Input: An $m \times n$ matrix $\mathbf{A}$ and a target rank k.

Output: Rank-k factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{U}\mathbf{D}\mathbf{V}^*$.

(1) Draw an $n \times k$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^*\,\mathbf{A}$.
(2) Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{A}\mathbf{R}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = \|\mathbf{A} - \mathbf{U}\mathbf{D}\mathbf{V}^*\|$? (Recall that $e_k = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|$.)

Input: An $m \times n$ matrix $\mathbf{A}$ and a target rank k.	
Output: Rank-k factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{U}\mathbf{D}\mathbf{V}^*$.	
(1) Draw an $n \times k$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^*\,\mathbf{A}$.
(2) Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{A}\mathbf{R}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{D}\mathbf{V}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = \|\mathbf{A} - \mathbf{U}\mathbf{D}\mathbf{V}^*\|$? (Recall that $e_k = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|$.)

Eckart-Young theorem: e_k is bounded from below by the singular value σ_{k+1} of $\mathbf{A}$.

Input: An $m \times n$ matrix $\mathbf{A}$ and a target rank k.

Output: Rank-k factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.

(1) Draw an $n \times k$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AR}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = \|\mathbf{A} - \mathbf{UDV}^*\|$? (Recall that $e_k = \|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|$.)

Eckart-Young theorem: e_k is bounded from below by the singular value σ_{k+1} of $\mathbf{A}$.

Question: Is e_k close to σ_{k+1}?

Input: An $m \times n$ matrix $\mathbf{A}$ and a target rank k.

Output: Rank-k factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.

(1) Draw an $n \times k$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AR}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = \|\mathbf{A} - \mathbf{UDV}^*\|$? (Recall that $e_k = \|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|$.)

Eckart-Young theorem: e_k is bounded from below by the singular value σ_{k+1} of $\mathbf{A}$.

Question: Is e_k close to σ_{k+1}?

Answer: Lamentably, no. The expectation of $\frac{e_k}{\sigma_{k+1}}$ is large, and has very large variance.

Input: An $m \times n$ matrix $\mathbf{A}$ and a target rank k. *Output:* Rank-k factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times k$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{AR}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = \|\mathbf{A} - \mathbf{UDV}^*\|$? (Recall that $e_k = \|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|$.)

Eckart-Young theorem: e_k is bounded from below by the singular value σ_{k+1} of $\mathbf{A}$.

Question: Is e_k close to σ_{k+1}?

Answer: Lamentably, no. The expectation of $\frac{e_k}{\sigma_{k+1}}$ is large, and has very large variance.

Remedy: Over-sample *slightly*. Compute $k+p$ samples from the range of $\mathbf{A}$.
It turns out that $p = 5$ or 10 is often sufficient. $p = k$ is almost always more than enough.

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank k, *and an over-sampling parameter p (say $p = 5$).* *Output:* Rank-$(k+p)$ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times (k+p)$ sample matrix $\mathbf{Y} = \mathbf{AR}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Bound on the expectation of the error for Gaussian test matrices

Let $\mathbf{A}$ denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let $\mathbf{R}$ denote an $n \times (k+p)$ Gaussian matrix.

Let $\mathbf{Q}$ denote the $m \times (k+p)$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{AR})$.

If $p \geq 2$, then

$$\mathbb{E}\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Frob}} \leq \left(1 + \frac{k}{p-1}\right)^{1/2} \left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2\right)^{1/2},$$

and

$$\mathbb{E}\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2\right)^{1/2}.$$

Ref: Halko, Martinsson, Tropp, 2009 & 2011

Large deviation bound for the error for Gaussian test matrices

Let $\mathbf{A}$ denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let $\mathbf{R}$ denote an $n \times (k+p)$ Gaussian matrix.

Let $\mathbf{Q}$ denote the $m \times (k+p)$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{AR})$.

If $p \geq 4$, and u and t are such that $u \geq 1$ and $t \geq 1$, then

$$\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \left(1 + t\sqrt{\frac{3k}{p+1}} + u\,t\,\frac{e\,\sqrt{k+p}}{p+1}\right)\sigma_{k+1} + \frac{t\,e\,\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}$$

except with probability at most $2\,t^{-p} + e^{-u^2/2}$.

Ref: Halko, Martinsson, Tropp, 2009 & 2011; Martinsson, Rokhlin, Tygert (2006)

---

u and t parameterize "bad" events — large u, t is bad, but unlikely.

Certain choices of t and u lead to simpler results. For instance,

$$\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \left(1 + 16\sqrt{1 + \frac{k}{p+1}}\right)\sigma_{k+1} + 8\frac{\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2},$$

except with probability at most $3\,e^{-p}$.

Large deviation bound for the error for Gaussian test matrices

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let **R** denote an $n \times (k+p)$ Gaussian matrix.

Let **Q** denote the $m \times (k+p)$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{AR})$.

If $p \geq 4$, and u and t are such that $u \geq 1$ and $t \geq 1$, then

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\| \leq \left(1 + t\sqrt{\frac{3k}{p+1}} + u\,t\,\frac{e\sqrt{k+p}}{p+1}\right)\sigma_{k+1} + \frac{t\,e\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}$$

except with probability at most $2\,t^{-p} + e^{-u^2/2}$.

Ref: Halko, Martinsson, Tropp, 2009 & 2011; Martinsson, Rokhlin, Tygert (2006)

---

u and t parameterize "bad" events — large u, t is bad, but unlikely.

Certain choices of t and u lead to simpler results. For instance,

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\| \leq \left(1 + 6\sqrt{(k+p)\cdot p\log p}\right)\sigma_{k+1} + 3\sqrt{k+p}\left(\sum_{j>k}\sigma_j^2\right)^{1/2},$$

except with probability at most $3\,p^{-p}$.

Proofs — Overview:

Let $\mathbf{A}$ denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.
Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.
Let $\mathbf{R}$ denote an $n \times \ell$ "test matrix", and let $\mathbf{Q}$ denote the $m \times \ell$ matrix $\mathbf{Q} = \mathtt{orth}(\mathbf{AR})$.

We seek to bound the error $e_k = e_k(\mathbf{A}, \mathbf{R}) = \|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|$, which is a random variable.

1. Make no assumption on $\mathbf{R}$. Construct a deterministic bound of the form

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\| \leq \cdots \mathbf{A} \cdots \mathbf{R} \cdots$$

2. Assume that $\mathbf{R}$ is drawn from a normal Gaussian distribution.
   Take expectations of the deterministic bound to attain a bound of the form

$$\mathbb{E}\big[\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|\big] \leq \cdots \mathbf{A} \cdots$$

3. Assume that $\mathbf{R}$ is drawn from a normal Gaussian distribution.
   Take expectations of the deterministic bound conditioned on "bad behavior" in $\mathbf{R}$ to get that

$$\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\| \leq \cdots \mathbf{A} \cdots$$

   holds with probability at least $\cdots$.

Part 1 (out of 3) — deterministic bound:

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.

Let **R** denote an $n \times \ell$ "test matrix", and let **Q** denote the $m \times \ell$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{AR})$.

Partition the SVD of **A** as follows:

$$\mathbf{A} = \mathbf{U} \begin{array}{c} \begin{array}{cc} k & n-k \end{array} \\ \begin{bmatrix} \mathbf{D}_1 & \\ & \mathbf{D}_2 \end{bmatrix} \end{array} \begin{bmatrix} \mathbf{V}_1^* \\ \mathbf{V}_2^* \end{bmatrix} \begin{array}{c} k \\ n-k \end{array}$$

Define $\mathbf{R}_1$ and $\mathbf{R}_2$ via

$$\underset{k \times (k+p)}{\mathbf{R}_1} = \underset{k \times n}{\mathbf{V}_1^*} \underset{n \times (k+p)}{\mathbf{R}} \quad \text{and} \quad \underset{(n-k) \times (k+p)}{\mathbf{R}_2} = \underset{(n-k) \times n}{\mathbf{V}_2^*} \underset{n \times (k+p)}{\mathbf{R}}.$$

Theorem: [HMT2009,HMT2011] Assuming that $\mathbf{R}_1$ is not singular, it holds that

$$|||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|||^2 \leq \underbrace{|||\mathbf{D}_2|||^2}_{\text{theoretically minimal error}} + |||\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger|||^2.$$

Here, $|||\cdot|||$ represents either ℓ^2-operator norm, or the Frobenius norm.

Note: A similar (but weaker) result appears in Boutsidis, Mahoney, Drineas (2009).

Recall: $\mathbf{A} = \mathbf{U}\begin{bmatrix} \mathbf{D}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{D}_2 \end{bmatrix}\begin{bmatrix} \mathbf{V}_1^* \\ \mathbf{V}_2^* \end{bmatrix}$, $\begin{bmatrix} \mathbf{R}_1 \\ \mathbf{R}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{V}_1^*\mathbf{R} \\ \mathbf{V}_2^*\mathbf{R} \end{bmatrix}$, $\mathbf{Y} = \mathbf{AR}$, $\mathbf{P}$ proj$^{\mathrm{n}}$ onto Ran($\mathbf{Y}$).

Thm: Suppose $\mathbf{D}_1\mathbf{R}_1$ has full rank. Then $\|\mathbf{A} - \mathbf{PA}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|^2$.

Proof: The problem is rotationally invariant $\Rightarrow$ We can assume $\mathbf{U} = \mathbf{I}$ and so $\mathbf{A} = \mathbf{DV}^*$.

Simple calculation: $\|(\mathbf{I} - \mathbf{P})\mathbf{A}\|^2 = \|\mathbf{A}^*(\mathbf{I} - \mathbf{P})^2\mathbf{A}\| = \|\mathbf{D}(\mathbf{I} - \mathbf{P})\mathbf{D}\|$.

$$\mathrm{Ran}(\mathbf{Y}) = \mathrm{Ran}\left(\begin{bmatrix} \mathbf{D}_1\mathbf{R}_1 \\ \mathbf{D}_2\mathbf{R}_2 \end{bmatrix}\right) = \mathrm{Ran}\left(\begin{bmatrix} \mathbf{I} \\ \mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\mathbf{D}_1 \end{bmatrix}\mathbf{D}_1\mathbf{R}_1\right) = \mathrm{Ran}\left(\begin{bmatrix} \mathbf{I} \\ \mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\mathbf{D}_1 \end{bmatrix}\right)$$

Set $\mathbf{F} = \mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\mathbf{D}_1^{-1}$. Then $\mathbf{P} = \begin{bmatrix} \mathbf{I} \\ \mathbf{F} \end{bmatrix}(\mathbf{I} + \mathbf{F}^*\mathbf{F})^{-1}[\mathbf{I}\ \mathbf{F}^*]$. (Compare to $\mathbf{P}_{\mathrm{ideal}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$.)

Use properties of psd matrices: $\mathbf{I} - \mathbf{P} \preccurlyeq \cdots \preccurlyeq \begin{bmatrix} \mathbf{F}^*\mathbf{F} & -(\mathbf{I} + \mathbf{F}^*\mathbf{F})^{-1}\mathbf{F}^* \\ -\mathbf{F}(\mathbf{I} + \mathbf{F}^*\mathbf{F})^{-1} & \mathbf{I} \end{bmatrix}$

Conjugate by $\mathbf{D}$ to get $\mathbf{D}(\mathbf{I} - \mathbf{P})\mathbf{D} \preccurlyeq \begin{bmatrix} \mathbf{D}_1\mathbf{F}^*\mathbf{F}\mathbf{D}_1 & -\mathbf{D}_1(\mathbf{I} + \mathbf{F}^*\mathbf{F})^{-1}\mathbf{F}^*\mathbf{D}_2 \\ -\mathbf{D}_2\mathbf{F}(\mathbf{I} + \mathbf{F}^*\mathbf{F})^{-1}\mathbf{D}_1 & \mathbf{D}_2^2 \end{bmatrix}$

Diagonal dominance: $\|\mathbf{D}(\mathbf{I} - \mathbf{P})\mathbf{D}\| \leq \|\mathbf{D}_1\mathbf{F}^*\mathbf{F}\mathbf{D}_1\| + \|\mathbf{D}_2^2\| = \|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|^2 + \|\mathbf{D}_2\|^2$.

Part 2 (out of 3) — bound on expectation of error when R is Gaussian:

Let $\mathbf{A}$ denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.

Let $\mathbf{R}$ denote an $n \times \ell$ "test matrix", and let $\mathbf{Q}$ denote the $m \times \ell$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{AR})$.

Recall: $|||\mathbf{A} - \mathbf{QQ}^*\mathbf{A}|||^2 \leq |||\mathbf{D}_2|||^2 + |||\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger|||^2$, where $\mathbf{R}_1 = \mathbf{V}_1^*\,\mathbf{R}$ and $\mathbf{R}_2 = \mathbf{V}_2^*\,\mathbf{R}$.

Assumption: $\mathbf{R}$ is drawn from a normal Gaussian distribution.

Since the Gaussian distribution is rotationally invariant, the matrices $\mathbf{R}_1$ and $\mathbf{R}_2$ also have a Gaussian distribution. (As a consequence, the matrices $\mathbf{U}$ and $\mathbf{V}$ do not enter the analysis and one could simply assume that $\mathbf{A}$ is diagonal, $\mathbf{A} = \text{diag}(\sigma_1, \sigma_2, \dots)$.)

What is the distribution of $\mathbf{R}_1^\dagger$ when $\mathbf{R}_1$ is a $k \times (k+p)$ Gaussian matrix?

If $p = 0$, then $\|\mathbf{R}_1^\dagger\|$ is typically large, and is very unstable.

Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. $p = 0$

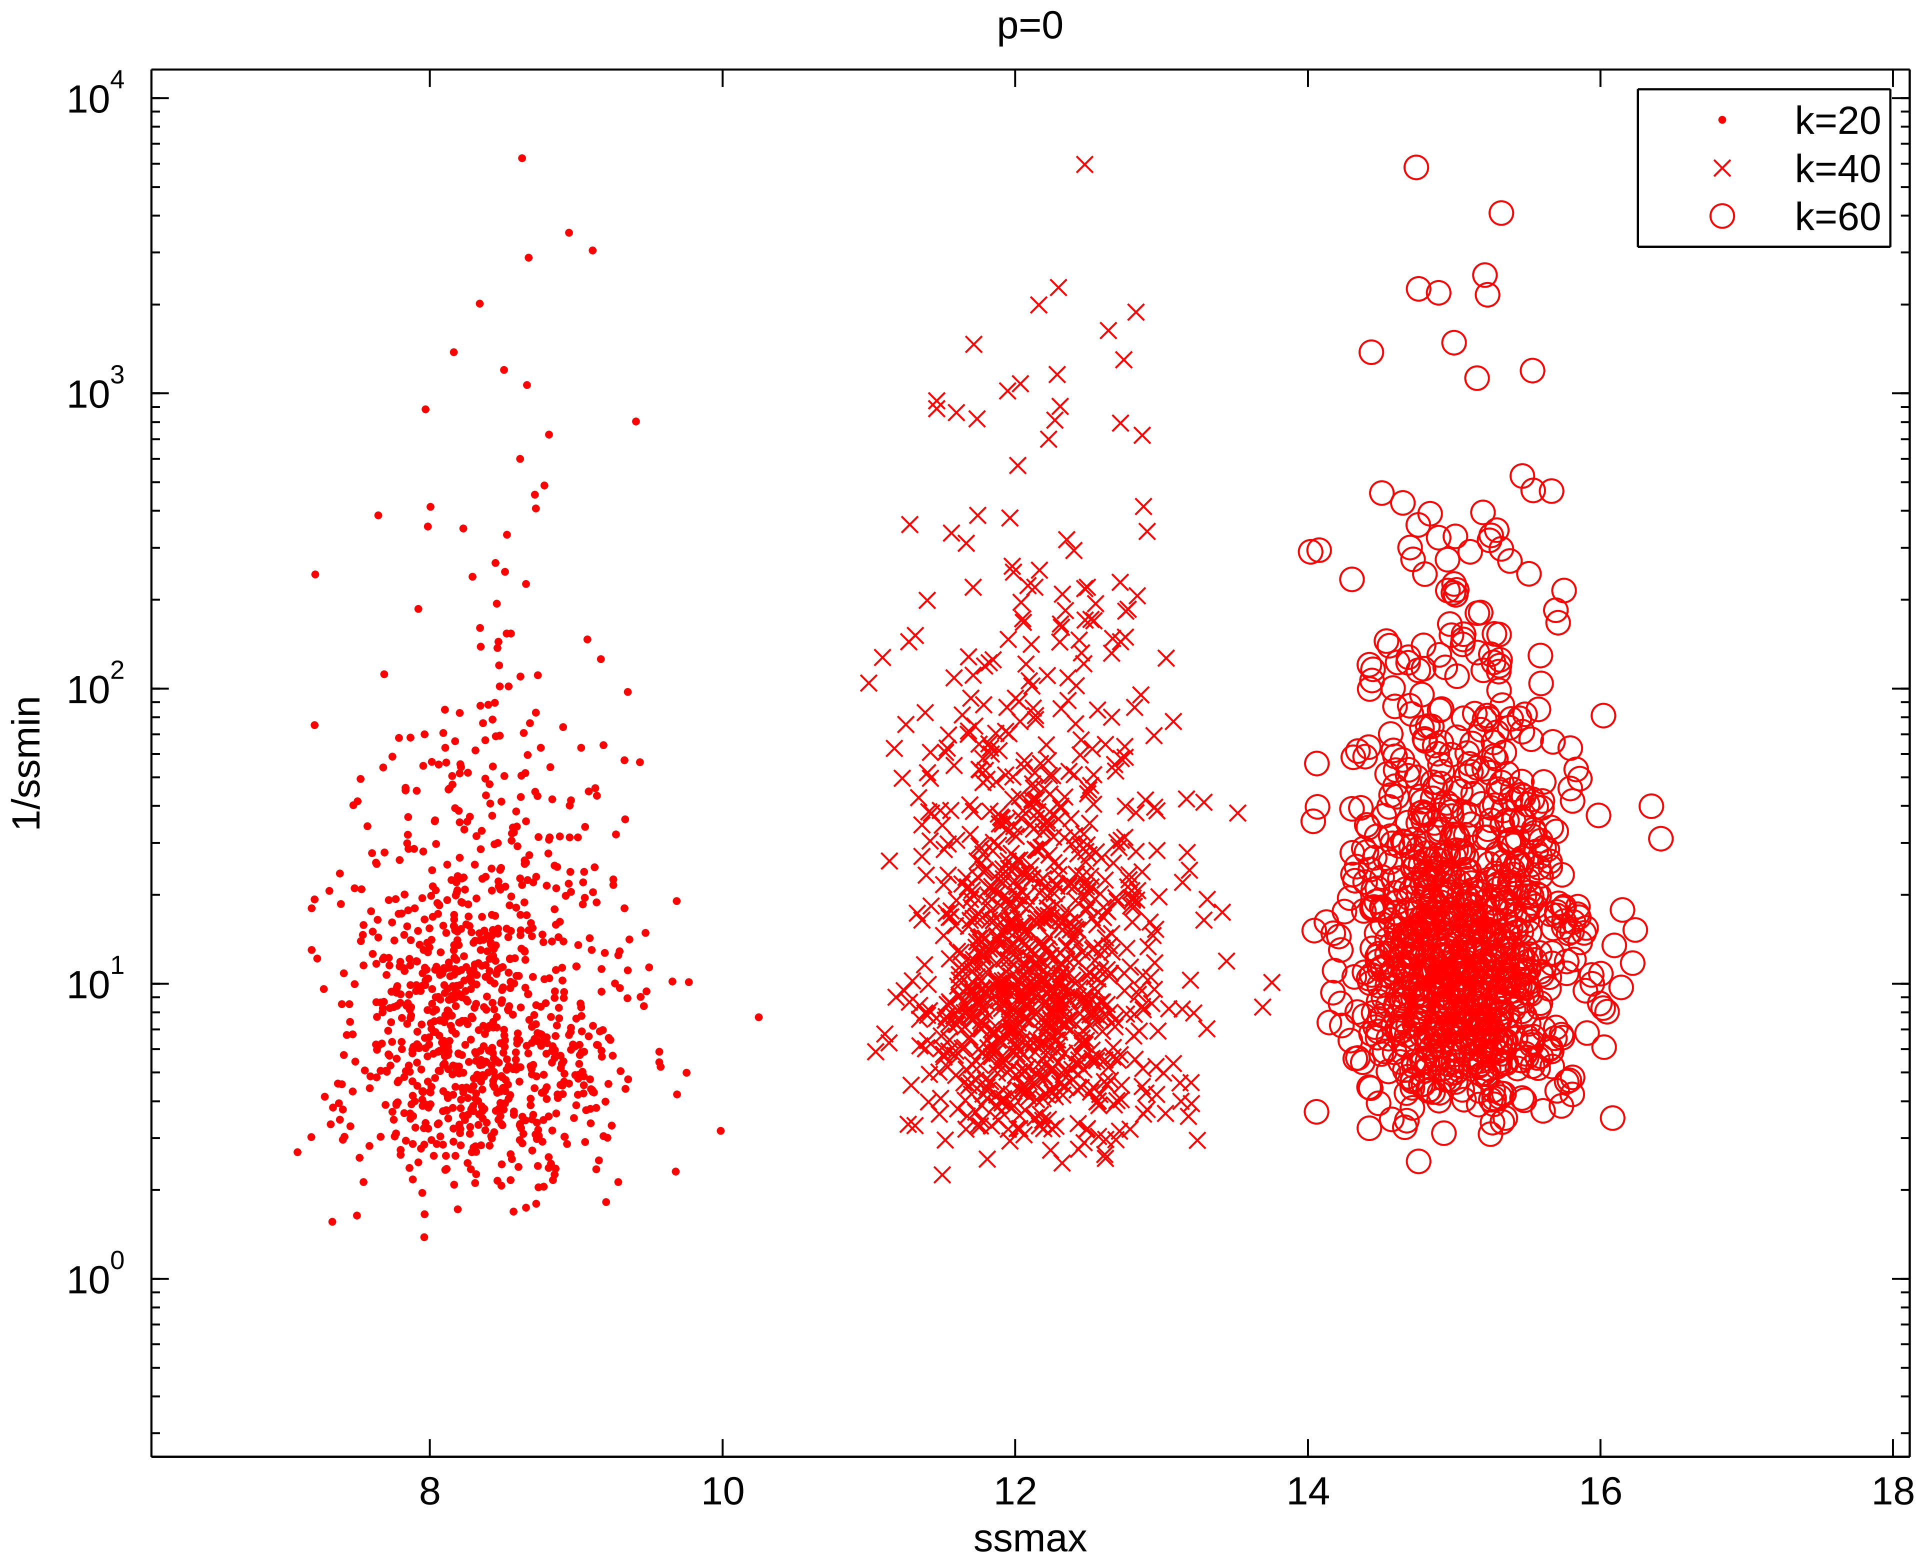

$1/\sigma_{\min}$ *is plotted against* $\sigma_{\max}$.

Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. $p = 2$

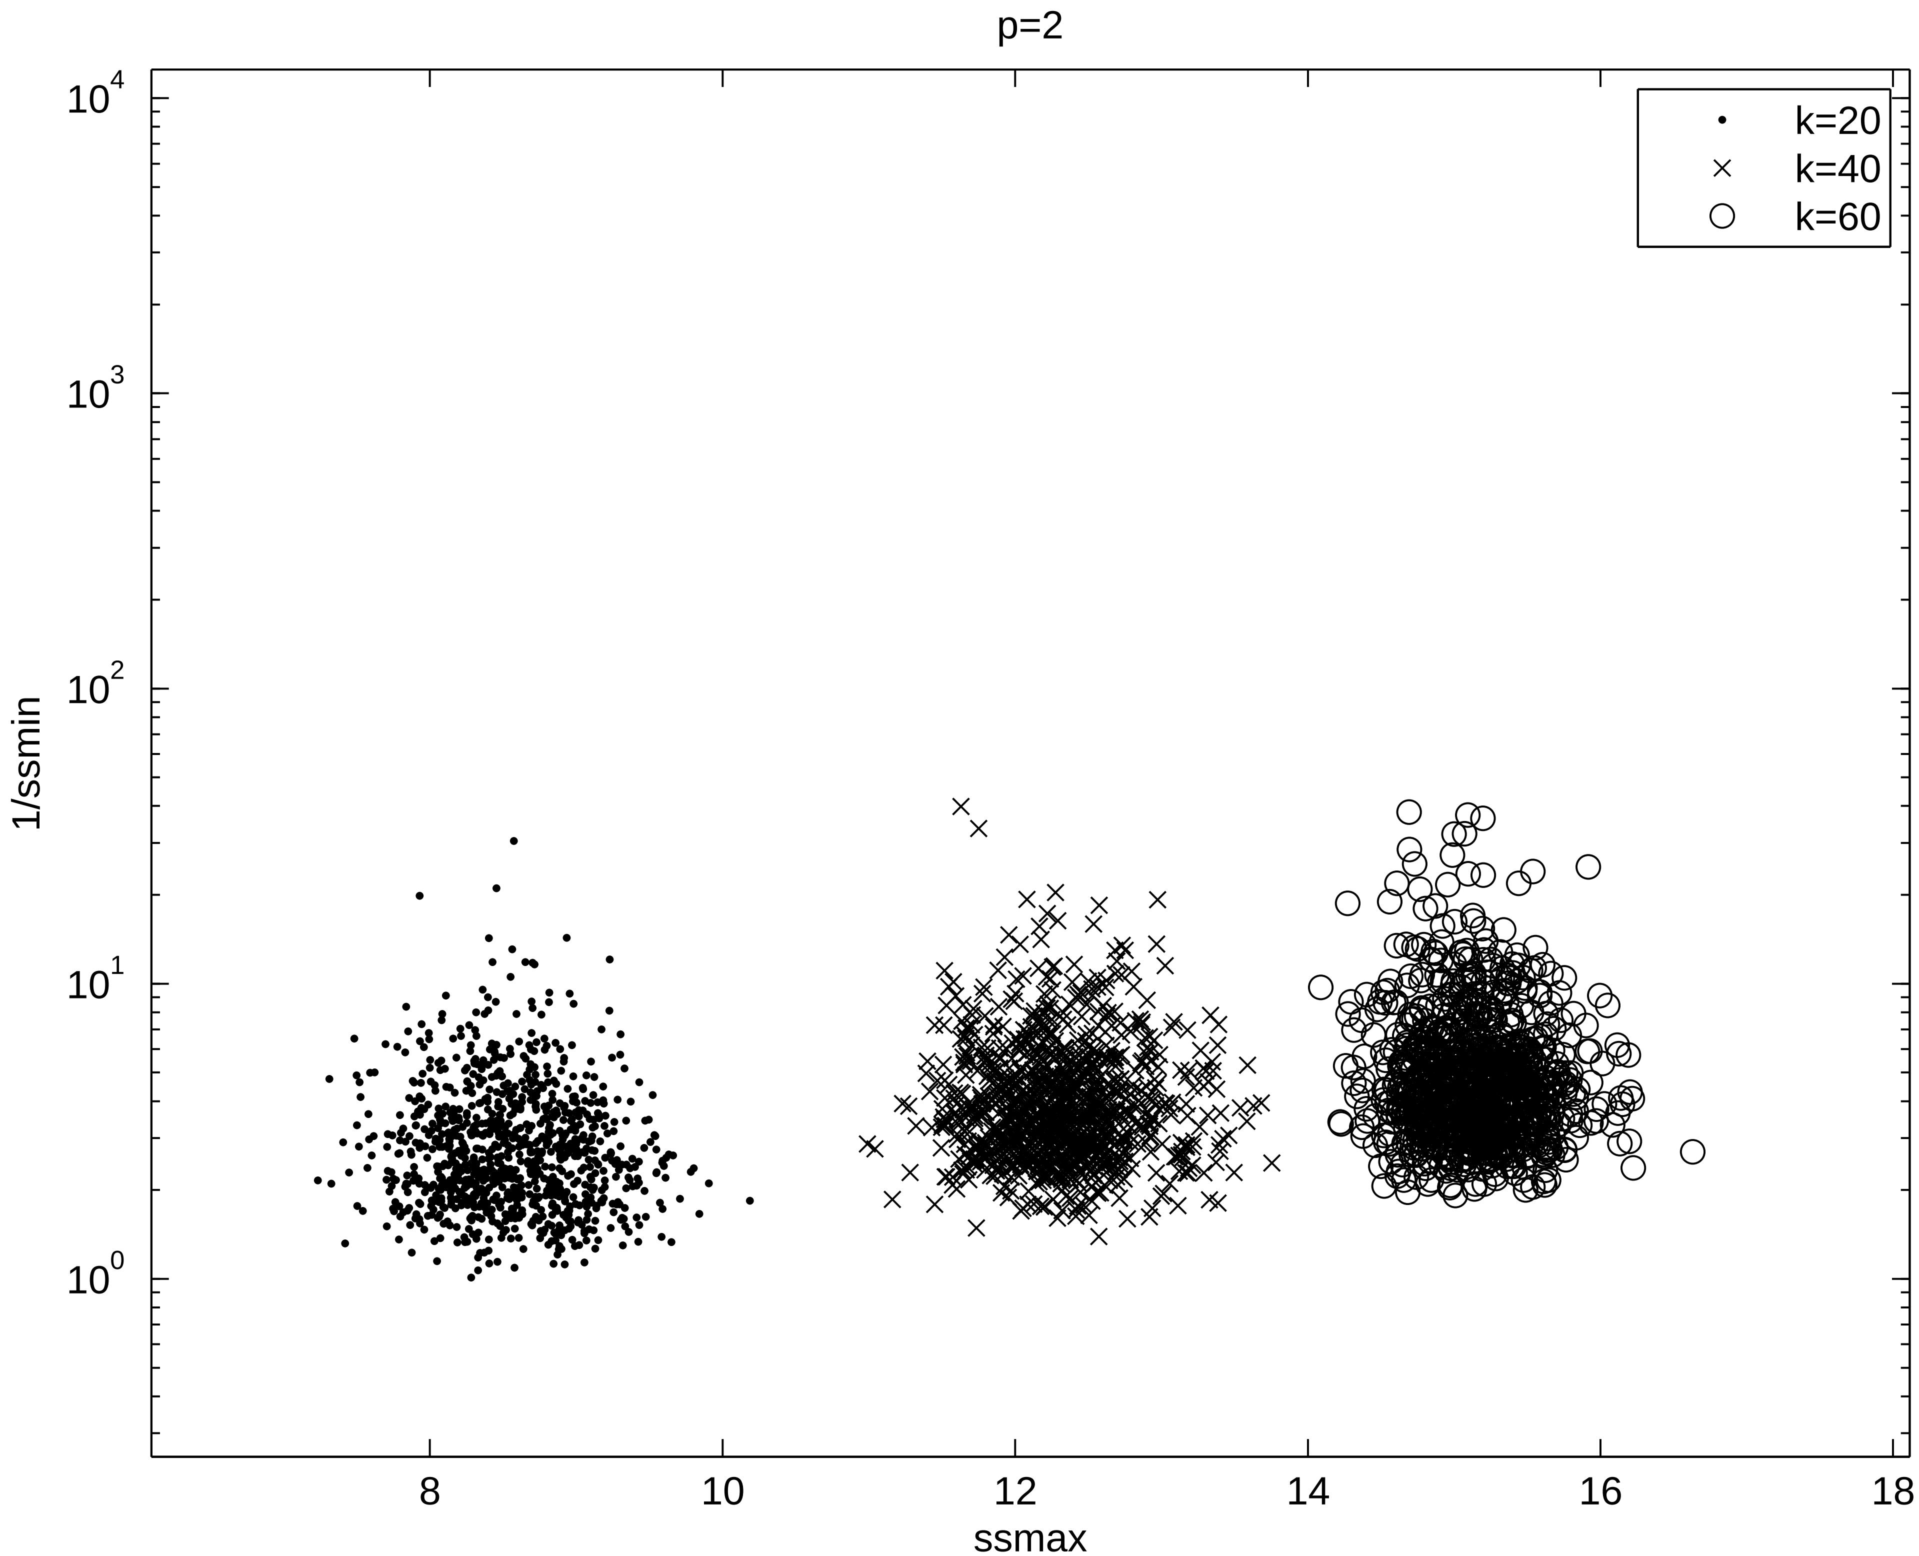

$1/\sigma_{\min}$ *is plotted against* $\sigma_{\max}$.

Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. $p = 5$

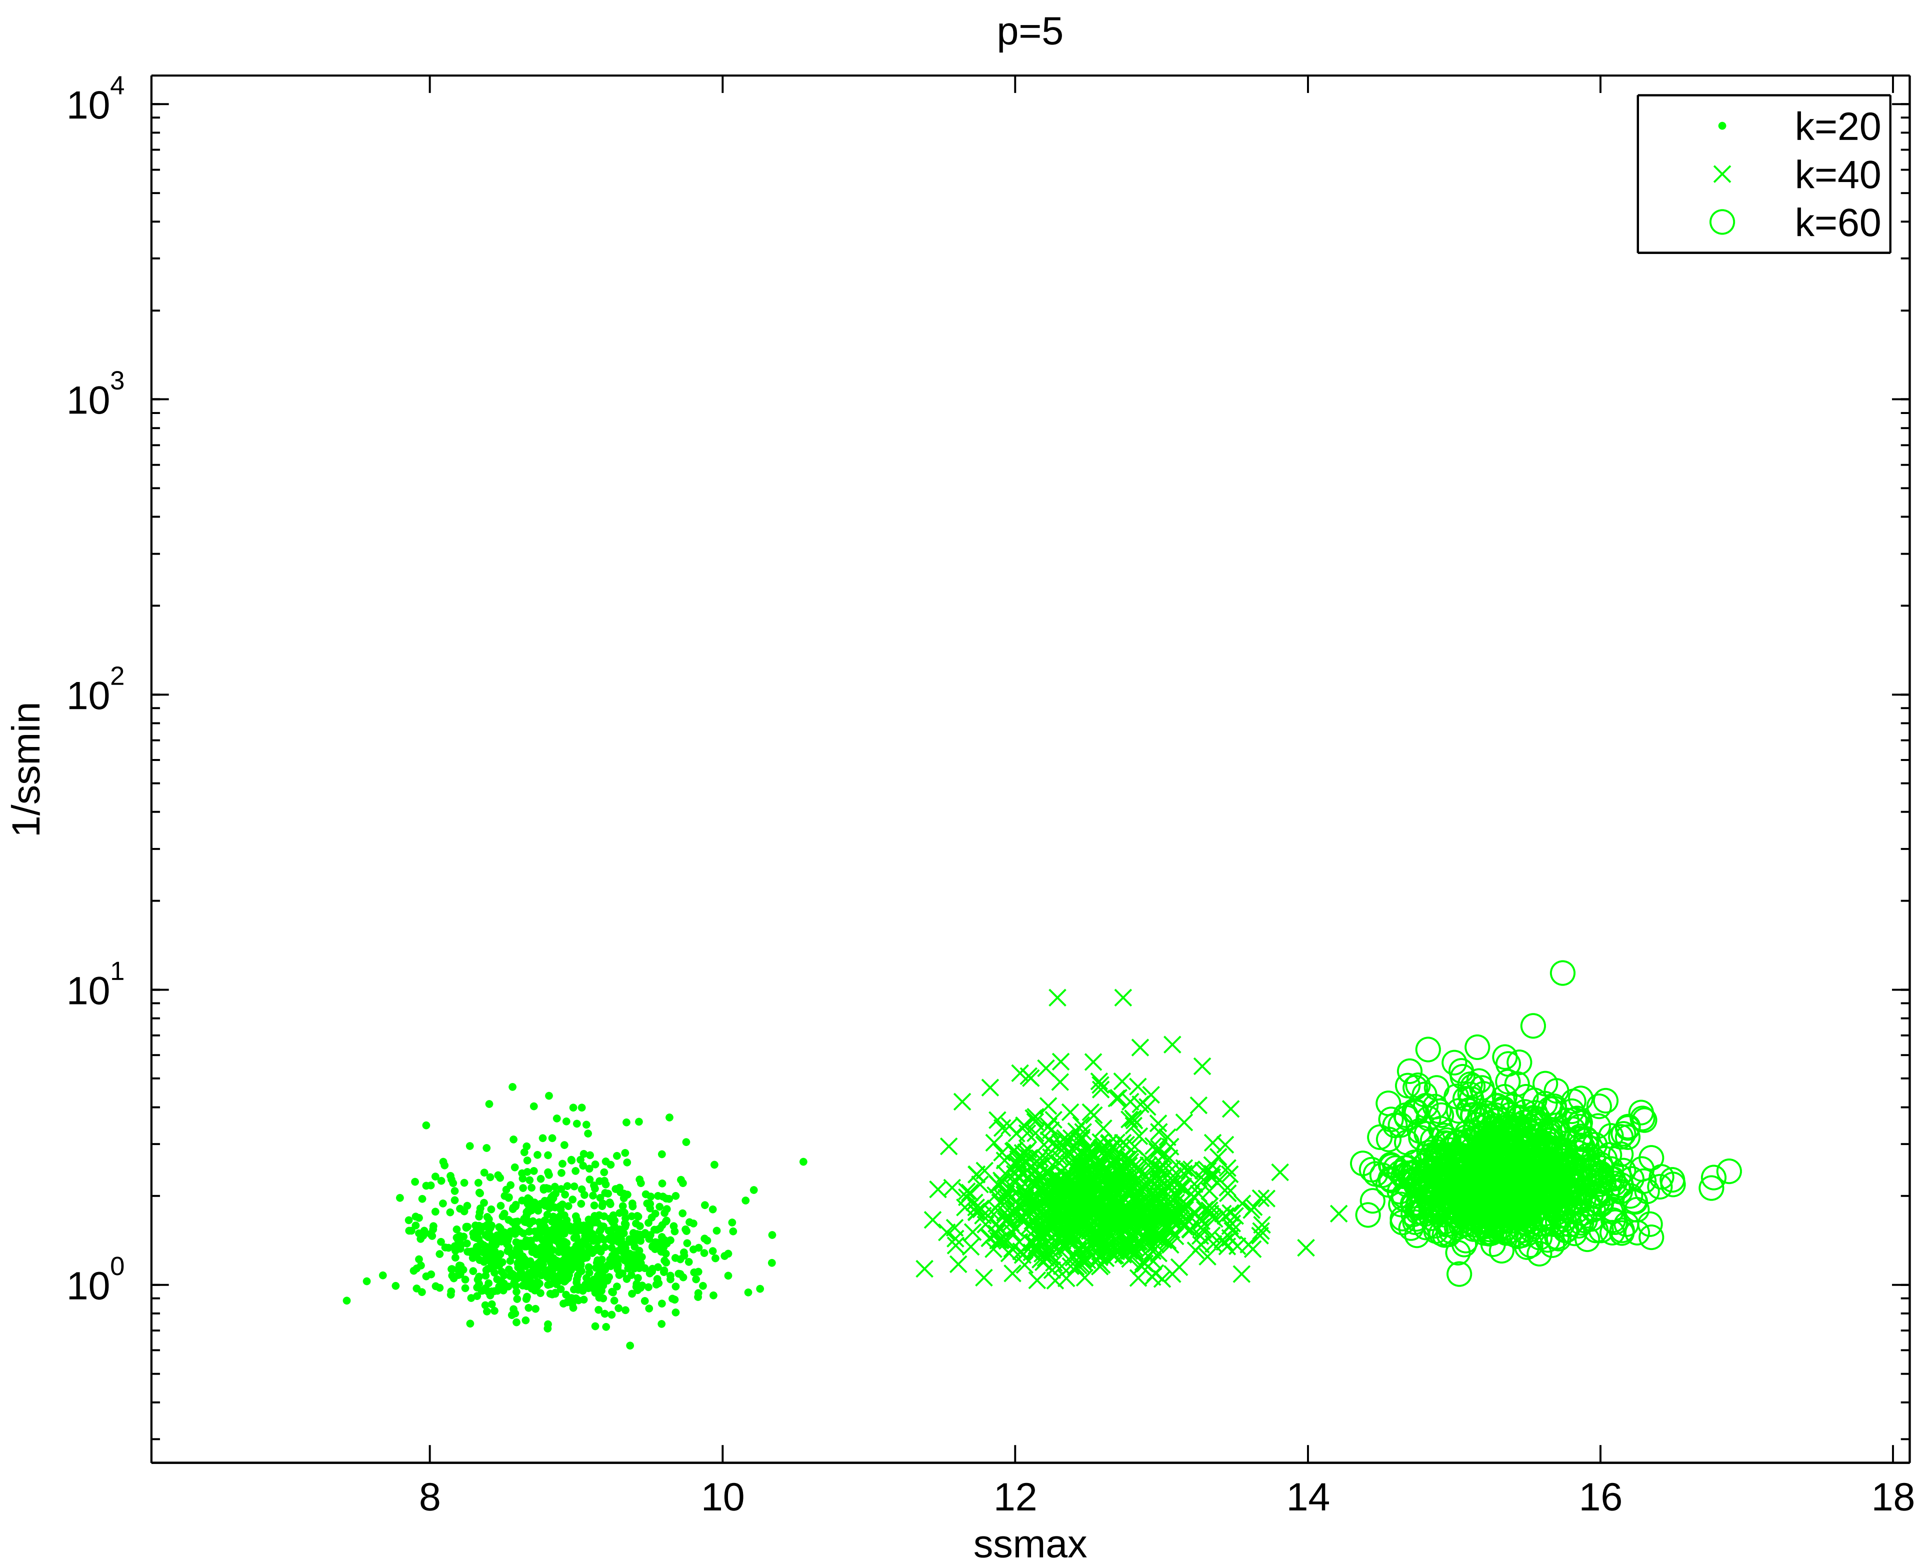

$1/\sigma_{\min}$ is plotted against $\sigma_{\max}$.

Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. $p = 10$

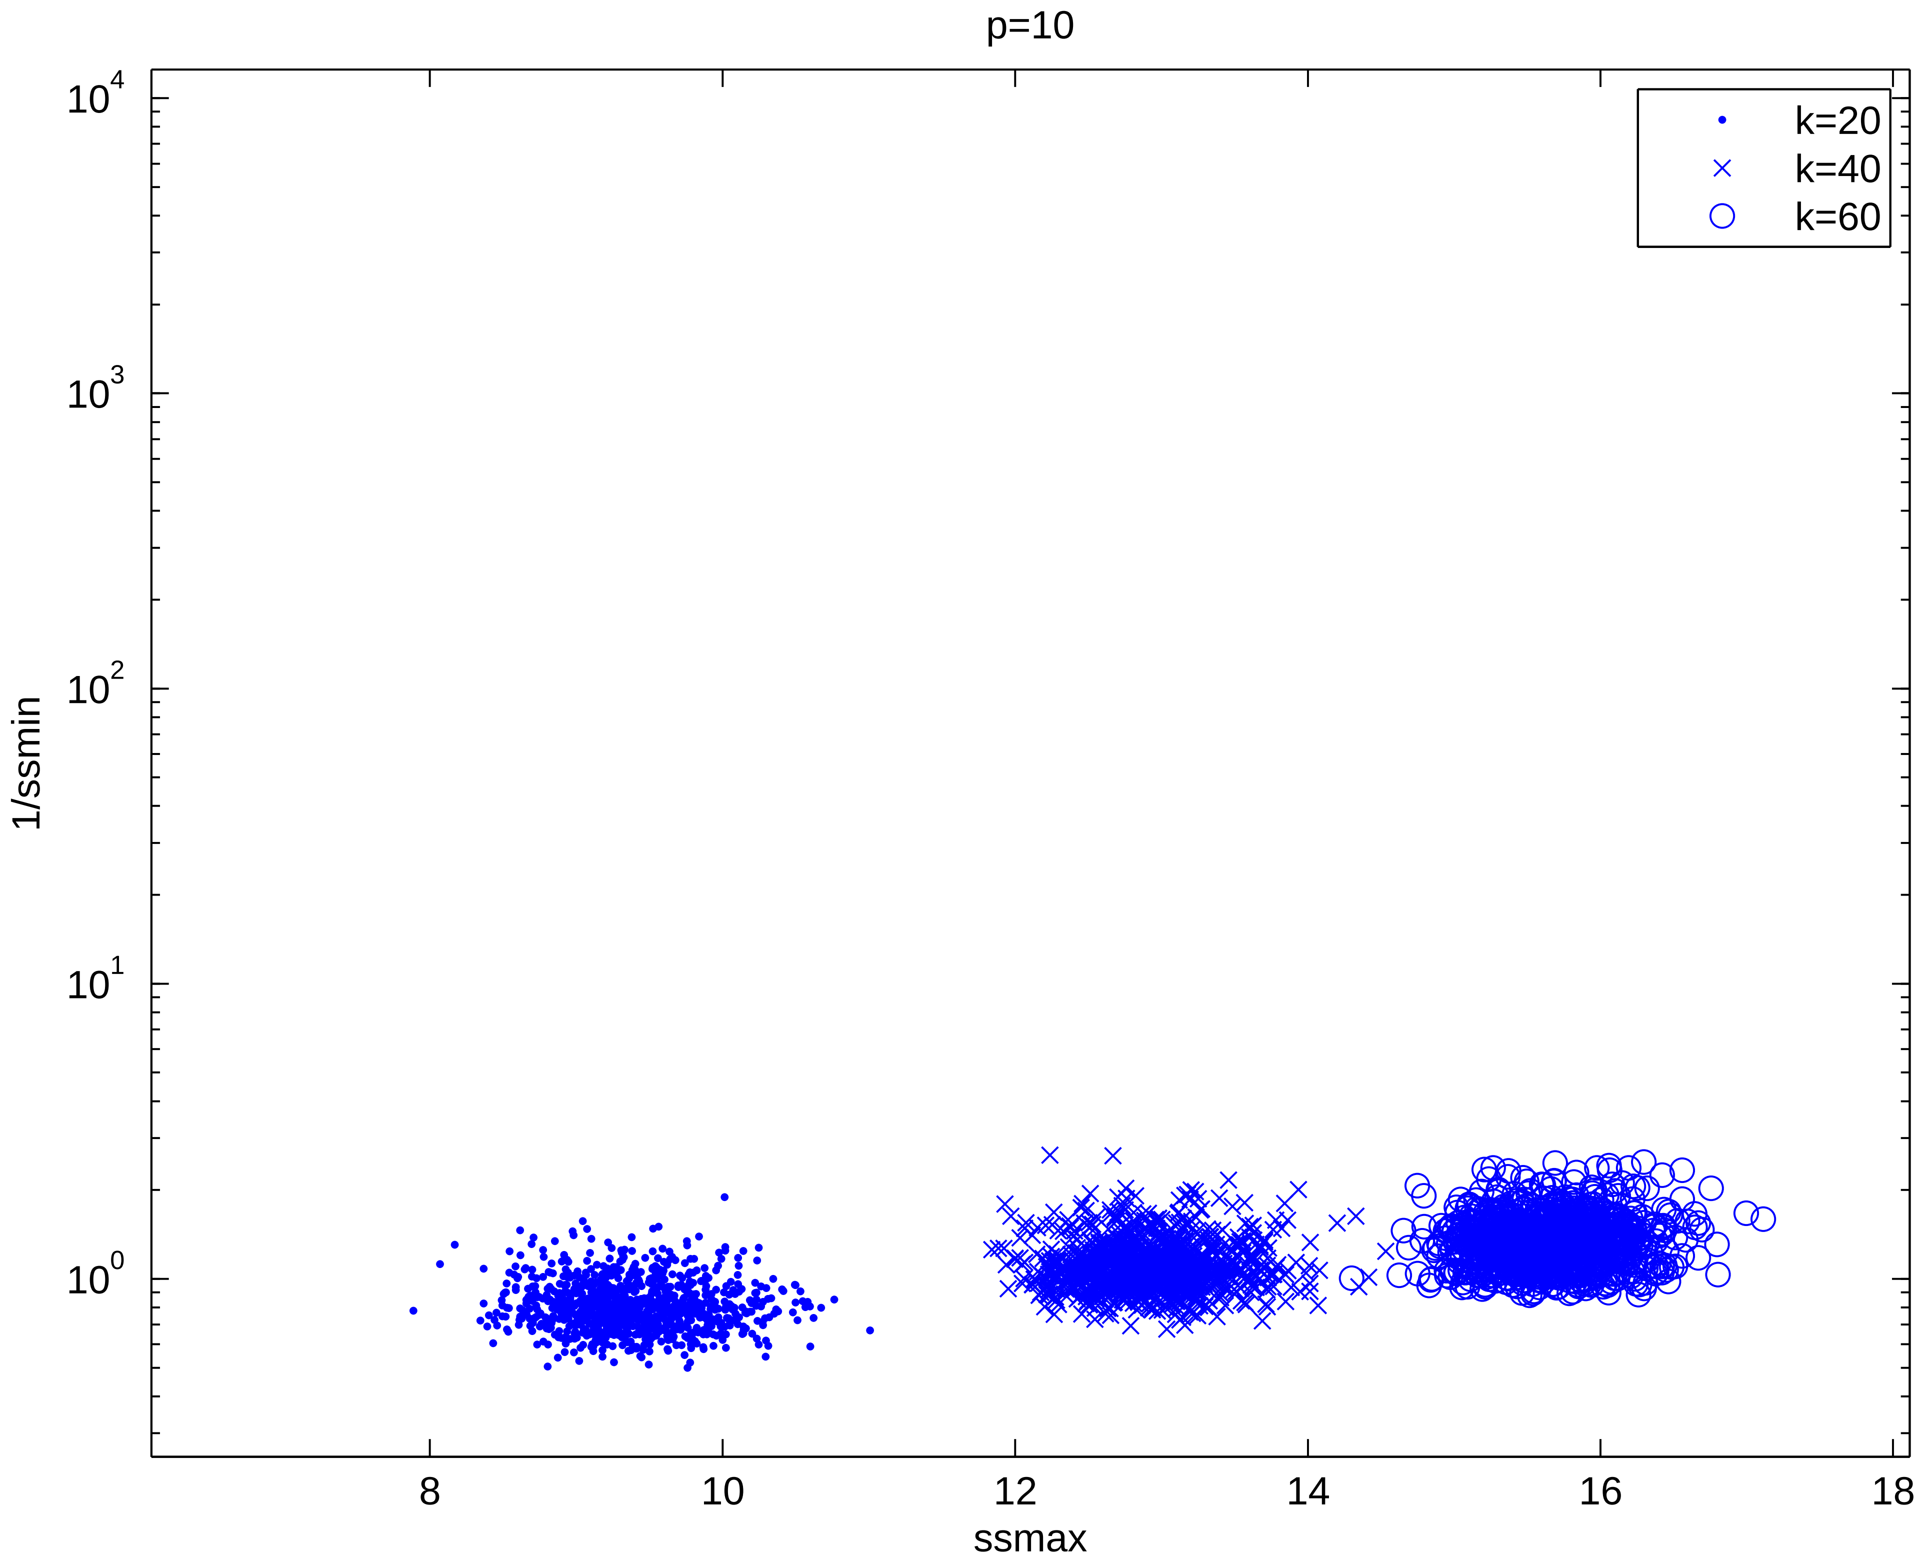

$1/\sigma_{\min}$ *is plotted against* $\sigma_{\max}$.

