hence this A has rank 1 (since we can write it as $A = V\Lambda^t$ with $\operatorname{dom} V = \mathbb{F}^1$, but we couldn't do it with $\operatorname{dom} V = \mathbb{F}^0$). To calculate Ax, we merely need to calculate the number $\alpha := (1,2,3,4,5)^t x$, and then obtain Ax as the particular scalar multiple $y\alpha$ of the vector $y := (1,1,1,1)$. That is much cheaper than computing the matrix product of the 4×5-matrix A with the 1-column matrix $[x]$.

As the example illustrates, any matrix

$$A := [v][w]^t = vw^t$$

with $v \in \mathbb{F}^m$ and $w \in \mathbb{F}^n$ has rank 1 unless it is trivial, i.e., unless either v or w is the zero vector. This explains why an *elementary* matrix is also called a **rank-one perturbation of the identity**.

The only linear map of rank 0 is the zero map. If A is not the zero map, then its range contains some nonzero vector, hence so must the range of any V for which $A = V\Lambda^t$ with $\operatorname{dom} V = \mathbb{F}^r$, therefore such r must be > 0.

As another *example*, for any vector space X,

$$\dim X = \operatorname{rank} \operatorname{id}_X.$$

Indeed, if $n = \dim X$, then, for any basis $V \in L(\mathbb{F}^n, X)$ for X, $\operatorname{id}_X = VV^{-1}$, therefore $\operatorname{rank} \operatorname{id}_X \le n$, while, for any factorization $\operatorname{id}_X = V\Lambda^t$ for some $V \in L(\mathbb{F}^r, X)$, V must necessarily be onto, hence $\dim X \le r$, by (4.6)Proposition, and therefore $\dim X \le \operatorname{rank} \operatorname{id}_X$. In fact, it is possible to make the rank concept the primary one and *define* $\dim X$ as the rank of id_X.

When A is an $m \times n$-matrix, then, trivially, $A = A \operatorname{id}_n = \operatorname{id}_m A$, hence $\operatorname{rank} A \le \min\{m, n\}$.

At times, particularly when A is a matrix, it is convenient to write the factorization $A = V\Lambda^t$ more explicitly as

$$A =: [v_1, v_2, \ldots, v_r][\lambda_1, \lambda_2, \ldots, \lambda_r]^t = \sum_{j=1}^{r} [v_j]\lambda_j. \tag{8.1}$$

Since each of the maps

$$v_j\lambda_j := [v_j]\lambda_j = [v_j] \circ \lambda_j : x \mapsto (\lambda_j x)v_j$$

has rank ≤ 1, this shows that *the rank of A gives the smallest number of terms necessary to write A as a sum of rank-one maps.*

(8.2) Proposition: $A = V\Lambda^t$ is minimal if and only if V is a basis for $\operatorname{ran} A$. In particular,

$$\operatorname{rank} A = \dim \operatorname{ran} A.$$

Proof: Let $A = V\Lambda^t$. Then $\operatorname{ran} A \subset \operatorname{ran} V$, hence

$$\dim \operatorname{ran} A \le \dim \operatorname{ran} V \le \#V,$$

with equality in the first $\le$ iff $\operatorname{ran} A = \operatorname{ran} V$ (by (4.13)Proposition), and in the second $\le$ iff V is 1-1. Thus, $\dim \operatorname{ran} A \le \#V$, with equality iff V is a basis for $\operatorname{ran} A$. □

One can prove in a similar way that $A = V\Lambda^t$ is minimal if and only if Λ^t is onto and $\operatorname{null} A = \operatorname{null} \Lambda^t$.

(8.3) Corollary: The factorization $A = A(:, \texttt{bound})\operatorname{rrref}(A)$ provided by elimination (see (3.11)) is minimal.

(8.4) Corollary: If $A = V\Lambda^t$ is minimal and A is invertible, then also V and Λ^t are invertible.

Proof: By (8.2)Proposition, $V \in L(\mathbb{F}^r, Y)$ is a basis for $\operatorname{ran} A$, while $\operatorname{ran} A = Y$ since A is invertible. Hence, V is invertible. Therefore, also $\Lambda^t = V^{-1}A$ is invertible. □

But note that the matrix $[\,1\,] = [\,1 \quad 0\,] \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ is invertible, even though neither of its two factors is.

