

Chapter 1

Introduction by examples

All happy families resemble one another, each unhappy family is unhappy in its own way.

(Lev Tolstoy, 1877)

Analogous to the quote, simple matrix nearness problems tend to look alike, while difficult ones are each difficult in their own distinct way. In this short introductory chapter, we present some of the matrix nearness problems that will be considered in more detail later in this book. We formulate the problems, give some background, and sketch ideas for numerical approaches that will be worked out in later chapters.

Nearest singular matrix, distance to singularity

It will be shown that if the least-squares criterion of approximation be adopted, this problem has a general solution which is relatively simple in a theoretical sense, though the amount of numerical work involved in applications may be prohibitive.

(Carl Eckart and Gale Young, 1936)

We begin with the matrix nearness problem that may be most basic and have the longest history. The problem is atypical in that its solution can be directly read off from a matrix factorization, which here is the singular value decomposition.

Problem. *Given an invertible matrix, find a nearest singular matrix.*

The term “nearest” refers to a metric given by a matrix norm. The problem can thus be rephrased as follows when general complex $n \times n$ matrices are considered: Given an invertible matrix $A \in \mathbb{C}^{n,n}$, find a perturbation $\Delta \in \mathbb{C}^{n,n}$ of minimal norm such that $A + \Delta$ is not invertible. The norm of the minimal perturbation Δ is called the *distance to singularity* of the matrix A .

The solution to the problem depends on the choice of matrix norm. Here we consider, for a matrix $M \in \mathbb{C}^{n,n}$ with singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0$, the matrix 2-norm

$$\|M\|_2 = \sup_{0 \neq x \in \mathbb{C}^n} \frac{\|Mx\|_2}{\|x\|_2} = \sigma_1$$

(where the 2-norm of vectors is the Euclidean norm) and the Frobenius norm, which is the Euclidean norm of the vector of matrix entries,

$$\|M\|_F = \left(\sum_{i,j=1}^n |m_{ij}|^2 \right)^{1/2} = \left(\sum_{i=1}^n \sigma_i^2 \right)^{1/2}.$$

For these two norms, the nearest singular matrix to the given invertible matrix $A \in \mathbb{C}^{n,n}$ is obtained directly from the singular value decomposition of A ,

$$A = \sum_{i=1}^n \sigma_i u_i v_i^*$$

(with the conjugate transpose $v^* = \bar{v}^\top$), where $u_i \in \mathbb{C}^n$ and $v_i \in \mathbb{C}^n$ are the left and right singular vectors, respectively, to the i th singular value $\sigma_i > 0$ of A . A nearest singular matrix to A is found by truncating the last term in the sum, i.e., by adding the rank-1 perturbation matrix $\Delta = -\sigma_n u_n v_n^*$ to A , and the distance to singularity thus equals the smallest singular value σ_n , both for the matrix 2-norm and the Frobenius norm. This result is the special case of rank $r = n - 1$ of the Eckart–Young theorem (1936), going back in its origins to Schmidt (1907). The theorem states that a nearest matrix to A of rank r is obtained by truncating the singular value decomposition after r terms. This holds true for both the matrix 2-norm and the Frobenius norm, and in fact for every unitarily invariant matrix norm, as shown by Mirsky (1960). For a general matrix norm induced by a vector norm, the distance to singularity equals $1/\|A^{-1}\|$; i.e., the relative distance is the reciprocal of the condition number $\text{cond}(A) = \|A\| \cdot \|A^{-1}\|$; moreover, a minimal perturbation matrix is still of rank 1. Kahan (1966) attributes this result to Gastinel; see also Wilkinson (1986).

Structured problem. While the distance to singularity with respect to a unitarily invariant matrix norm is obtained directly from the singular value decomposition for a general complex or real invertible matrix A , the situation changes when A is a *structured* invertible matrix and we ask for the nearest singular matrix with the same given linear structure, for example, having the same sparsity pattern or belonging to the class of Hamiltonian or Toeplitz or Hankel matrices. The nearest singular structured matrix cannot be determined from the singular value decomposition of A , and the structured distances to singularity are in general not the same for the matrix 2-norm and the Frobenius norm. The minimal structured perturbation Δ is no longer a rank-1 matrix in general, but for the Frobenius norm we will show in Section 5.8 that a minimal Δ is the orthogonal projection of a rank-1 matrix onto the structure.