Scatter plot showing distribution of $k \times (k+p)$ Gaussian matrices.

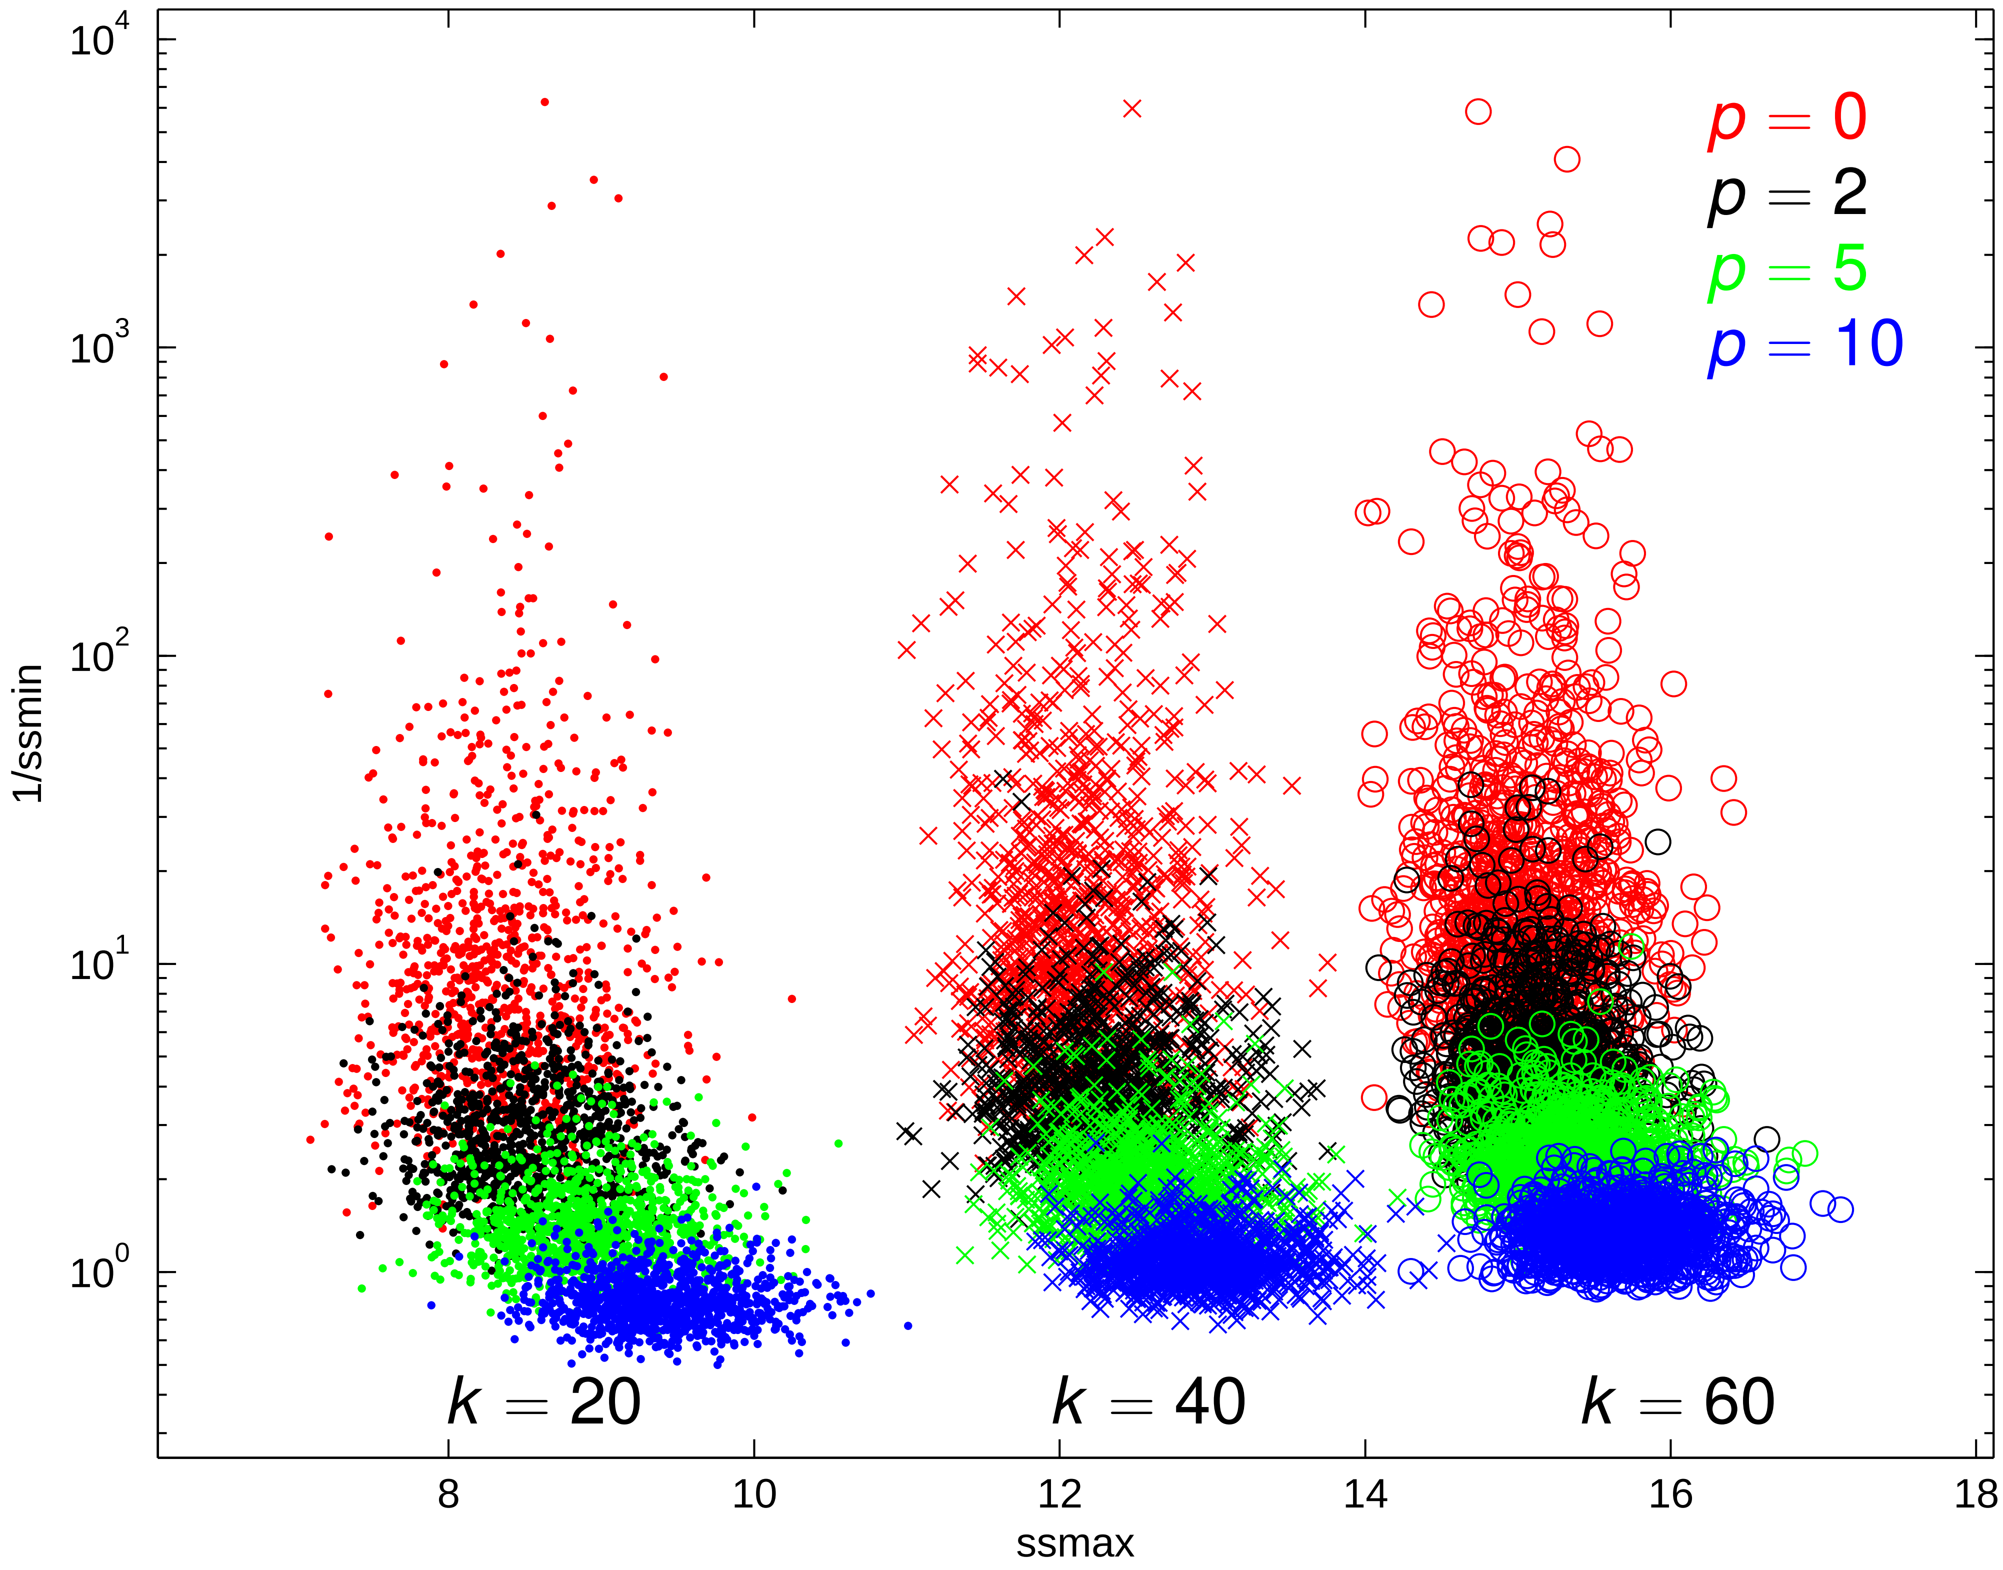

$1/\sigma_{\min}$ *is plotted against* $\sigma_{\max}$.

Simplistic proof that a rectangular Gaussian matrix is well-conditioned:

Let **G** denote a $k \times \ell$ Gaussian matrix where $k < \ell$.
Let "g" denote a generic $\mathcal{N}(0,1)$ variable and "r_j^2" denote a generic χ_j^2 variable. Then

$$\mathbf{G} \sim \begin{bmatrix} g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots & \end{bmatrix} \sim \begin{bmatrix} r_\ell & 0 & 0 & 0 & 0 & 0 & \cdots \\ g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots & \end{bmatrix}$$

$$\sim \begin{bmatrix} r_\ell & 0 & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & g & g & g & g & g & \cdots \\ 0 & g & g & g & g & g & \cdots \\ 0 & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots & \end{bmatrix} \sim \begin{bmatrix} r_\ell & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & g & g & g & g & \cdots \\ 0 & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix}$$

$$\sim \begin{bmatrix} r_\ell & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & r_{k-2} & g & g & g & \cdots \\ 0 & 0 & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix} \sim \cdots \sim \begin{bmatrix} r_\ell & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & r_{k-2} & r_{\ell-2} & 0 & 0 & \cdots \\ 0 & 0 & r_{k-3} & r_{\ell-3} & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \end{bmatrix}$$

Gershgorin's circle theorem will now show that **G** is well-conditioned if, e.g., $\ell = 2k$.

More sophisticated methods are required to get to $\ell = k + 2$.

Some results on Gaussian matrices. *Adapted from HMT 2009/2011; Gordon (1985,1988) for Proposition 1; Chen & Dongarra (2005) for Propositions 2 and 4; Bogdanov (1998) for Proposition 3.*

Proposition 1: Let **G** be a Gaussian matrix. Then

$$\begin{aligned}\left(\mathbb{E}\left[\|\mathbf{SGT}\|_{\mathrm{F}}^2\right]\right)^{1/2} &\leq \|\mathbf{S}\|_{\mathrm{F}}\,\|\mathbf{T}\|_{\mathrm{F}}\\ \mathbb{E}\left[\|\mathbf{SGT}\|\right] &\leq \|\mathbf{S}\|\,\|\mathbf{T}\|_{\mathrm{F}} + \|\mathbf{S}\|_{\mathrm{F}}\,\|\mathbf{T}\|\end{aligned}$$

Proposition 2: Let **G** be a Gaussian matrix of size $k \times k + p$ where $p \geq 2$. Then

$$\begin{aligned}\left(\mathbb{E}\left[\|\mathbf{G}^\dagger\|_{\mathrm{F}}^2\right]\right)^{1/2} &\leq \sqrt{\frac{k}{p-1}}\\ \mathbb{E}\left[\|\mathbf{G}^\dagger\|\right] &\leq \frac{e\sqrt{k+p}}{p}.\end{aligned}$$

Proposition 3: Suppose h is Lipschitz $|h(\mathbf{X}) - h(\mathbf{Y})| \leq L\|\mathbf{X}-\mathbf{Y}\|_{\mathrm{F}}$ and **G** is Gaussian. Then

$$\mathbb{P}\left[h(\mathbf{G}) > \mathbb{E}[h(\mathbf{G})] + L\,u\right] \leq e^{-u^2/2}.$$

Proposition 4: Suppose **G** is Gaussian of size $k \times k + p$ with $p \geq 4$. Then for $t \geq 1$:

$$\begin{aligned}\mathbb{P}\left[\|\mathbf{G}^\dagger\|_{\mathrm{F}} \geq \sqrt{\frac{3k}{p+1}}\,t\right] &\leq t^{-p}\\ \mathbb{P}\left[\|\mathbf{G}^\dagger\| \geq \frac{e\sqrt{k+p}}{p+1}\,t\right] &\leq t^{-(p+1)}\end{aligned}$$

Recall: $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|^2$, where $\mathbf{R}_1$ and $\mathbf{R}_2$ are Gaussian and $\mathbf{R}_1$ is $k \times k + p$.

Theorem: $\mathbb{E}\big[\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|\big] \leq \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j=k+1}^{\min(m,n)}\sigma_j^2\right)^{1/2}$.

Proof: First observe that

$$\mathbb{E}\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| = \mathbb{E}\big(\|\mathbf{D}_2\|^2 + \|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|^2\big)^{1/2} \leq \|\mathbf{D}_2\| + \mathbb{E}\|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|.$$

Condition on $\mathbf{R}_1$ and use Proposition 1:

$$\mathbb{E}\|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\| \leq \mathbb{E}\big[\|\mathbf{D}_2\|\,\|\mathbf{R}_1^\dagger\|_{\mathrm{F}} + \|\mathbf{D}_2\|_{\mathrm{F}}\,\|\mathbf{R}_1^\dagger\|\big]$$
$$\leq \{\text{Hölder}\} \leq \|\mathbf{D}_2\|\,\big(\mathbb{E}\|\mathbf{R}_1^\dagger\|_{\mathrm{F}}^2\big)^{1/2} + \|\mathbf{D}_2\|_{\mathrm{F}}\,\mathbb{E}\|\mathbf{R}_1^\dagger\|.$$

Proposition 2 now provides bounds for $\mathbb{E}\|\mathbf{R}_1^\dagger\|_{\mathrm{F}}^2$ and $\mathbb{E}\|\mathbf{R}_1^\dagger\|$ and we get

$$\mathbb{E}\|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\| \leq \sqrt{\frac{k}{p-1}}\|\mathbf{D}_2\| + \frac{e\sqrt{k+p}}{p}\|\mathbf{D}_2\|_{\mathrm{F}} = \sqrt{\frac{k}{p-1}}\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}.$$

Some results on Gaussian matrices. *Adapted from HMT2009/2011; Gordon (1985,1988) for Proposition 1; Chen & Dongarra (2005) for Propositions 2 and 4; Bogdanov (1998) for Proposition 3.*

Proposition 1: Let $\mathbf{G}$ be a Gaussian matrix. Then

$$\left(\mathbb{E}\left[\|\mathbf{SGT}\|_{\mathrm{F}}^2\right]\right)^{1/2} \leq \|\mathbf{S}\|_{\mathrm{F}}\, \|\mathbf{T}\|_{\mathrm{F}}$$

$$\mathbb{E}\left[\|\mathbf{SGT}\|\right] \leq \|\mathbf{S}\|\, \|\mathbf{T}\|_{\mathrm{F}} + \|\mathbf{S}\|_{\mathrm{F}}\, \|\mathbf{T}\|$$

Proposition 2: Let $\mathbf{G}$ be a Gaussian matrix of size $k \times k + p$ where $p \geq 2$. Then

$$\left(\mathbb{E}\left[\|\mathbf{G}^{\dagger}\|_{\mathrm{F}}^2\right]\right)^{1/2} \leq \sqrt{\frac{k}{p-1}}$$

$$\mathbb{E}\left[\|\mathbf{G}^{\dagger}\|\right] \leq \frac{e\sqrt{k+p}}{p}.$$

Proposition 3: Suppose h is Lipschitz $|h(\mathbf{X}) - h(\mathbf{Y})| \leq L\|\mathbf{X} - \mathbf{Y}\|_{\mathrm{F}}$ and $\mathbf{G}$ is Gaussian. Then

$$\mathbb{P}\left[h(\mathbf{G}) > \mathbb{E}[h(\mathbf{G})] + L\,u\right] \leq e^{-u^2/2}.$$

Proposition 4: Suppose $\mathbf{G}$ is Gaussian of size $k \times k + p$ with $p \geq 4$. Then for $t \geq 1$:

$$\mathbb{P}\left[\|\mathbf{G}^{\dagger}\|_{\mathrm{F}} \geq \sqrt{\frac{3k}{p+1}}\,t\right] \leq t^{-p}$$

$$\mathbb{P}\left[\|\mathbf{G}^{\dagger}\| \geq \frac{e\sqrt{k+p}}{p+1}\,t\right] \leq t^{-(p+1)}$$

Recall: $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|^2 \leq \|\mathbf{D}_2\|^2 + \|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|^2$, where $\mathbf{R}_1$ and $\mathbf{R}_2$ are Gaussian and $\mathbf{R}_1$ is $k \times k + p$.

Theorem: With probability at least $1 - 2\,t^{-p} - e^{-u^2/2}$ it holds that

$$\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \left(1 + t\sqrt{\frac{3k}{p+1}} + u\,t\,\frac{e\,\sqrt{k+p}}{p+1}\right)\sigma_{k+1} + \frac{t\,e\,\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}.$$

Proof: Set $E_t = \left\{\|\mathbf{R}_1\| \leq \frac{e\sqrt{k+p}}{p+1}t \quad \text{and} \quad \|\mathbf{R}_1^\dagger\|_{\mathrm{F}} \leq \sqrt{\frac{3k}{p+1}}\,t\right\}$. By Proposition 4: $\mathbb{P}(E_t^{\mathrm{c}}) \leq 2\,t^{-p}$.

Set $h(\mathbf{X}) = \|\mathbf{D}_2\mathbf{X}\mathbf{R}_1^\dagger\|$. A direct calculation shows

$$|h(\mathbf{X}) - h(\mathbf{Y})| \leq \|\mathbf{D}_2\|\,\|\mathbf{R}_1^\dagger\|\,\|\mathbf{X} - \mathbf{y}\|_{\mathrm{F}}.$$

Hold $\mathbf{R}_1$ fixed and take the expectation on $\mathbf{R}_2$. Then Proposition 1 applies and so

$$\mathbb{E}\big[h(\mathbf{R}_2)\,\big|\,\mathbf{R}_1\big] \leq \|\mathbf{D}_2\|\,\|\mathbf{R}_1^\dagger\|_{\mathrm{F}} + \|\mathbf{D}_2\|_{\mathrm{F}}\,\|\mathbf{R}_1^\dagger\|.$$

Now use Proposition 3 (concentration of measure)

$$\mathbb{P}\Big[\underbrace{\|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\|}_{=h(\mathbf{R}_2)} > \underbrace{\|\mathbf{D}_2\|\,\|\mathbf{R}_1^\dagger\|_{\mathrm{F}} + \|\mathbf{D}_2\|_{\mathrm{F}}\,\|\mathbf{R}_1^\dagger\|}_{=\mathbb{E}[h(\mathbf{R}_2)]} + \underbrace{\|\mathbf{D}_2\|\,\|\mathbf{R}_1^\dagger\|}_{=L}\,u\,\Big|\,E_t\Big] < e^{-u^2/2}.$$

When E_t holds true, we have bounds on the "badness" of $\mathbf{R}_1^\dagger$:

$$\mathbb{P}\Big[\|\mathbf{D}_2\mathbf{R}_2\mathbf{R}_1^\dagger\| > \|\mathbf{D}_2\|\sqrt{\frac{3k}{p+1}}t + \|\mathbf{D}_2\|_{\mathrm{F}}\frac{e\sqrt{k+p}}{p+1}t + \|\mathbf{D}_2\|\frac{e\sqrt{k+p}}{p+1}ut\,\Big|\,E_t\Big] < e^{-u^2/2}.$$

The theorem is obtained by using $\mathbb{P}(E_t^{\mathrm{c}}) \leq 2\,t^{-p}$ to remove the conditioning of E_t.

