

Least squares, the singular value decomposition, and linear inverse problems

Marcel Oliver

Perspectives of Mathematics II – Spring 2010

1 Simple least squares

The *simple least squares* is the following special case of a *linear regression* problem: Find the equation of a line which is “closest” to a given set of points in the plane. More precisely, given tuples of real numbers $(x_1, y_1), \dots, (x_n, y_n)$, find numbers a and b such that

$$f(a, b) \equiv \sum_{i=1}^n (a x_i + b - y_i)^2 \quad (1)$$

is minimal.

As f is a smooth function defined for all $(a, b) \in \mathbb{R}^2$, calculus tells us that it can only have a minimum provided its partial derivatives vanish. I.e.,

$$0 = \frac{\partial f}{\partial b} = 2 \sum_{i=1}^n a x_i + b - y_i \quad (2a)$$

and

$$0 = \frac{\partial f}{\partial a} = 2 \sum_{i=1}^n x_i (a x_i + b - y_i). \quad (2b)$$

Defining the *mean* of the $\{x_i\}$ and $\{y_i\}$ as

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad \text{and} \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i, \quad (3)$$

equation (2a) can be written

$$b = \bar{y} - a \bar{x}. \quad (4a)$$

Substituting into (2b) and solving for a , we find

$$a = \frac{\sum_{i=1}^n x_i y_i - n \bar{x} \bar{y}}{\sum_{i=1}^n x_i^2 - n \bar{x}^2} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2}. \quad (4b)$$

Thus, we can first compute a via (4b) and then b via (4a).

Two remarks are in order. First, vanishing of the partial derivatives is, in general, only a necessary condition for the existence of a minimum. Here, however, it is not difficult to show that $f(a, b) \rightarrow \infty$ as $a, b \rightarrow \infty$ (in any way) provided there are at least two distinct data points. Thus, since the necessary condition yields a unique solution, this solution must correspond to a minimum of f .

Second, the minimization of the Euclidean norm of the error as opposed to some other error measure is certainly convenient as it yields a smooth and even linear relationship between measured y -values and estimates for the coefficients. However, it is often also preferred for statistical reasons: When the errors are assumed Gaussian, the least-squares fit equals the so-called maximum-likelihood estimator. When the errors are all from the same probability distribution, least-squares is the so-called best unbiased linear estimator. We will not dwell on these issues for the time being.

2 Calculus on normed vector spaces

In the following, we give a simple framework for doing calculus on arbitrary normed¹ vector spaces.² A typical first year course on multivariate calcu-

¹A normed vector space \mathcal{V} is a vector space together with a map $\|\cdot\|: \mathcal{V} \rightarrow [0, \infty)$ such that

- (i) $\|v\| = 0$ if and only if $v = 0$,
- (ii) $\|\lambda v\| = |\lambda| \|v\|$ for all $\lambda \in \mathbb{R}$ and $v \in \mathcal{V}$, and
- (iii) $\|v + w\| \leq \|v\| + \|w\|$ for all $v, w \in \mathcal{V}$.

The last property is called the *triangle inequality*. When $\mathcal{V} \in \mathbb{R}^n$, we take without further discussion the *Euclidean Norm* defined via

$$\|\mathbf{v}\|^2 = \sum_{i=1}^n v_i^2 = \mathbf{v}^T \mathbf{v}$$

for $\mathbf{v} = (v_1, \dots, v_n)^T \in \mathbb{R}^n$.

²For simplicity, we only consider vector spaces over \mathbb{R} . In case that the space is infinite dimensional, we would typically assume completeness, i.e., that we have a Banach space.

lus typically considers the vector space \mathbb{R}^n only. However, looking at more general vector spaces is in some sense easier, and naturally includes differentiation with respect to matrices as well as the calculus of variations, both of which occur frequently in applications. In this basic introduction, we make no attempt to give a rigorous definition of differentiability nor do we discuss any of the analytical subtleties of doing calculus on infinite dimensional vector spaces such as function spaces.

Let \mathcal{V} and \mathcal{W} be vector spaces and $f: \mathcal{V} \rightarrow \mathcal{W}$ a continuous function. Central to the understanding of differentiation in this setting is the notion of *directional derivative*. Namely, fix $v \in \mathcal{V}$ and choose an arbitrary $\delta v \in \mathcal{V}$.³ We define the directional derivative as the instantaneous rate of change of f when changing v by one unit of δv . We write

$$df(v)[\delta v] = \lim_{h \rightarrow 0} \frac{f(v + h \delta v) - f(v)}{h} = \left. \frac{df(v + h \delta v)}{dh} \right|_{h=0}. \quad (5)$$

If df is continuous on $\mathcal{V} \times \mathcal{V}$, for each $v \in \mathcal{V}$, the map $df(v)[\cdot]: \mathcal{V} \rightarrow \mathcal{W}$, the called the *total derivative*, is linear.⁴ I.e., for fixed $v \in \mathcal{V}$,

$$df(v)[\lambda \delta v] = \lambda df(v)[\delta v], \quad (6a)$$

$$df(v)[\delta v + \delta w] = df(v)[\delta v] + df(v)[\delta w] \quad (6b)$$

for all $\delta v, \delta w \in \mathcal{V}$ and $\lambda \in \mathbb{R}$.

Practical computation The following trick often simplifies the computation of the derivative. Writing

$$\delta f \equiv df(v)[\delta v], \quad (7)$$

the δ symbol formally behaves like differentiation,⁵ so we can apply the product rule, chain rule, etc. as usual.

Example 1. On $\mathcal{V} = \mathbb{R}^n$, fix an $A \in M(n \times n)$ and define

$$f(\mathbf{v}) = \mathbf{v}^T A \mathbf{v}. \quad (8)$$

³We must issue an important warning: In typical generalizations, v and δv are taken from different sets. For example, \mathcal{V} may be an affine space $a + \mathcal{U}$. Then the “admissible variations” δv must be taken from \mathcal{U} . Or \mathcal{V} may be a manifold, then δv is an element from its tangent space $T\mathcal{V}$.

⁴In the context of infinite dimensional Banach spaces, the directional derivative is called the Gâteaux derivative and the total derivative leads to so-called Fréchet derivative. For Gâteaux differentiability to imply Fréchet differentiability, one needs, moreover, continuity of $df: \mathcal{V} \rightarrow \mathcal{L}(\mathcal{V}, \mathcal{W})$.