This rank-1 property leads us to propose a two-level iterative algorithm, where we solve a system of differential equations for two vectors depending on a distance parameter in the inner iteration – which requires only matrix-vector products with

structured matrices and inner products of vectors – and solve a scalar nonlinear equation for the structured distance to singularity in the outer iteration.

In a numerical example in Section 5.8, we use this algorithm to compute the distance of a given pair of polynomials to a pair of polynomials with a common zero. This distance equals the structured distance to singularity of a Sylvester matrix, and the coefficients of the nearest pair of polynomials with a common zero can be read off from the nearest singular Sylvester matrix.

Nearest unstable matrix, distance to instability

How near is a stable matrix to an unstable matrix?

(Charles Van Loan, 1985)

Many matrix nearness problems considered in this book are motivated by dynamical systems. In particular, the question posed by Van Loan in the title of his 1985 paper has received much attention in the literature; see the Notes at the end of Chapters 3 and 4. A linear differential equation $\dot{x}(t) = Ax(t)$ is asymptotically stable – i.e., every solution $x(t)$ tends to 0 as $t \rightarrow \infty$ – if and only if all eigenvalues of the matrix A have negative real part. Van Loan's question addresses the robustness of asymptotic stability under perturbations of A .

Problem. *Given a matrix having all eigenvalues in the open complex left half-plane, find a nearest matrix having some eigenvalue on the imaginary axis.*

Various algorithms have been proposed in the literature to solve this problem for general complex matrices (see again the Notes of Chapters 3 and 4). The problem can be rephrased in terms of the ε -pseudospectrum of $A \in \mathbb{C}^{n,n}$, which for $\varepsilon > 0$ is defined as

$$\Lambda_\varepsilon(A) = \{\lambda \in \mathbb{C} : \lambda \in \Lambda(A + \Delta) \text{ for some } \Delta \in \mathbb{C}^{n,n} \text{ with } \|\Delta\| \leq \varepsilon\},$$

where $\Lambda(A)$ is the spectrum of A , that is, the set of eigenvalues of A . Equivalently, $\Lambda_\varepsilon(A)$ is the set of all complex λ for which the distance to singularity of $A - \lambda I$ is at most ε . In the following, the norm is chosen as the Frobenius norm or the matrix 2-norm. Both yield the same ε -pseudospectrum since the distance to singularity is the same for both norms.

The problem is then to find the smallest $\varepsilon > 0$ with $\Lambda_\varepsilon(A) \cap i\mathbb{R} \neq \emptyset$. This optimal ε is called the *distance to instability* or *stability radius* of A . A minimal perturbation $\Delta \in \mathbb{C}^{n,n}$ with $\Lambda(A + \Delta) \cap i\mathbb{R} \neq \emptyset$ is known to have rank 1. This fact has been used to derive numerical algorithms for the above matrix nearness problem; see Section 3.2 and the Notes of Chapters 3 and 4.

Structured problem. Consider now the same problem posed for real matrices, i.e., finding $\Delta \in \mathbb{R}^{n,n}$ of minimal norm having $\Lambda(A + \Delta) \cap i\mathbb{R} \neq \emptyset$, or, more generally, posed for perturbation matrices Δ in a complex-linear or real-linear subspace $\mathcal{S} \subset \mathbb{C}^{n,n}$ (the *structure space*), e.g., complex or real matrices with a given sparsity pattern or special matrix classes, often named after 19th-century mathematicians. The $\mathcal{S}$ -structured ε -pseudospectrum $\Lambda_{\varepsilon}^{\mathcal{S}}(A)$, where $\mathbb{C}^{n,n}$ in the definition of $\Lambda_{\varepsilon}(A)$ is replaced by the subspace $\mathcal{S}$, is usually not the same for the Frobenius norm and the matrix 2-norm. With respect to both norms, the $\mathcal{S}$ -structured distance to instability, which is the smallest $\varepsilon > 0$ such that $\Lambda_{\varepsilon}^{\mathcal{S}}(A) \cap i\mathbb{R} \neq \emptyset$, is in general strictly larger than the complex-unstructured distance to instability, including the real case $\mathcal{S} = \mathbb{R}^{n,n}$.

A minimal structured perturbation $\Delta \in \mathcal{S}$ is no longer a rank-1 matrix in general, but for the Frobenius norm we will show in Section 5.2 that a minimal $\Delta \in \mathcal{S}$ is the orthogonal projection of a rank-1 matrix onto the structure space $\mathcal{S}$. (For the real case $\mathcal{S} = \mathbb{R}^{n,n}$, this is the real part of a complex rank-1 matrix.) Based on this rank-1 property, we propose a two-level algorithm, where we solve a system of differential equations for two vectors depending on a distance parameter in the inner iteration – which requires computing eigenvalues and eigenvectors of perturbed matrices – and solve a scalar nonlinear equation for the structured distance to instability in the outer iteration.