Example 1:

We consider a $1\,000 \times 1\,000$ matrix $\mathbf{A}$ whose singular values are shown below:

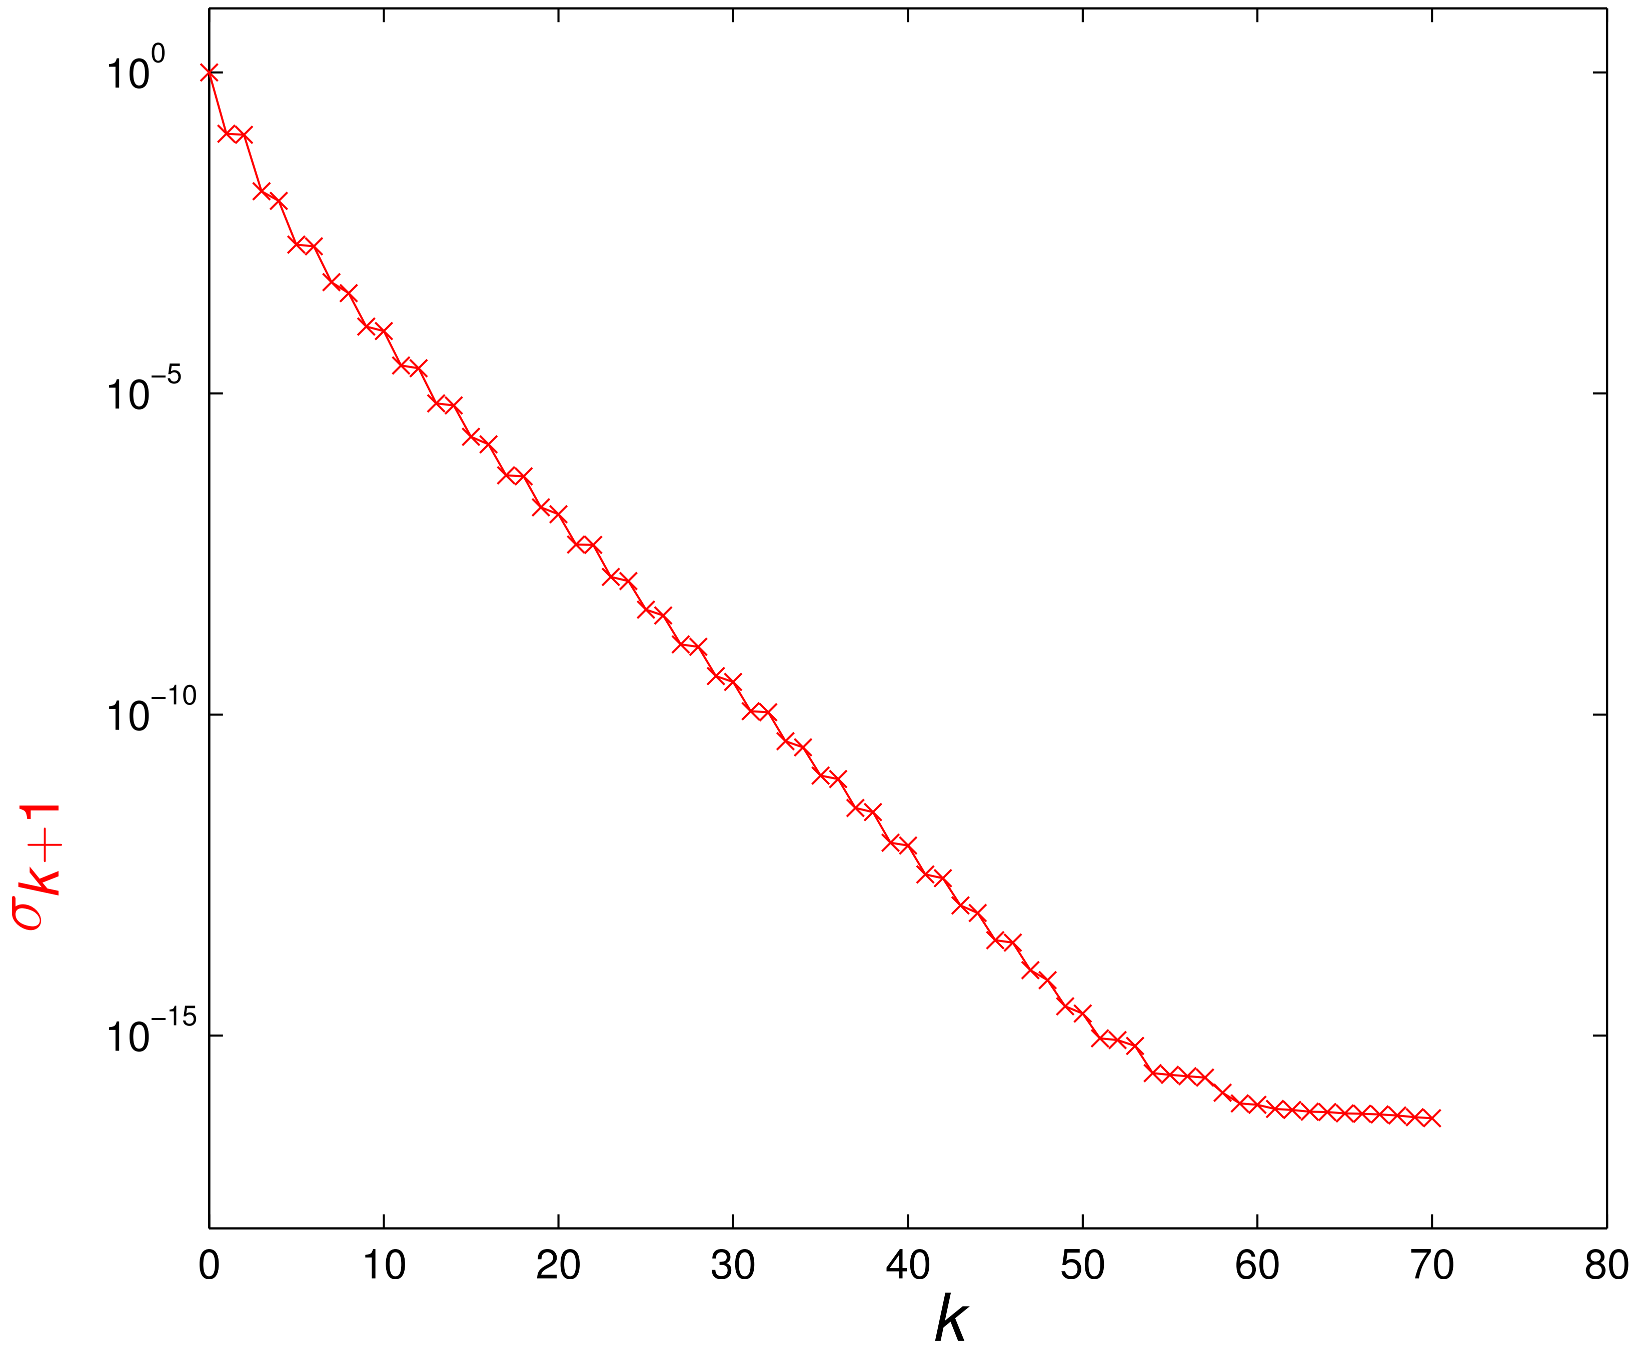

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 1:

We consider a $1\,000 \times 1\,000$ matrix $\mathbf{A}$ whose singular values are shown below:

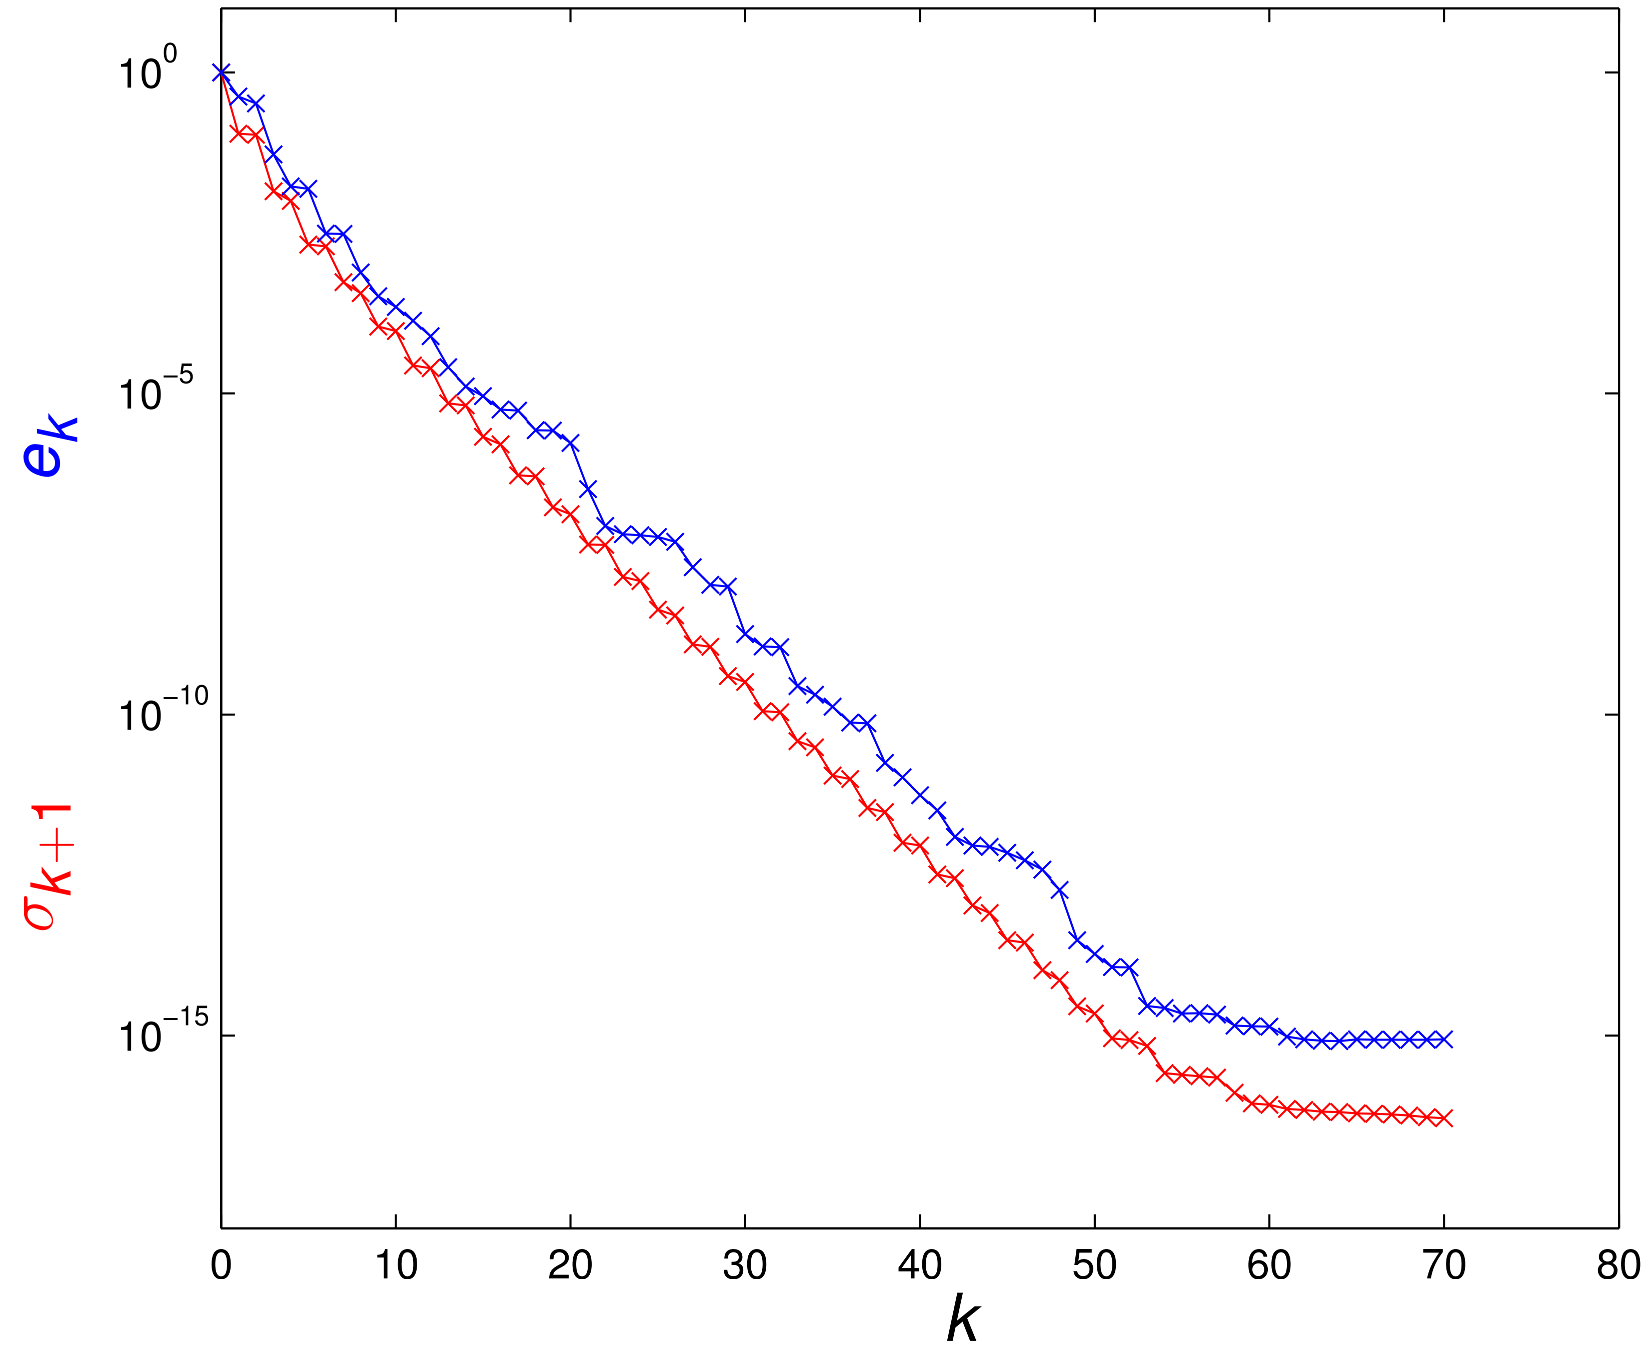

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

The blue line indicates the actual errors e_k incurred by one instantiation of the proposed method.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 1:

We consider a $1\,000 \times 1\,000$ matrix $\mathbf{A}$ whose singular values are shown below:

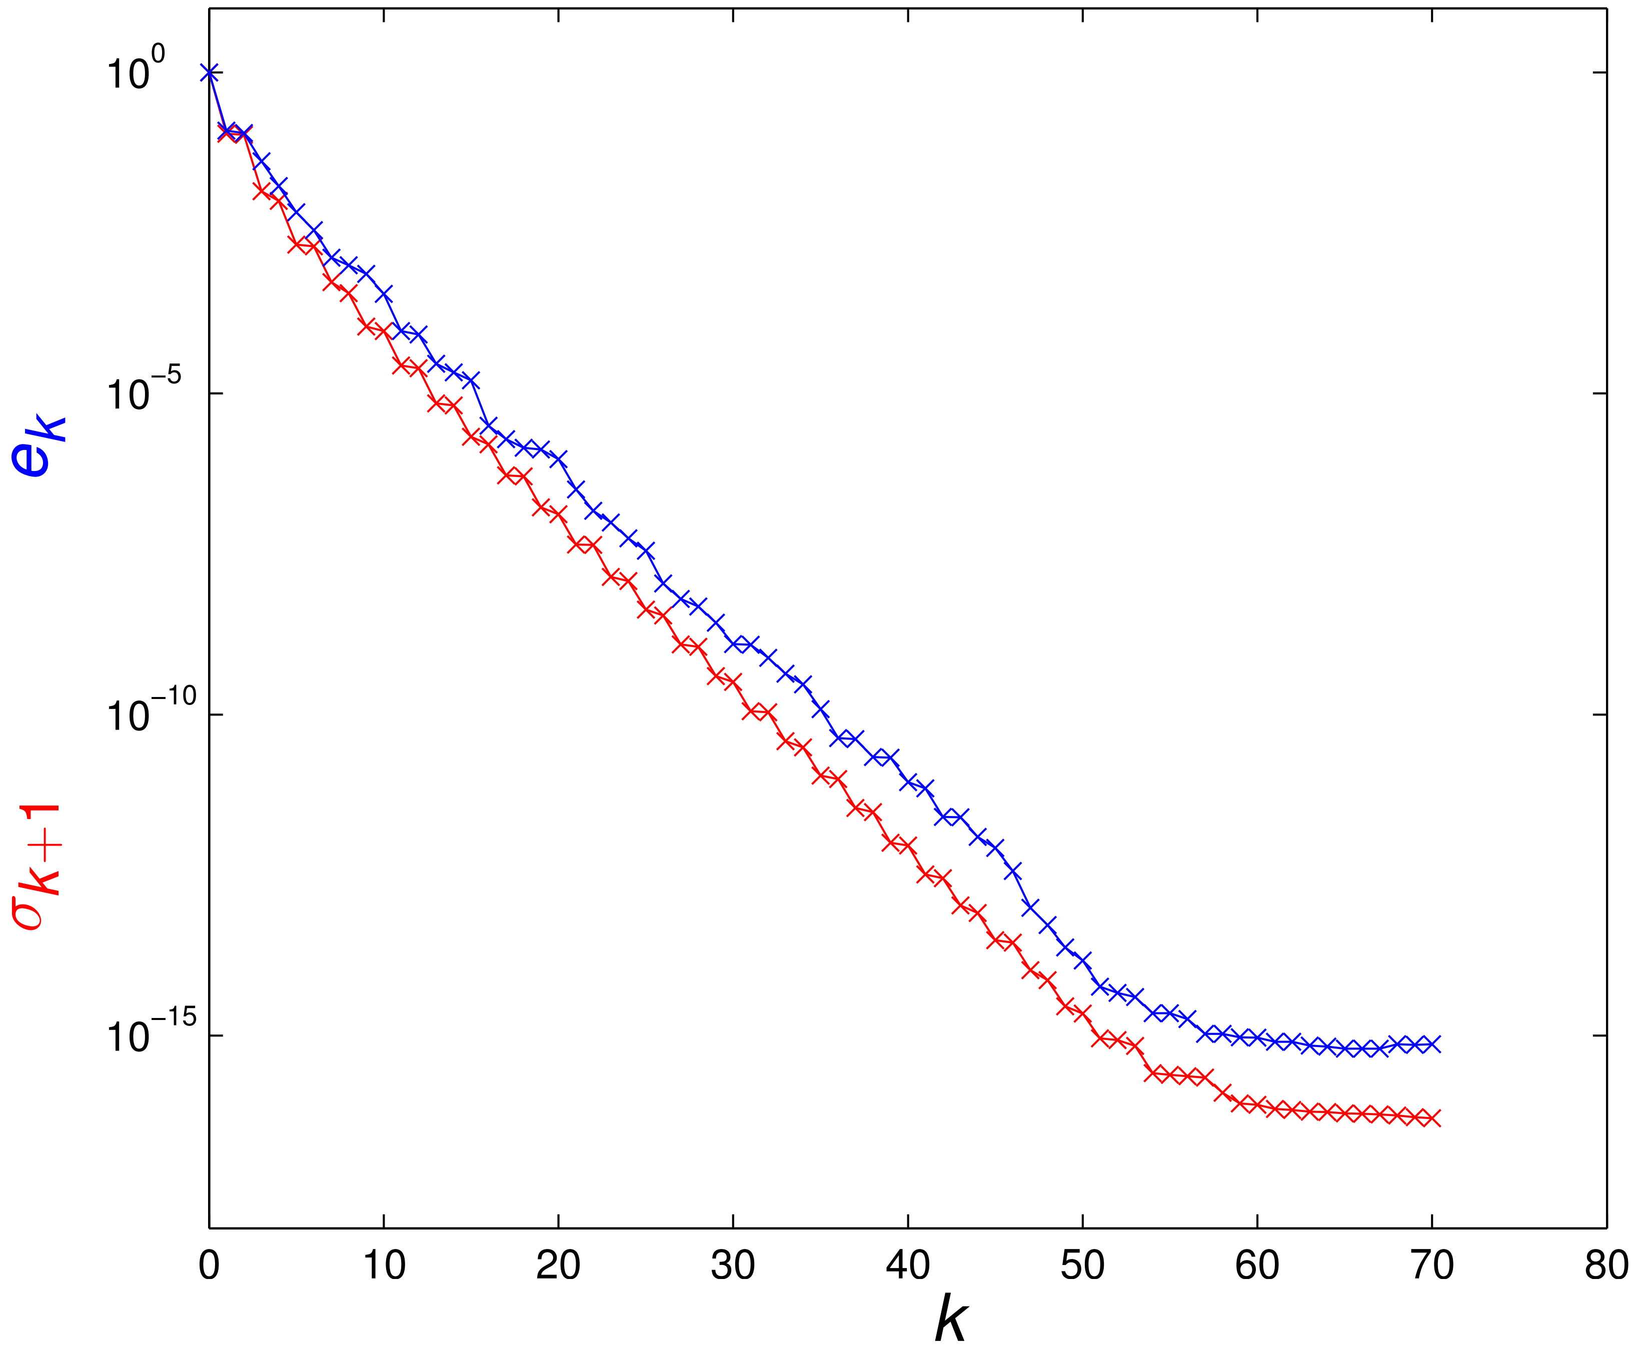

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

The blue line indicates the actual errors e_k incurred by a different instantiation of the proposed method.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 1:

We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:

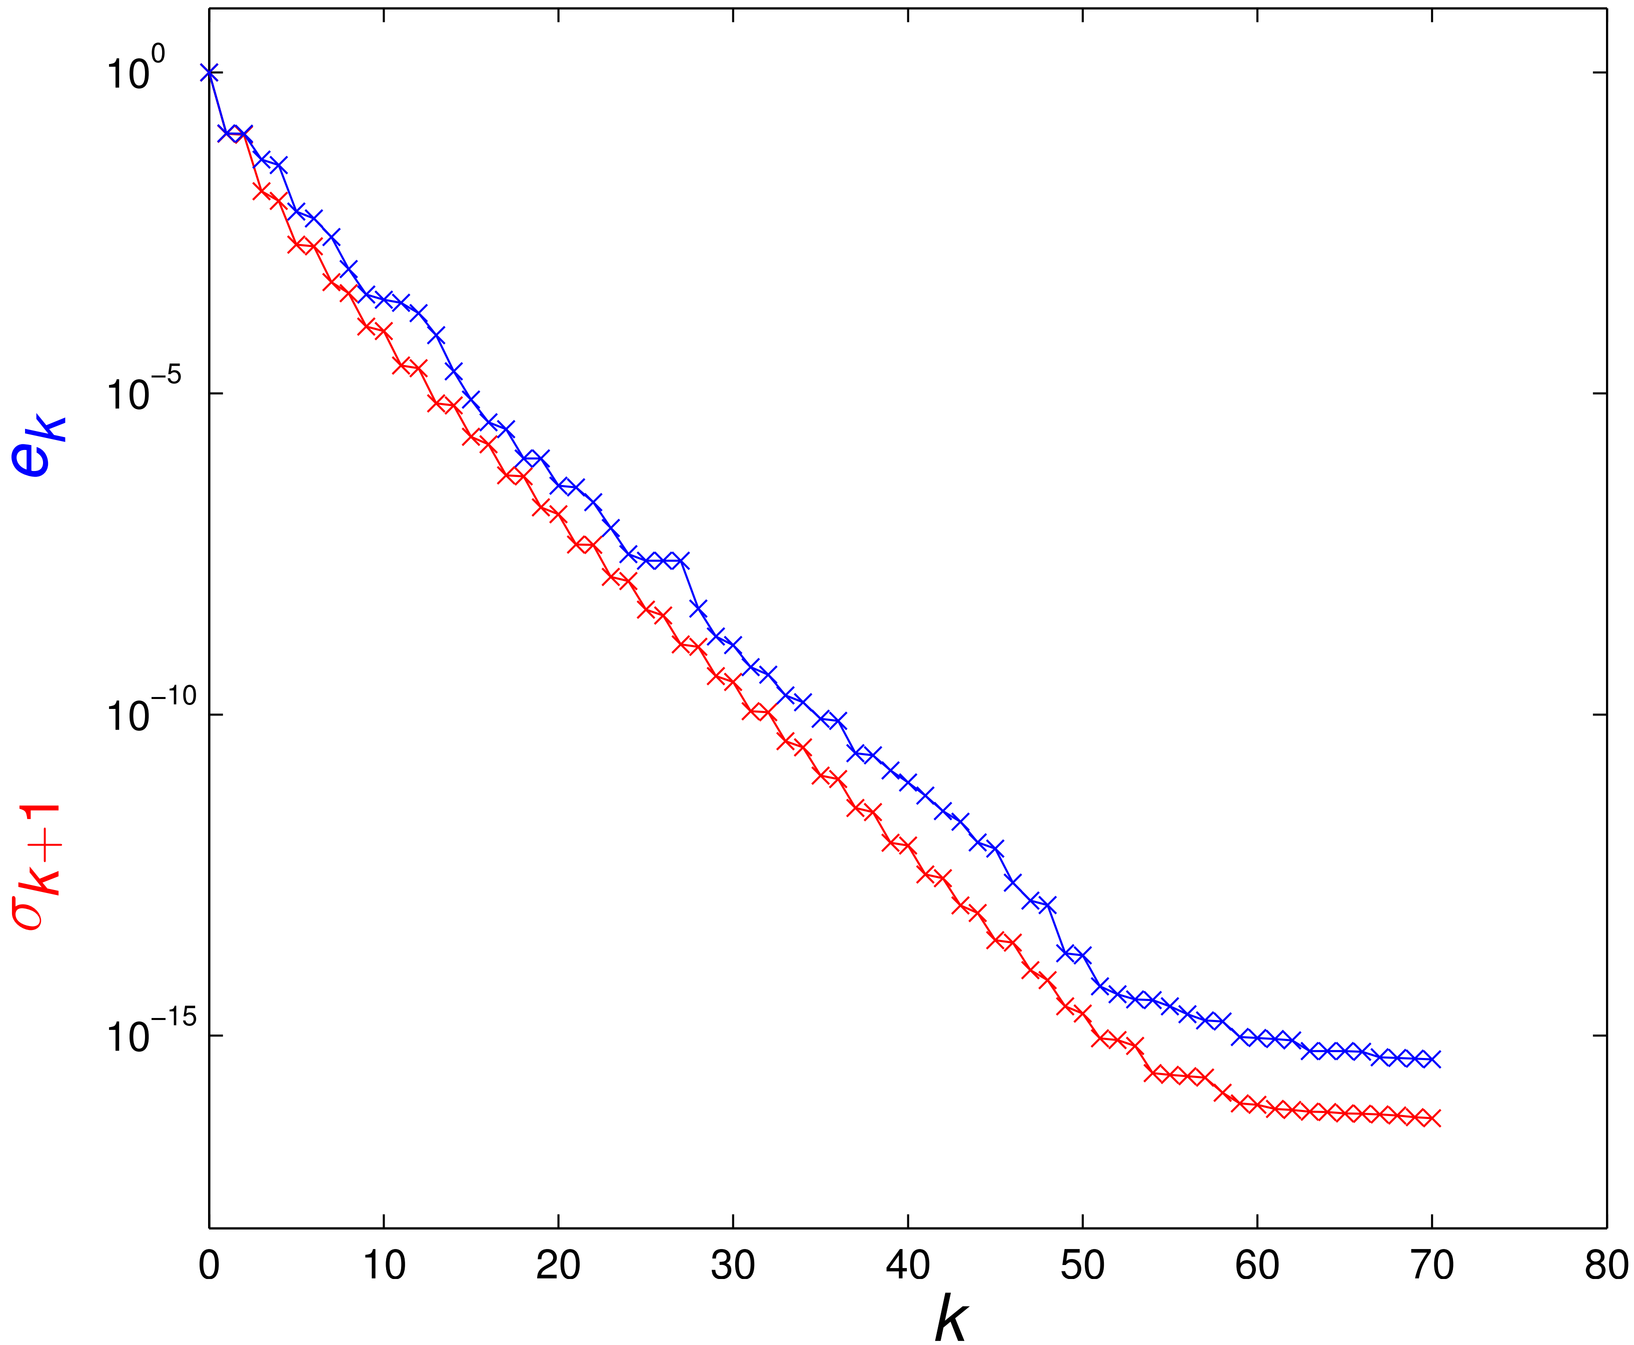

The red line indicates the singular values σ_{k+1} of **A**. These indicate the theoretically minimal approximation error.