⁵In fact, with only a change of language we could call this a differential.

Then, using the product rule,

$$\delta f = (\delta \mathbf{v})^T A \mathbf{v} + \mathbf{v}^T A \delta \mathbf{v} = \mathbf{v}^T A^T \delta \mathbf{v} + \mathbf{v}^T A \delta \mathbf{v} \quad (9)$$

so that $df(\mathbf{v})[\delta \mathbf{v}] = \mathbf{v}^T (A + A^T) \delta \mathbf{v}$ or $df(\mathbf{v}) = \mathbf{v}^T (A + A^T)$. You probably know that the derivative of a real-valued function on \mathbb{R}^n can be computed via its Jacobian, which here is the $1 \times n$ matrix

$$df(\mathbf{v}) = (\partial_1 f, \dots, \partial_n f). \quad (10)$$

Sure enough, using this formula, you'll find the same answer, but with a lot more pain. Try it!

Example 2. Let

$$\mathcal{V} = \{\phi \in C^2([0, 1]): \phi(0) = \phi(1) = 0\} \quad (11)$$

and fix $g \in C([0, 1])$. Define

$$f(\phi) = \int_0^1 \left(\frac{1}{2} (\phi'(x))^2 + g(x) \phi(x) \right) dx. \quad (12)$$

Then

$$\delta f = \int_0^1 (\phi'(x) \delta \phi'(x) + g(x) \delta \phi(x)) dx = \int_0^1 (-\phi''(x) + g(x)) \delta \phi(x) dx \quad (13)$$

where, in the last equality, we have used integration by parts noting that the boundary terms vanish due to the definition of \mathcal{V} . Thus, (13) is an expression for $df(\phi)[\delta \phi]$; we note that there is no elementary way to separate out $\delta \phi$ in this expression.

Extrema We now specifically consider maps $f: \mathcal{V} \rightarrow \mathbb{R}$, as in the examples above, and ask for conditions under which such maps assume local minima or maxima. Throughout, we assume that $f \in C^1(\mathcal{V}, \mathbb{R})$, i.e., f is at least once continuously differentiable. In this case, $df(v)$ is a linear functional (i.e., a linear map from \mathcal{V} to \mathbb{R}) called the *gradient* of f . We say that f has a local maximum at $v \in \mathcal{V}$ if there is $\varepsilon > 0$ such that $f(v) \geq f(w)$ for all $\|v - w\| < \varepsilon$; f has a local minimum if $-f$ has a local maximum. Then we have the following necessary condition for the existence of a local extreme point:

f has a local minimum or maximum at $v \in \mathcal{V}$ only if V is a *stationary point* of f , i.e., if $df(v) = 0$.

The proof is a direct extension of the proof of the “first derivative test” from single-variable calculus and shall not be given here.

Note that if f is only defined on a domain $D \subsetneq V$, then it may assume a local or global extremum on the boundary (with obvious modification to the definition of the local extremum). This case may have to be considered separately. The technique of Lagrange multipliers below can be useful if the boundary can be described by a constraint.

Lagrange Multipliers Sometimes we need to find extreme values of f where v is subject to additional constraints. In the simplest case, assume that the constraint is given by the equation $g(v) = 0$ for some $g \in C^1(\mathcal{V}, \mathbb{R})$. We have the following necessary condition.

f has a local minimum or maximum on the set $\{v \in V : g(v) = 0\}$ only if there exists $\lambda \in \mathbb{R}$ such that

$$df(v) = \lambda dg(v). \tag{14}$$

We will not give a proof, but shall give the following geometric motivation. For the purposes of discussion, suppose that $\mathcal{V} = \mathbb{R}^n$. First, we notice that the gradient of a function is a vector pointing in the direction of steepest ascent.⁶ Thus, if we vary \mathbf{v} in the constraint set $\{v \in V : g(v) = 0\}$ in the direction $\delta\mathbf{v} = dg(\mathbf{v})^T$ leave the constraint set. On the other hand, any variation $\delta\mathbf{v} \perp dg(\mathbf{v})^T$ is tangent to the constraint set and therefore permitted. Thus, when “probing” f for a local extremum, we can allow f to ascent or decent in the “forbidden” direction $dg(\mathbf{v})^T$. On the other hand, when f ’s steepest ascent direction $df(\mathbf{v})$ aligns with the “forbidden” direction $dg(\mathbf{v})^T$, then the instantaneous rate of change of f is zero in any direction tangent to the constraint set. This alignment is represented by the Lagrange multiplier condition (14).

This argument also points out how to generalize to K constraints: Each constraint function g_1, \dots, g_K induces a “forbidden” direction $dg_i(\mathbf{v})^T$ which must be excluded from the “search” when probing f for an extremum. Thus, the ascent direction of f may lie in the “forbidden” space spanned by all

⁶To see this, recall that in Euclidean geometry,

$$\mathbf{u}^T \mathbf{w} = \|\mathbf{u}\| \|\mathbf{w}\| \cos \angle(\mathbf{u}, \mathbf{v}).$$

Thus, if $\mathbf{u} = df(\mathbf{v})$, the instantaneous rate of change in the direction of $\delta\mathbf{v}$ can be written $\mathbf{u}^T \delta\mathbf{v}$ and is therefore maximal among all vectors of length $\|\mathbf{u}\|$ if $\delta\mathbf{v} = \mathbf{u}$.

the $dg_i(\mathbf{v})^T$. Thus, a necessary condition for a local extremum under K constraints is the existence of a Lagrange multiplier $\boldsymbol{\lambda} \in \mathbb{R}^K$ such that

$$df(v) = \boldsymbol{\lambda}^T dg(v). \quad (15)$$

3 General linear least squares

Recall that in Section 1, the model for our data is the linear equation

$$y = ax + b. \quad (16)$$

We generalize as follows. Assume $\mathbf{x} \in \mathbb{R}^m$ with $m \in \mathbb{N}$ while we still take $y \in \mathbb{R}$. (The situation discussed here easily generalizes to vector-valued “output data” y by considering each of its components separately.) So our data is again a set of tuples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$. A *linear model* for our data is a linear combination with unknown coefficients of a fixed finite set of possibly nonlinear “ansatz functions” $h_i: \mathbb{R}^m \rightarrow \mathbb{R}$ for $i = 1, \dots, k$. I.e., we seek to find $\mathbf{v} \in \mathbb{R}^k$ such that

$$y = \mathbf{h}(\mathbf{x})^T \mathbf{v} \quad (17)$$

with $\mathbf{h}(\mathbf{x}) = (h_1(\mathbf{x}), \dots, h_k(\mathbf{x}))^T$ reproduces the data with minimal mean-square error

$$f(\mathbf{v}) = \sum_{i=1}^n (\mathbf{h}(\mathbf{x}_i)^T \mathbf{v} - y_i)^2. \quad (18)$$