Robustness of transient bounds

This phenomenon has traditionally been investigated by linearizing the equations of flow and testing for unstable eigenvalues of the linearized problem, but the results of such investigations agree poorly in many cases with experiments. Nevertheless, linear effects play a central role in hydrodynamic instability.

(Lloyd N. Trefethen, Anne E. Trefethen, Satish C. Reddy, and Tobin A. Driscoll, 1993)

As has been emphasized by Trefethen et al. (1993, 2005), eigenvalues of non-normal matrices A give no indication of transient bounds of solutions of linear differential equations with A , whereas pseudospectra do. This relies on the characterization of the ε -pseudospectrum in terms of resolvent norms,

$$\begin{aligned}\Lambda_{\varepsilon}(A) &= \{\lambda \in \mathbb{C} : \lambda \in \Lambda(A + \Delta) \text{ for some } \Delta \in \mathbb{C}^{n,n} \text{ with } \|\Delta\| \leq \varepsilon\} \\ &= \{\lambda \in \mathbb{C} : \|(A - \lambda I)^{-1}\|_2 \geq \varepsilon^{-1}\},\end{aligned}$$

which is valid for both the matrix 2-norm and the Frobenius norm as matrix norm $\|\cdot\|$. This fundamental equality was apparently first noted by Wilkinson (1986), who “would like to emphasize the sheer economy of this theorem”; see also the Notes of Chapter 3. The result follows directly from the fact that the distance to singularity of the shifted matrix $A - \lambda I$ is equal to $1/\|(A - \lambda I)^{-1}\|_2$ in the cases of both the matrix 2-norm and the Frobenius-norm pseudospectra.

If $\Lambda_\varepsilon(A)$ lies in the closed complex left half-plane, resolvent norms can be used to derive bounds reciprocal to ε of solutions of homogeneous and inhomogeneous linear differential equations $\dot{x}(t) = Ax(t) + f(t)$ that are valid for all times $t \geq 0$; see the Notes of Chapter 3. To assess the robustness of such transient bounds under perturbations of A , we therefore consider the following pseudospectral version of the previous problem. Here, ε -pseudoeigenvalues are the points in the ε -pseudospectrum $\Lambda_\varepsilon(A)$.

Problem. *Given a matrix having all ε -pseudoeigenvalues in the open complex left half-plane, find a nearest matrix having some ε -pseudoeigenvalue on the imaginary axis.*

The task here is to find $\Delta \in \mathbb{C}^{n,n}$ of minimal norm having $\Lambda_\varepsilon(A + \Delta) \cap i\mathbb{R} \neq \emptyset$. This problem is readily solved in terms of the distance to instability $\varepsilon_\star$ of A : for $\varepsilon < \varepsilon_\star$, the distance is simply $\varepsilon_\star - \varepsilon$ and the minimum perturbation matrix Δ equals $(1 - \varepsilon/\varepsilon_\star)$ times the minimum perturbation matrix in the nearest unstable matrix problem.

Structured problem

The figure shows that the unstructured spectral value set can be a misleading indicator of the robustness of stability.

(Diederich Hinrichsen and Anthony J. Pritchard, 2005)

By contrast, the eigenvalues that arise from structured perturbations do not bear as close a relation to the resolvent norm and may not provide much information about matrix behaviour.

(Lloyd N. Trefethen and Mark Embree, 2005)

Consider now the same problem posed for real matrices, i.e., finding $\Delta \in \mathbb{R}^{n,n}$ of minimal norm having $\Lambda_\varepsilon(A + \Delta) \cap i\mathbb{R} \neq \emptyset$, or, more generally, posed for perturbation matrices Δ in a complex-linear or real-linear structure space $\mathcal{S} \subset \mathbb{C}^{n,n}$. Here, the situation is more intricate. The key observation is that pseudoeigenvalues that arise from structured perturbations do bear a close relation to the resolvent norm. Given $\varepsilon < \varepsilon_\star$, we introduce the joint unstructured-structured pseudospectrum as a means to bridge the apparent discrepancy between the two quotes,

