



Vector Spaces

Linear Transformations
and Matrices

Properties of Matrices

Singular Value
Decomposition

Chapter 1

Mathematical Background: Linear Algebra

Multiple View Geometry
Summer 2017

Prof. Daniel Cremers
Chair for Computer Vision and Pattern Recognition
Department of Computer Science
Technische Universität München



Vector Spaces

Linear Transformations
and Matrices

Properties of Matrices

Singular Value
Decomposition

1 Vector Spaces

2 Linear Transformations and Matrices

3 Properties of Matrices

4 Singular Value Decomposition

Vector Space (Vektorraum)



A set V is called a **linear space** or a **vector space over the field** \mathbb{R} if it is closed under vector summation

$$+ : V \times V \rightarrow V$$

and under scalar multiplication

$$\cdot : \mathbb{R} \times V \rightarrow V,$$

i.e. $\alpha v_1 + \beta v_2 \in V \forall v_1, v_2 \in V, \forall \alpha, \beta \in \mathbb{R}$. With respect to addition (+) it forms a commutative group (existence of neutral element 0, inverse element $-v$). Scalar multiplication respects the structure of \mathbb{R} : $\alpha(\beta u) = (\alpha\beta)u$. Multiplication and addition respect the distributive law:

$$(\alpha + \beta)v = \alpha v + \beta v \text{ and } \alpha(v + u) = \alpha v + \alpha u.$$

Example: $V = \mathbb{R}^n$, $v = (x_1, \dots, x_n)^\top$.

A subset $W \subset V$ of a vector space V is called **subspace** if $0 \in W$ and W is closed under $+$ and \cdot (for all $\alpha \in \mathbb{R}$).

Linear Independence and Basis



The spanned subspace of a set of vectors

$S = \{v_1, \dots, v_k\} \subset V$ is the subspace formed by all linear combinations of these vectors:

$$\text{span}(S) = \left\{ v \in V \mid v = \sum_{i=1}^k \alpha_i v_i \right\}$$

The set S is called **linearly independent** if:

$$\sum_{i=1}^k \alpha_i v_i = 0 \Rightarrow \alpha_i = 0 \quad \forall i,$$

in other words: if none of the vectors can be expressed as a linear combination of the remaining vectors. Otherwise the set is called **linearly dependent**.

A set of vectors $B = \{v_1, \dots, v_n\}$ is called a **basis of V** if it is linearly independent and if it spans the vector space V . A basis is a maximal set of linearly independent vectors.

Properties of a Basis

Let B and B' be two bases of a linear space V .

- 1 B and B' contain the same number of vectors. This number n is called the **dimension of the space V** .
- 2 Any vector $v \in V$ can be uniquely expressed as a linear combination of the basis vectors in $B = \{b_1, \dots, b_n\}$:

$$v = \sum_{i=1}^n \alpha_i b_i.$$

- 3 In particular, all vectors of B can be expressed as linear combinations of vectors of another basis $b'_i \in B'$:

$$b'_i = \sum_{j=1}^n \alpha_{ji} b_j$$

The coefficients α_{ji} for this **basis transform** can be combined in a matrix A . Setting $B \equiv (b_1, \dots, b_n)$ and $B' \equiv (b'_1, \dots, b'_n)$ as the matrices of basis vectors, we can write: $B' = BA \Leftrightarrow B = B'A^{-1}$.



Inner Product

On a vector space one can define an **inner product** (dot product, dt.: Skalarprodukt \neq skalare Multiplikation):

$$\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$$

which is defined by three properties:

- ① $\langle u, \alpha v + \beta w \rangle = \alpha \langle u, v \rangle + \beta \langle u, w \rangle$ (linear)
- ② $\langle u, v \rangle = \langle v, u \rangle$ (symmetric)
- ③ $\langle v, v \rangle \geq 0$ and $\langle v, v \rangle = 0 \Leftrightarrow v = 0$ (positive definite)

The scalar product induces a **norm**

$$|\cdot| : V \rightarrow \mathbb{R}, \quad |v| = \sqrt{\langle v, v \rangle}$$

and a **metric**

$$d : V \times V \rightarrow \mathbb{R}, \quad d(v, w) = |v - w| = \sqrt{\langle v - w, v - w \rangle}$$

for measuring lengths and distances, making V a **metric space**.
Since the metric is induced by a scalar product V is called a **Hilbert space**.