8.1 Determine a minimal factorization for the matrix

$$\begin{bmatrix} 1 & 2 & 0 & 3 & 4 \\ 4 & 3 & 0 & 2 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 8 & 7 & 0 & 6 & 5 \end{bmatrix}$$

8.2 With A the matrix of the previous problem, give a basis for $\operatorname{ran} A$ and a basis for $\operatorname{ran} A^{\mathrm t}$.

8.3 Give an example of a pair of matrices, A and B, of order 4, each of rank 2, yet $\operatorname{ran} A \cap \operatorname{ran} B = \{0\}$.

8.4 Prove: *For any two linear maps A and B for which AB is defined,* $\operatorname{rank}(AB) \le \min\{\operatorname{rank} A, \operatorname{rank} B\}$. (Hint: If $A = V_A \Lambda_A{}^{\mathrm t}$ and $B = V_B \Lambda_B{}^{\mathrm t}$, then $AB = V_A(\Lambda_A{}^{\mathrm t} V_B \Lambda_B{}^{\mathrm t}) = (V_A \Lambda_A{}^{\mathrm t} V_B)\Lambda_B{}^{\mathrm t}$. Totally different hint: Use the Dimension Formula together with the fact that $\operatorname{rank} C = \dim \operatorname{ran} C$.)

The trace of a linear map

Each $A \in L(X)$ can be factored in possibly many different ways as

$$A = V\Lambda^{\mathrm t} = [v_1, \ldots, v_n][\lambda_1, \ldots, \lambda_n]^{\mathrm t}$$

for some n (necessarily $\ge \operatorname{rank} A$). It may therefore be surprising that, nevertheless, the number

$$\sum_j \lambda_j v_j$$

only depends on A. For the proof of this claim, we notice that

$$\sum_j \lambda_j v_j = \operatorname{trace}(\Lambda^{\mathrm t} V).$$

Now, let W be a basis for X, with dual basis $\mathrm{M} := W^{-1}$. Then

$$\hat{A} := \mathrm{M}^{\mathrm t} A W = \mathrm{M}^{\mathrm t} V \Lambda^{\mathrm t} W,$$

while

$$\Lambda^{\mathrm t} W \mathrm{M}^{\mathrm t} V = \Lambda^{\mathrm t} V.$$

Hence, by (6.27),

$$\operatorname{trace}(\hat{A}) = \operatorname{trace}(\mathrm{M}^{\mathrm t} V \Lambda^{\mathrm t} W) = \operatorname{trace}(\Lambda^{\mathrm t} W \mathrm{M}^{\mathrm t} V) = \operatorname{trace}(\Lambda^{\mathrm t} V).$$

By holding our factorization $A = V\Lambda^{\mathrm t}$ fixed, this implies that $\operatorname{trace}(\hat{A})$ does not depend on the particular basis W for X we happen to use here, hence only depends on the linear map A. With that, holding now this linear map A fixed, we see that also $\operatorname{trace}(\Lambda^{\mathrm t} V)$ does not depend on the particular factorization $A = V\Lambda^{\mathrm t}$ we picked, but only depends on A. This number is called the trace of A, written

$$\operatorname{trace}(A).$$

The problems provide the basic properties of the trace of a linear map.

8.5 $\operatorname{trace}(\operatorname{id}_X) = \dim X$.

8.6 If $P \in L(X)$ is a projector (i.e., $P^2 = P$), then $\operatorname{trace}(P) = \dim \operatorname{ran} P$.

8.7 $A \mapsto \operatorname{trace}(A)$ is the unique scalar-valued linear map on $L(X)$ for which $\operatorname{trace}([x]\lambda) = \lambda x$ for all $x \in X$ and $\lambda \in X'$.

8.8 If $A \in L(X, Y)$ and $B \in L(Y, X)$, then (both AB and BA are defined) $\operatorname{trace}(AB) = \operatorname{trace}(BA)$.

8.9 Prove that, for column maps V, W into X, and row maps $\Lambda^{\mathrm t}$, $\mathrm{M}^{\mathrm t}$ from X, $V\Lambda^{\mathrm t} = W\mathrm{M}^{\mathrm t}$ implies that $\operatorname{trace}(\Lambda^{\mathrm t} V) = \operatorname{trace}(\mathrm{M}^{\mathrm t} W)$ even if X is not finite-dimensional.

