

Chapter 3

Inverse Problems

Mathematicians often cannot help but cringe when physicists are doing math, and even more so whenever physicists claim to be ‘rigorous’ with their math. This chapter is not written to satisfy the mathematicians for two reasons. First, I am not a mathematician, so I am quite certain that they will cringe despite my best efforts. More importantly though, this is meant to be accessible for engineers and physicists, and often what mathematicians consider ‘special cases’ are the only ones we happen to care about. So with apologies to any mathematicians reading this, the goals of this chapter are threefold: First, we want to help the reader develop an appreciation for what inverse problems are and what makes them difficult. Second, we want to introduce the specialized tools that are used to solve these inverse problems. Finally, we bring the focus back to our particular application, and fine tune the ideas developed for the purpose of photonic device design.

There are many excellent references on inverse problems. A standard reference is the textbook by Engl [21] which gives a thorough overview of the subject, but the mathematics is quite formal. A very nice introduction to the subject for physicists can be found in a series of lecture notes by Sze Tan and Colin Fox at the University of Auckland [22]. The work by Per Christian Hansen is more focused on discrete and finite dimensional problems, and hence particularly suitable to our application. He has also written a package of matlab functions for inverse problems available for download as well [23]. The ideas presented in this chapter are mostly taken from his work, although the discussion of the role of noise in distinguishing between a forward

and inverse problem has to our knowledge not been articulated elsewhere. Arnold Neumaier also provides a concise treatment similar to Hansen's, but bridges the gap to the infinite dimensional treatment of inverse problems with more mathematical rigor [24].

We begin the chapter by attempting to define what an inverse problem is through some examples of simple physical problems. We introduce the concept of an *ill-posed* problem to distinguish between the *forward* or *direct* problem vs. the inverse problem. In section 3.2, we restrict our discussion to finite dimensional linear operators, allowing us to illustrate the pathologies of inverse problems in the linear algebra formalism. A numerical example is provided to help illustrate the effects of ill-conditioning. We make use of the singular value decomposition (SVD) to explain why standard techniques will fail to solve inverse problems. The SVD also allows us to utilize the condition number as a quantifying metric for how ill-posed a particular problem is. In section 3.3 we introduce regularization as a tool for solving inverse problems. We conclude with a glimpse of the difficulties we expect to encounter for the purpose of PBG device design.

3.1 Introduction

At first glance, the meaning of the term ‘inverse problem’ seems obvious. It is the complement of some other problem, one that presumably preceded the inverse problem, and is more well known. To a physicist though, such a ‘definition’ is rather unsavory, for if that were the case, then the distinction between a forward problem and an inverse problem seems rather arbitrary. Our obsession with symmetries in natural laws lead us naturally to wonder why one problem formulation is more ‘privileged’ than the other. A good example of this that we have already encountered in this thesis is the Fourier transform and the inverse Fourier transform. The two operations are simply labeled that way by convention, and nothing would have been lost had we reversed the labels. It becomes a question of semantics, rather than a matter of any fundamental significance. By the end of this chapter, we will see that

the inverse Fourier transform in fact does not fit our definition of an inverse problem.

The distinction is in reality more than just semantics or there would not be an entire journal devoted to inverse problems. One's first exposure to inverse problems is typically accompanied by some claim that inverse problems are difficult to solve, with the implication being that it is *more* difficult than the associated forward problem. We give a more formal definition in the next section, but first, we review a few well-known examples of inverse problems to develop some intuition.

3.1.1 Examples

Our first example is found in medical imaging, such as computerized tomography (CT) scans. The forward problem is a form of scattering or diffraction problem, such that for some radiation incident upon a given material distribution, we determine the scattered radiation in the far field. For medical applications, the goal is to non-invasively determine the internal structure of a patient's body. This is accomplished by measuring the scattered field at various angles given some incident radiation, and solving the inverse scattering problem for the scatterer distribution. A related inverse problem is found in geophysics, where the internal structure of the earth is determined based on surface measurements of seismic waves.