$$\begin{aligned} \Lambda_{\delta,\varepsilon}^{\mathcal{S}}(A) &= \{\lambda \in \mathbb{C} : \lambda \in \Lambda_\varepsilon(A + \Delta) \text{ for some } \Delta \in \mathcal{S} \text{ with } \|\Delta\|_F \leq \delta\} \\ &= \{\lambda \in \mathbb{C} : \lambda \in \Lambda(A + \Delta + \Theta) \text{ for some } \Delta \in \mathcal{S} \text{ with } \|\Delta\|_F \leq \delta \\ &\quad \text{and } \Theta \in \mathbb{C}^{n,n} \text{ with } \|\Theta\|_F \leq \varepsilon\} \\ &= \{\lambda \in \mathbb{C} : \|(A + \Delta - \lambda I)^{-1}\|_2 \geq \varepsilon^{-1} \text{ for some } \Delta \in \mathcal{S} \text{ with } \|\Delta\|_F \leq \delta\}. \end{aligned}$$

We ask for the smallest $\delta > 0$ such that $\Lambda_{\delta,\varepsilon}^{\mathcal{S}}(A) \cap i\mathbb{R} \neq \emptyset$. We call this minimal δ the $\mathcal{S}$ -structured ε -stability radius and denote it δ_ε . Then, transient bounds reciprocal

to ε are valid uniformly for all structured perturbations $\Delta \in \mathcal{S}$ with $\|\Delta\|_F \leq \delta_\varepsilon$; see Section 5.5. A minimal perturbation $\Delta \in \mathcal{S}$ again turns out to be the orthogonal projection onto $\mathcal{S}$ of a rank-1 matrix. A two-level algorithm for computing the $\mathcal{S}$ -structured ε -stability radius is based on a rank-1 matrix differential equation in the inner iteration and is given in Section 5.6. It takes essentially the same computational steps as the rank-1 algorithm for computing the unstructured stability radius. See also the Notes of Chapter 5.

Matrix stabilization

The question is to find the smallest perturbation that stabilizes a given unstable matrix A , or, equivalently, to find the closest stable matrix X to a given unstable matrix A . This kind of problem occurs in system identification when one needs to identify a stable system from observations.

(François-Xavier Orbandexivry, Yurii Nesterov, and Paul Van Dooren, 2013)

How near is an unstable matrix to a stable matrix? This formal converse of Van Loan's question has turned out to be challenging from a computational perspective. The problem arises in system identification as well as in model reduction of a stable system, which can yield a smaller but unstable matrix. The question then is how to stabilize it.

Problem. *Given a matrix having some eigenvalues of positive real part, find a nearest matrix having all eigenvalues of nonpositive real part (or below a negative threshold).*

A variety of numerical algorithms for this problem in the general complex case have been proposed in the literature; see the Notes of Chapter 6. For the Frobenius-norm setting, we describe in Section 6.1 two complementary new approaches, which can also be used for real and structured versions of the problem. The problem can be made more meaningful for transient behaviour – and potentially less computationally challenging – by replacing eigenvalues with ε -pseudoeigenvalues in the problem statement.

Se non è vero, è molto ben trovato.

(Giordano Bruno, 1585)

Contrary to computing the distance to instability considered before, different algorithms tend to compute different stabilized matrices with different distances to the given unstable matrix; see the numerical example in Section 6.1. Rigorously speaking, the algorithms only yield upper bounds of the distance to stability via locally minimal stabilizing perturbations. Nonetheless, providing a fairly tight bound together with a stabilized matrix can be useful when one considers that the problem

posed is a nonsmooth, nonconvex optimization problem that can have many local minima.

Nearest correlation matrix

Given a symmetric matrix what is the nearest correlation matrix, that is, the nearest symmetric positive semidefinite matrix with unit diagonal? This problem arises in the finance industry, where the correlations are between stocks.

(Nicholas J. Higham, 2002)

An important special case of the structured matrix stabilization problem is to find, for a given symmetric real matrix, a nearest symmetric positive semidefinite matrix with a prescribed additional structure. A case of particular interest is the following.