Canonical and Induced Inner Product

On $V = \mathbb{R}^n$, one can define the canonical inner product for the canonical basis $B = I_n$ as

$$\langle x, y \rangle = x^\top y = \sum_{i=1}^n x_i y_i$$

which induces the standard L_2 -norm or Euclidean norm

$$|x|_2 = \sqrt{x^\top x} = \sqrt{x_1^2 + \dots + x_n^2}$$

With a basis transform A to the new basis B' given by $I = B' A^{-1}$ the canonical inner product in the new coordinates x', y' is given by:

$$\langle x, y \rangle = x^\top y = (Ax')^\top (Ay') = x'^\top A^\top A y' \equiv \langle x', y' \rangle_{A^\top A}$$

The latter product is called the **induced inner product** from the matrix A .

Two vectors v and w are **orthogonal** iff $\langle v, w \rangle = 0$.



Kronecker Product and Stack of a Matrix

Given two matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{k \times l}$, one can define their **Kronecker product** $A \otimes B$ by:

$$A \otimes B \equiv \begin{pmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \cdots & a_{mn}B \end{pmatrix} \in \mathbb{R}^{mk \times nl}.$$

In Matlab this can be implemented by `C=kron(A,B)`.

Given a matrix $A \in \mathbb{R}^{m \times n}$, its **stack** A^s is obtained by stacking its n column vectors $a_1, \dots, a_n \in \mathbb{R}^m$:

$$A^s \equiv \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \in \mathbb{R}^{mn}.$$

These notations allow to rewrite algebraic expressions, for example:

$$u^\top A v = (v \otimes u)^\top A^s.$$



Linear Transformations and Matrices

Linear algebra studies the properties of linear transformations between linear spaces. Since these can be represented by matrices, linear algebra studies the properties of matrices.

A **linear transformation** L between two linear spaces V and W is a map $L : V \rightarrow W$ such that:

- $L(x + y) = L(x) + L(y) \quad \forall x, y \in V$
- $L(\alpha x) = \alpha L(x) \quad \forall x \in V, \alpha \in \mathbb{R}.$

Due to the linearity, the action of L on the space V is uniquely defined by its action on the basis vectors of V . In the canonical basis $\{e_1, \dots, e_n\}$ we have:

$$L(x) = Ax \quad \forall x \in V,$$

where

$$A = (L(e_1), \dots, L(e_n)) \in \mathbb{R}^{m \times n}.$$

The set of all real $m \times n$ -matrices is denoted by $\mathcal{M}(m, n)$. In the case that $m = n$, the set $\mathcal{M}(m, n) \equiv \mathcal{M}(n)$ forms a **ring** over the field \mathbb{R} , i.e. it is closed under matrix multiplication and summation.



The Linear Groups $GL(n)$ and $SL(n)$



There exist certain sets of linear transformations which form a group.

A **group** is a set G with an operation $\circ : G \times G \rightarrow G$ such that:

- 1 closed: $g_1 \circ g_2 \in G \quad \forall g_1, g_2 \in G$,
- 2 assoc.: $(g_1 \circ g_2) \circ g_3 = g_1 \circ (g_2 \circ g_3) \quad \forall g_1, g_2, g_3 \in G$,
- 3 neutral: $\exists e \in G : e \circ g = g \circ e = g \quad \forall g \in G$,
- 4 inverse: $\exists g^{-1} \in G : g \circ g^{-1} = g^{-1} \circ g = e \quad \forall g \in G$.