The rank of a matrix and of its (conjugate) transpose

In this section, let A' denote either the transpose or the conjugate transpose of the matrix A. Then, either way, $A = VW'$ iff $A' = WV'$. This trivial observation implies all kinds of things about the relationship between a matrix and its (conjugate) transpose.

As a starter, it says that $A = VW'$ is minimal if and only if $A' = WV'$ is minimal. Therefore:

Proposition: $\operatorname{rank} A = \operatorname{rank} A^c = \operatorname{rank} A^t$.

(8.5) Corollary: If A is a matrix, then $\dim \operatorname{ran} A = \dim \operatorname{ran} A^c = \dim \operatorname{ran} A^t$.

(8.6) Corollary: For any matrix A, A' is 1-1 (onto) if and only if A is onto (1-1).

Proof: If $A \in \mathbb{F}^{m\times n}$, then A is onto iff $\operatorname{rank} A = m$ iff $\operatorname{rank} A' = m$ iff the natural factorization $A' = A' \operatorname{id}_m$ is minimal, i.e., iff A' is 1-1.

The other equivalence follows from this since $(A')' = A$. □

For a different proof of these results, see the comments that follow (6.17)Corollary and (6.18)Corollary.

Elimination as factorization

The description (3.2) of elimination does not rely on any particular ordering of the rows of the given $(m \times n)$-matrix A. At any stage, it only distinguishes between pivot rows and those rows not yet used as pivot rows. We may therefore imagine that we initially place the rows of A into the workarray B in exactly the order in which they are going to be used as pivot rows, followed, in any order whatsoever, by those rows (if any) that are never going to be used as pivot rows.

In terms of the n-vector `p` provided by the (3.2)Elimination Algorithm, this means that we start with $B = A(\mathtt{q},:)$, with `q` obtained from `p` by

```
q = p(find(p>0)); 1:m; ans(q) = []; q = [q, ans];
```

Indeed, to recall, `p(j)` is positive if and only if the `j`th unknown is bound, in which case row `p(j)` is the pivot row for that unknown. Thus the assignment `q = p(find(p>0))` initializes `q` so that `A(q,:)` contains the pivot rows in order of their use. With that, `1:m; ans(q) = []` leaves, in `ans`, the indices of all rows not used as pivot rows.

Note that `q` is a permutation of order m. Hence $B = QA$, with Q the corresponding permutation matrix, meaning the matrix obtained from the identity matrix by the very same reordering, $Q =$`eye(m)(q,:)`.

We prefer to write this as $A = PB$, with P the inverse of Q, hence obtainable from `q` by

```
P = eye(m); P(q,:) = P;
```

□

With that done, we have, at the beginning of the algorithm,

$$B = P^{-1}A$$

for some permutation matrix P, and all the work in the algorithm consists of repeatedly subtracting some multiple α of some row h of B from some *later* row, i.e., some row i with $i > h$. In terms of matrices, this means the repeated replacement

$$B \leftarrow E_{i,h}(-\alpha)B$$

with $i > h$. Since, by (2.19), $E_{i,h}(-\alpha)^{-1} = E_{i,h}(\alpha)$, this implies that

$$A = PLU,$$

with L the product of all those elementary matrices $E_{i,h}(\alpha)$ (in the appropriate order), and U the final state of the workarray B. Specifically, U is in row-echelon form (as defined in (3.7)); in particular, U is upper triangular.

Each $E_{i,h}(\alpha)$ is **unit lower triangular**, i.e., of the form $\mathrm{id} + N$ with N **strictly lower triangular**, i.e.,

$$N(r,s) \neq 0 \quad \Longrightarrow \quad r > s.$$

For, because of the initial ordering of the rows in B, only $E_{i,h}(\alpha)$ with $i > h$ appear. This implies that L, as the product of unit lower triangular matrices, is itself unit lower triangular.

If we apply the elimination algorithm to the matrix $[A, C]$, with $A \in \mathbb{F}^{m\times m}$ invertible, then the first m columns are bound, hence the remaining columns are free. In particular, both P and L in the resulting factorization depend only on A and not at all on C.