Problem. *Given a symmetric real matrix, find a nearest correlation matrix.*

Here we look for a minimal-norm perturbation of the form $D + \Delta$ with $D = I - \text{diag}(A)$ and $\Delta \in \mathcal{S}_{\text{cor}}$ for $\mathcal{S}_{\text{cor}} = \{\Delta \in \mathbb{R}^{n,n} : \Delta \text{ is symmetric with zero diagonal}\}$ such that all eigenvalues of $A + D + \Delta$ are nonnegative. In the Frobenius-norm setting, we observe $\|D + \Delta\|_F^2 = \|D\|_F^2 + \|\Delta\|_F^2$, and hence $\|\Delta\|_F$ (instead of $\|D + \Delta\|_F$) needs to be minimal. The problem can thus be viewed as the previous matrix stabilization problem for the matrix $-(A + D)$ with structured perturbations $\Delta \in \mathcal{S}_{\text{cor}}$ (unless $A + D$ is already positive semidefinite). A variety of dedicated algorithms, which are not based on this interpretation, have been given in the literature; see Higham (2002, 2025) and references therein. On the other hand, the structured matrix stabilization algorithms of Section 6.1 offer an alternative that can be used for the structured problem with perturbation matrices $\Delta \in \mathcal{S}_{\text{cor}}$ having a prescribed sparsity pattern.

Nearest defective matrix

Summary. Gives a bound for the distance of a matrix having an ill-conditioned eigenvalue problem from a matrix having a multiple eigenvalue which is generally sharper than that which has been published hitherto.

(James H. Wilkinson, 1972, complete abstract)

For many problems of numerical analysis, there is a simple relationship between the condition number of a problem and the shortest distance from that problem to an ill-posed one: the shortest distance is proportional to the reciprocal of the condition number (or bounded by the reciprocal of the condition number). This is true for matrix inversion, computing eigenvalues and eigenvectors, finding zeros of polynomials, and pole assignment in linear control systems.

(James W. Demmel, 1987, from the abstract)

A matrix is called *defective* if its Jordan canonical form has a Jordan block of dimension greater than 1. For the corresponding defective eigenvalues (and only for them), the eigenvalue condition number becomes infinite. Here, the reciprocal of the eigenvalue condition number is the absolute value of the inner product of left and right normalized eigenvectors.

Problem. *Given a matrix having distinct eigenvalues, find a nearest defective matrix.*

The problem was originally motivated by ill-conditioned eigenvalue problems in numerical linear algebra; see Alam, Bora, Byers, and Overton (2011) and the Notes of Chapter 6 for references to the history. It has recently found unexpected practical interest in non-Hermitian optics and photonics, where the desired “exceptional points” are nearest defective matrices, possibly with a linear structure and under further constraints; see the references in the Notes of Chapter 6.

For the general complex case, there exist various algorithms based on a characterization by Alam and Bora (2005), which reduces the problem to finding the smallest $\varepsilon > 0$ for which two components of $\Lambda_\varepsilon(A)$ coalesce in a point; see again the Notes of Chapter 6.

Structured problem. In the unstructured complex case, every matrix with multiple eigenvalues is arbitrarily close to a matrix with a defective eigenvalue. This need not hold true in structured cases, as is obvious when the structure space consists of normal matrices. We are not aware of any algorithm in the literature for finding nearest real or structured defective matrices. A two-level algorithm, which is suited also for such cases, is given in Section 6.3. The algorithm minimizes the reciprocal of the eigenvalue condition number of perturbed matrices with a fixed perturbation size via a rank-2 matrix differential equation in the inner iteration. It determines the distance to defectivity in the outer iteration by requiring that the reciprocal of the eigenvalue condition number become 0.

Nearest singular matrix pencil

This is a first step toward the goal of finding a way to calculate a smallest norm de-regularizing perturbation of a given square matrix pencil.

(Ralph Byers, Chunyang He, and Volker Mehrmann, 1998)

A matrix pencil $\{A - \mu B : \mu \in \mathbb{C}\}$ or more briefly (A, B) (with square matrices A and B) is called *singular* if $A - \mu B$ is a singular matrix for all $\mu \in \mathbb{C}$ (and is called a regular matrix pencil otherwise). This notion is fundamental in the theory of linear differential-algebraic equations $B\dot{x}(t) = Ax(t)$, which arise in descriptor systems and as linearizations of nonlinear differential-algebraic equations $F(x, \dot{x}) = 0$ near a stationary point. The matrix pencil (A, B) is singular if and only if no initial value $x(0)$ exists such that the corresponding initial value problem has a unique solution.

Problem. *Given a regular matrix pencil, find a nearest singular matrix pencil.*

Byers, He, and Mehrmann (1998) derived characterizations of minimal perturbations and upper and lower bounds for the distance. Algorithms for this problem appeared later; see the Notes of Chapter 6. In Section 6.4, we extend the algorithm for computing a nearest (structured) singular matrix, as given in Section 5.8, to compute the nearest singular matrix pencil to (A, B) under (possibly structured) perturbations of A when B is a fixed singular $n \times n$ matrix. The algorithm extends further to the case where B can also be perturbed. It still uses only matrix-vector products and vector inner products, but now $(n + 1)$ times more of them in every iteration. Such an extra factor does not occur for the simpler problem of finding the nearest matrix pencil for which the two matrices have a common null-vector.