Another example is in image processing, or image restoration, where the ideal image must pass through non-ideal optics, leading to blurring and other distortion to the captured image. The forward problem of blurring is typically modeled as a convolution of the original image $i_o(x)$ with a point spread function $h(x)$. Sharp features are smeared out, leading to a loss of resolution. Formally, our captured image $i_c(x)$ becomes:

$$i_c(x) = \int i_o(\xi)h(x - \xi)d\xi, \text{ or} \tag{3.1}$$

$$I_c(k) = I_o(k)H(k) \tag{3.2}$$

The inverse problem becomes a deconvolution, which can be performed as a simple division of the captured image with the point spread function in their respective Fourier

representations. One can determine $h(x)$ by characterizing the optical elements carefully. It turns out that the image cannot be reconstructed with a straightforward application of the deconvolution theorem. Section 1.6 in [22] provides a nice pictorial example of this problem.

A final example is the heat conduction problem. The forward problem is of course a standard undergraduate-level problem, and can be described by some variant of the following:

$$\frac{\partial u(x, t)}{\partial t} = \frac{\partial^2 u(x, t)}{\partial x^2}, \quad x \in [0, \pi], t \geq 0, \quad (3.3)$$

$$\text{where } u(x, 0) = f(x), \quad (3.4)$$

$$u(0, t) = u(\pi, t) = 0. \quad (3.5)$$

This is solved in the usual way by separation of variables and then an eigenfunction expansion for the spatial dependence, with the set of normalized sine functions $\{\phi_n(x)\}$ forming a complete orthonormal set. Expressing the initial distribution in terms of a superposition of the eigenfunctions $f(x) = \sum_n c_n \phi_n(x)$, we obtain the heat distribution $u(x, t)$ as

$$u(x, t) = \sum_n c_n e^{-n^2 t} \phi_n(x). \quad (3.6)$$

There is often some remark about our inability to solve the backwards heat conduction problem, namely given some final distribution $u(x, t_f)$, we generally cannot go backwards in time and deduce an initial distribution. Typically this is attributed to the exponential factor, and we see that it blows up if we go backwards in time.

Based on these examples, we can make some observations that will prove to be helpful. First, the heat conduction example makes explicit a common theme among the examples provided: that of *cause* and *effect*. Whereas forward problems in physics tend to study the unknown effects of known causes (in order to derive a model for predicting the effects), inverse problems seek the unknown cause of measured effects. The backwards heat conduction equation makes this transparent because of the explicit time dependence, but the rest of the examples all seek an explanation of some final

observed phenomenon, given some well-characterized, presumably accurate model of the forward problem.

The other observation is that at least in some of these forward problems, there appears to be a ‘smoothing’ process. For example, in the heat conduction and the blurring examples, features present in the initial condition seems to get lost or smoothed out as the state evolves forward. Often, we arrive at steady state solutions of dynamical systems, therefore independent of initial conditions: The system forgets or loses information about the initial state. In such a case, we can certainly see just by physical principles alone why an inverse problem would be ‘difficult.’ A more precise way to look at this might be how ‘solvable’ a given problem is, which leads to the notion of well-posed and ill-posed problems proposed by Hadamard.

3.1.2 Well-posedness

Hadamard proposed three properties that a problem must possess in order to be classified as well-posed [21]:

1. For all admissible data, a solution exists.
2. For all admissible data, the solution is unique.
3. The solution depends continuously on the data.

What constitutes ‘admissible data,’ ‘solution’ and ‘continuous’ will of course depend on the nature of a specific problem. For our purposes, considering only finite dimensional linear operators, we can think of data and solution as the *input* and *output* vectors of some linear transformation. The first two properties seem rather obvious, as it is not much of a linear mapping if we, for some given input, cannot get an output, or get multiple or non-unique outputs. The third property is a question of stability, requiring that small changes to the input does not produce arbitrarily large changes to the output. A problem that lacks any one of these properties is by definition *ill-posed*.

We can apply our intuition to the heat conduction example and readily see that indeed a solution to the inverse problem does not always exist. A simple example is if our final field distribution corresponds to a state with minimal entropy. Since entropy must increase with time, we know that there is no way to go backwards in time, since we are ‘starting’ in a minimum entropy state. As for the uniqueness of the solution to the inverse problem, we already addressed the problem of the steady state fields, which means *any* initial state would reach the same steady state. The final property of stability relates to the smoothing behavior of the forward problem. If we perturb the initial conditions by a small amount, the perturbations will be smoothed out over time to yield similar output fields. By extension then, small perturbations at the final time must have corresponded to large changes in the initial condition. We observed this effect quantified by the exponential term in the heat conduction equation. We now express these ideas in a more formal mathematical footing in the context of finite dimensional linear operators which can be represented by matrices.