In particular, in solving $A? = y$, there is no need to subject all of $[A, y]$ to the elimination algorithm. If elimination just applied to A gives the factorization

$$A = PLU \tag{8.7}$$

for an invertible A, then we can find the unique solution x to the equation $A? = y$ by the two-step process:

$$c \leftarrow L^{-1}P^{-1}y$$
$$x \leftarrow U^{-1}c$$

and these two steps are easily carried out. The first step amounts to subjecting the rows of the matrix $[y]$ to all the row operations (including reordering) used during elimination applied to A. The second step is handled by the Backsubstitution Algorithm (3.3), with input $B = [U, c]$, $p = (1, 2, \ldots, m, 0)$, and $z = (0, \ldots, 0, -1)$.

Once it is understood that the purpose of elimination for solving $A? = y$ is the factorization of A into a product of "easily" invertible factors, then it is possible to seek factorizations that might serve the same goal in a better way. The best-known alternative is the QR factorization, in which one obtains

$$A = QR,$$

with R upper triangular and Q o.n., i.e., $Q^{\mathrm{c}}Q = \mathrm{id}$. Such a factorization is obtained by doing elimination a column at a time, usually with the aid of **Householder matrices**. These are elementary matrices of the form

$$H_w := E_{w,w}(-2/w^{\mathrm{c}}w) = \mathrm{id} - \frac{2}{w^{\mathrm{c}}w}ww^{\mathrm{c}},$$

and are easily seen to be **self-inverse** or **involutory** (i.e., $H_wH_w = \mathrm{id}$), **hermitian** (i.e., ${H_w}^{\mathrm{c}} = H_w$), hence **unitary** (i.e., ${H_w}^{\mathrm{c}}H_w = \mathrm{id} = H_w{H_w}^{\mathrm{c}}$).

While the computational cost of constructing the QR factorization is roughly double that needed for the PLU factorization, the QR factorization has the advantage of being more impervious to the effects of rounding errors. Precisely, the relative rounding error effects in both a PLU factorization $A = PLU$ and in a QR factorization $A = QR$ can be shown to be proportional to the condition numbers of the factors. Since Q is o.n., $\kappa(Q) = 1$ and $\kappa(R) = \kappa(A)$, while, for a PLU factorization $A = PLU$, only the permutation matrix, P, is o.n., and $\kappa(L)$ and $\kappa(U)$ can be quite large.

SVD

Let $A = VW^c$ be a minimal factorization for the $m \times n$-matrix A of rank r. Then $A^c = WV^c$ is a minimal factorization for A^c. By (8.2), this implies that V is a basis for $\operatorname{ran} A$ and W is a basis for $\operatorname{ran} A^c$.

Can we choose both these bases to be o.n.?

Well, if both V and W are o.n., then, for any x, $\|Ax\| = \|VW^c x\| = \|W^c x\|$, while, for $x \in \operatorname{ran} A^c$, $x = WW^c x$, hence $\|x\| = \|W^c x\|$. Therefore, altogether, in such a case, A is an isometry on $\operatorname{ran} A^c$, a very special situation.

Nevertheless and, perhaps, surprisingly, there is an o.n. basis W for $\operatorname{ran} A^c$ for which the columns of AW are *orthogonal*, i.e., $AW = V\Sigma$ with V o.n. and Σ diagonal, hence $A = V\Sigma W^c$ with also V o.n.

(8.8) Theorem: For every $A \in \mathbb{F}^{m\times n}$, there exist o.n. bases V and W for $\operatorname{ran} A$ and $\operatorname{ran} A^c$, respectively, and a diagonal matrix Σ with positive diagonal entries, so that

$$A = V\Sigma W^c. \tag{8.9}$$

Proof: For efficiency, the proof given here uses results, concerning the eigenstructure of hermitian positive definite matrices, that are established only later in these notes. This may help to motivate the study to come of the eigenstructure of matrices.

For motivation of the proof, assume for the moment that $A = V\Sigma W^c$ is a factorization of the kind we claim to exist. Then, with $\Sigma =: \operatorname{diag}(\sigma_1, \ldots, \sigma_r)$, it follows that

$$A^c A = W\Sigma^c V^c\, V\Sigma W^c = W\Sigma^c\Sigma W^c,$$

hence

$$A^c A W = W\mathrm{T}, \quad \text{with} \;\; \mathrm{T} := \operatorname{diag}(\tau_1, \tau_2, \ldots, \tau_r) \tag{8.10}$$

and W o.n., and the $\tau_j = \overline{\sigma_j}\sigma_j = |\sigma_j|^2$ all positive.