The blue line indicates the actual errors e_k incurred by a different instantiation of the proposed method.

A is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of **A** is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 1:

We consider a $1\,000 \times 1\,000$ matrix $\mathbf{A}$ whose singular values are shown below:

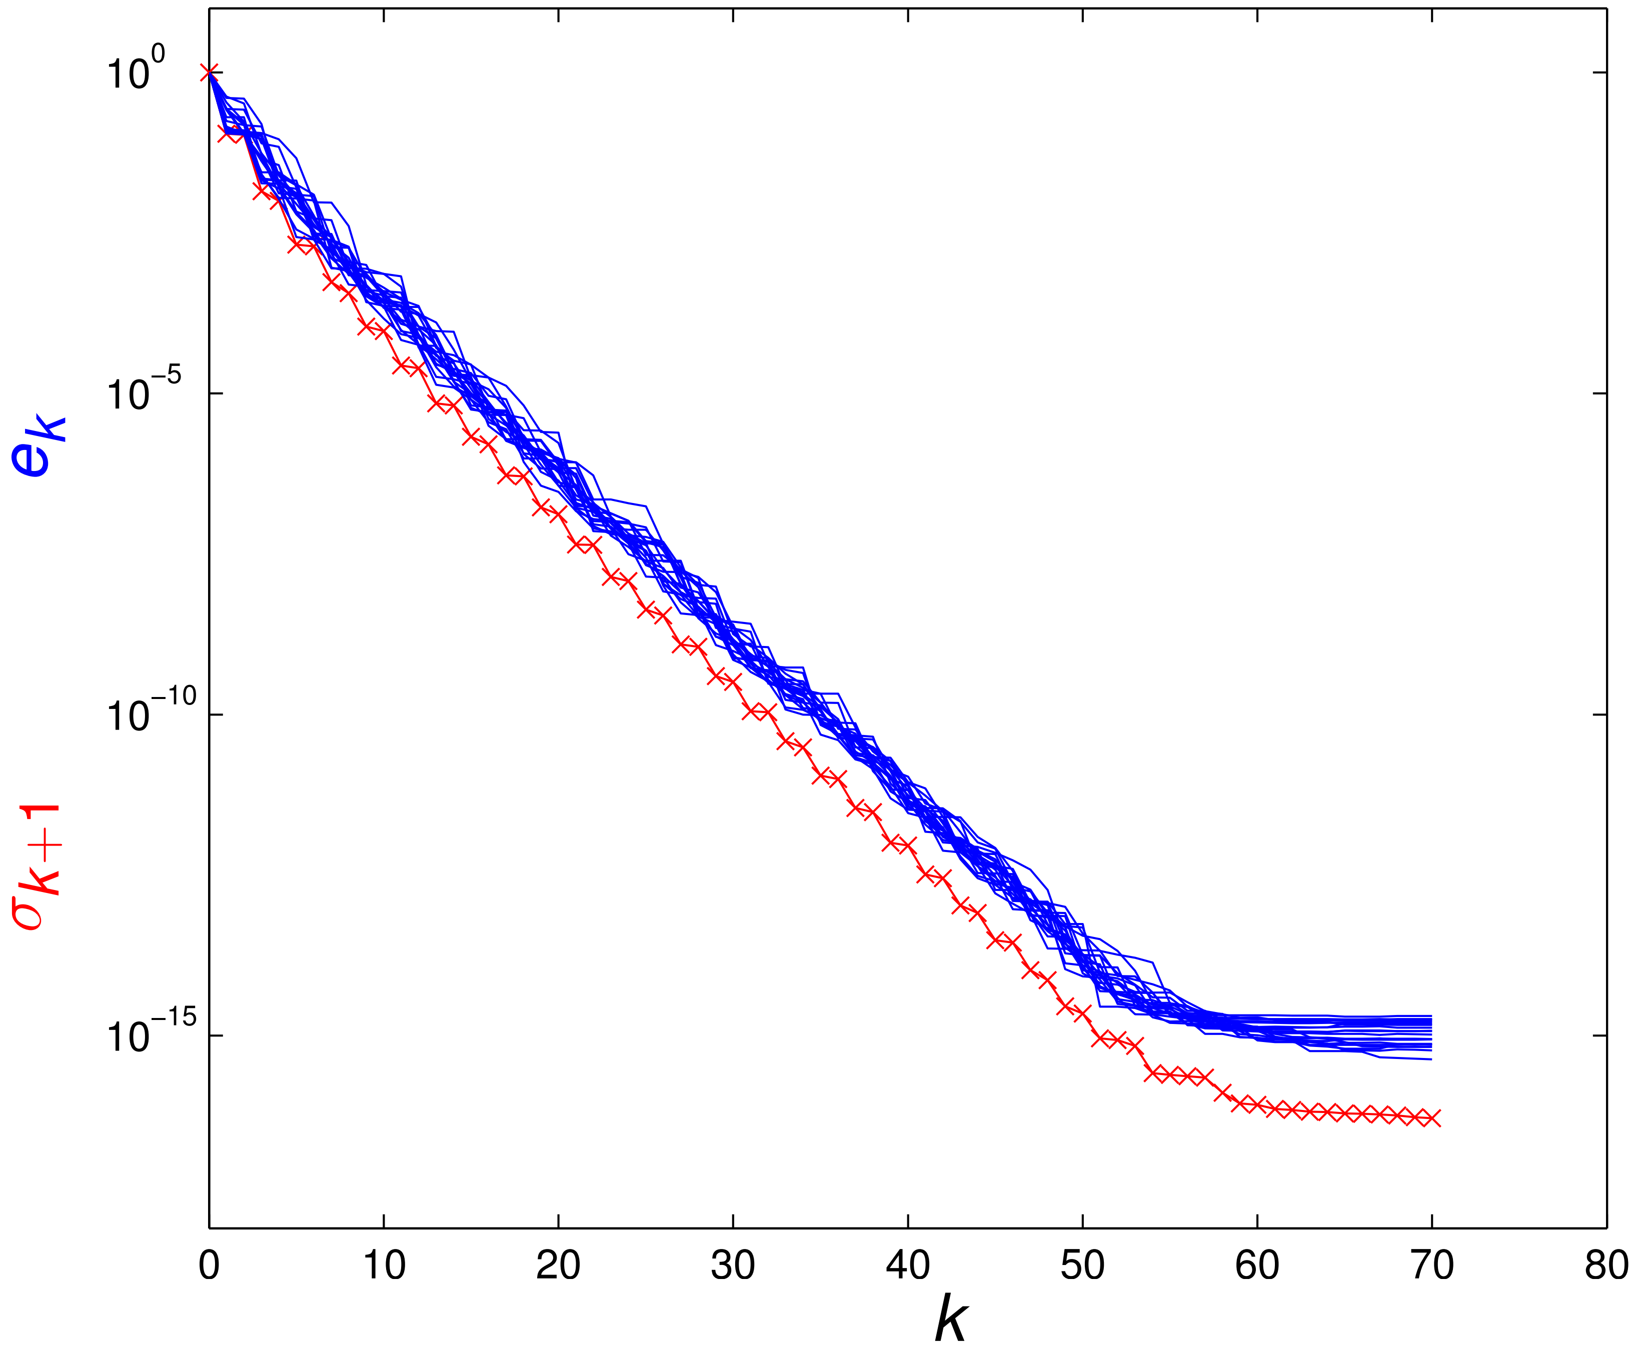

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

The blue lines indicate the actual errors e_k incurred by 20 instantiations of the proposed method.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 2:

We consider a 1 000 $\times$ 1 000 matrix $\mathbf{A}$ whose singular values are shown below:

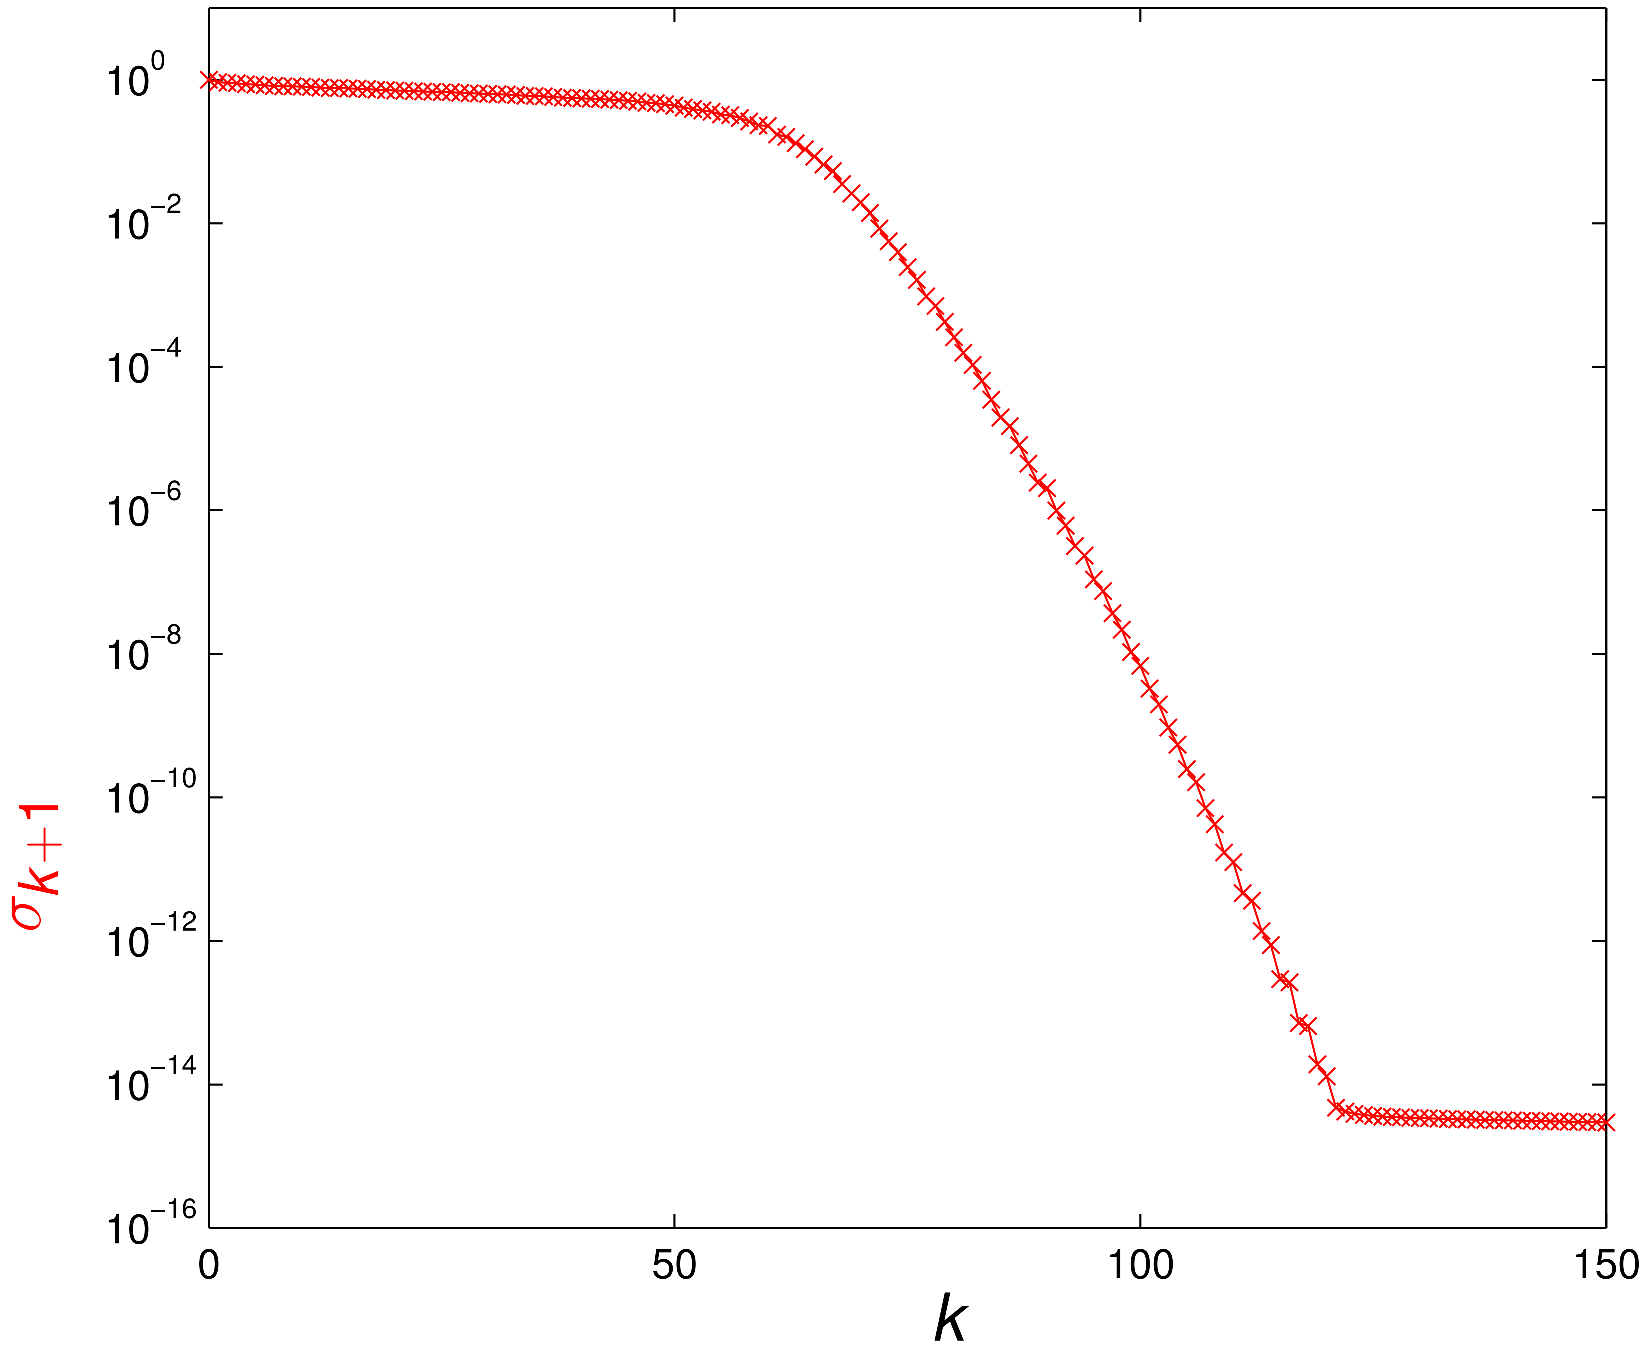

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 2:

We consider a $1\,000 \times 1\,000$ matrix $\mathbf{A}$ whose singular values are shown below:

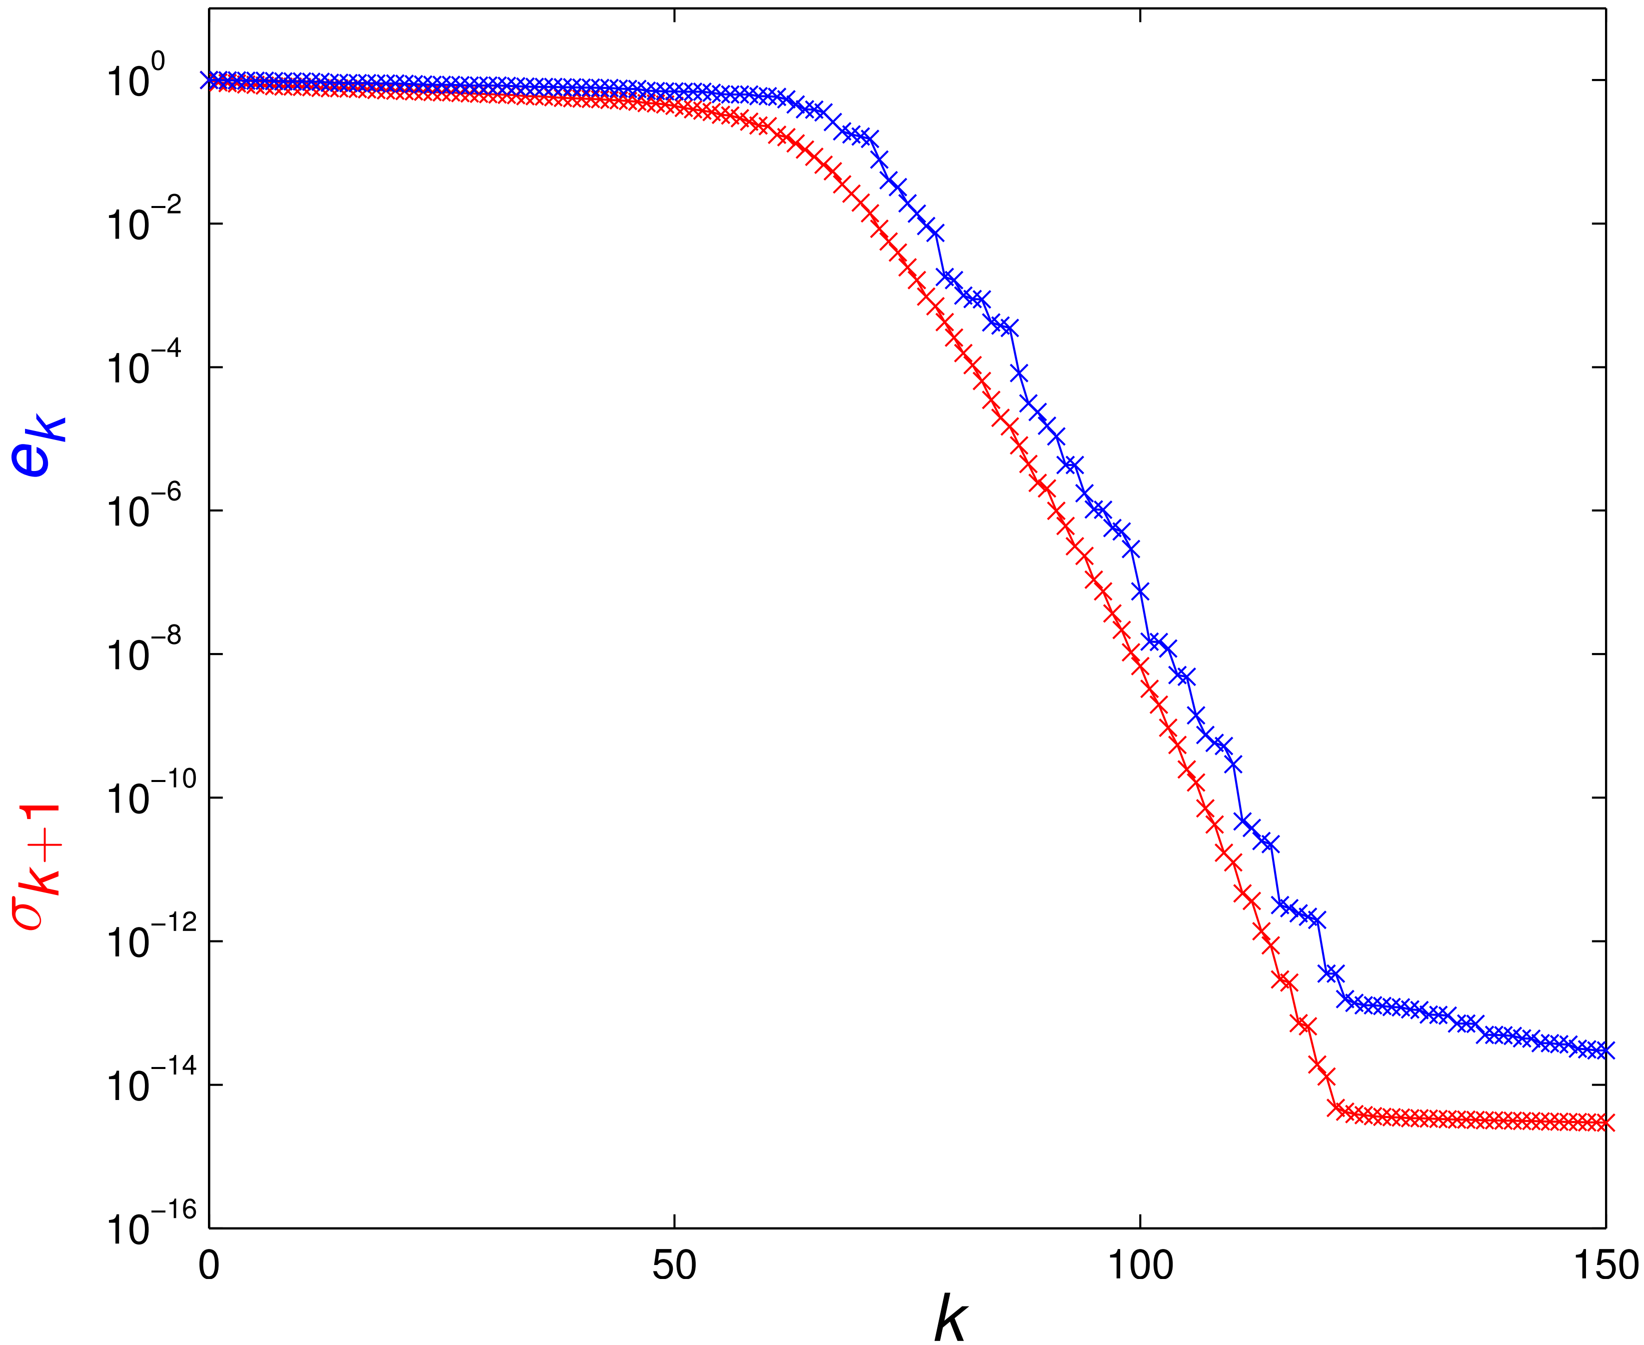

The red line indicates the singular values σ_{k+1} of $\mathbf{A}$. These indicate the theoretically minimal approximation error.

The blue line indicates the actual errors e_k incurred by one instantiation of the proposed method.