There are many more interesting nearness problems for matrix pencils and, more generally, matrix-valued polynomials and even more general nonlinear matrix-valued functions of a complex parameter, not least those related to robust stability of differential-algebraic equations and delay differential equations; see Section 6.5.

Matrix nearness problems in robust control

Problems of uncertainty and robustness, which had been forgotten for some time in “modern control”, gradually re-emerged and came to the foreground of control theory.

(Diederich Hinrichsen and Anthony J. Pritchard, 2005)

Consider a linear time-invariant system with input $u(t) \in \mathbb{C}^p$, output $y(t) \in \mathbb{C}^m$, and state vector $x(t) \in \mathbb{C}^n$,

$$\begin{aligned}\dot{x}(t) &= Ax(t) + Bu(t), \\ y(t) &= Cx(t) + Du(t),\end{aligned}$$

where all eigenvalues of A have negative real part.

In Chapter 7, we discuss basic problems of robust control of such systems. We rephrase them as eigenvalue optimization problems and matrix (or operator) nearness problems.

Problems. *The problems considered from this perspective include the following:*

- *Compute the $\mathcal{H}_\infty$ -norm of the system, which is the operator norm of the input-to-output map $u \in L^2(0, \infty; \mathbb{C}^p) \mapsto y \in L^2(0, \infty; \mathbb{C}^m)$.*
- *Find an $\mathcal{H}_\infty$ -nearest uncontrollable system to a given controllable system under perturbations of the matrix B .*
- *Find an $\mathcal{H}_\infty$ -nearest passive system to a given non-passive system under perturbations of the matrix C .*

- *Find a non-contractive system by a minimal structured perturbation to the state matrix A of a given contractive system.*

The distances from systems with undesired properties are important robustness measures of a given control system, whereas algorithms for finding a nearby system with prescribed desired properties are important design tools; see the Notes at the end of Chapter 7. In that chapter, the algorithmic approach to eigenvalue optimization via low-rank matrix differential equations and the two-level approach to matrix nearness problems of previous chapters is extended to problems of robust control and is shown to yield a versatile approach to novel algorithms in that field.

Graph problems as matrix nearness problems

There is no question that eigenvalues play a central role in our fundamental understanding of graphs.

(Fan Chung, 1997)

Eigenvalues or eigenvectors associated with matrices such as the graph Laplacian or the weighted adjacency matrix characterize basic properties of graphs such as connectivity and centrality. In Chapter 8, we study problems of clustering and ranking nodes of undirected weighted graphs, which can be viewed as matrix nearness problems that refer to eigenvalues and eigenvectors of sparsity-structured symmetric real matrices under nonnegativity constraints for the weights. We consider in detail the following two exemplary problems, which will be reformulated as matrix nearness problems.

Problem. *Given a connected weighted undirected graph, find a minimum cut subject to constraints such as must-link, cannot-link, and cardinality constraints.*

Problem. *Given a connected weighted undirected graph where node i is ranked higher than node j , find a minimal perturbation of the weights such that the ranking is reversed.*

For both problems we propose efficient algorithms that perform eigenvalue optimization without eigenvalues, but instead use Rayleigh quotients. This requires only vector inner products and matrix-vector products with matrices having the sparsity pattern of the adjacency matrix of the graph.

Outline

The matrix nearness problems presented above (and some more) will be discussed in detail in the following chapters. For their numerical solution, a common framework throughout this book is provided by a two-level approach that combines eigenvalue

optimization and root-finding in a nested or alternating way: an eigenvalue optimization problem for a fixed perturbation size is to be solved via a gradient-based low-rank matrix differential equation, and the root-finder adjusts the perturbation size by solving a nonlinear scalar equation, typically by a Newton-bisection method. In the next two chapters, we concentrate on the basic eigenvalue optimization problem, and in Chapter 4 on its combination with scalar root-finding. Chapter 5 extends the programme to structured matrix nearness problems. Chapter 6 discusses diverse matrix nearness problems that go beyond those of previous chapters. Chapters 7 and 8 are about matrix nearness problems in the areas of robust control and graph theory, respectively. Appendix A presents known basic results on derivatives of eigenvalues and eigenvectors, which are often used in this book.