Example: All invertible (non-singular) real $n \times n$ -matrices form a group with respect to matrix multiplication. This group is called the **general linear group** $GL(n)$. It consists of all $A \in \mathcal{M}(n)$ for which

$$\det(A) \neq 0$$

All matrices $A \in GL(n)$ for which $\det(A) = 1$ form a group called the **special linear group** $SL(n)$. The inverse of A is also in this group, as $\det(A^{-1}) = \det(A)^{-1}$

Matrix Representation of Groups



A group G has a **matrix representation** (dt.: Darstellung) or can be realized as a matrix group if there exists an injective transformation:

$$R : G \rightarrow GL(n)$$

which **preserves the group structure** of G , that is inverse and composition are preserved by the map:

$$R(e) = I_{n \times n}, \quad R(g \circ h) = R(g)R(h) \quad \forall g, h \in G$$

Such a map R is called a **group homomorphism** (dt. Homomorphismus).

The idea of matrix representations of a group is that they allow to analyze more abstract groups by looking at the properties of the respective matrix group. Example: The rotations of an object form a group, as there exists a neutral element (no rotation) and an inverse (the inverse rotation) and any concatenation of rotations is again a rotation (around a different axis). Studying the properties of the rotation group is easier if rotations are represented by respective matrices.

The Affine Group $A(n)$



An affine transformation $L : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined by a matrix $A \in GL(n)$ and a vector $b \in \mathbb{R}^n$ such that:

$$L(x) = Ax + b$$

The set of all such affine transformations is called the **affine group of dimension n** , denoted by $A(n)$.

L defined above is not a linear map unless $b = 0$. By introducing **homogeneous coordinates** to represent $x \in \mathbb{R}^n$ by $\begin{pmatrix} x \\ 1 \end{pmatrix} \in \mathbb{R}^{n+1}$, L becomes a linear mapping from

$$L : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}; \quad \begin{pmatrix} x \\ 1 \end{pmatrix} \mapsto \begin{pmatrix} A & b \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ 1 \end{pmatrix}.$$

A matrix $\begin{pmatrix} A & b \\ 0 & 1 \end{pmatrix}$ with $A \in GL(n)$ and $b \in \mathbb{R}^n$ is called an **affine matrix**. It is an element of $GL(n+1)$. The affine matrices form a subgroup of $GL(n+1)$. Why?

The Orthogonal Group $O(n)$

A matrix $A \in \mathcal{M}(n)$ is called **orthogonal** if it preserves the inner product, i.e.:

$$\langle Ax, Ay \rangle = \langle x, y \rangle, \quad \forall x, y \in \mathbb{R}^n.$$

The set of all orthogonal matrices forms the **orthogonal group** $O(n)$, which is a subgroup of $GL(n)$. For an orthogonal matrix R we have

$$\langle Rx, Ry \rangle = x^\top R^\top Ry = x^\top y, \quad \forall x, y \in \mathbb{R}^n$$

Therefore we must have $R^\top R = RR^\top = I$, in other words:

$$O(n) = \{R \in GL(n) \mid R^\top R = I\}$$

The above identity shows that for any orthogonal matrix R , we have $\det(R^\top R) = (\det(R))^2 = \det(I) = 1$, such that $\det(R) \in \{\pm 1\}$.

The subgroup of $O(n)$ with $\det(R) = +1$ is called the **special orthogonal group** $SO(n)$. $SO(n) = O(n) \cap SL(n)$. In particular, $SO(3)$ is the group of all 3-dimensional **rotation matrices**.



The Euclidean Group $E(n)$

A Euclidean transformation L from \mathbb{R}^n to \mathbb{R}^n is defined by an orthogonal matrix $R \in O(n)$ and a vector $T \in \mathbb{R}^n$:

$$L : \mathbb{R}^n \rightarrow \mathbb{R}^n; \quad x \mapsto Rx + T.$$

The set of all such transformations is called the **Euclidean group $E(n)$** . It is a subgroup of the affine group $A(n)$. Embedded by homogeneous coordinates, we get:

$$E(n) = \left\{ \begin{pmatrix} R & T \\ 0 & 1 \end{pmatrix} \mid R \in O(n), T \in \mathbb{R}^n \right\}.$$

If $R \in SO(n)$ (i.e. $\det(R) = 1$), then we have the **special Euclidean group $SE(n)$** . In particular, $SE(3)$ represents the **rigid-body motions** (dt.: Starrkörpertransformationen) in \mathbb{R}^3 .