3.2 Matrices as Linear Operators

3.2.1 A numerical example

We first provide a numerical example to give a concrete illustration of the ideas presented in the previous section. The matlab code used to generate this example is included in appendix A. Consider the $N \times N$ Hilbert matrix.

$$A(i, j) = \frac{1}{i + j - 1} \tag{3.7}$$

$$\text{with } Ax_{in} = x_{out} \tag{3.8}$$

For this example, we will choose $N = 5$, and choose a relatively simple x_{in} .

$$A = \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} \\ \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} \\ \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} \end{bmatrix}, \quad x_{in} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad (3.9)$$

Evaluating x_{out} gives:

$$x_{out} = \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} \\ \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} \\ \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 2.2833 \\ 1.4500 \\ 1.0929 \\ 0.8845 \\ 0.7456 \end{bmatrix}. \quad (3.10)$$

Clearly, for any x_{in} , we can evaluate a unique x_{out} . Therefore the first two Hadamard conditions are satisfied. To test the stability condition, we can define an additive noise vector n that is sufficiently small to form x'_{in} . Evaluating x'_{out} gives:

$$x'_{out} = \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} \\ \frac{1}{2} & \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} \\ \frac{1}{4} & \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} \\ \frac{1}{5} & \frac{1}{6} & \frac{1}{7} & \frac{1}{8} & \frac{1}{9} \end{bmatrix} \begin{bmatrix} 0.9982 \\ 0.9992 \\ 0.9997 \\ 1.0003 \\ 1.0010 \end{bmatrix} = \begin{bmatrix} 2.2813 \\ 1.4490 \\ 1.0922 \\ 0.8840 \\ 0.7452 \end{bmatrix}. \quad (3.11)$$

We see that x'_{out} is close to the nominal solution x_{out} . Formally we can define the relative magnitude of the input and output error to provide a measure of the stability




Figure 3.1: Distribution of stability values for the Hilbert matrix operator

\mathcal{S} :

$$e_{in} = \frac{|x'_{in} - x_{in}|}{|x_{in}|} \quad (3.12)$$

$$e_{out} = \frac{|x'_{out} - x_{out}|}{|x_{out}|} \quad (3.13)$$

$$\mathcal{S} \equiv \frac{e_{out}}{e_{in}} = 0.7846 \quad (3.14)$$

where $|\cdot|$ denotes the 2-norm (i.e., $|x| = (\sum_i x_i^2)^{1/2}$). We repeat this with 10,000 different noise vectors and show the distribution of \mathcal{S} in figure 3.1. Most of the values fall between 0 and 1, with the maximum value of about 1.1. Therefore, we see that this problem is stable against perturbations to the input vector, i.e., errors remain small.

We now look at the ‘reverse’ problem of finding x_{in} given x_{out} . We will look at the stability again, but this time, we add the noise to the nominal x_{out} . We solve for

$x'_{in} = A^{-1}x_{out}$. We use matlab to find the inverse of A .

$$A^{-1} = \begin{bmatrix} 25 & -300 & 1050 & -1400 & 630 \\ -300 & 4800 & -18900 & 26880 & -12600 \\ 1050 & -18900 & 79380 & -117600 & 56700 \\ -1400 & 26880 & -117600 & 179200 & -88200 \\ 630 & -12600 & 56700 & -88200 & 44100 \end{bmatrix} \quad (3.15)$$

Again, using 10,000 different noise vectors, we obtain the distribution of \mathcal{S} for this reverse problem (as shown in 3.2). Notice the x-axis is scaled by 10^5 , meaning the relative error is greatly amplified. To illustrate, suppose we rounded x_{out} to 3 decimal places and then evaluated x'_{in} .

$$x'_{out} = \begin{bmatrix} 2.283 & 1.450 & 1.093 & 0.885 & 0.746 \end{bmatrix}^T \quad (3.16)$$

$$x'_{in} = A^{-1}x'_{out} = \begin{bmatrix} 2.105 & -20.28 & 94.29 & -141.4 & 71.19 \end{bmatrix}^T \quad (3.17)$$