$\mathbf{A}$ is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of $\mathbf{A}$ is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 2:

We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:

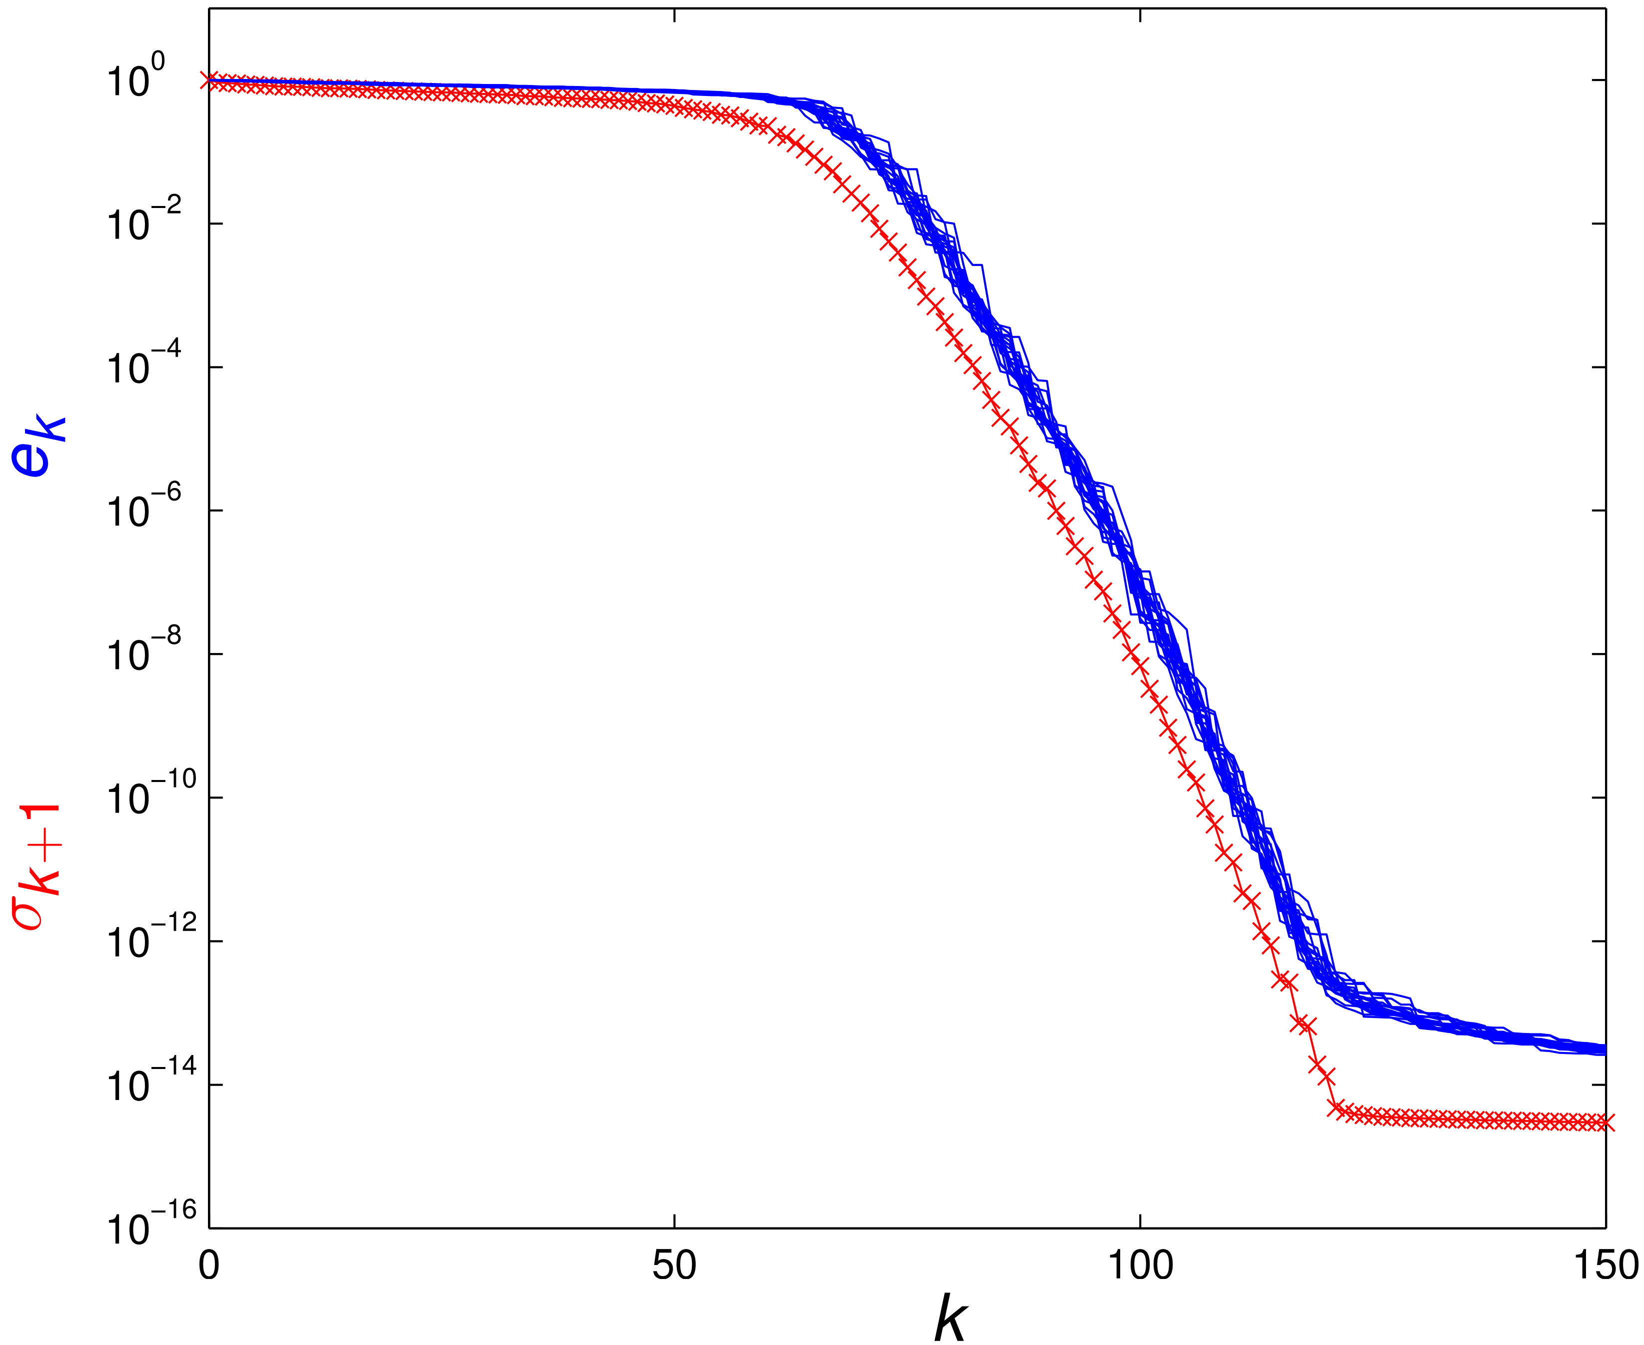

The red line indicates the singular values σ_{k+1} of **A**. These indicate the theoretically minimal approximation error.

The blue lines indicate the actual errors e_k incurred by 20 instantiations of the proposed method.

A is a discrete approximation of a certain compact integral operator normalized so that $\|\mathbf{A}\| = 1$.

Curiously, the nature of **A** is in a strong sense irrelevant: the error distribution depends only on $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Example 3:

The matrix **A** being analyzed is a 9025×9025 matrix arising in a diffusion geometry approach to image processing.

To be precise, **A** is a graph Laplacian on the manifold of 3×3 patches.

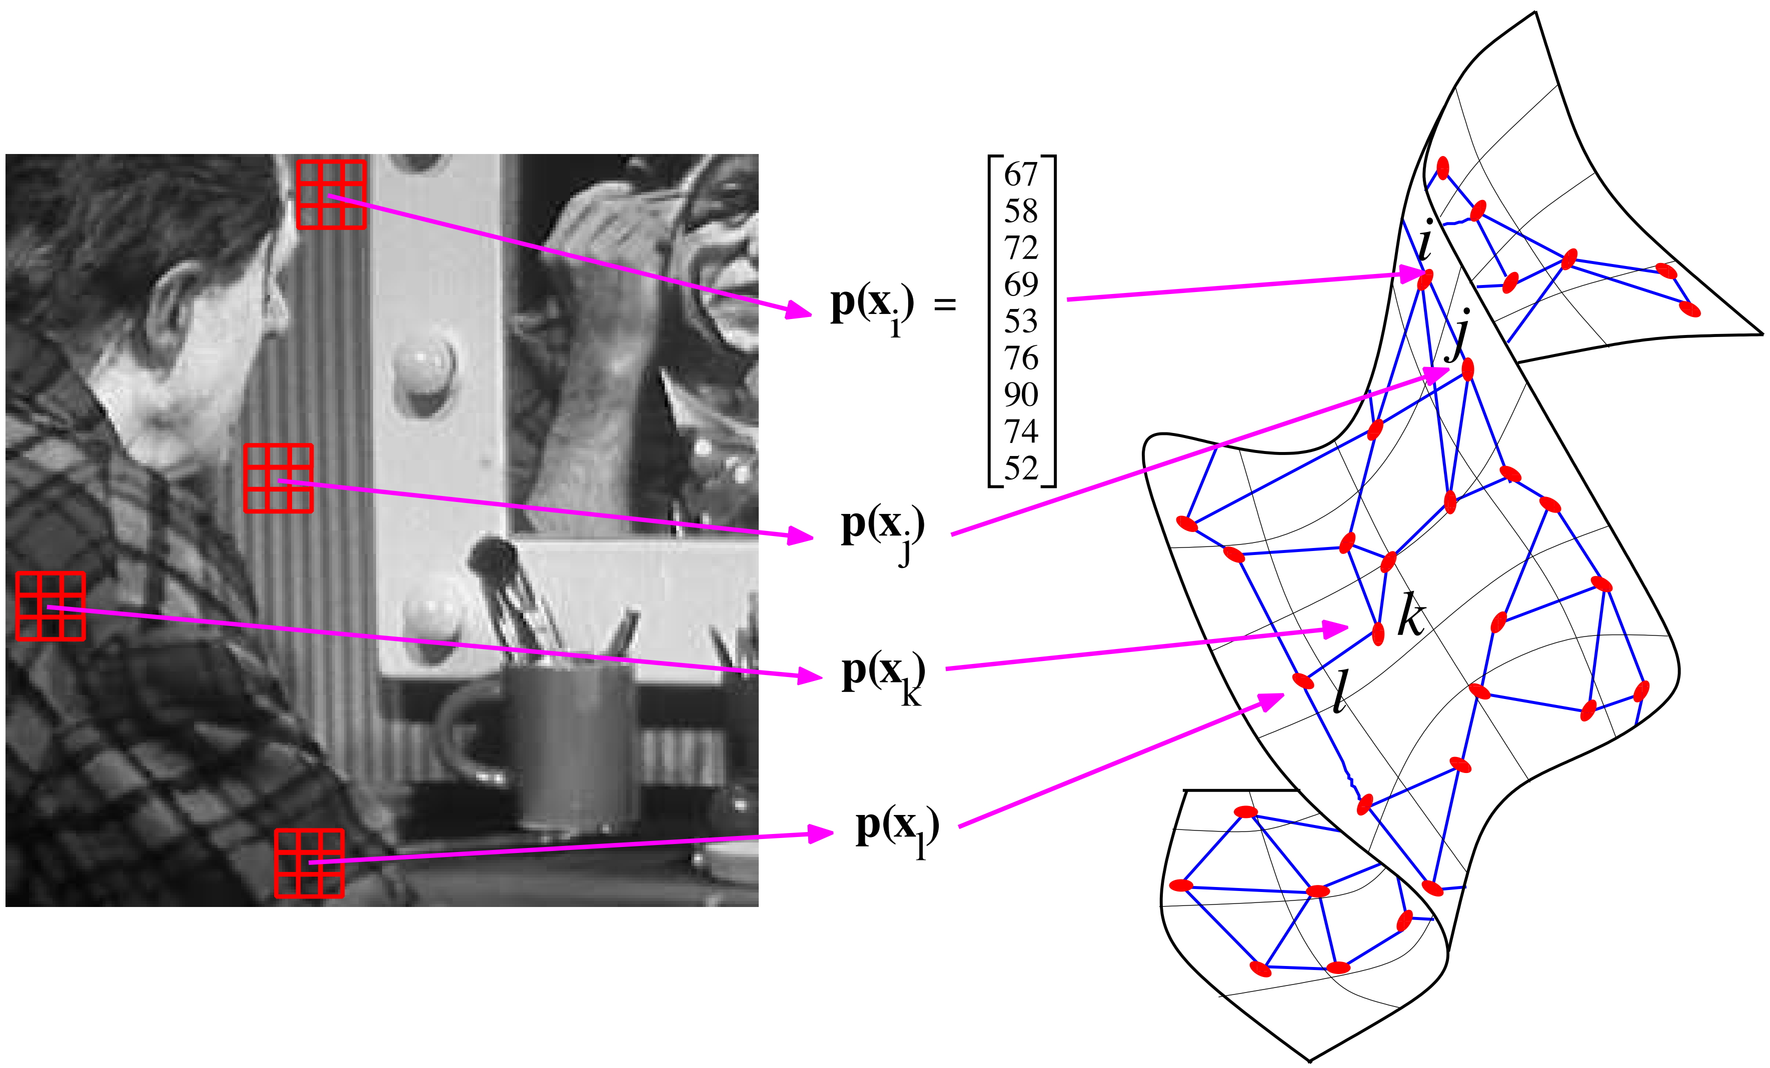

Joint work with François Meyer of the University of Colorado at Boulder.

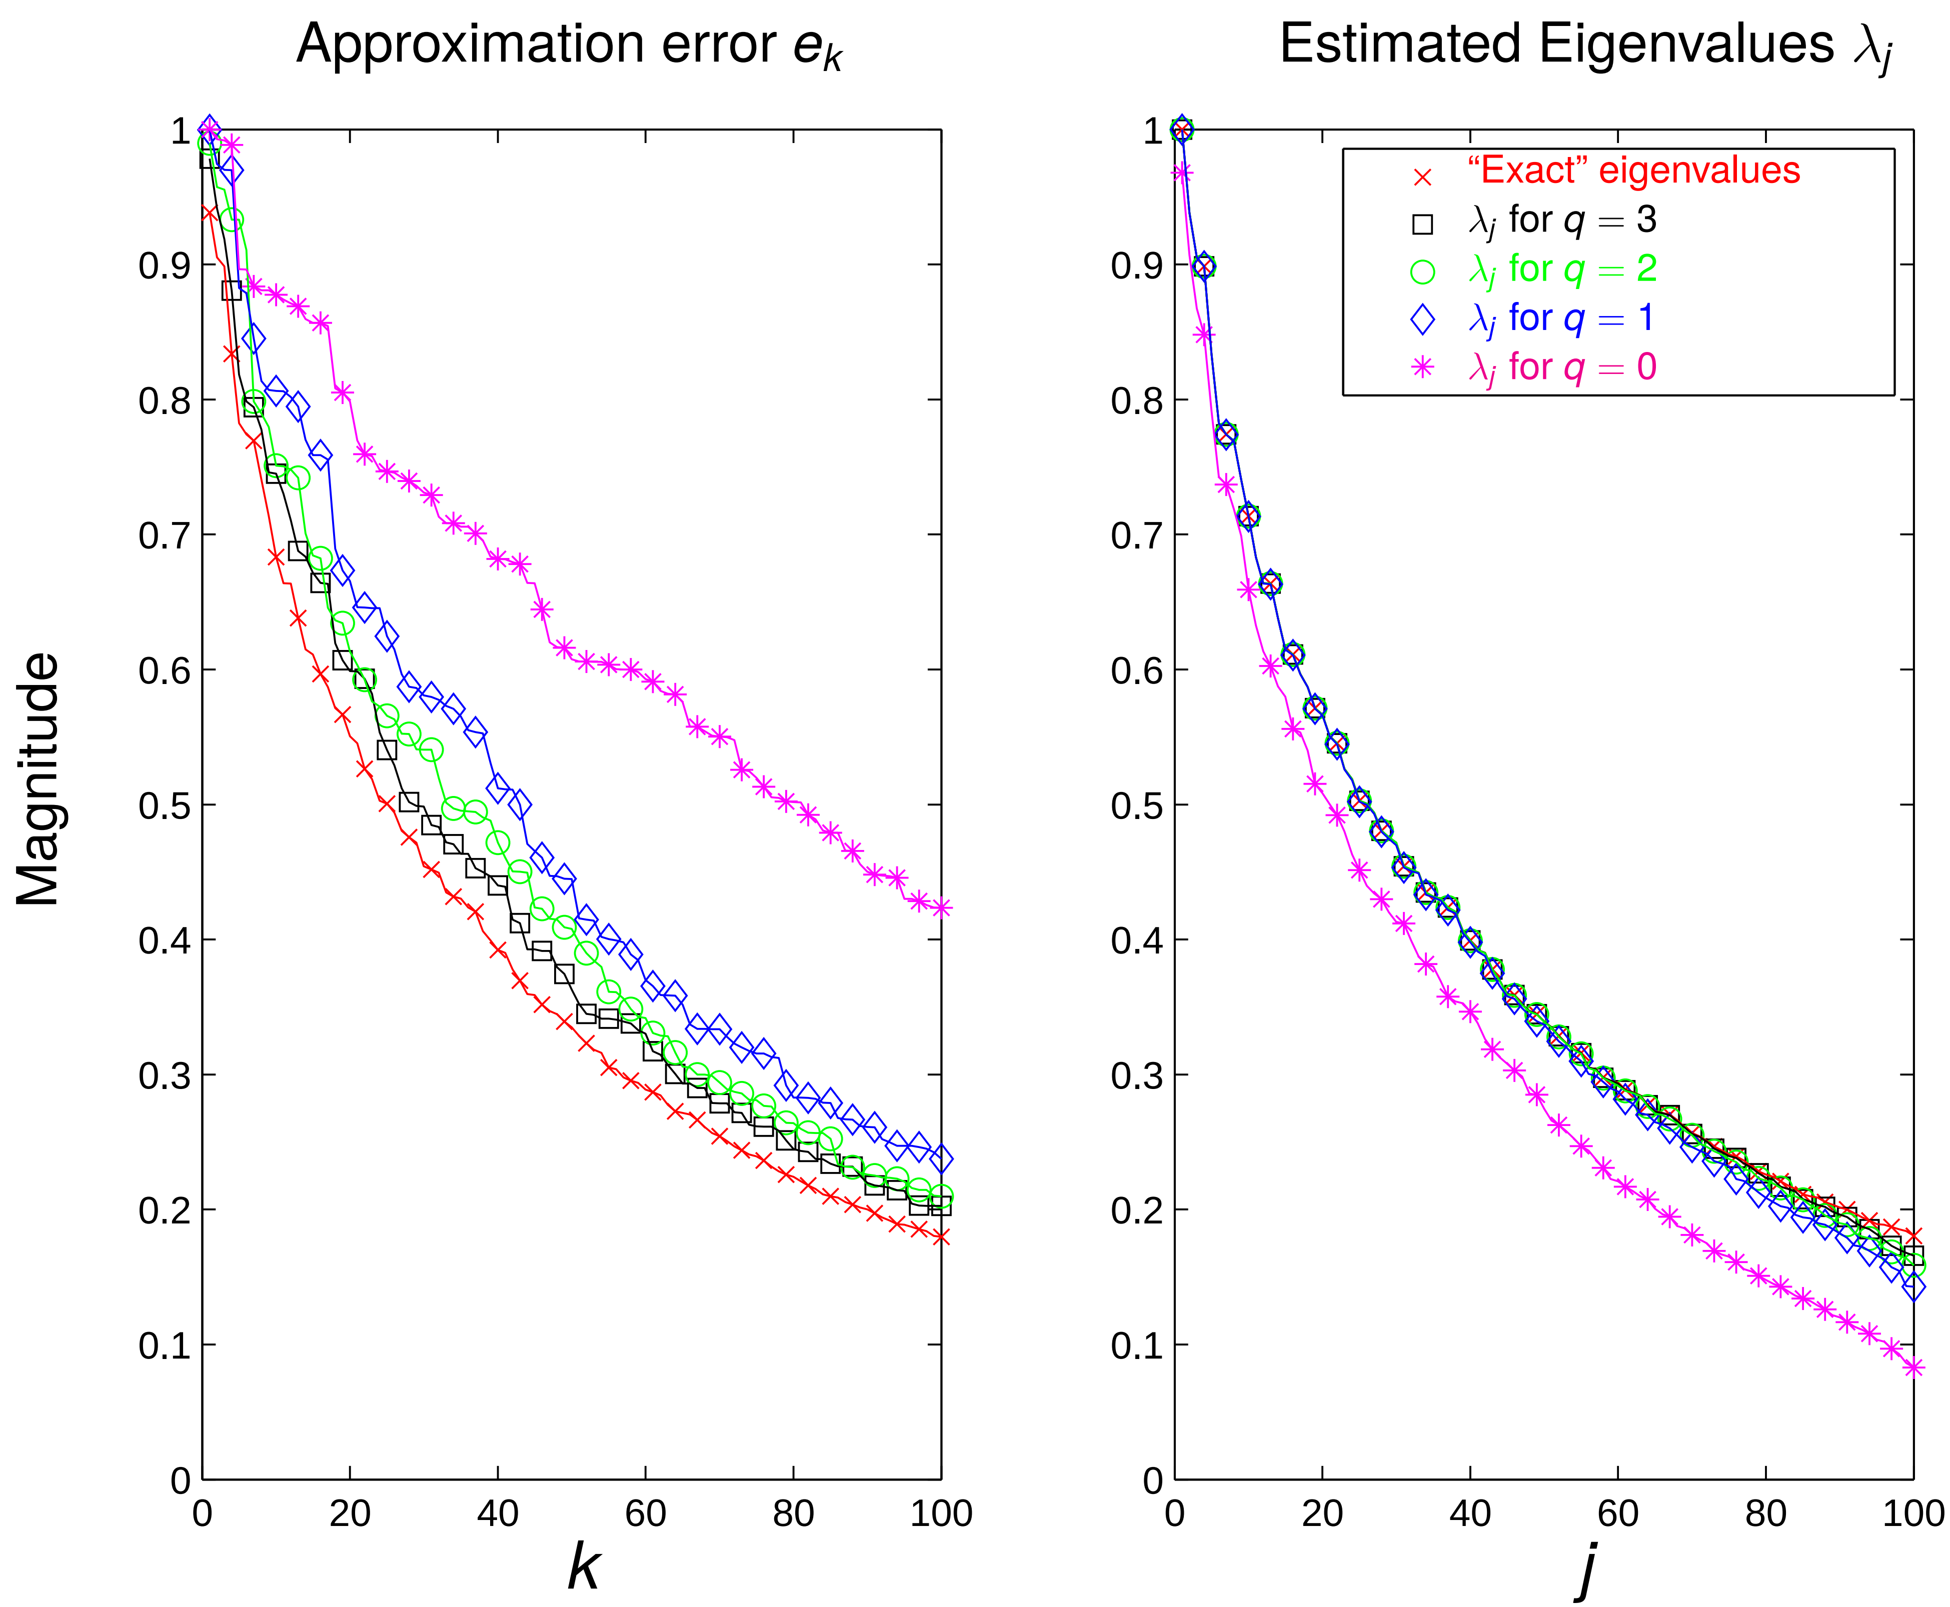

The pink lines illustrates the performance of the basic random sampling scheme.

The errors are huge, and the estimated eigenvalues are much too small.

Example 4: "Eigenfaces"

We next process a data base containing $m = 7\,254$ pictures of faces

Each image consists of $n = 384 \times 256 = 98\,304$ gray scale pixels.

We center and scale the pixels in each image, and let the resulting values form a column of a $98\,304 \times 7\,254$ data matrix $\mathbf{A}$.

The left singular vectors of $\mathbf{A}$ are the so called *eigenfaces* of the data base.