Just such an o.n. $W \in \mathbb{F}^{n\times r}$ and positive scalars τ_j do exist by (12.2)Corollary and (15.2)Proposition, since the matrix $A^c A$ is **hermitian** (i.e., $(A^c A)^c = A^c A$) and **positive semidefinite** (i.e., $\langle A^c Ax, x\rangle \ge 0$ for all x) and has rank r.

With W and the τ_j so chosen, it follows that W is an o.n. basis for $\operatorname{ran} A^c$, since (8.10) implies that $\operatorname{ran} W \subset \operatorname{ran} A^c$, and W is a 1-1 column map of order $r = \dim \operatorname{ran} A^c$. Further, $U := AW$ satisfies $U^c U = W^c A^c A W = W^c W \mathrm{T} = \mathrm{T}$, hence

$$V := AW\Sigma^{-1}, \quad \text{with} \;\; \Sigma := \mathrm{T}^{1/2} := \operatorname{diag}(\sqrt{\tau_j} : j = 1{:}r),$$

is o.n., and so $V\Sigma W^c = A$, because $WW^c = P := P_{\operatorname{ran} A^c}$, hence $\operatorname{ran}(\operatorname{id} - P) = \operatorname{null} P = \operatorname{ran} A^{c\perp} = \operatorname{null} A$, and so $AWW^c = AP = A(P + (\operatorname{id} - P)) = A$. □

It is customary to order the numbers

$$\sigma_j := \sqrt{\tau_j}, \quad j = 1{:}r.$$

Specifically, one assumes that

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_r.$$

These numbers σ_j are called the (nonzero) **singular value**s of A, and with this ordering, the factorization

$$A = \sum_{j=1}^{\operatorname{rank} A} v_j \sigma_j w_j{}^c$$

is called a (**reduced**) **singular value decomposition** or **svd** for A.

Offhand, a svd is *not* unique. E.g., *any* o.n. basis V for $\mathbb{F}^n$ provides the svd $V\,\mathrm{id}_n V^c$ for id_n.

Some prefer to have a factorization $A = \tilde{V}\tilde{\Sigma}\tilde{W}^c$ in which both $\tilde{V}$ and $\tilde{W}$ are o.n. bases for all of $\mathbb{F}^m$ and $\mathbb{F}^n$, respectively (rather than just for $\operatorname{ran} A$ and $\operatorname{ran} A^c$, respectively). This can always be achieved by extending V and W from (8.9) in any manner whatsoever to o.n. bases $\tilde{V} := [V, V_1]$ and $\tilde{W} := [W, W_1]$ and, correspondingly, extending Σ to

$$\tilde{\Sigma} := \operatorname{diag}(\Sigma, 0) = \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{F}^{m\times n}$$

by the adjunction of blocks of 0 of appropriate size. With this, we have

$$A = \tilde{V}\tilde{\Sigma}\tilde{W}^c = \sum_{j=1}^{\min\{m,n\}} v_j \sigma_j w_j{}^c, \qquad (8.11)$$

and the diagonal entries

$$\sigma_1 \ge \sigma_2 \ge \cdots \ge \sigma_r > 0 = \sigma_{r+1} = \cdots = \sigma_{\min\{m,n\}}$$

of $\tilde{\Sigma}$ are altogether referred to as the **singular value**s of A. Note that this sequence is still ordered. We will refer to (8.11) as a **Singular Value Decomposition** or **SVD**.

The `MATLAB` command `svd(A)` returns the SVD rather than the svd of `A` when issued in the form `[V,S,W] = svd(A)`. Specifically, `A = V*S*W'`, with `V` and `W` both unitary, of order m and n, respectively, if `A` is an $m \times n$-matrix. By itself, `svd(A)` returns, in a one-column matrix, the (ordered) sequence of singular values of `A`.

□

The Pseudo-inverse

Here is a first of many uses to which the svd has been put. It concerns the solution of the equation

$$A? = y$$

in case A is not invertible (for whatever reason). In a previous chapter (see page 69), we looked in this case for a solution of the 'projected' problem

$$A? = P_{\operatorname{ran} A}\, y =: \widehat{y} \qquad (8.12)$$

for the simple reason that any solution x of this equation makes the **residual** $\|Ax - y\|_2$ as small as it can be made by any x. For this reason, any solution of (8.12) is called a **least-squares solution** for $A? = y$.