In summary:

$$SO(n) \subset O(n) \subset GL(n), \quad SE(n) \subset E(n) \subset A(n) \subset GL(n+1).$$



Range, Span, Null Space and Kernel

Let $A \in \mathbb{R}^{m \times n}$ be a matrix defining a linear map from \mathbb{R}^n to \mathbb{R}^m . The **range** or **span** of A (dt.: Bild) is defined as the subspace of \mathbb{R}^m which can be 'reached' by A :

$$\text{range}(A) = \{y \in \mathbb{R}^m \mid \exists x \in \mathbb{R}^n : Ax = y\}.$$

The range of a matrix A is given by the span of its column vectors.

The **null space** or **kernel** of a matrix A (dt.: Kern) is given by the subset of vectors $x \in \mathbb{R}^n$ which are mapped to zero:

$$\text{null}(A) \equiv \ker(A) = \{x \in \mathbb{R}^n \mid Ax = 0\}.$$

The null space of a matrix A is given by the vectors orthogonal to its row vectors. **Matlab:** `z=null(A)`.

The concepts of range and null space are useful when studying the **solution of linear equations**. The system $Ax = b$ will have a solution $x \in \mathbb{R}^n$ if and only if $b \in \text{range}(A)$. Moreover, this solution will be unique only if $\ker(A) = \{0\}$. Indeed, if x_s is a solution of $Ax = b$ and $x_o \in \ker(A)$, then $x_s + x_o$ is also a solution: $A(x_s + x_o) = Ax_s + Ax_o = b$.



Rank of a Matrix

The **rank** of a matrix (dt. Rang) is the dimension of its range:

$$\text{rank}(A) = \dim(\text{range}(A)).$$

The rank of a matrix $A \in \mathbb{R}^{m \times n}$ has the following properties:

- 1 $\text{rank}(A) = n - \dim(\ker(A)).$
- 2 $0 \leq \text{rank}(A) \leq \min\{m, n\}.$
- 3 $\text{rank}(A)$ is equal to the maximum number of linearly independent row (or column) vectors of A .
- 4 $\text{rank}(A)$ is the highest order of a nonzero minor of A , where a **minor of order k** is the determinant of a $k \times k$ submatrix of A .
- 5 **Sylvester's inequality:** Let $B \in \mathbb{R}^{n \times k}$.
Then $AB \in \mathbb{R}^{m \times k}$ and
 $\text{rank}(A) + \text{rank}(B) - n \leq \text{rank}(AB) \leq \min\{\text{rank}(A), \text{rank}(B)\}.$
- 6 For any nonsingular matrices $C \in \mathbb{R}^{m \times m}$ and $D \in \mathbb{R}^{n \times n}$, we have: $\text{rank}(A) = \text{rank}(CAD).$ **Matlab:** `d=rank(A).`



Eigenvalues and Eigenvectors



Let $A \in \mathbb{C}^{n \times n}$ be a complex matrix. A non-zero vector $v \in \mathbb{C}^n$ is called a (right) eigenvector of A if:

$$Av = \lambda v, \quad \text{with } \lambda \in \mathbb{C}.$$

λ is called an eigenvalue of A . Similarly v is called a left eigenvector of A , if $v^\top A = \lambda v^\top$ for some $\lambda \in \mathbb{C}$.

The spectrum $\sigma(A)$ of a matrix A is the set of all its eigenvalues.

Matlab:

$$[V, D] = \text{eig}(A);$$

where D is a diagonal matrix containing the eigenvalues and V is a matrix whose columns are the corresponding eigenvectors, such that $AV = VD$.