In the reverse problem, we cannot even tolerate rounding errors as x'_{in} bears no resemblance to x_{in} at all. Therefore, this problem fails to satisfy Hadamard's third condition and is therefore ill-posed. Because of the ill-posedness, this reverse problem is the one that is defined to be the inverse problem. Therefore, it is not simply a question of semantics, but there are fundamental distinctions between a forward and its inverse problem. Even if we had first defined an operator $B = A^{-1}$ and went through this same analysis, we would still conclude that B is the 'inverse problem,' and B^{-1} is the 'forward problem,' objectively based on the stability criterion.

In chapter 6, when we derive the inverse Helmholtz equation, we will encounter a more severe manifestation of this problem, where we cannot (even without the additive noise) recover the input with the computed output. However, if we understand the analysis in this chapter, it will no longer be surprising when we get to chapter 6. To understand the origins of this ill-conditioning, we need to take a closer look at the properties of a general linear operator A in terms of the singular value decomposition.

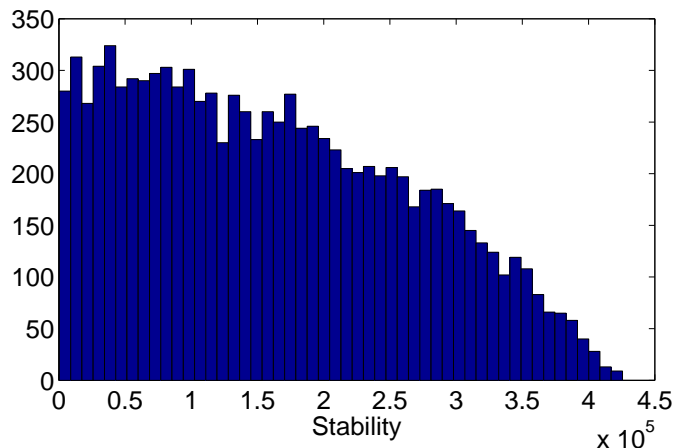


Figure 3.2: Distribution of stability values for the inverse Hilbert matrix operator. Notice the scale on the x-axis is in increments of 10^5 .

3.2.2 Singular value decomposition

Consider a general linear transformation (or linear mapping) $A : \mathcal{C}^n \rightarrow \mathcal{C}^m$, such that

$$Ax = b \quad (3.18)$$

with A an $m \times n$ matrix, $x \in \mathcal{C}^n$ and $b \in \mathcal{C}^m$.

Any matrix A can be decomposed by the singular value decomposition (SVD) such that

$$A = U\Sigma V^\dagger, \quad (3.19)$$

where the $m \times m$ matrix U and the $n \times n$ matrix V are unitary, and Σ is an $m \times n$ matrix whose only non-zero elements are along the diagonal with $\{\sigma_i \geq \sigma_{i+1} \geq 0\}$ called the singular values. The columns of U and V are known as the left $\{u_i\}$ and right $\{v_i\}$ singular vectors. This is a generalization of the eigenvalue decomposition. In fact (although one would not actually do so in practice), one can get the SVD by performing an eigenvalue decomposition of AA^\dagger and $A^\dagger A$. The eigenvectors of AA^\dagger and $A^\dagger A$ are the left and right singular vectors of A respectively, and the eigenvalues are the singular values squared. Since both AA^\dagger and $A^\dagger A$ are Hermitian and positive semi-definite, we are guaranteed real non-negative eigenvalues, thus ensuring $\sigma \geq 0$.




Figure 3.3: The vector x to be transformed is decomposed into v_1 and v_2 and then mapped onto corresponding u_1 and u_2 , with each component stretched or compressed by the respective σ . In this example, $\sigma_2 < 1$.