If now A is 1-1, then (8.12) has exactly one solution. The question is what to do in the contrary case. One proposal is to get the **best least-squares solution**, i.e., the solution of minimal norm. The svd for A makes it easy to find this particular solution.

If $A = V\Sigma W^c$ is a svd for A, then V is an o.n. basis for $\operatorname{ran} A$, hence

$$\widehat{b} = P_{\operatorname{ran} A}\, b = V V^c b.$$

Therefore, (8.12) is equivalent to the equation

$$V\Sigma W^c ? = V V^c b.$$

Since V is o.n., hence 1-1, and Σ is invertible, this equation is, in turn, equivalent to

$$W^c ? = \Sigma^{-1} V^c b,$$

hence to

$$WW^c? = W\Sigma^{-1}V^c b. \tag{8.13}$$

Since W is also o.n., $WW^c = P_W$ is an o.n. projector, hence, by (6.14)Proposition, strictly reduces norms unless it is applied to something in its range. Since the right-hand side of (8.13) is in $\operatorname{ran} W$, it follows that the solution of smallest norm of (8.13), i.e., the best least-squares solution of $A? = y$, is that right-hand side, i.e., the vector

$$\widehat{x} := A^+ y,$$

with the matrix

$$A^+ := W\Sigma^{-1}V^c$$

the **Moore-Penrose pseudo-inverse** of A.

Note that

$$A^+A = W\Sigma^{-1}V^c V\Sigma W^c = WW^c,$$

hence A^+ is a left inverse for A in case W is square, i.e., in case $\operatorname{rank} A = \#A$. Similarly,

$$AA^+ = V\Sigma W^c W\Sigma^{-1}V^c = VV^c,$$

hence A^+ is a right inverse for A in case V is square, i.e., in case $\operatorname{rank} A = \#A^c$. In any case,

$$A^+A = P_{\operatorname{ran} A^c}, \qquad AA^+ = P_{\operatorname{ran} A},$$

therefore, in particular,

$$AA^+A = A.$$

2-norm and 2-condition of a matrix

Recall from (6.23) that o.n. matrices are 2-norm-preserving, i.e.,

$$\|x\|_2 = \|Ux\|_2, \qquad \forall x \in \mathbb{F}^n, \text{o.n. } U \in \mathbb{F}^{m\times n}.$$

This implies that

$$\|TB\|_2 = \|B\|_2 = \|BU^c\|_2, \qquad \forall \text{ o.n. } T \in \mathbb{F}^{r\times m},\ B \in \mathbb{F}^{m\times n},\ \text{o.n. } U \in \mathbb{F}^{r\times n}.$$

Indeed,

$$\|TB\|_2 = \max_{x\neq 0} \frac{\|TBx\|_2}{\|x\|_2} = \max_{x\neq 0} \frac{\|Bx\|_2}{\|x\|_2} = \|B\|_2.$$

By (7.21), this implies that also

$$\|BU^c\|_2 = \|UB^c\|_2 = \|B^c\|_2 = \|B\|_2.$$

It follows that, with $A = V\Sigma W^c \in \mathbb{F}^{m\times n}$ a svd for A,

$$\|A\|_2 = \|\Sigma\|_2 = \sigma_1, \tag{8.14}$$

the last equality because of the fact that $\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_r)$ with $\sigma_1 \geq \sigma_2 \geq \cdots \geq 0$.

Assume that, in addition, A is invertible, therefore $r = \operatorname{rank} A = n = m$, making also V and W square, hence A^+ is both a left and a right inverse for A, therefore necessarily $A^{-1} = A^+ = V\Sigma^{-1}W^c$. It follows that $\|A^{-1}\|_2 = 1/\sigma_n$. Hence, the 2-condition of $A \in \mathbb{F}^{n\times n}$ is

$$\kappa_2(A) = \|A\|_2\|A^{-1}\|_2 = \sigma_1/\sigma_n,$$

and this is how this condition number is frequently *defined*.

The effective rank of a noisy matrix

The problem to be addressed here is the following. If we construct a matrix in the computer, we have to deal with the fact that the entries of the constructed matrix are not quite exact; rounding errors during the calculations may have added some noise. This is even true for a matrix merely entered into the computer, in case some of its entries cannot be represented exactly by the floating point arithmetic used (as is the case, e.g., for the number .1 or the number 1/3 in any of the standard binary-based floatingpoint arithmetics).