Properties of Eigenvalues and Eigenvectors

Let $A \in \mathbb{R}^{n \times n}$ be a square matrix. Then:

- 1 If $Av = \lambda v$ for some $\lambda \in \mathbb{R}$, then there also exists a **left-eigenvector** $\eta \in \mathbb{R}^n : \eta^\top A = \lambda \eta^\top$. Hence $\sigma(A) = \sigma(A^\top)$.
- 2 The eigenvectors of a matrix A associated with different eigenvalues are linearly independent.
- 3 All eigenvalues $\sigma(A)$ are the roots of the **characteristic polynomial** $\det(\lambda I - A) = 0$. Therefore $\det(A)$ is equal to the product of all eigenvalues (some of which may appear multiple times).
- 4 If $B = PAP^{-1}$ for some nonsingular matrix P , then $\sigma(B) = \sigma(A)$.
- 5 If $\lambda \in \mathbb{C}$ is an eigenvalue, then its conjugate $\bar{\lambda}$ is also an eigenvalue. Thus $\sigma(A) = \overline{\sigma(A)}$ for real matrices A .



Symmetric Matrices

A matrix $S \in \mathbb{R}^{n \times n}$ is called **symmetric** if $S^T = S$. A symmetric matrix S is called **positive semi-definite** (denoted by $S \geq 0$ or $S \succeq 0$) if $x^T S x \geq 0$. S is called **positive definite** (denoted by $S > 0$ or $S \succ 0$) if $x^T S x > 0 \forall x \neq 0$.

Let $S \in \mathbb{R}^{n \times n}$ be a real symmetric matrix. Then:

- 1 All eigenvalues of S are real, i.e. $\sigma(S) \subset \mathbb{R}$.
- 2 Eigenvectors v_i and v_j of S corresponding to distinct eigenvalues $\lambda_i \neq \lambda_j$ are orthogonal.
- 3 There always exist n orthonormal eigenvectors of S which form a basis of \mathbb{R}^n . Let $V = (v_1, \dots, v_n) \in O(n)$ be the orthogonal matrix of these eigenvectors, and $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_n\}$ the diagonal matrix of eigenvalues. Then we have $S = V \Lambda V^T$.
- 4 S is positive (semi-)definite, if all eigenvalues are positive (nonnegative).
- 5 Let S be positive semi-definite and λ_1, λ_n the largest and smallest eigenvalue. Then $\lambda_1 = \max_{|x|=1} \langle x, Sx \rangle$ and $\lambda_n = \min_{|x|=1} \langle x, Sx \rangle$.



Norms of Matrices

There are many ways to define norms on the space of matrices $A \in \mathbb{R}^{m \times n}$. They can be defined based on norms on the domain or codomain spaces on which A operates. In particular, the **induced 2-norm of a matrix A** is defined as

$$\|A\|_2 \equiv \max_{|x|_2=1} |Ax|_2 = \max_{|x|_2=1} \sqrt{\langle x, A^\top A x \rangle}.$$

Alternatively, one can define the **Frobenius norm of A** as:

$$\|A\|_f \equiv \sqrt{\sum_{i,j} a_{ij}^2} = \sqrt{\text{trace}(A^\top A)}.$$

Note that these norms are in general not the same. Since the matrix $A^\top A$ is symmetric and pos. semi-definite, we can diagonalize it as: $A^\top A = V \text{diag} \{ \sigma_1^2, \dots, \sigma_n^2 \} V^\top$ with $\sigma_1^2 \geq \sigma_i^2 \geq 0$. This leads to:

$$\|A\|_2 = \sigma_1, \quad \text{and} \quad \|A\|_f = \sqrt{\text{trace}(A^\top A)} = \sqrt{\sigma_1^2 + \dots + \sigma_n^2}.$$



Skew-symmetric Matrices



A matrix $A \in \mathbb{R}^{n \times n}$ is called **skew-symmetric** or **anti-symmetric** (dt. schiefsymmetrisch) if $A^\top = -A$.

If A is a real skew-symmetric matrix, then:

- 1 All eigenvalues of A are either zero or purely imaginary, i.e. of the form $i\omega$ with $i^2 = -1$, $\omega \in \mathbb{R}$.
- 2 There exists an orthogonal matrix V such that

$$A = V \Lambda V^\top,$$

where Λ is a block-diagonal matrix

$\Lambda = \text{diag}\{A_1, \dots, A_m, 0, \dots, 0\}$, with real skew-symmetric matrices A_i of the form:

$$A_i = \begin{pmatrix} 0 & a_i \\ -a_i & 0 \end{pmatrix} \in \mathbb{R}^{2 \times 2}, \quad i = 1, \dots, m.$$

In particular, the rank of any skew-symmetric matrix is even.