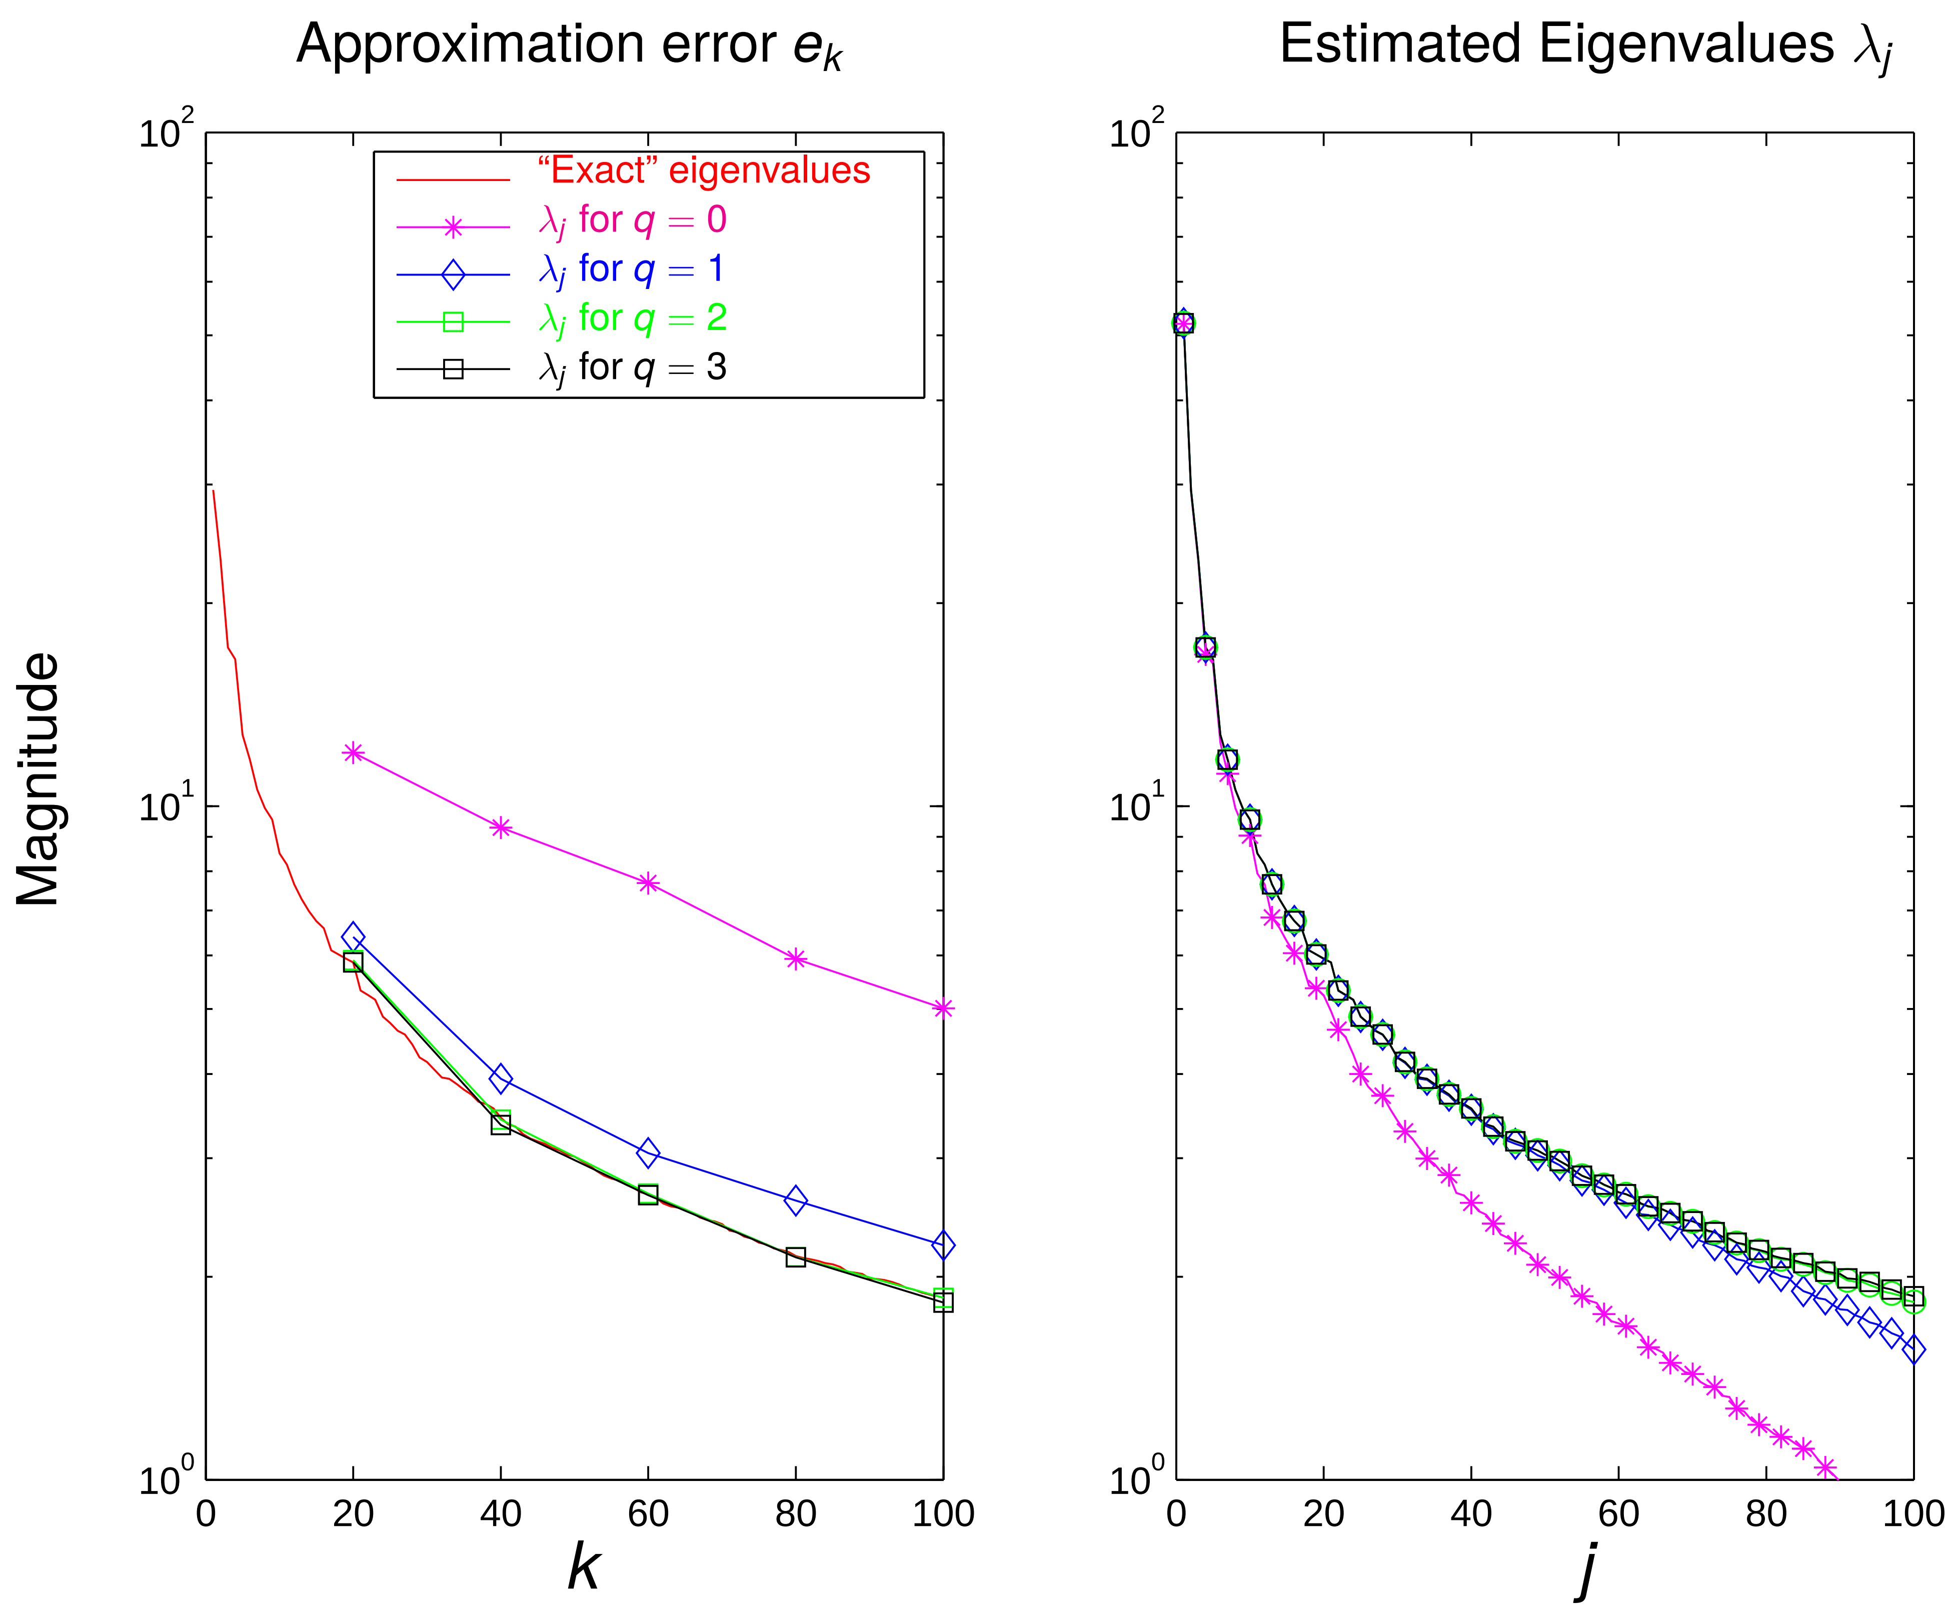

The pink lines illustrates the performance of the basic random sampling scheme.
Again, the errors are huge, and the estimated eigenvalues are much too small.

Power method for improving accuracy:

The error depends on how quickly the singular values decay. Recall that

$$\mathbb{E}\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\| \leq \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\,\sqrt{k+p}}{p}\left(\sum_{j=k+1}^{\min(m,n)} \sigma_j^2\right)^{1/2}.$$

The faster the singular values decay — the stronger the relative weight of the dominant modes in the samples.

Idea: The matrix $(\mathbf{A}\,\mathbf{A}^*)^q\,\mathbf{A}$ has the same left singular vectors as $\mathbf{A}$, and singular values

$$\sigma_j((\mathbf{A}\,\mathbf{A}^*)^q\,\mathbf{A}) = (\sigma_j(\mathbf{A}))^{2\,q+1}.$$

Much faster decay — so let us use the sample matrix

$$\mathbf{Y} = (\mathbf{A}\,\mathbf{A}^*)^q\,\mathbf{A}\,\mathbf{R}$$

instead of

$$\mathbf{Y} = \mathbf{A}\,\mathbf{R}.$$

References: Paper by Rokhlin, Szlam, Tygert (2008). Suggestions by Ming Gu. Also similar to “block power method,” “block Lanczos,” “subspace iteration.”

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank ℓ, and a small integer q.	
Output: Rank-ℓ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{UDV}^*$.	
(1) Draw an $n \times \ell$ random matrix $\mathbf{R}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $m \times \ell$ sample matrix $\mathbf{Y} = (\mathbf{A}\,\mathbf{A}^*)^q\,\mathbf{AR}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}}\mathbf{DV}^*$.
(3) Compute an ON matrix $\mathbf{Q}$ s.t. $\mathbf{Y} = \mathbf{QQ}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

- Detailed (and, we believe, close to sharp) error bounds have been proven.
  For instance, with $\mathbf{A}^{\text{computed}} = \mathbf{UDV}^*$, the expectation of the error satisfies:

$$(1) \qquad \mathbb{E}\left[\|\mathbf{A} - \mathbf{A}^{\text{computed}}\|\right] \le \left(1 + 4\sqrt{\frac{2\,\min(m,n)}{k-1}}\right)^{1/(2q+1)} \sigma_{k+1}(\mathbf{A}).$$

  Reference: Halko, Martinsson, Tropp (2011).

- The improved accuracy from the modified scheme comes at a cost;
  $2q + 1$ passes over the matrix are required instead of 1.
  However, q can often be chosen quite small in practice, $q = 2$ or $q = 3$, say.

- The bound (1) assumes exact arithmetic.
  To handle round-off errors, variations of subspace iterations can be used.
  These are entirely numerically stable and achieve the same error bound.

A numerically stable version of the "power method":

Input: An $m \times n$ matrix $\mathbf{A}$, a target rank ℓ, and a small integer q.

Output: Rank-ℓ factors $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ in an approximate SVD $\mathbf{A} \approx \mathbf{U}\mathbf{D}\mathbf{V}^*$.

Draw an $n \times \ell$ Gaussian random matrix $\mathbf{R}$.
Set $\mathbf{Q} = \texttt{orth}(\mathbf{A}\mathbf{R})$
for $i = 1, 2, \ldots, q$
  $\mathbf{W} = \texttt{orth}(\mathbf{A}^*\mathbf{Q})$
  $\mathbf{Q} = \texttt{orth}(\mathbf{A}\mathbf{W})$
end for
$\mathbf{B} = \mathbf{Q}^*\mathbf{A}$
$[\hat{\mathbf{U}}, \mathbf{D}, \mathbf{V}] = \texttt{svd}(\mathbf{B})$
$\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Note: Algebraically, the method with orthogonalizations is identical to the "original" method where $\mathbf{Q} = \texttt{orth}((\mathbf{A}\mathbf{A}^*)^q\mathbf{A}\mathbf{R})$.

Note: This is a classic subspace iteration.

The novelty is the error analysis, and the finding that using a very small q is often fine.

(In fact, our analysis allows q to be zero...)

Adaptive rank determination

How to proceed when the rank of a matrix is not known in advance.

Adaptive rank determination — vector-by-vector technique

Let us again start by considering the simplistic case where $\mathbf{A}$ is *exactly* rank-deficient.

Let $\mathbf{A}$ be an $m \times n$ matrix of exact rank k, where k is unknown.

We seek an $m \times k$ matrix $\mathbf{Q}$ whose columns form an ON basis for col($\mathbf{A}$).

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots, ???$
>
> Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
> Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
> Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
> Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
> **end for**

Adaptive rank determination — vector-by-vector technique

Let us again start by considering the simplistic case where $\mathbf{A}$ is *exactly* rank-deficient.

Let $\mathbf{A}$ be an $m \times n$ matrix of exact rank k, where k is unknown.

We seek an $m \times k$ matrix $\mathbf{Q}$ whose columns form an ON basis for col($\mathbf{A}$).

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \ldots, ???$
>
> Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
> Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
> Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
> Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
> **end for**

Observation 1: While $i \leq k$, we know that $\mathbf{z}_i \neq \mathbf{0}$ with probability 1.

Adaptive rank determination — vector-by-vector technique

Let us again start by considering the simplistic case where $\mathbf{A}$ is *exactly* rank-deficient.

Let $\mathbf{A}$ be an $m \times n$ matrix of exact rank k, where k is unknown.

We seek an $m \times k$ matrix $\mathbf{Q}$ whose columns form an ON basis for col($\mathbf{A}$).

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots, ???$
>
> Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
> Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
> Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
> Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
> **end for**

Observation 1: While $i \leq k$, we know that $\mathbf{z}_i \neq \mathbf{0}$ with probability 1.

Observation 2: Once you come to step $i = k + 1$, the vector $\mathbf{z}_{k+1}$ must be zero!

Adaptive rank determination — vector-by-vector technique

Let us again start by considering the simplistic case where **A** is *exactly* rank-deficient.

Let **A** be an $m \times n$ matrix of exact rank k, where k is unknown.

We seek an $m \times k$ matrix **Q** whose columns form an ON basis for col(**A**).

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots$
>
>   Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
>   Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
>   Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
>   **if** $\left[\mathbf{z}_i = \mathbf{0}\right]$ **then**
>
>   The rank is $k = i - 1$.
>
>   **break**
>
>   **else**
>
>   Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
>   **end if**
>
> **end for**

Adaptive rank determination — vector-by-vector technique

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance. We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots, ???$
>
>   Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
>   Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
>   Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
>   Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
> **end for**

Adaptive rank determination — vector-by-vector technique

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance. We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

$\mathbf{Q} = [\,]$;

for $i = 1, 2, 3, \ldots, ???$

Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.

Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.

Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.

Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.

end for

Observe that

$$\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i = \mathbf{A}\mathbf{r}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\mathbf{r}_i = \big(\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\big)\mathbf{r}_i.$$

Adaptive rank determination — vector-by-vector technique

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance. We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^{*}\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots, ???$
>
> Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
> Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
> Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^{*}\mathbf{y}_i$.
>
> Add the new element to the basis $\mathbf{Q} = \left[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\right]$.
>
> **end for**

Observe that

$$\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^{*}\mathbf{y}_i = \mathbf{A}\mathbf{r}_i - \mathbf{Q}\mathbf{Q}^{*}\mathbf{A}\mathbf{r}_i = \big(\mathbf{A} - \mathbf{Q}\mathbf{Q}^{*}\mathbf{A}\big)\mathbf{r}_i.$$

In consequence, since $\mathbf{r}_i$ is Gaussian,

$$\mathbb{E}\big[\|\mathbf{z}_i\|^2\big] = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^{*}\mathbf{A}\|_{\mathrm{Fro}}^2.$$

Adaptive rank determination — vector-by-vector technique

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance. We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\rm Fro} \leq \varepsilon$.

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1,\,2,\,3,\,\dots,\,???$
>
> Draw an $n \times 1$ Gaussian random vector $\mathbf{r}_i$.
>
> Compute an $m \times 1$ sample vector $\mathbf{y}_i = \mathbf{A}\mathbf{r}_i$.
>
> Project the sample vector away from the basis computed $\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i$.
>
> Add the new element to the basis $\mathbf{Q} = \big[\mathbf{Q}\ \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}\big]$.
>
> **end for**

Observe that

$$\mathbf{z}_i = \mathbf{y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{y}_i = \mathbf{A}\mathbf{r}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\mathbf{r}_i = \big(\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\big)\mathbf{r}_i.$$

In consequence, since $\mathbf{r}_i$ is Gaussian,

$$\mathbb{E}\big[\|\mathbf{z}_i\|^2\big] = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\rm Fro}^2.$$

Observation 1: Once you observe several consecutive $\mathbf{z}_i$ such that, say, $\|\mathbf{z}_i\| \leq \varepsilon/2$, it will "likely" be the case that $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}\mathbf{A}\|_{\rm Fro} \leq \varepsilon$.

Observation 2: You need to block the algorithm for computational efficiency.

Adaptive rank determination

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance, and let b be a "block size."

We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

> $\mathbf{Q} = [\,]$;
>
> **for** $i = 1, 2, 3, \dots$
>
> Draw an $n \times b$ Gaussian random matrix $\mathbf{R}_i$.
>
> Compute an $m \times b$ sample matrix $\mathbf{Y}_i = \mathbf{A}\mathbf{R}_i$.
>
> Project the sample columns away from the basis computed $\mathbf{Z}_i = \mathbf{Y}_i - \mathbf{Q}\mathbf{Q}^*\mathbf{Y}_i$.
>
> Orthonormalize the samples $[\mathbf{Q}_i,\ \mathbf{R}_i] = \mathtt{qr}(\mathbf{Z}_i, 0)$. (*Unpivoted* QR factorization!)
>
> **if** $\big[$"several consecutive columns of $\mathbf{R}_i$ are small"$\big]$ **then**
>
> Add the appropriate number of columns of $\mathbf{Q}_i$ to $\mathbf{Q}$.
>
> **break**
>
> **else**
>
> Add the new element to the basis $\mathbf{Q} = \big[\mathbf{Q}\ \mathbf{Q}_i\big]$.
>
> **end if**
>
> **end for**

Warning: Re-orthogonalization is often needed to combat floating point errors.

Adaptive rank determination — with updating

Consider the special case that **A** *can be updated, e.g. if it is dense and stored in RAM.*

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance, and let b be a "block size."

We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Fro}} \le \varepsilon$.

(1) $\mathbf{Q} = [\,]$; $\mathbf{B} = [\,]$;
(2) **while** $\|\mathbf{A}\| > \varepsilon$
(3) Draw an $n \times b$ Gaussian matrix $\mathbf{R}_i$.
(4) Compute the $m \times b$ matrix $[\mathbf{Q}_i, \sim] = \mathtt{qr}(\mathbf{A}\mathbf{R}_i, 0)$.
(5) $\mathbf{B}_i = \mathbf{Q}_i^*\mathbf{A}$
(6) $\mathbf{Q} = [\mathbf{Q}\ \mathbf{Q}_i]$
(7) $\mathbf{B} = \begin{bmatrix} \mathbf{B} \\ \mathbf{B}_i \end{bmatrix}$
(8) $\mathbf{A} = \mathbf{A} - \mathbf{Q}_i\mathbf{B}_i$
(9) **end while**

A blocked and randomized variation of the classical "modified Gram-Schmidt" algorithm.

Warning: Re-orthogonalization is often needed to combat floating point errors.

Adaptive rank determination — with updating

*Consider the special case that **A** can be updated, e.g. if it is dense and stored in RAM.*

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance, and let b be a "block size."

We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

(1) $\mathbf{Q} = [\ ]$; $\mathbf{B} = [\ ]$;
(2) **while** $\|\mathbf{A}\| > \varepsilon$
(3) Draw an $n \times b$ Gaussian matrix $\mathbf{R}_i$.
(4) Compute the $m \times b$ matrix $\mathbf{Q}_i = \mathtt{qr}(\mathbf{A}\mathbf{R}_i, 0)$.
(5) $\mathbf{B}_i = \mathbf{Q}_i^*\mathbf{A}$
(6) $\mathbf{Q} = [\mathbf{Q}\ \mathbf{Q}_i]$
(7) $\mathbf{B} = \begin{bmatrix} \mathbf{B} \\ \mathbf{B}_i \end{bmatrix}$
(8) $\mathbf{A} = \mathbf{A} - \mathbf{Q}_i\mathbf{B}_i$
(9) **end while**

Observation: Almost all the work is done by matrix-matrix multiplies.

Adaptive rank determination — with updating

Consider the special case that **A** *can be updated, e.g. if it is dense and stored in RAM.*

Let $\mathbf{A}$ be an $m \times n$ matrix whose singular values decay, but do not hit zero.

Let $\varepsilon > 0$ be a given tolerance, and let b be a "block size."

We seek an $m \times k$ ON matrix $\mathbf{Q}$ s.t. $\|\mathbf{A} - \mathbf{QQ}^*\mathbf{A}\|_{\mathrm{Fro}} \leq \varepsilon$.

(1) $\mathbf{Q} = [\,]$; $\mathbf{B} = [\,]$;
(2) **while** $\|\mathbf{A}\| > \varepsilon$
(3) Draw an $n \times b$ Gaussian matrix $\mathbf{R}_i$.
(4) Compute the $m \times b$ matrix $\mathbf{Q}_i = \mathtt{qr}(\mathbf{AR}_i, 0)$.
(5) $\mathbf{B}_i = \mathbf{Q}_i^*\mathbf{A}$
(6) $\mathbf{Q} = [\mathbf{Q}\ \mathbf{Q}_i]$
(7) $\mathbf{B} = \begin{bmatrix} \mathbf{B} \\ \mathbf{B}_i \end{bmatrix}$
(8) $\mathbf{A} = \mathbf{A} - \mathbf{Q}_i\mathbf{B}_i$
(9) **end while**

Observation: Almost all the work is done by matrix-matrix multiplies.

This algorithm is ideal for running on modern CPUs and GPUs!