This makes it impossible to use, e.g., the rref algorithm to determine the rank of the underlying matrix. However, if one has some notion of the size of the noise involved, then one can use the svd to determine a sharp *lower* bound on the rank of the underlying matrix, because of the following.

(8.15) Proposition: If $A = V\Sigma W^c$ is a svd for A and $\text{rank}(A) > k$, then $\min\{\|A - B\|_2 : \text{rank}(B) \le k\} = \sigma_{k+1} = \|A - A_k\|_2$, with

$$A_k := \sum_{j=1}^{k} v_j \sigma_j {w_j}^c.$$

Proof: If $B \in \mathbb{F}^{m\times n}$ with $\text{rank}(B) \le k$, then $\dim \text{null}(B) > n - (k+1) = \dim \mathbb{F}^n - \dim \text{ran}\, W_{k+1}$, with

$$W_{k+1} := [w_1, w_2, \ldots, w_{k+1}].$$

Therefore, by (4.21)Corollary, the intersection $\text{null}(B) \cap \text{ran}\, W_{k+1}$ contains a vector z of norm 1. Then $Bz = 0$, and $W^c z = {W_{k+1}}^c z$, and $\|{W_{k+1}}^c z\|_2 = \|z\|_2 = 1$. Therefore, $Az = V\Sigma W^c z = V_{k+1}\Sigma_{k+1}{W_{k+1}}^c z$, hence

$$\|A - B\|_2 \ge \|Az - Bz\|_2 = \|Az\|_2 = \|\Sigma_{k+1}{W_{k+1}}^c z\|_2 \ge \sigma_{k+1}\|{W_{k+1}}^c z\|_2 = \sigma_{k+1}.$$

On the other hand, for the specific choice $B = A_k$, we get $\|A - A_k\|_2 = \sigma_{k+1}$ by (8.14), since $A - A_k = \sum_{j>k} v_j\sigma_j {w_j}^c$ is a svd for it, hence its largest singular value is σ_{k+1}. □

In particular, if we have in hand a svd

$$A + E = V\,\text{diag}(\hat\sigma_1, \hat\sigma_2, \ldots, \hat\sigma_{\hat r})W^c$$

for the *perturbed* matrix $A + E$, and know (or believe) that $\|E\|_2 \le \varepsilon$, then the best we can say about the rank of A is that it must be at least

$$r_\varepsilon := \max\{j : \hat\sigma_j > \varepsilon\}.$$

For example, the matrix

$$A = \begin{bmatrix} 2/3 & 1 & 1/3 \\ 4/3 & 2 & 2/3 \\ 1 & 1 & 1 \end{bmatrix}$$

is readily transformed by elimination into the matrix

$$B = \begin{bmatrix} 0 & 1/3 & -1/3 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix},$$

hence has rank 2. However, on entering A into a computer correct to four decimal places after the decimal point, we get (more or less) the matrix

$$A_c = \begin{bmatrix} .6667 & 1 & .3333 \\ 1.3333 & 2 & .6667 \\ 1 & 1 & 1 \end{bmatrix},$$

and for it, `MATLAB` correctly returns id_3 as its `rref`. However, the singular values of A_c, as returned by `svd`, are

```
(3.2340..., 0.5645..., 0.000054...)
```

indicating that there is a rank-2 matrix B with $\|A_c - B\|_2 < .000055$. Since entries of A_c are only accurate to within 0.00005, the safe conclusion is that A has rank ≥ 2; it happens to have rank 2 in this particular example.

The polar decomposition

The svd can also be very helpful in establishing results of a more theoretical flavor, as the following discussion is intended to illustrate.

This discussion concerns a useful extension to square matrices of the polar form (see Backgrounder)

$$z = |z| \exp(\mathrm{i}\varphi)$$

of a complex number z, i.e., a factorization of z into a nonnegative number $|z| = \sqrt{z\overline{z}}$ (its modulus or absolute value) and a number whose absolute value is equal to 1, a socalled **unimodular** number.

There is, for any $A \in \mathbb{C}^{n\times n}$, a corresponding decomposition

$$A = \sqrt{AA^{\mathrm{c}}}\,E, \qquad (8.16)$$

called a **polar decomposition**, with $\sqrt{AA^{\mathrm{c}}}$ 'nonnegative' in the sense that it is hermitian and positive semidefinite, and E 'unimodular' in the sense that it is unitary, hence norm-preserving, i.e., an isometry.