Examples of Skew-symmetric Matrices



In Computer Vision, a common skew-symmetric matrix is given by the **hat operator** of a vector $u \in \mathbb{R}^3$ is:

$$\hat{u} = \begin{pmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{pmatrix} \in \mathbb{R}^{3 \times 3}.$$

This is a linear operator from the space of vectors \mathbb{R}^3 to the space of skew-symmetric matrices in $\mathbb{R}^{3 \times 3}$.

In particular, the matrix \hat{u} has the property that

$$\hat{u}v = u \times v,$$

where \times denotes the standard vector cross product in \mathbb{R}^3 . For $u \neq 0$, we have $\text{rank}(\hat{u}) = 2$ and the null space of \hat{u} is spanned by u , because $\hat{u}u = u^\top \hat{u} = 0$.

The Singular Value Decomposition (SVD)



In the last slides, we have studied many properties of matrices, such as rank, range, null space, and induced norms of matrices. Many of these properties can be captured by the so-called **singular value decomposition (SVD)**.

SVD can be seen as a generalization of eigenvalues and eigenvectors to non-square matrices. The computation of SVD is numerically well-conditioned. It is very useful for solving linear-algebraic problems such as matrix inversion, rank computation, linear least-squares estimation, projections, and fixed-rank approximations.

In practice, both singular value decomposition and eigenvalue decompositions are used quite extensively.

Algebraic Derivation of SVD

Let $A \in \mathbb{R}^{m \times n}$ with $m \geq n$ be a matrix of $\text{rank}(A) = p$. Then there exist

- $U \in \mathbb{R}^{m \times p}$ whose columns are orthonormal
- $V \in \mathbb{R}^{n \times p}$ whose columns are orthonormal, and
- $\Sigma \in \mathbb{R}^{p \times p}$, $\Sigma = \text{diag}\{\sigma_1, \dots, \sigma_p\}$, with $\sigma_1 \geq \dots \geq \sigma_p$,

such that

$$A = U \Sigma V^\top.$$

Note that this **generalizes the eigenvalue decomposition**. While the latter decomposes a symmetric square matrix A with an orthogonal transformation V as:

$$A = V \Lambda V^\top, \quad \text{with } V \in O(n), \Lambda = \text{diag}\{\lambda_1, \dots, \lambda_n\},$$

SVD allows to decompose an arbitrary (non-square) matrix A of rank p with two transformations U and V with orthonormal columns as shown above. Nevertheless, we will see that SVD is based on the eigenvalue decomposition of symmetric square matrices.



Proof of SVD Decomposition 1

Given a matrix $A \in \mathbb{R}^{m \times n}$ with $m \geq n$ and $\text{rank}(A) = p$, the matrix

$$A^\top A \in \mathbb{R}^{n \times n}$$

is symmetric and positive semi-definite. Therefore it can be decomposed with non-negative eigenvalues $\sigma_1^2 \geq \dots \geq \sigma_n^2 \geq 0$ with orthonormal eigenvectors v_1, \dots, v_n . The σ_i are called **singular values**. Since

$$\ker(A^\top A) = \ker(A) \quad \text{and} \quad \text{range}(A^\top A) = \text{range}(A^\top),$$

we have $\text{span}\{v_1, \dots, v_p\} = \text{range}(A^\top)$ and $\text{span}\{v_{p+1}, \dots, v_n\} = \ker(A)$. Let

$$u_i \equiv \frac{1}{\sigma_i} A v_i \Leftrightarrow A v_i = \sigma_i u_i, \quad i = 1, \dots, p$$

then the $u_i \in \mathbb{R}_m$ are orthonormal:

$$\langle u_i, u_j \rangle = \frac{1}{\sigma_i \sigma_j} \langle A v_i, A v_j \rangle = \frac{1}{\sigma_i \sigma_j} \langle v_i, A^\top A v_j \rangle = \delta_{ij}.$$