Setting

$$H = \begin{pmatrix} h_1(\mathbf{x}_1) & \dots & h_k(\mathbf{x}_1) \\ \vdots & & \vdots \\ h_1(\mathbf{x}_n) & \dots & h_k(\mathbf{x}_n) \end{pmatrix} \quad \text{and} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, \quad (19)$$

we can interpret this problem as the minimization of the so-called residual $H\mathbf{v} - \mathbf{y}$ corresponding to the potentially inconsistent linear system $H\mathbf{v} = \mathbf{y}$. (In the following, nothing depends on the fact that these equations come from a general linear fitting problem; we are generally asking for the least-square solution of some inconsistent system of linear equations.) In matrix notation, we ask for the minimum of

$$f(\mathbf{v}) = \|H\mathbf{v} - \mathbf{y}\|^2 = (H\mathbf{v} - \mathbf{y})^T (H\mathbf{v} - \mathbf{y}). \quad (20)$$

The necessary condition for a minimum of f is

$$0 = \delta f = (H\delta\mathbf{v})^T(H\mathbf{v} - \mathbf{y}) + (H\mathbf{v} - \mathbf{y})^T H\delta\mathbf{v} = 2(H\mathbf{v} - \mathbf{y})^T H\delta\mathbf{v} \quad (21)$$

for any $\delta\mathbf{v} \in \mathbb{R}^k$, which implies $(H\mathbf{v})^T H = \mathbf{y}^T H$ or

$$(\mathbf{y} - H\mathbf{v})^T H = \mathbf{0}. \quad (22)$$

Geometric interpretation In the language of Linear Algebra, equation (22) can be seen as an orthogonality condition, namely⁷

$$\mathbf{y} - H\mathbf{v} \perp \text{Range } H, \quad (23)$$

i.e., the process of minimizing $f(\mathbf{v})$ yields a decomposition

$$\mathbf{y} = \underbrace{\mathbf{y} - H\mathbf{v}}_{\perp \text{Range } H} + \underbrace{H\mathbf{v}}_{\in \text{Range } H}. \quad (24)$$

Thus, the point $H\mathbf{v}$ is the point on the hyperplane $\text{Range } H$ of closest Euclidean distance to \mathbf{y} .

Solvability In practical terms, we solve (22) by solving the linear system

$$H^T H\mathbf{v} = H^T \mathbf{y}, \quad (25)$$

called the *normal equations*. It is an easy exercise in Linear Algebra to verify that $\text{Range } H^T = \text{Range } H^T H$, therefore (25) has a solution \mathbf{v} for any $\mathbf{y} \in \mathbb{R}^n$.

In the typical case, the $k \times k$ matrix $H^T H$ is invertible, so that we can write

$$\mathbf{v} = (H^T H)^{-1} H^T \mathbf{y}. \quad (26)$$

This equation can be solved directly. (For big problems, matrix inversion is expensive and should be replaced by more efficient computational methods, for details see [1].)

Example 3. We can recover the result of Section 1 by setting $m = 1$, $k = 2$, $h_1(x) = x$, $h_2(x) = 1$, and $\mathbf{v} = (a, b)^T$. Equation (26) is thus a compact way of writing (4b) and (4a).

⁷We define, as usual,

$$\text{Range } H = \{H\mathbf{u} : \mathbf{u} \in \mathbb{R}^k\}$$

and

$$\text{Ker } H = \{\mathbf{u} \in \mathbb{R}^k : H\mathbf{u} = \mathbf{0}\}.$$

Example 4. Suppose we would like to fit a polynomial of degree $k-1$ to a set of data points taken at integer points $x_1 = 0, \dots, x_n = n-1$. If we choose as ansatz functions $h_1(x) = 1, h_2(x) = x, \dots, h_k(x) = x^{k-1}$, the matrix H reads

$$H = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 2^2 & \cdots & 2^{k-1} \\ \vdots & & & & \vdots \\ 1 & n-1 & (n-1)^2 & \cdots & (n-1)^{k-1} \end{pmatrix} \quad (27)$$

This matrix is known as the *Vandermonde matrix*. It is known to be poorly conditioned [to be discussed later in the course] for even moderate values of k ; as a result, polynomial least squares with a degree higher than about 3–6 are almost never a good idea.

When $H^T H$ is not invertible, i.e., when $\text{Ker } H \neq \{\mathbf{0}\}$, there is linear dependence in the data points or in the ansatz functions. In this case, the effective number of independent observations is less than k . In this case, we may wish to compute a least-norm solution of (25), discussed in the next section.

4 Least norm solutions

When a system of linear equations

$$A\mathbf{v} = \mathbf{b}, \quad (28)$$

where $A \in M(n \times k)$ is consistent, but does not have a unique solution (we can express this situation in technical terms by saying that $\mathbf{b} \in \text{Range } A$ and $\text{Ker } A \neq \{\mathbf{0}\}$), we may ask for a solution \mathbf{v} which has the least norm amongst the family of solutions. (In the context of the general linear least squares problem of Section 3, $A = H^H$ and $\mathbf{b} = H^T \mathbf{y}$.)

Geometric interpretation If \mathbf{w} is any solution of $A\mathbf{w} = \mathbf{b}$, then it can be decomposed as

$$\mathbf{w} = \underbrace{\mathbf{w} - \mathbf{v}}_{\in \text{Ker } A} + \underbrace{\mathbf{v}}_{\perp \text{Ker } A}, \quad (29)$$

where \mathbf{v} is the minimum norm solution. Thus, $\|\mathbf{v}\|$ is the distance of \mathbf{w} to the hyperplane $\text{Ker } A$. This decomposition is the domain-side analog to the range-side decomposition (24) that comes up when computing least-square solutions.