Having obtained the SVD of A , we can write down the linear mapping in a more suggestive form:

$$A\mathbf{x} = U\Sigma(V^\dagger\mathbf{x}) \quad (3.20)$$

$$A\mathbf{x} = \sum_{i=1}^{\min(m,n)} (\mathbf{v}_i^\dagger\mathbf{x}) \sigma_i \mathbf{u}_i, \quad (3.21)$$

where for clarity we have written the vectors in boldface. Looking at eqn. (3.21), we see that any linear mapping can be viewed as a transformation of the input vector x into the right singular vector basis, then stretching each component by the associated singular value, and finally mapping these components to the corresponding left singular vectors. A pictorial representation for a simple 2D mapping is shown in figure 3.3. For special cases of A that has an eigenvalue decomposition (i.e. diagonalizable), the left singular vectors are the same as the right singular vectors, so in the diagonalized basis, the linear transformation is particularly simple (just stretch each component by the eigenvalues; this is, of course, why we prefer to work in a diagonalized basis). The singular values of A play an important role since they determine how much gain is in a particular component of the linear map. For the time being, let us consider the problem of finding b (given A and x) to be the forward problem, while the problem of finding x (given that Ax produces b) is the inverse problem.

Revisiting Hadamard's conditions then, in the forward problem, the first two conditions are automatically satisfied if we can express the problem in this form. As for the third condition, we can think of it as requiring reasonable gains (i.e. not too large) for the system. Stability can also be achieved if random perturbations are spectrally decomposed to singular vectors that have relatively small singular values. In other words, singular vectors associated with small singular values should look like noise. For physical systems the relevant physics are embodied by the linear operator A .

If we now attempt to solve the inverse problem, we need to do the following:

$$Ax = b \quad (3.22)$$

$$x = A^{-1}b \quad (3.23)$$

$$= V\Sigma^{-1}U^\dagger b \quad (3.24)$$

$$= \sum_i^{\min(m,n)} \left(\frac{1}{\sigma_i} \mathbf{u}_i^\dagger \mathbf{b} \right) \mathbf{v}_i. \quad (3.25)$$

In eqn. (3.25) above, we have expressed the inverse of A using the SVD expansion. Even when A is singular (i.e. not strictly invertible), the expression can be used and interpreted as a generalized inverse or 'pseudo-inverse,' although there are of course limitations associated with a singular A . Now A is obviously singular when $m \neq n$, but even when A is square it can still be singular if $\sigma_i = 0$. Singularity of A implies that A does not have full rank, i.e. A is *rank deficient*, or A has a non-trivial nullspace:

$$\text{For } \sigma_i = 0, \quad (3.26)$$

$$Av_i = 0 \quad (3.27)$$

$$\therefore A(x + \alpha v_i) = b, \text{ and furthermore,} \quad (3.28)$$

$$\nexists y \mid Ay = \beta u_i \quad (3.29)$$

Eqn. (3.28) shows that we fail the uniqueness test, and eqn. (3.29) shows that we fail the existence test for Hadamard's condition for well-posedness. In most physical

problems, the singular values may not be identically zero, as it would be impractical to numerically evaluate them to that level of precision. Based on the importance of the singular value spectrum though, we can define a *condition number*:

$$C \equiv \frac{\max \sigma}{\min \sigma} \quad (3.30)$$

As C becomes larger, the problem becomes more ill-conditioned, and for a strictly singular matrix, $C \rightarrow \infty$. Even though we now have this quantity defined, the boundary between what is considered well-conditioned and poor-conditioned is not a sharp one. A generally acceptable figure is $C \leq 10^3$.

We now return to our numerical example of the Hilbert matrix. The singular values are $\{1.567, 0.2085, 0.0114, 0.0003, 0.000003\}$. The condition number is 4.766×10^5 , so as suspected, the problem is ill-conditioned. Specifically, let us examine eqn. (3.25), especially the factor σ_i^{-1} . In the forward problem, the small singular values damp out the contributions from the additive noise. In the inverse problem, however, they become an amplification for the noise components, drowning out the original signal x_{in} . This amplification picture is consistent with our result above, as we found $|A^{-1}x'_{out} = x'_{in}| \gg |x_{in}|$. If a problem is ill-conditioned, any standard matrix inversion algorithm will fail to recover the desired solution x_{in} . Having understood the origins of the difficulties, we can now discuss strategies for overcoming these difficulties.

3.3 Regularization and the L-curve

The specialized technique that is used to solve inverse problems is called *regularization*. There are many regularization schemes that have been developed, and Engl's text is a good starting point. Here, we will only discuss the most common regularization scheme, known as Tikhonov regularization, that works well in many situations.