The polar decomposition is almost immediate, given that we already have a SVD $A = \tilde{V}\tilde{\Sigma}\tilde{W}^{\mathrm{c}}$ for A (see (8.11)) in hand. Indeed, from that,

$$A = \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}\ \tilde{V}\tilde{W}^{\mathrm{c}},$$

with $P := \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}$ evidently hermitian, and also positive semidefinite since

$$\langle Px, x\rangle = x^{\mathrm{c}}\tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}x = \sum_j \tilde{\sigma}_j |(\tilde{V}^{\mathrm{c}}x)_j|^2$$

is nonnegative for all x, given that $\tilde{\sigma}_j \ge 0$ for all j; and

$$P^2 = \tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}}\tilde{V}\tilde{\Sigma}\tilde{V}^{\mathrm{c}} = \tilde{V}\tilde{\Sigma}\tilde{\Sigma}^{\mathrm{c}}\tilde{V}^{\mathrm{c}} = \tilde{V}\tilde{\Sigma}\tilde{W}^{\mathrm{c}}\tilde{W}\tilde{\Sigma}^{\mathrm{c}}\tilde{V}^{\mathrm{c}} = AA^{\mathrm{c}};$$

and, finally, $E := \tilde{V}\tilde{W}^{\mathrm{c}}$ unitary as the product of unitary maps.

Equivalence and similarity

The SVD provides a particularly useful example of *equivalence*. The linear maps A and $\widehat{A}$ are called **equivalent** if there are *invertible* linear maps V and W so that

$$A \ = \ V\widehat{A}W^{-1}.$$

Since both V and W are invertible, such equivalent linear maps share all essential properties, such as their rank, being 1-1, or onto, or invertible.

Equivalence is particularly useful when the domains of V and W are coordinate spaces, i.e., when V and W are *bases*, and, correspondingly, $\widehat{A}$ is a matrix, as in the following diagram:

$$\begin{array}{ccc} X & \xrightarrow{A} & Y \\ W\uparrow & & \uparrow V \\ \mathbb{F}^n & \xrightarrow[\widehat{A}]{} & \mathbb{F}^m \end{array}$$

In this situation, $\widehat{A} = V^{-1}AW$ is called a **matrix representation for** A.

For example, we noted earlier that the matrix

$$\widehat{D}_k := \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 2 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & k \end{bmatrix}$$

is the standard matrix representation used in Calculus for the linear map $D : \Pi_k \to \Pi_{k-1}$ of differentiation of polynomials of degree $\le k$.

In practice, one looks, for given $A \in L(X, Y)$, for matrix representations $\widehat{A}$ that are as simple as possible. If that means a matrix with as many zero entries as possible and, moreover, all the nonzero entries the same, say equal to 1, then a simplest such matrix representation is of the form

$$\widehat{A} = \operatorname{diag}(\operatorname{id}_{\operatorname{rank} A}, 0) = \begin{bmatrix} \operatorname{id}_{\operatorname{rank} A} & 0 \\ 0 & 0 \end{bmatrix},$$

with 0 indicating zero matrices of the appropriate size to make $\widehat{A}$ of size $\dim \operatorname{tar} A \times \dim \operatorname{dom} A$. It can be obtained from any minimal factorization $A = \tilde{V}\Lambda^{\mathrm{t}}$ by extending $\tilde{V}$ to a basis $V = [\tilde{V}, V_1]$ of $\operatorname{tar} A$ and extending Λ^{t} to the inverse W^{-1} of a basis W for $\operatorname{dom} A$.

The situation becomes much more interesting and challenging when $\operatorname{dom} A = \operatorname{tar} A$ and, correspondingly, we insist that also $V = W$. Linear maps A and $\widehat{A}$ for which there exists an invertible linear map V with

$$A = V\widehat{A}V^{-1}$$

are called *similar*. Such similarity will drive much of the rest of these notes.

8.10 T/F

() If A, B, M are matrices such that $\operatorname{rank} AM = \operatorname{rank} B$, then M is invertible.

() If M is invertible and $AM = B$, then $\operatorname{rank} AM = \operatorname{rank} B$.

() If M is invertible and $MA = B$, then $\operatorname{rank} MA = \operatorname{rank} B$.