Formulation as constrained minimization The least norm problem can be formulated as the following constrained minimization problem.

Find \mathbf{v} minimizing $f(\mathbf{v}) = \mathbf{v}^T \mathbf{v}$ subject to $\mathbf{g}(\mathbf{v}) = A\mathbf{v} - \mathbf{b} = \mathbf{0}$.

Since $df(\mathbf{v}) = 2\mathbf{v}^T$ and $d\mathbf{g}(\mathbf{v}) = A$, the method of Lagrange multipliers (15) yields the necessary condition $2\mathbf{v}^T = \boldsymbol{\lambda}^T A$ or, setting $\boldsymbol{\mu} = -2\boldsymbol{\lambda}$,

$$I\mathbf{v} + A^T \boldsymbol{\mu} = \mathbf{0}. \quad (30)$$

Together with the constraint $A\mathbf{v} = \mathbf{b}$, we can write the necessary condition in block-matrix form,

$$\begin{pmatrix} I & A^T \\ A & 0 \end{pmatrix} \begin{pmatrix} \mathbf{v} \\ \boldsymbol{\mu} \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ \mathbf{b} \end{pmatrix}. \quad (31)$$

It is easy to check that this system is consistent as long as (28) is consistent. Moreover,

$$\text{Ker} \begin{pmatrix} I & A^T \\ A & 0 \end{pmatrix} = \left\{ \begin{pmatrix} \mathbf{0} \\ \mathbf{u} \end{pmatrix} \in \mathbb{R}^{k+n} : \mathbf{u} \in \text{Ker } A^T \right\}. \quad (32)$$

In other words, the solution of (31) exists and, if not unique, the non-uniqueness affects only the Lagrange multiplier $\boldsymbol{\mu}$; the least-norm vector \mathbf{v} is uniquely determined.

5 Singular value decomposition (SVD)

Recall from your linear algebra class that if $S \in M(k \times k)$ is a symmetric matrix, it has an orthogonal diagonalization. More precisely, there exists an orthogonal matrix⁸ $V \in M(k \times k)$ such that

$$D = V^T S V \quad (33)$$

is diagonal. The orthogonal diagonalizability of symmetric matrices can be used to construct a decomposition of *any* matrix into a product of orthogonal and diagonal matrices. We proceed as follows.

Let $A \in M(n \times k)$. Without loss of generality, we may assume that $n \geq k$. Then $S = A^T A \in M(k \times k)$ is symmetric and has an orthogonal

⁸A matrix $V \in M(k \times k)$ is orthogonal if $V^{-1} = V^T$. This is equivalent to $V^T V = I$ (the columns of V are an orthogonal basis of the column space) which is equivalent to $V V^T = I$ (the rows of V are an orthogonal basis of the row space of V).

diagonalization of the form (33). Moreover, S is positive semidefinite.⁹ By permuting the columns of V , we can always achieve that the diagonal entries d_{ii} of D appear in decreasing order so that, in particular, $d_{11}, \dots, d_{mm} > 0$ and $d_{m+1,m+1}, \dots, d_{kk} = 0$ for some $m \leq k$. We now define

$$\Sigma_0 = \sqrt{D} \equiv \begin{pmatrix} \sqrt{d_{11}} & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ & & \sqrt{d_{mm}} & \\ \vdots & & 0 & \ddots \\ 0 & \dots & & 0 \end{pmatrix} \quad (34)$$

and¹⁰

$$\Sigma_0^\dagger = \begin{pmatrix} d_{11}^{-1/2} & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ & & d_{mm}^{-1/2} & \\ \vdots & & 0 & \ddots \\ 0 & \dots & & 0 \end{pmatrix}. \quad (35)$$

and

$$U_0 = AV\Sigma_0^\dagger. \quad (36)$$

First, we note that the first m columns of U_0 are an orthonormal set in the column space, as

$$U_0^T U_0 = \Sigma_0^\dagger V^T A^T AV \Sigma_0^\dagger = \Sigma_0^\dagger V^T S V \Sigma_0^\dagger = \Sigma_0^\dagger D \Sigma_0^\dagger = \begin{pmatrix} I_m & 0 \\ 0 & 0 \end{pmatrix}. \quad (37)$$

Second,

$$U_0 \Sigma_0 V^T = AV \Sigma_0^\dagger \Sigma_0 V^T = A. \quad (38)$$

⁹A matrix $S \in M(k \times k)$ is positive definite if $\mathbf{v}^T S \mathbf{v} > 0$ for all nonzero $\mathbf{v} \in \mathbb{R}^k$; it is positive semi-definite if $\mathbf{v}^T S \mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^k$. It is easy to see that all eigenvalues of a positive definite matrix must be positive; all eigenvalues of a positive semidefinite matrix must be nonnegative.

¹⁰ Σ_0^\dagger is called the *pseudo-inverse* of Σ_0 . It is characterized by

$$\Sigma_0^\dagger \Sigma_0 = \begin{pmatrix} I_m & 0 \\ 0 & 0 \end{pmatrix},$$

where I_m denotes the $m \times m$ identity matrix.

The last identity is obvious if A is nonsingular so that $m = k$ and $\Sigma_0^\dagger \Sigma_0 = I$. In general, note that $A^T A = V D V^T$ implies

$$\text{Ker } A^\perp = \text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_m\} \quad (39a)$$

and

$$\text{Ker } A = \text{Span}\{\mathbf{v}_{m+1}, \dots, \mathbf{v}_k\}, \quad (39b)$$

where $\mathbf{v}_1, \dots, \mathbf{v}_k$ denote the columns of V . The last step of (38) is then easily verified by restricting to $\text{Ker } A$ and $\text{Ker } A^\perp$, respectively. Equation (38) is referred to as the *restricted singular value decomposition* of A .