Proof of SVD Decomposition 2

Complete $\{u_i\}_{i=1}^p$ to a basis $\{u_i\}_{i=1}^m$ of \mathbb{R}_m . Since $Av_i = \sigma_i u_i$, we have

$$A(v_1, \dots, v_n) = (u_1, \dots, u_m) \begin{pmatrix} \sigma_1 & 0 & 0 & \dots & 0 \\ 0 & \ddots & 0 & \vdots & 0 \\ 0 & \dots & \sigma_p & \vdots & 0 \\ \vdots & \dots & \dots & \vdots & 0 \\ 0 & \dots & \dots & 0 & 0 \end{pmatrix},$$

which is of the form $A\tilde{V} = \tilde{U}\tilde{\Sigma}$, thus

$$A = \tilde{U}\tilde{\Sigma}\tilde{V}^\top.$$

Now simply delete all columns of \tilde{U} and the rows of \tilde{V}^\top which are multiplied by zero singular values and we obtain the form $A = U\Sigma V^\top$, with $U \in \mathbb{R}^{m \times p}$ and $V \in \mathbb{R}^{n \times p}$.

In Matlab: `[U, S, V] = svd(A).`



Vector Spaces

Linear Transformations
and Matrices

Properties of Matrices

Singular Value
Decomposition

A Geometric Interpretation of SVD

For $A \in \mathbb{R}^{n \times n}$, the singular value decomposition $A = U \Sigma V^T$ is such that the columns $U = (u_1, \dots, u_n)$ and $V = (v_1, \dots, v_n)$ form orthonormal bases of \mathbb{R}^n . If a point $x \in \mathbb{R}^n$ is mapped to a point $y \in \mathbb{R}^n$ by the transformation A , then the coordinates of y in basis U are related to the coordinates of x in basis V by the diagonal matrix Σ : each coordinate is merely scaled by the corresponding singular value:

$$y = Ax = U \Sigma V^T x \quad \Leftrightarrow \quad U^T y = \Sigma V^T x.$$

The matrix A maps the unit sphere into an ellipsoid with semi-axes $\sigma_i u_i$. To see this, we call $\alpha \equiv V^T x$ the coefficients of the point x in the basis V and those of y in basis U shall be called $\beta \equiv U^T y$. All points of the circle fulfill $|x|_2^2 = \sum_i \alpha_i^2 = 1$. The above statement says that $\beta_i = \sigma_i \alpha_i$. Thus for the points on the sphere we have

$$\sum_i \alpha_i^2 = \sum_i \beta_i^2 / \sigma_i^2 = 1,$$

which states that the transformed points lie on an ellipsoid oriented along the axes of the basis U .



The Generalized (Moore Penrose) Inverse



For certain quadratic matrices one can define an inverse matrix, if $\det(A) \neq 0$. The set of all invertible matrices forms the group $GL(n)$. One can also define a (generalized) inverse (also called pseudo inverse) (dt.: Pseudoinverse) for an arbitrary (non-quadratic) matrix $A \in \mathbb{R}^{m \times n}$. If its SVD is $A = U \Sigma V^T$ the pseudo inverse is defined as:

$$A^\dagger = V \Sigma^\dagger U^T, \text{ where } \Sigma^\dagger = \begin{pmatrix} \Sigma_1^{-1} & 0 \\ 0 & 0 \end{pmatrix}_{n \times m},$$

where Σ_1 is the diagonal matrix of non-zero singular values. In **Matlab**: `X=pinv(A)`. In particular, the pseudo inverse can be employed in a similar fashion as the inverse of quadratic invertible matrices:

$$AA^\dagger A = A, \quad A^\dagger AA^\dagger = A^\dagger.$$

The linear system $Ax = b$ with $A \in \mathbb{R}^{m \times n}$ of rank $r \leq \min(m, n)$ can have multiple or no solutions. $x_{min} = A^\dagger b$ is among all minimizers of $|Ax - b|^2$ the one with the smallest norm $|x|$.