching for the Kaczmarz method
... E. Rebrova, D. Needell
... Numerical algorithms (NUMA), 2019

Stochastic Gradient Descent Methods for Corrupted Systems of Linear Equations
... J. Haddock, D. Needell, E. Rebrova, W. Swartworth
... Proc. Conference on Information Sciences and Systems, 2020

Thanks for your attention!