Now let $\mathbf{u}_1, \dots, \mathbf{u}_m$ denote the first m columns of U_0 . Note that (37) shows that

$$\text{Range } A = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\}. \quad (39c)$$

It is often convenient to extend the orthonormal set $\mathbf{u}_1, \dots, \mathbf{u}_m$ to an orthonormal basis $\mathbf{u}_1, \dots, \mathbf{u}_n$ of \mathbb{R}^n , so that

$$\text{Range } A^\perp = \text{Span}\{\mathbf{u}_{m+1}, \dots, \mathbf{u}_n\}. \quad (39d)$$

Then, setting

$$U = \begin{pmatrix} | & & | \\ \mathbf{u}_1 & \cdots & \mathbf{u}_n \\ | & & | \end{pmatrix} \in M(n \times n) \quad \text{and} \quad \Sigma = \begin{pmatrix} \Sigma_0 \\ 0 \end{pmatrix} \in M(n \times k), \quad (40)$$

we can write

$$A = U \Sigma V^T \quad (41)$$

This is the *singular value decomposition* of A ; the diagonal entries of Σ , denoted $\sigma_1, \dots, \sigma_k$, are the *singular values* of A , the vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ are the *right singular vectors*, and the vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$ are the *left singular vectors*.

Variational characterization Recall, from Homework 1, that the principal unit eigenvector of a symmetric matrix $S \in M(k \times k)$ maximizes the function

$$f(\mathbf{v}) = \mathbf{v}^T S \mathbf{v} \quad (42)$$

among all unit vectors $\mathbf{v} \in \mathbb{R}^k$. The singular value decomposition of a matrix $A \in M(n \times k)$ has a similar variational characterization. Consider the function

$$f(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T A \mathbf{v}. \quad (43)$$

Then the maximum of f over the set of unit vectors $\mathbf{u} \in \mathbb{R}^n$ and unit vectors $\mathbf{v} \in \mathbb{R}^k$ is the largest singular value of A ; vectors \mathbf{u} and \mathbf{v} at which the maximum is attained can be taken as the first left singular vector \mathbf{u}_1 and the first right singular vector \mathbf{v}_1 , respectively, in the notation of (39).

To establish our claim, we proceed as follows. Our task can be formulated as an optimization problem for f on \mathbb{R}^{n+k} with the two scalar constraints

$$g_1(\mathbf{u}, \mathbf{v}) \equiv \mathbf{u}^T \mathbf{u} - 1 = 0 \quad \text{and} \quad g_2(\mathbf{u}, \mathbf{v}) \equiv \mathbf{v}^T \mathbf{v} - 1 = 0. \quad (44)$$

Clearly,

$$df(\mathbf{u}, \mathbf{v})[\delta\mathbf{u}, \delta\mathbf{v}] = \mathbf{v}^T A^T \delta\mathbf{u} + \mathbf{u}^T A \delta\mathbf{v}, \quad (45a)$$

$$dg_1(\mathbf{u}, \mathbf{v})[\delta\mathbf{u}, \delta\mathbf{v}] = 2 \mathbf{u}^T \delta\mathbf{u}, \quad (45b)$$

and

$$dg_2(\mathbf{u}, \mathbf{v})[\delta\mathbf{u}, \delta\mathbf{v}] = 2 \mathbf{v}^T \delta\mathbf{v}. \quad (45c)$$

Then, the necessary condition for a local maximum is given by the Lagrange multiplier principle (15) which reads, after substituting in (45) and rearranging terms,

$$(\mathbf{v}^T A^T \delta\mathbf{u} - 2 \lambda_1 \mathbf{u}^T \delta\mathbf{u}) + (\mathbf{u}^T A \delta\mathbf{v} - 2 \lambda_2 \mathbf{v}^T \delta\mathbf{v}) = 0 \quad (46)$$

which, since $\delta\mathbf{u}$ and $\delta\mathbf{v}$ can be chosen independently, implies

$$\mathbf{v}^T A^T = 2 \lambda_1 \mathbf{u}^T \quad \text{and} \quad \mathbf{u}^T A = 2 \lambda_2 \mathbf{v}^T. \quad (47)$$

Right-multiplying the first equality with \mathbf{u} and the second equality with \mathbf{v} , we conclude that

$$\mathbf{u}^T A \mathbf{v} = 2 \lambda_1 = 2 \lambda_2 \equiv \sigma. \quad (48)$$

The existence of a solution is clear as we are maximizing a smooth function over a compact constraint set; hence the necessary condition must be satisfied somewhere. Denote some such solution triple \mathbf{u}_1 , \mathbf{v}_1 , and σ_1 .¹¹ Now, restricting \mathbf{u} and \mathbf{v} to the respective orthogonal complements of the vectors already established, we can iterate this argument k times to obtain a full

¹¹We are not claiming uniqueness. In fact, it can be shown that the solution is unique up to a choice of sign if the singular values of A are nondegenerate; in the degenerate case, it is clear that we cannot expect uniqueness.

set of “singular values” $\sigma_1, \dots, \sigma_k$, “right singular vectors” $\mathbf{v}_1, \dots, \mathbf{v}_k$ and, completing the orthogonal set to a basis, “left singular vectors” $\mathbf{u}_1, \dots, \mathbf{u}_n$.

Two claims remain to be verified. First, that the result of the above construction is indeed indeed a singular value decomposition. Second, that the largest singular value is indeed the global (not just some local) maximum of f as claimed. The verification of both statements is not difficult and shall be left as an exercise.

Matrix norms Given any norm $\|\cdot\|$ on \mathbb{R}^n , we can define the norm of a matrix $A = M(n \times k)$ as the smallest number $\|A\|$ such that

$$\|A\mathbf{x}\| \leq \|A\| \|\mathbf{x}\| \quad (49)$$

holds true. In other words,

$$\|A\| = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|}. \quad (50)$$

This definition generalizes to operators between Banach spaces and is often referred to as the *operator norm*. Note that, if the coefficients of A are denoted a_{ij} ,

$$\|A\|_F^2 = \sum_{\substack{i \in \{1, \dots, n\} \\ j \in \{1, \dots, k\}}} |a_{ij}|^2 \quad (51)$$

defines a norm, called the *Frobenius norm*, which is is not an operator norm.

Consider, in particular, the operator norm induced by the Euclidean norm on \mathbb{R}^n , which we generally assume in these notes. Then $\|\mathbf{x}\|^2 = \mathbf{x}^T \mathbf{x}$, so that, using the singular value decomposition,

$$\begin{aligned} \|A\|^2 &= \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T V \Sigma^T U^T U \Sigma V^T \mathbf{x}}{\mathbf{x}^T \mathbf{x}} \\ &= \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T V \Sigma^T \Sigma V^T \mathbf{x}}{\mathbf{x}^T V V^T \mathbf{x}} \\ &= \max_{\mathbf{y} \neq \mathbf{0}} \frac{\mathbf{y}^T \Sigma^T \Sigma \mathbf{y}}{\mathbf{y}^T \mathbf{y}} \\ &= \sigma_1^2 \end{aligned} \quad (52)$$

In other words, the matrix norm induced by the Euclidean metric equals the largest singular value of the matrix.