First, we must address the three properties of ill-posedness. In a way, they are related, because any time we have a singular (or near-singular as defined by C)

matrix, any of the three can occur. The lack of an existence theorem is overcome by minimizing the residual $|Ax'_{in} - x'_{out}|$ in usual inverse problem applications, and is not considered too serious. One must give up on the notion of an exact solution to inverse problems. Rather, we just try to reconstruct a ‘sensible’ solution that satisfies our given equation ‘well enough.’ We will comment further on this issue in the final section of this chapter.

Non-uniqueness is considered much more serious. In our numerical example, having given up the notion of an exact solution, we know that $|Ax_{in} - x'_{out}|$ would have been nonzero but small. In fact, it is exactly the norm of the small noise term added to x_{out} . The problem becomes how to pick out the nice solution among all the many that would still give reasonably small residual norms. We observed at the end of the last section that the small singular values lead to large noise amplification. We note also that these bad solutions do tend to blow up and have large norms, much larger than the desired solution. Therefore, one strategy would be to restrict the size of the solution. This additional constraint allows us to choose systematically a unique solution that at least allows a ‘sufficiently small’ residual. Rather than minimizing only the residual, we can include the solution norm as well to formulate the following regularized inverse problem:

$$x_{out}^{(\lambda)} = \min_{x_{out}} \{|Ax_{out} - x'_{in}|^2 - \lambda|x_{out}|^2\} \quad (3.31)$$

The scalar λ is known as the regularization parameter, and determines the relative weight between minimizing the residual norm versus the solution norm. The standard graphical tool that assists in choosing the regularization parameter is the L-curve. The L-curve plots the solution norm on the y-axis and the residual norm on the x-axis for a series of λ 's. The L-curve for our numerical example is shown in figure 3.4, and it takes on its name because of the characteristic L shape. The corner is usually taken as the appropriate regularization parameter that indicates optimal compromise between the two norms. In practice however, the corner rarely gives the optimal solution, so it is best to use it as a rough guide. Constructing the L-curve for our problem, we




Figure 3.4: L-curve for the Hilbert operator. The curve is a parametric log-log plot of the solution norm vs. the residual norm for different regularization parameter λ in the direction shown. Most references suggest the optimal point is at the corner as shown in red, but physical problems often have the ‘best’ solution elsewhere.

find a value of $\lambda = 6.7 \times 10^{-5}$ at the corner. The solution then is:

$$x'_{\lambda=6.7 \times 10^{-5}} = \begin{bmatrix} 0.9298 & 1.5388 & 0.2284 & 0.8474 & 1.4818 \end{bmatrix}^T \quad (3.32)$$

$$x'_{\lambda=1.6 \times 10^{-3}} = \begin{bmatrix} 0.9800 & 1.0661 & 1.0103 & 0.9773 & 0.9474 \end{bmatrix}^T \quad (3.33)$$

$$\cong x_{in}. \quad (3.34)$$

By looking at some more values near the corner, we find that the solution closer to our ‘true’ solution actually has $\lambda = 1.6 \times 10^{-3}$. So we see that we can in fact recover sensible results even for badly conditioned problems.

3.3.1 An alternate interpretation

There is an alternative picture to justify the Tikhonov regularization scheme. Recognizing that it is the small singular values that cause the difficulties, we can imagine

applying a filter on the singular values when we construct the inverse in eqn. (3.25). Applying a Lorentzian filter to the reciprocal singular values, we get:

$$\frac{1}{\sigma_i} \rightarrow \left(\frac{\sigma_i^2}{\sigma_i^2 + \lambda^2} \right) \sigma_i^{-1} \quad (3.35)$$