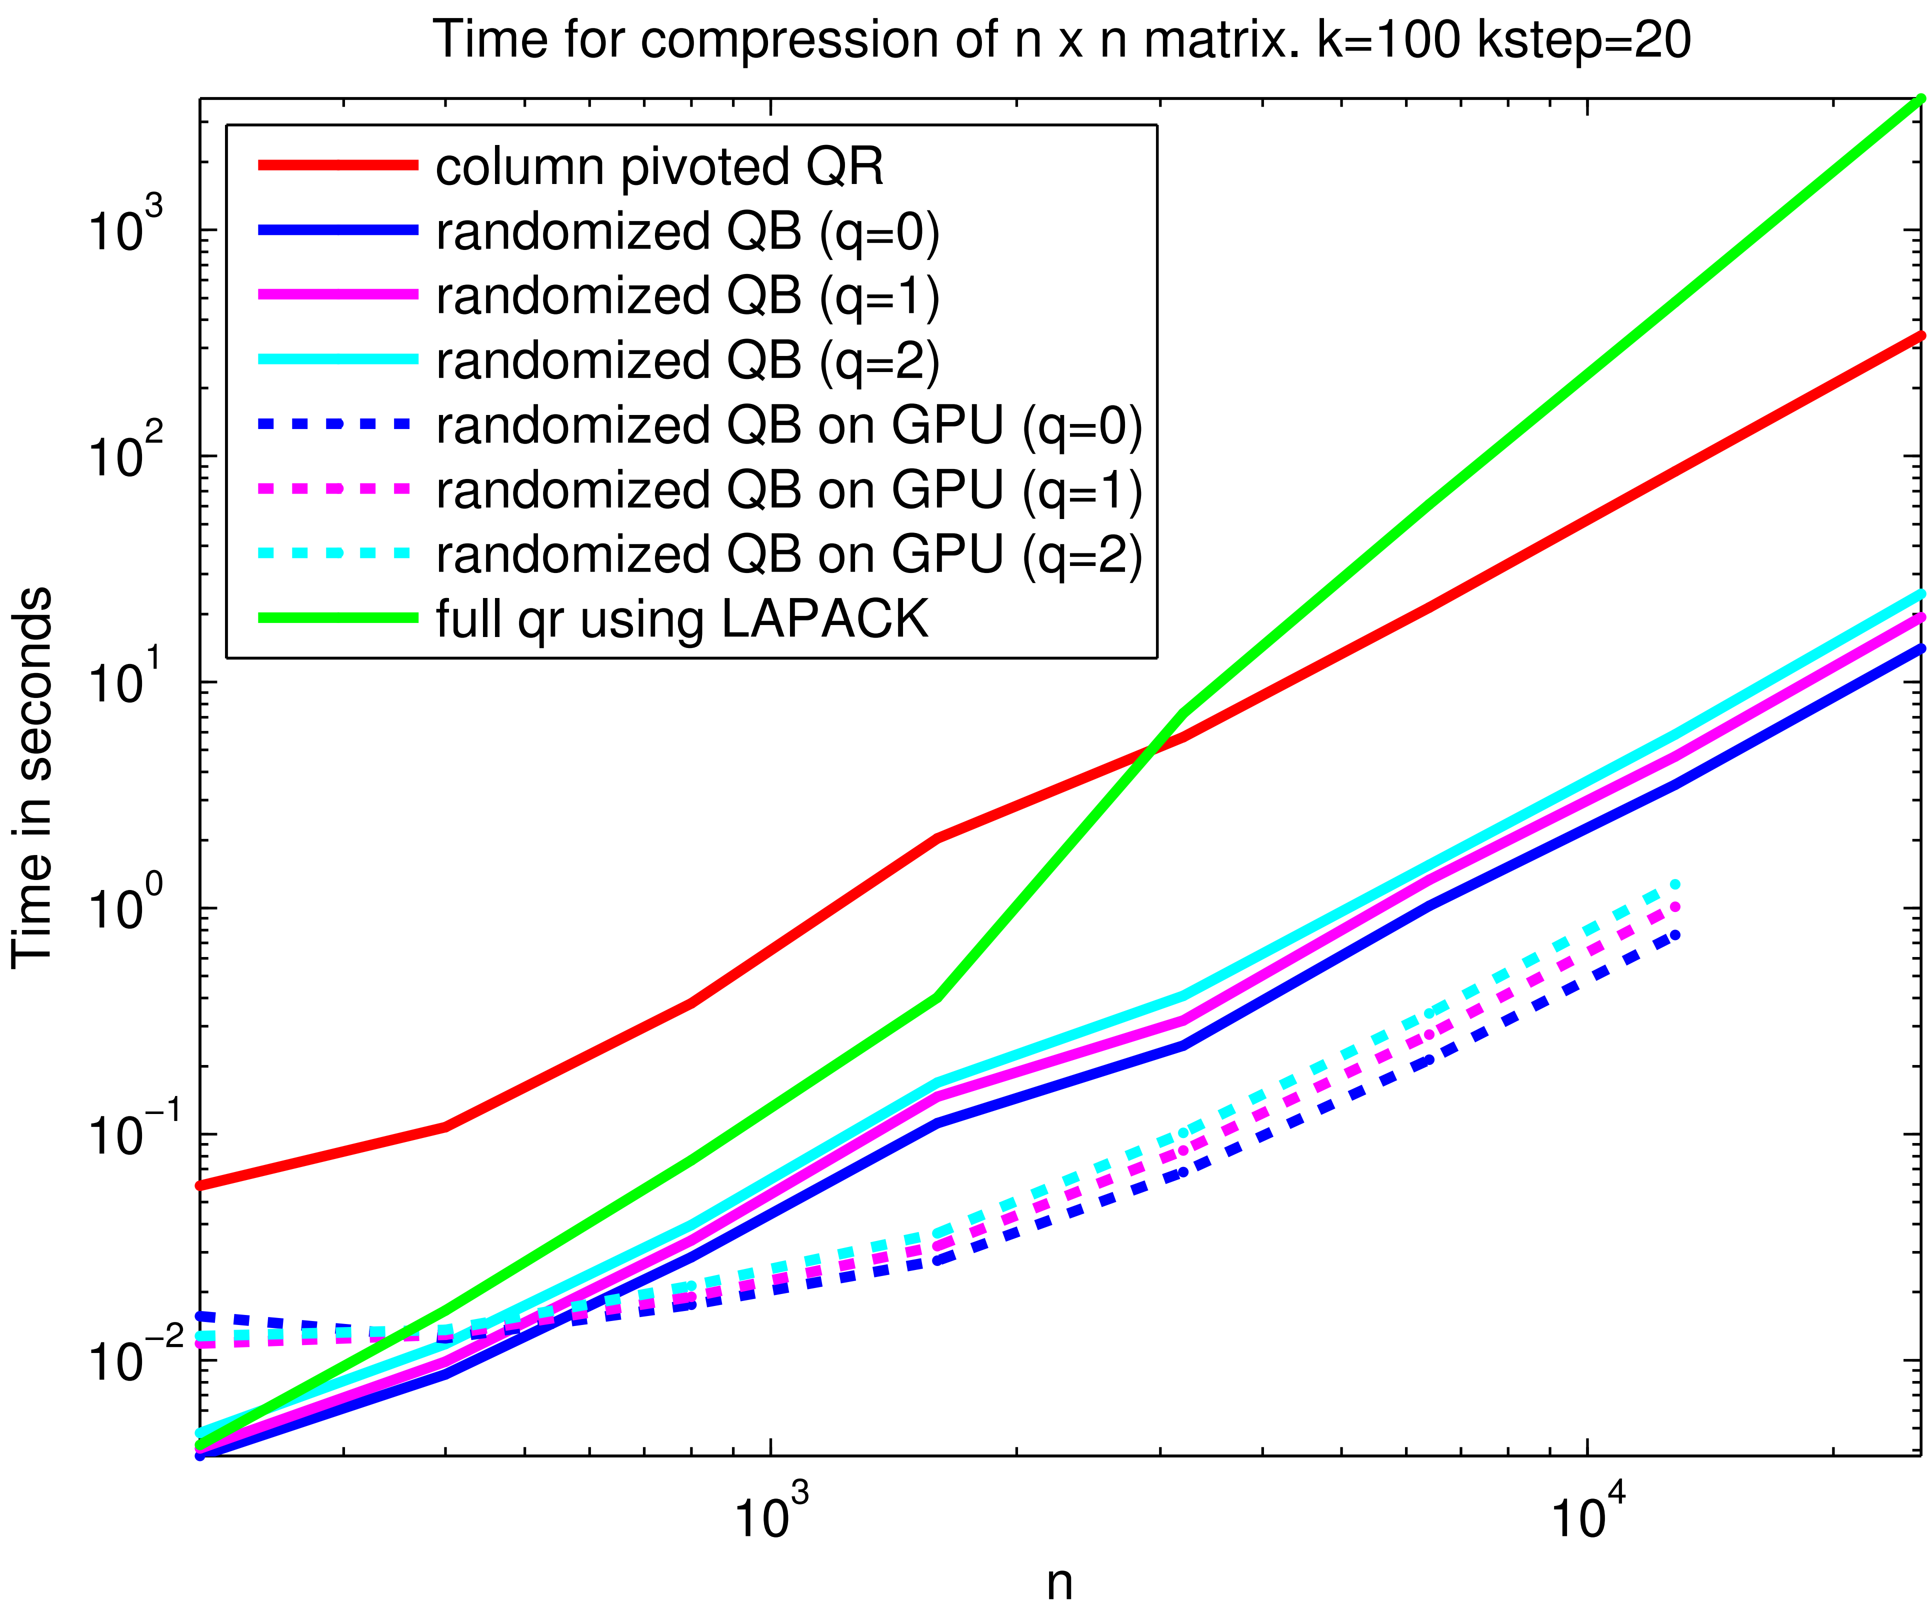

Everything is implemented in Matlab. The “full qr” line refers to Matlab built in qr.

CPU = Intel Xeon E-1660 (6 cores, 3.3GHz). GPU = Tesla K40c (2880 cores, 12GB).

Caveat: Matlab overhead makes column-pivoted QR slower than it could be.

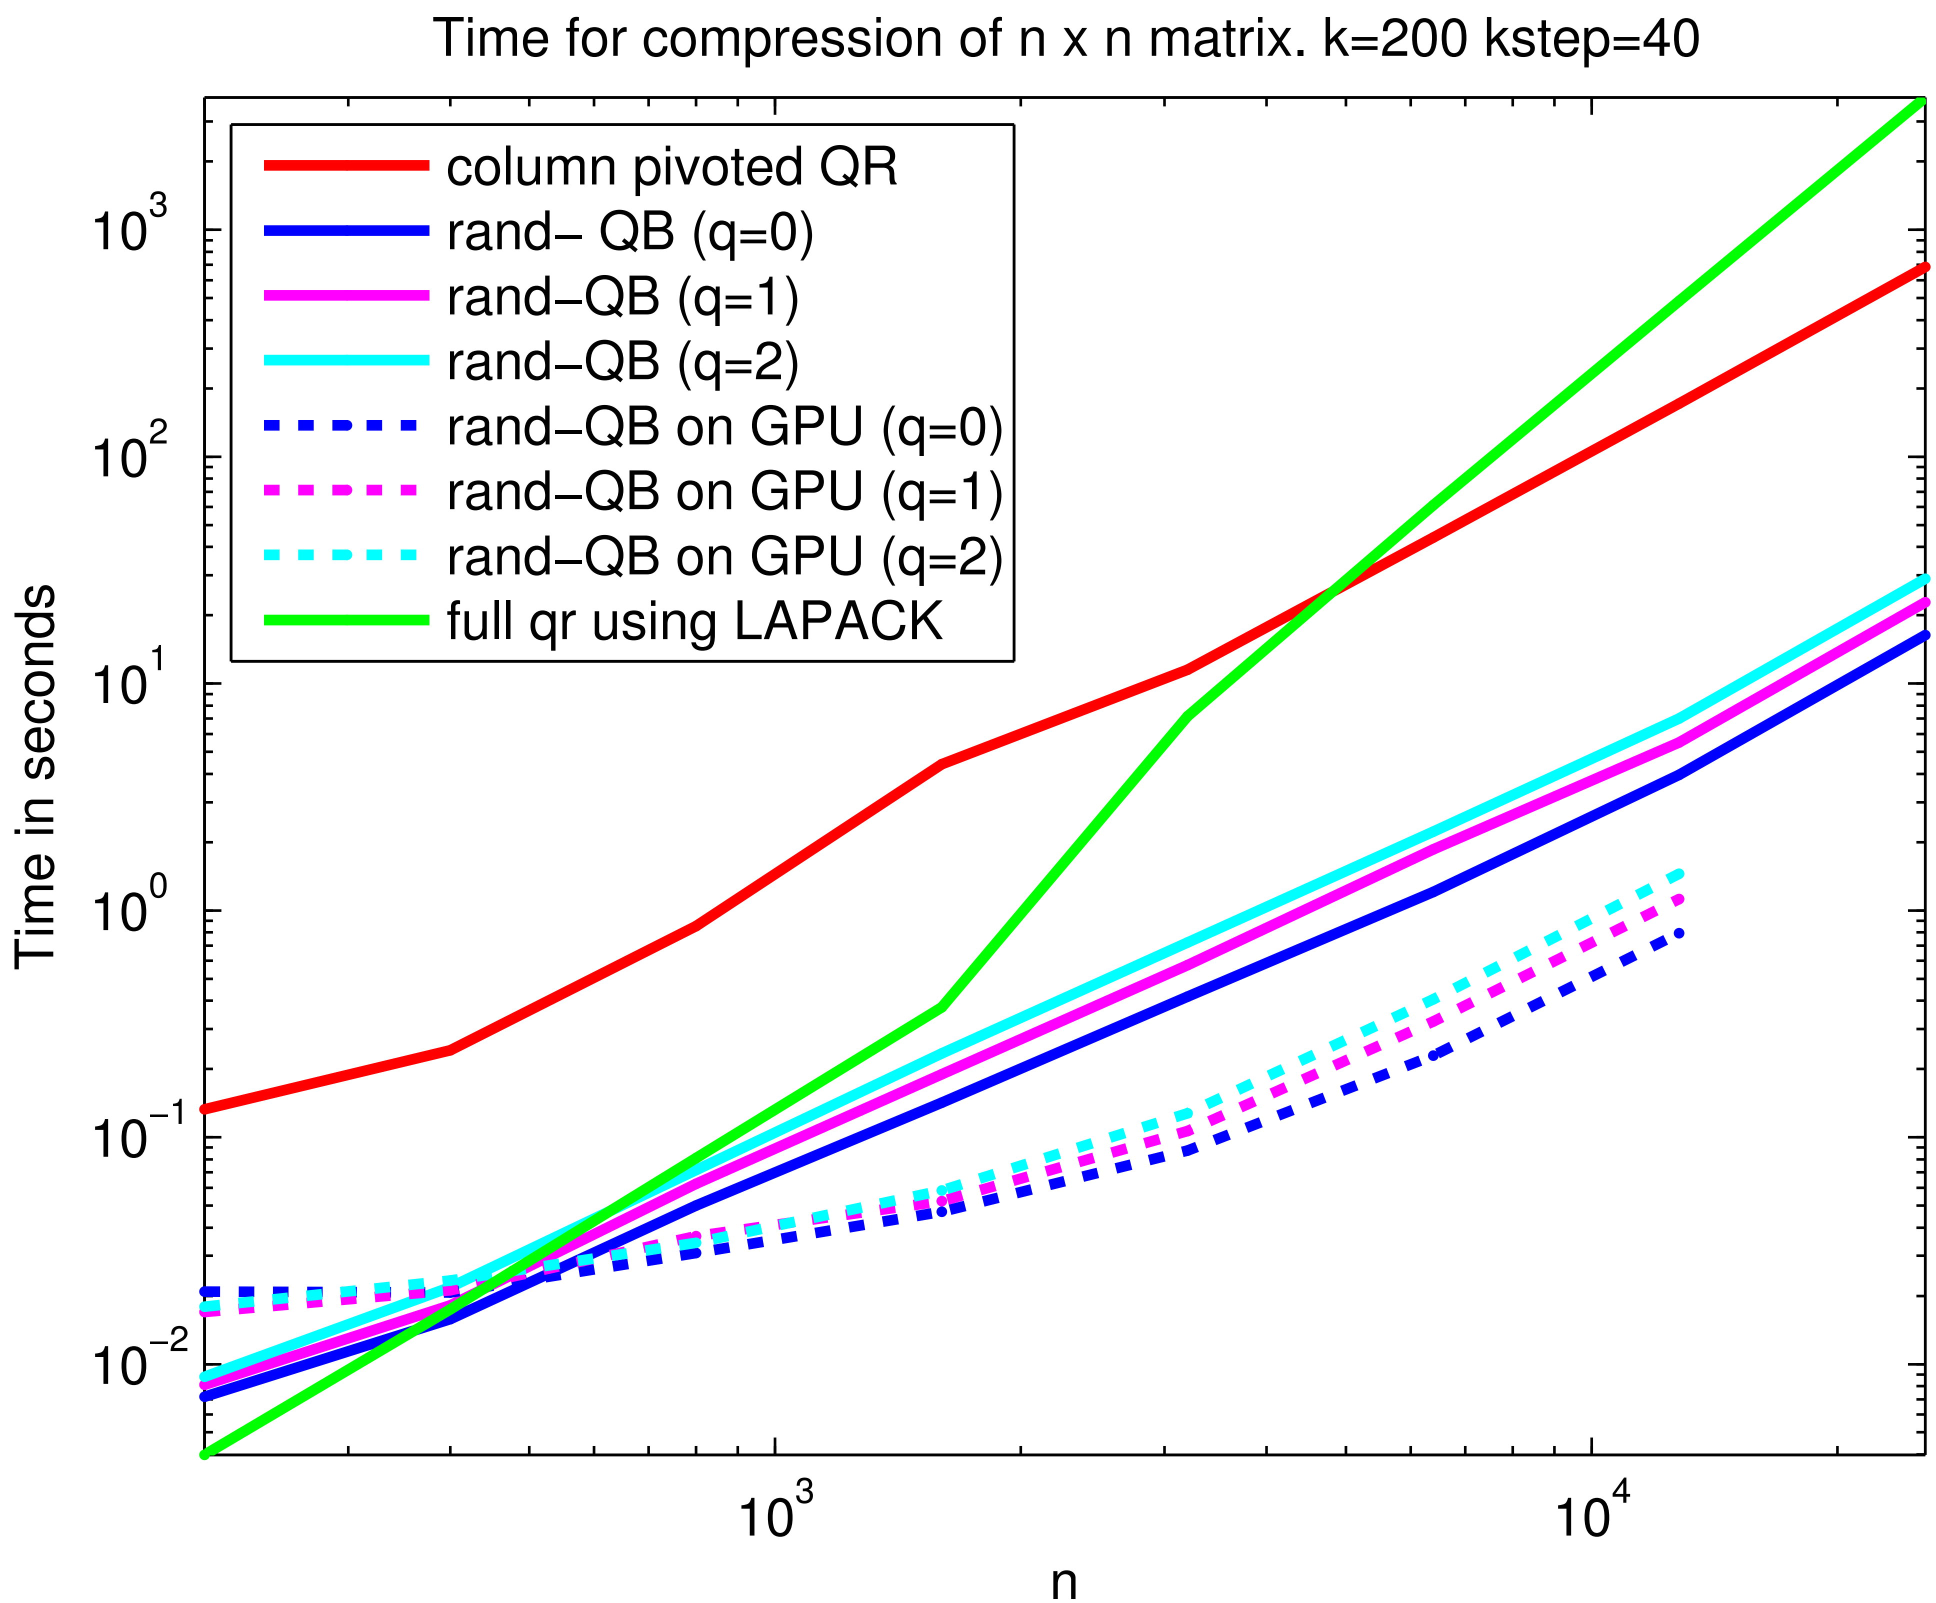

Everything is implemented in Matlab. The “full qr” line refers to Matlab built in qr.

CPU = Intel Xeon E-1660 (6 cores, 3.3GHz). GPU = Tesla K40c (2880 cores, 12GB).

Caveat: Matlab overhead makes column-pivoted QR slower than it could be.

Randomized algorithms for FULL factorizations: Column pivoted QR

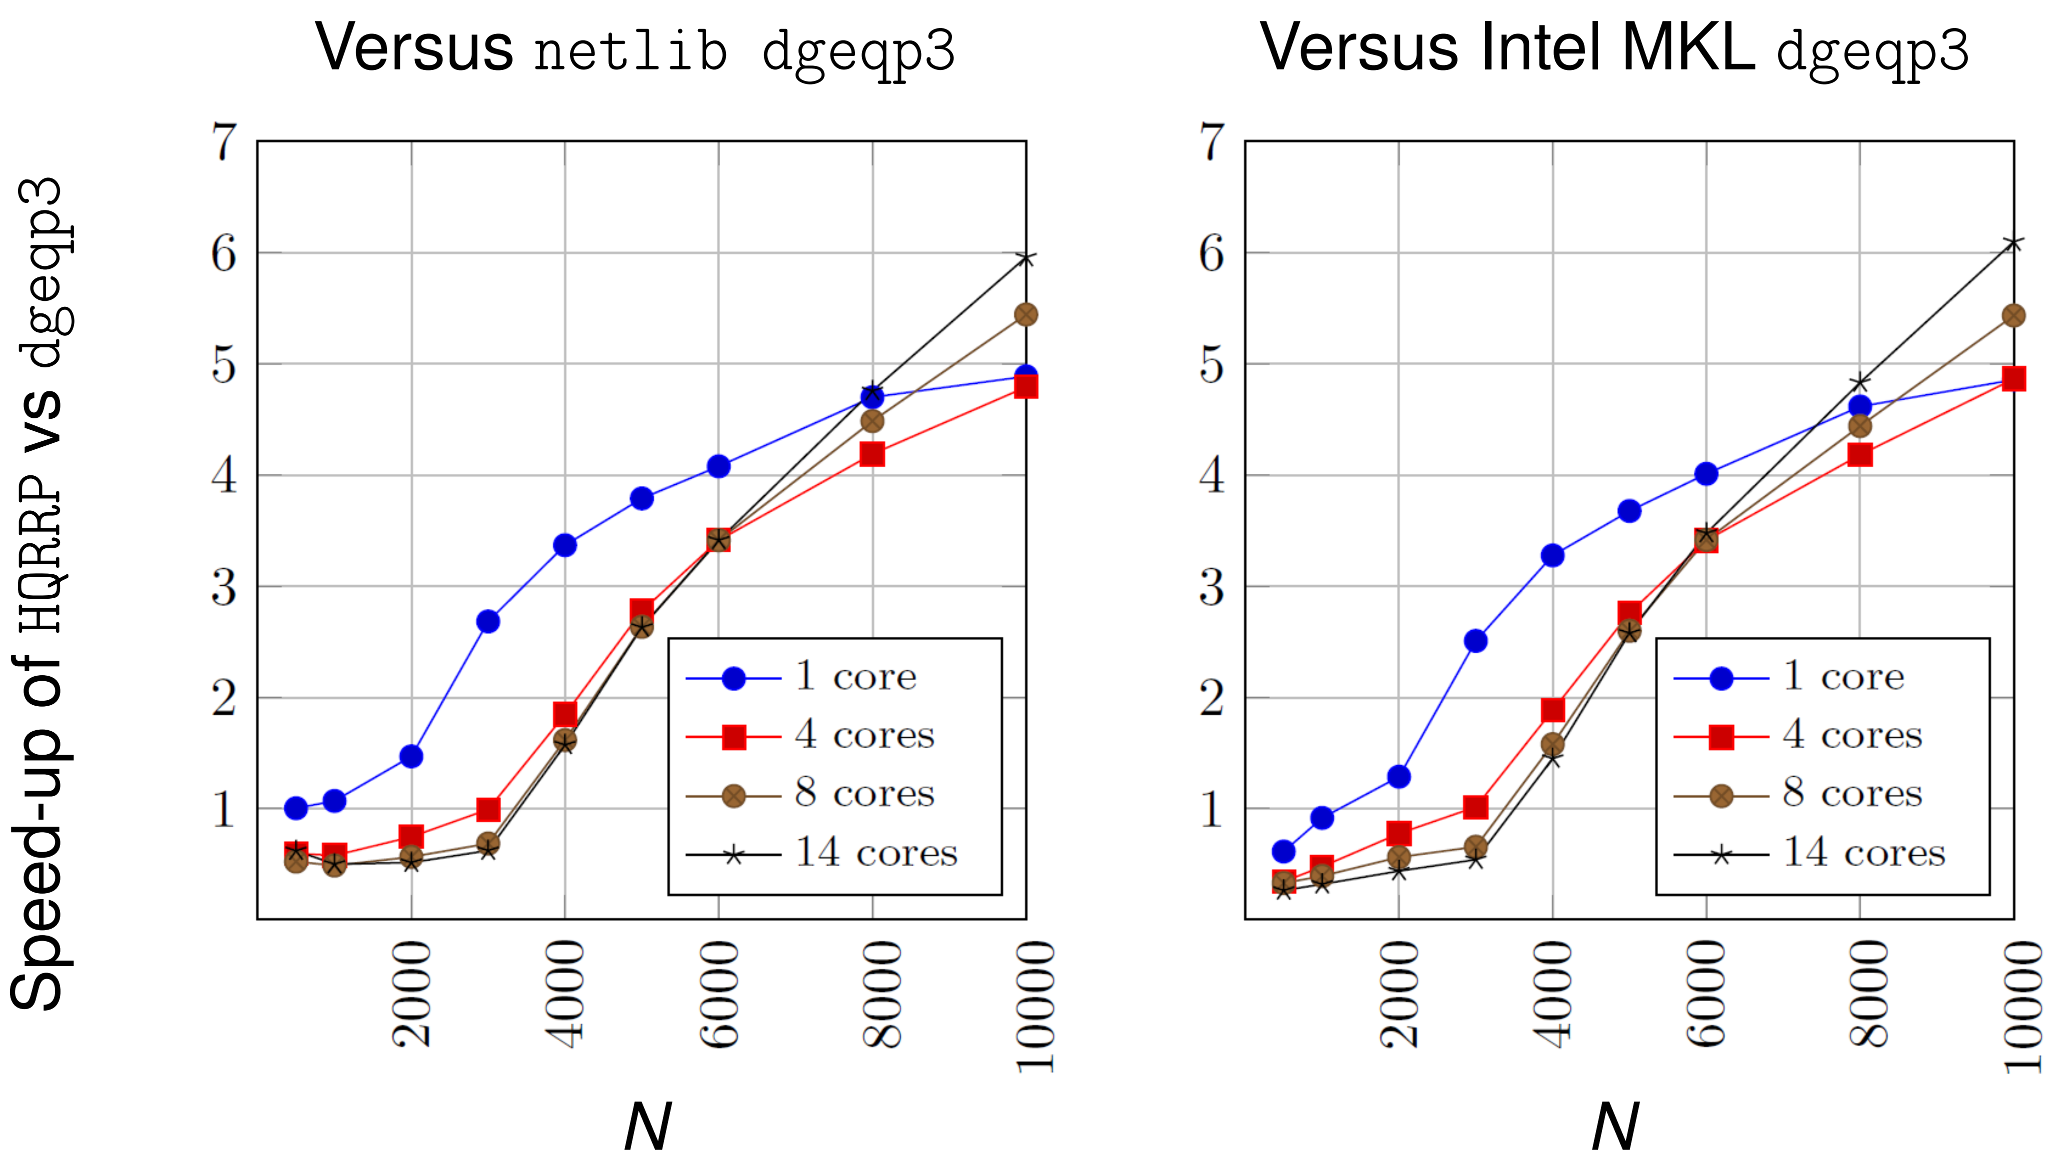

Speedup attained by our randomized algorithm `HQRRP` *for computing a full column pivoted QR factorization of an* $N \times N$ *matrix. The speed-up is measured versus LAPACK's faster routine* `dgeqp3` *as implemented in Netlib (left) and Intel's MKL (right). Our implementation was done in C, and was executed on an Intel Xeon E5-2695. Joint work with G. Quintana-Ortí, N. Heavner, and R. van de Geijn. Available at:* https://github.com/flame/hqrrp/