In contrast, it is not difficult to show that

$$\|A\|_F^2 = \sum_{i=1}^k \sigma_i^2. \quad (53)$$

Condition numbers Suppose, for simplicity, that $A \in M(n \times n)$ is a non-singular matrix. We would like to characterize the stability of the solution \mathbf{x} to the system of linear equations

$$A\mathbf{x} = \mathbf{b} \tag{54}$$

under perturbation of the right hand side. By linearity, $\Delta\mathbf{b} = A\Delta\mathbf{x}$, where $\Delta\mathbf{b}$ denotes an *absolute* change of the vector \mathbf{b} , and $\Delta\mathbf{x}$ the corresponding absolute change in \mathbf{x} . Thus, by (49), we obtain $\|\Delta\mathbf{x}\| \leq \|A^{-1}\| \|\Delta\mathbf{b}\|$ and $\|\mathbf{b}\| \leq \|A\| \|\mathbf{x}\|$, so that

$$\frac{\|\Delta\mathbf{x}\|}{\|\mathbf{x}\|} \leq \|A\| \|A^{-1}\| \frac{\|\Delta\mathbf{b}\|}{\|\mathbf{b}\|}. \tag{55}$$

The worst-case amplification factor from a *relative* change in \mathbf{b} to the relative change in \mathbf{x} ,

$$\kappa(A) = \|A\| \|A^{-1}\|, \tag{56}$$

is called the *condition number* of A . Due to (52),

$$\kappa(A) = \frac{\sigma_1}{\sigma_n}. \tag{57}$$

This notion of condition number generalizes naturally to non-square matrices A of full rank. When the condition number is large, the solution \mathbf{x} becomes very sensitive to error (e.g. measurement error) in the data \mathbf{b} ; we speak of an *ill-conditioned* problem.

6 Approximate solution of ill-conditioned linear equations

Let us consider the following prototypical version of an ill-conditioned system of linear equations. Our presentation closely follows [2]. Let $\mathbf{x} \in \mathbb{R}^n$ represent the true pixel values of a, for simplicity, one-dimensional image. The image is taken by a device which, due to lens imperfections, say, records a blurred set of pixel values $\mathbf{b} \in \mathbb{R}^n$. We assume a simple linear model for the device imperfections, namely that

$$b_i = \sum_{j=1}^n a_{i-j} x_j, \tag{58}$$

where b_1, \dots, b_n and x_1, \dots, x_n denote the components of \mathbf{b} and \mathbf{x} , respectively. This type of expression is referred to as a *discrete convolution* with *kernel* a_k ; typical is a bell-like shape of its graph, e.g.,

$$a_k = c e^{-\gamma k^2} \quad (59)$$

for some given constants c and γ . Equation (58) can be written as a system of linear equations

$$\mathbf{b} = A\mathbf{x} \quad (60)$$

with $n \times n$ matrix

$$A = \begin{pmatrix} a_0 & a_{-1} & \cdots & & & \\ a_1 & \ddots & \ddots & & & \\ \vdots & \ddots & & \ddots & \vdots & \\ & & \ddots & \ddots & a_{-1} & \\ & & \cdots & a_1 & a_0 & \end{pmatrix}. \quad (61)$$

Matrices of this type are known as *Toeplitz matrices*. They are typically very ill-conditioned¹², so that, even in the case that A is non-singular, the direct solution of (60) may be practically meaningless.

Regularization The ill conditioning of (60) can be understood in terms of the singular values of A as follows. Suppose \mathbf{x}_{true} refers to the true pixel values of the image; due to measurement error,

$$A\mathbf{x}_{\text{true}} = \mathbf{b} + \boldsymbol{\eta}, \quad (62)$$

where $\boldsymbol{\eta}$ is the measurement error vector with norm $\|\boldsymbol{\eta}\| > 0$. For simplicity, we suppose that $A \in M(n \times n)$ is nonsingular. Thus, when attempting to compute the pixel values by direct inversion of AS , we obtain

$$\mathbf{x}_{\text{naive}} = A^{-1}\mathbf{b} = \mathbf{x}_{\text{true}} + V\Sigma^{-1}U^T\boldsymbol{\eta}. \quad (63)$$

¹²In fact, (58) can be interpreted as a discrete approximation of the integral equation

$$b(x) = \int_a^b k(x-y)f(y)dy$$

where, with a smooth kernel function k , the right hand side defines a compact, hence non-invertible operator on a Hilbert space. Such problem are referred to as *ill-posed*. Direct discretization of an ill-posed continuum problem always leads to an ill-conditioned finite dimensional system. In this sense, the ill-conditioning in our example is a generic phenomenon.

The ill-conditioning manifests itself in the occurrence of small singular values for A , so that Σ^{-1} in (63) has small denominators—some components of $\boldsymbol{\eta}$ with respect to the orthonormal basis $\{\boldsymbol{v}_i\}$ are being multiplied with large gains. As a result, we cannot expect $\boldsymbol{x}_{\text{naive}}$ and $\boldsymbol{x}_{\text{true}}$ to be close.

Since the problem is due to small denominators, let us define a filtered solution

$$\boldsymbol{x}_\alpha = V\Sigma_\alpha^{-1}U^T\boldsymbol{b} \quad (64)$$

where Σ_α^{-1} is the filtered inverse of Σ , defined by

$$\Sigma_\alpha^{-1} = \text{diag}\{w_\alpha(\sigma_i^2)\sigma_i^{-1} : i = 1, \dots, n\} \quad (65)$$

with $w(\sigma^2)$ converging to zero at least linearly as $\sigma \rightarrow 0$ to avoid unbounded singular values of the filtered inverse in this limit. Among the possible choices for w_α , the following are particularly natural.