For $\sigma_i \gg \lambda$, the filter has little effect, whereas if $\sigma_i \ll \lambda$, then $\sigma_i^{-1} \rightarrow \lambda^{-1}$, limiting the unstable spectrum. We see that the two views are equivalent since we can analytically solve the Tikhonov minimization (eqn. (3.31)). For a given λ , the function is minimized if x'_{out} is constructed using filtered coefficients of eqn. (3.35) instead of the reciprocal singular values σ_i^{-1} . Different regularization schemes effectively change how we evaluate the filtering coefficients. For example, if we take our Hilbert operator and increase N to 100, we find the spectrum of singular values as shown in figure 3.5. Because of the distinctive corner in the spectrum, we might use an aggressive strategy here and simply truncate beyond the 20th singular value. This is known as the truncated singular value decomposition (TSVD) regularization scheme. (Note: this scheme alone would not work for the Hilbert problem because the remaining σ 's would still give a condition number of 10^{17} .) Most physical inverse problems do not have these obvious clusters, making hard truncation more difficult, so Tikhonov is really a good general strategy to use. In figure 3.6 we show the spectrum for the inverse photonic problem (see chapter 6) using a Gaussian output mode. This spectrum is more representative of real world inverse problems.

3.4 Conclusion

In this chapter, we explored the reasons why there is a real fundamental distinction between a forward problem and an inverse problem. In particular, the notion of stability against random perturbations is what sets the two apart, and we gave the condition number as a quantity that helps us identify ill-conditioning. For completeness, we now elaborate on a subtle point that we commented on earlier in passing. We motivated the need to find a fundamental distinction between forward and inverse




Figure 3.5: The spectrum of singular values for the 100×100 Hilbert operator. Note the distinct corner, showing an obvious point where we can perform a truncation.

problems due to the inherent symmetry of the two problems, i.e. one is the inverse of the other. It should now be clear that neither the Fourier transform nor its inverse can be considered an inverse problem, because they are both unitary, so $C = 1$. For an inverse or ill-posed problem, $C \gg 1$.

Of course, the condition number of the forward and inverse problem are the same as we have defined it (since it is just the ratio of the largest to smallest singular values), so the spectrum of singular values does not break the symmetry between the two problems. So what actually breaks the symmetry? It turns out to be the special status given to the random fluctuation or noise. The forward problem is defined as the one that is stable against changes *caused by random fluctuations*. However, given a large condition number, stability against noise necessarily implies it will be ‘unstable’ to a different form of perturbation. Of course, we usually do not use the term ‘instability,’ but instead we use the term ‘sensitivity’ in this context. Historically, this makes sense in how one studies physics. To model a physical system, we vary its parameters and measure its effects. A model is good if it makes good predictions about




Figure 3.6: The spectrum of singular values for the inverse photonic problem using a Gaussian shaped desired mode. In contrast to figure 3.5, we find no obvious place for a hard truncation.

the effects. Given a new physical system we are trying to model, if any small noise (i.e. a perturbation to the system the experimentalist cannot control) will create a large disturbance in the effect we can measure, it will be very difficult to come up with a model. What we need is a system that is sensitive to controllable and systematic variations to the input, so the effects can be readily observed with adequate signal-to-noise ratio. By its very nature, most problems studied are stable (in the sense given here) against most forms of noise. Any physical model derived based on experimental results will necessarily reflect this process. Therefore, we do expect ‘noise’ vectors to have large projections onto the ‘bad’ singular vectors in physical problems, using our linear algebra language.

3.4.1 Parameter estimation vs. design

We conclude this chapter by making an observation about the difficulties in transferring over from standard inverse problems to PBG device design. First, most standard inverse problems can assume implicitly the existence of a solution even if the prob-

lem formally does not guarantee you a solution [21]. Going back to the cause and effect picture, you are measuring a real effect from a cause that necessarily exists. The problem as we saw is that noisy measurements hide the underlying cause. If you really cannot reconstruct the cause from the measured effect, it is probably an indication that the model is wrong. As applied to a design paradigm, that is not the case. The ‘desired effect’ has not been observed or measured. It is a mere figment of our imagination, so to speak. We will see that this is a much more serious problem for design purposes. If we encounter a design problem where we encounter a non-existent solution situation, we would like to make strong claims to that effect, but we cannot do so because the Tikhonov regularization scheme is not really the most appropriate for the PBG problem. This brings us to our second observation. The solution we seek in the PBG inverse problem is the dielectric function, and their norms are not necessarily small, particularly if there are discontinuities. Physically, η cannot take on negative values, but Tikhonov would be happy accommodating negative values as long as they are small. Fortunately, we can use the insight developed in this chapter to implement a much more appropriate regularization scheme for the PBG problem. We expand on these ideas and develop the necessary tools in the next chapter.