Truncation The simplest choice is the cut-off filter

$$w_\alpha(\sigma^2) = \begin{cases} 1 & \text{if } \sigma^2 \geq \alpha \\ 0 & \text{if } \sigma^2 < \alpha. \end{cases} \quad (66)$$

Tychonov regularization With the choice

$$w_\alpha(\sigma^2) = \frac{\sigma^2}{\sigma^2 + \alpha}, \quad (67)$$

we obtain

$$w_\alpha(\sigma^2)\sigma^{-1} = \frac{\sigma}{\sigma^2 + \alpha} \quad (68)$$

or

$$\Sigma_\alpha^{-1} = (\Sigma^2 + \alpha I)^{-1}\Sigma, \quad (69)$$

thereby providing a smooth transition between the filtered and unfiltered regime. Substituting this expression into (64), we obtain

$$\begin{aligned} \boldsymbol{x}_\alpha &= V(\Sigma^2 + \alpha I)^{-1}\Sigma U^T\boldsymbol{b} \\ &= V(\Sigma^2 + \alpha I)^{-1}V^T V\Sigma U^T\boldsymbol{b} \\ &= (V(\Sigma^2 + \alpha I)V^T)^{-1}V\Sigma U^T\boldsymbol{b} \\ &= (A^T A + \alpha I)^{-1}A^T\boldsymbol{b}. \end{aligned} \quad (70)$$

Note that this expression formally resembles the normal equations (26) of the linear least squares problem, suggesting a close relationship between the two problems. This expectation will be substantiated later when we describe a variational formulation of the Tychonov regularization.

Error analysis We provide a deterministic error analysis, i.e., we assume that $\boldsymbol{\eta}$ is a fixed measurement error vector of known magnitude. We can then define

$$\begin{aligned}
\mathbf{e}_\alpha &= \mathbf{x}_{\text{true}} - \mathbf{x}_\alpha \\
&= \mathbf{x}_{\text{true}} - V\Sigma_\alpha^{-1}U^T\mathbf{b} \\
&= VIV^T\mathbf{x}_{\text{true}} - V\Sigma_\alpha^{-1}U^T(A\mathbf{x}_{\text{true}} - \boldsymbol{\eta}) \\
&= V(I - \Sigma_\alpha^{-1}\Sigma)V^T\mathbf{x}_{\text{true}} - V\Sigma_\alpha^{-1}U^T\boldsymbol{\eta} \\
&\equiv \mathbf{e}_{\text{trunc}} + \mathbf{e}_{\text{noise}}
\end{aligned} \tag{71}$$

Note that, in contrast to the noise term in (63), the expression for $\mathbf{e}_{\text{noise}}$ involves a filtered inverse, thus has bounded gains. We expect that $\|\mathbf{e}_{\text{noise}}\|$ decreases with increasing filter parameter α , while the truncation error $\|\mathbf{e}_{\text{trunc}}\|$ increases; thus, there should be an optimal filter parameter α . This is borne out by the following estimation.

First, due to (49) and the fact that the Euclidean matrix norm of an orthogonal matrix equals 1,

$$\|\mathbf{e}_{\text{noise}}\| \leq \|V\| \|\Sigma_\alpha^{-1}\| \|U^T\| \|\boldsymbol{\eta}\| \leq \max_{\sigma>0} \frac{w_\alpha(\sigma^2)}{\sigma} \|\boldsymbol{\eta}\| \leq \frac{1}{\sqrt{\alpha}} \|\boldsymbol{\eta}\|. \tag{72}$$

The last inequality above can be verified independently for the truncation filter and for the Tychonov regularization.

Second, we seek to derive estimates on the truncation error. Here we are faced with a fundamental dilemma. In order to conclude that truncation error—the error committed by filtering—is small, we need to ensure that the information which gets removed by the filter is in some sense unimportant for the true solution. In other words, we need to provide additional *a priori* information on the problem.¹³ In a standard analysis of this problem, one assumes the *range condition*¹⁴

$$\mathbf{x}_{\text{true}} = A^T\mathbf{z} \tag{73}$$

¹³In fact, in the absence of further information, we can still conclude, directly from the definition of the w_α considered, that $\|\mathbf{e}_{\text{trunc}}\| \rightarrow 0$ as $\alpha \rightarrow 0$. This conclusion, however, is practically meaningless because we have no control about the rate of convergence, so there is no way to ensure that we are doing substantially better than the naive, direct inversion (63).

¹⁴In an infinite dimensional Hilbert space setting, the range condition can be stated concisely as

$$\mathbf{x}_{\text{true}} \in \text{Range } A^T.$$

There, this condition is nontrivial.

such that $\|\mathbf{z}\|$ is “reasonable,” meaning of the same order of magnitude as $\|\mathbf{x}_{\text{true}}\|$. It is important to realize that this condition, from the mathematical point of view, is arbitrary. Rather, it represents additional knowledge about the specific problem providing additional restrictions on the the class of vectors in which we seek a solution. As such, it must be justified on physical grounds which necessarily lie outside of the analysis presented here.

Proceeding with our estimation, we obtain, inserting the range condition (73) into the expression for $\mathbf{e}_{\text{trunc}}$,

$$\begin{aligned}\|\mathbf{e}_{\text{trunc}}\|^2 &= \mathbf{e}_{\text{trunc}}^T \mathbf{e}_{\text{trunc}} = \mathbf{z}^T AV(I - \Sigma_\alpha^{-1}\Sigma)V^T V(I - \Sigma_\alpha^{-1}\Sigma)V^T A^T \mathbf{z} \\ &= \mathbf{z}^T U (I - \Sigma_\alpha^{-1}\Sigma)^2 \Sigma^2 U^T \mathbf{z} \\ &\leq \max_{\sigma>0}(1 - w_\alpha(\sigma^2))^2 \sigma^2 \|\mathbf{z}\|^2 \\ &\leq \alpha \|\mathbf{z}\|^2.\end{aligned}\tag{74}$$

Inserting the bounds for $\mathbf{e}_{\text{trunc}}$ and $\mathbf{e}_{\text{noise}}$ into (71), we obtain

$$\|\mathbf{e}_\alpha\| \leq \sqrt{\alpha} \|\mathbf{z}\| + \frac{1}{\sqrt{\alpha}} \|\boldsymbol{\eta}\|.\tag{75}$$

It is easy to show that the right hand side is minimized when the two term balance, i.e., when

$$\alpha = \frac{\|\boldsymbol{\eta}\|}{\|\mathbf{z}\|}\tag{76}$$

so that

$$\|\mathbf{e}_\alpha\| = 2\sqrt{\|\boldsymbol{\eta}\| \|\mathbf{z}\|}.\tag{77}$$

We call this an *a priori* estimate for the regularization parameter since knowledge of the true solution \mathbf{x}_{true} is required. In the following, we discuss methods of how to choose the regularization parameter in an optimal way knowing only the data.

Morozov discrepancy principle In contrast to what was done in (71), we can split the error into two components as follows,¹⁵

$$\mathbf{e}_\alpha^* \equiv A(\mathbf{x}_{\text{true}} - \mathbf{x}_\alpha) = (\mathbf{b} - A\mathbf{x}_\alpha) + \boldsymbol{\eta} \equiv \mathbf{e}_{\text{trunc}}^* + \mathbf{e}_{\text{noise}}^*.\tag{78}$$

¹⁵The definition $\mathbf{e}_\alpha^* = A\mathbf{e}_\alpha$ seems arbitrary, but is motivated by the need to avoid any inversion of the nonregularized A as this computation generally cannot be performed in a stable manner.

Since $\mathbf{e}_{\text{trunc}}^*$ now only depends on the measured data \mathbf{b} , we can compute an *a posteriori* estimate of α by balancing the norms of $\mathbf{e}_{\text{trunc}}^*$ and $\mathbf{e}_{\text{noise}}^*$ rather than those of $\mathbf{e}_{\text{trunc}}$ and $\mathbf{e}_{\text{noise}}$ as in (75). Notice that

$$\mathbf{e}_{\text{trunc}}^* = U(I - \Sigma\Sigma_\alpha^{-1})U^T\mathbf{b} \quad (79)$$

so that

$$\|\mathbf{e}_{\text{trunc}}^*\|^2 = \mathbf{b}^T U (I - \Sigma\Sigma_\alpha^{-1})^2 U^T \mathbf{b} \rightarrow \begin{cases} 0 & \text{as } \alpha \rightarrow 0 \\ \mathbf{b}^T \mathbf{b} & \text{as } \alpha \rightarrow \infty; \end{cases} \quad (80)$$

the left hand expression is, moreover, smooth and monotonic. Thus, whenever $\|\mathbf{b}\| \geq \|\boldsymbol{\eta}\|$, the intermediate value theorem guarantees the existence of an α_{opt} such that $\|\mathbf{e}_{\text{trunc}}^*(\alpha_{\text{opt}})\| = \|\boldsymbol{\eta}\|$. This regularization parameter selection criterion is called the *Morozov discrepancy principle*.

Error analysis of the Morozov principle We now ask the question how the error \mathbf{e}_α behaves under the Morozov discrepancy principle. This question is only easy to answer in the case of the Tychonov regularization where \mathbf{x}_α is known to minimize the function

$$f_\alpha(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \alpha \|\mathbf{x}\|^2 \quad (81)$$

over all $\mathbf{x} \in \mathbb{R}^n$; see Homework 2. Thus, in particular,

$$f_\alpha(\mathbf{x}_\alpha) \leq f_\alpha(\mathbf{x}_{\text{true}}) \quad (82)$$

which we can write

$$\|\mathbf{A}\mathbf{x}_\alpha - \mathbf{b}\|^2 + \alpha \|\mathbf{x}_\alpha\|^2 \leq \|\boldsymbol{\eta}\|^2 + \alpha \|\mathbf{x}_{\text{true}}\|^2. \quad (83)$$

When α is chosen according to the Morozov principle, $\|\mathbf{A}\mathbf{x}_\alpha - \mathbf{b}\| = \|\boldsymbol{\eta}\|$, so that the first terms on left and right of this inequality drop out and we conclude that

$$\|\mathbf{x}_\alpha\| \leq \|\mathbf{x}_{\text{true}}\|. \quad (84)$$

Then,

$$\begin{aligned} \|\mathbf{e}_\alpha\|^2 &= \|\mathbf{x}_{\text{true}}\|^2 - 2\mathbf{x}_\alpha^T \mathbf{x}_{\text{true}} + \|\mathbf{x}_\alpha\|^2 \\ &\leq 2\|\mathbf{x}_{\text{true}}\|^2 - 2\mathbf{x}_\alpha^T \mathbf{x}_{\text{true}} = 2\mathbf{e}_\alpha^T \mathbf{x}_{\text{true}}. \end{aligned} \quad (85)$$

As before, we assume the range condition with $\mathbf{x}_{\text{true}} = \mathbf{A}^T \mathbf{z}$, so that

$$\|\mathbf{e}_\alpha\|^2 \leq 2\mathbf{e}_\alpha^T \mathbf{A}^T \mathbf{z} = 2(\mathbf{e}_\alpha^*)^T \mathbf{z} \leq 2\|\mathbf{e}_\alpha^*\| \|\mathbf{z}\|. \quad (86)$$

When α is chosen according to the Morozov principle, $\|e_\alpha^*\| = 2\|\boldsymbol{\eta}\|$. Inserting this estimate back into (86) and taking the square root, we find that

$$\|e_\alpha\| = 2\sqrt{\|\boldsymbol{\eta}\|\|z\|} \quad (87)$$

which coincides with the *a priori* error estimate (77).

Final Remarks The range condition is crucial for getting good reconstructions. In the case when A represents a discrete convolution or the inverse of a discrete differential operator, the range condition qualitatively expresses that $\boldsymbol{x}_{\text{true}}$ consists of samples of a smooth function on a grid. A concise discussion is beyond the scope of these notes, but the consequences of this statement are easily visible in simple computational experiments. In particular, the filtered inverse will reconstruct or reproduce step functions only poorly—the edges will get eroded.

In the variational characterization (81), the term $\alpha\|\boldsymbol{x}\|^2$ can be seen as a “penalty term” which might be replaced with very different types of penalties, not all of which might have an easy description in terms of the singular value decomposition. One important example is to replace $\|\boldsymbol{x}\|^2$ by the so-called *total variation*. The total variation penalizes oscillatory behavior while allowing jumps. The book by Vogel [2] gives more complete, yet relatively accessible introduction to theory and computational methods for inverse problems.

References

- [1] Å. Björck, “Numerical Methods for Least Squares Problems,” SIAM, Philadelphia, 1996.
- [2] C. Vogel, “Computational Methods for Inverse Problems,” SIAM, Philadelphia, 2002.