University of Pennsylvania
ScholarlyCommons

Departmental Papers (BE)

Department of Bioengineering

August 2008

The Bad Truth about Laplace's Transform

Charles L. Epstein

University of Pennsylvania, cle@math.upenn.edu

John C. Schotland

University of Pennsylvania, schotland@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/be_papers

Recommended Citation

Epstein, C. L., & Schotland, J. C. (2008). The Bad Truth about Laplace's Transform. Retrieved from https://repository.upenn.edu/be_papers/123

Copyright SIAM 2008. Reprinted from *Siam Review*, Volume 50, Issue 3, pp. 504–520.
Publisher URL: <http://dx.doi.org/10.1137/060657273>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/be_papers/123
For more information, please contact repository@pobox.upenn.edu.

The Bad Truth about Laplace's Transform

Abstract

Inverting the Laplace transform is a paradigm for exponentially ill-posed problems. For a class of operators, including the Laplace transform, we give forward and inverse formulae that have fast implementations using the fast Fourier transform. These formulae lead easily to regularized inverses whose effects on noise and filtered data can be precisely described. Our results give cogent reasons for the general sense of dread most mathematicians feel about inverting the Laplace transform.

Keywords

finite Fourier transform, FFT, Laplace transform, spectral theory, SVD, regularized inverse

Comments

Copyright SIAM 2008. Reprinted from *Siam Review*, Volume 50, Issue 3, pp. 504–520.

Publisher URL: <http://dx.doi.org/10.1137/060657273>

The Bad Truth about Laplace's Transform*

Charles L. Epstein[†]
John Schotland[‡]

Abstract. Inverting the Laplace transform is a paradigm for exponentially ill-posed problems. For a class of operators, including the Laplace transform, we give forward and inverse formulae that have fast implementations using the fast Fourier transform. These formulae lead easily to regularized inverses whose effects on noise and filtered data can be precisely described. Our results give cogent reasons for the general sense of dread most mathematicians feel about inverting the Laplace transform.

Key words. finite Fourier transform, FFT, Laplace transform, spectral theory, SVD, regularized inverse

AMS subject classifications. 44A10, 15A12, 65R10, 65R30, 65T50, 62J10

DOI. 10.1137/060657273

1. Introduction. Inversion of the Laplace transform is the paradigmatic exponentially ill-posed problem. In many inverse scattering problems, the Laplace transform is, at least implicitly, a part of the forward model, and so the solution of the inverse scattering problem entails inverting the Laplace transform; see [12, 13, 9, 6]. While it is well understood that this inversion is problematic, to the best of our knowledge, no one has yet spelled out the details of why, where, and how things go wrong. In this note we introduce the harmonic analysis appropriate to this problem. On one hand, this leads to fast numerical forward and inverse algorithms for data which is log-uniformly sampled. On the other hand, we apply it to study regularized inverses of the Laplace transform. We analyze the consequences of passing noisy, filtered measurements through the approximate inverse. The picture that emerges is probably much worse than most people imagine.

We begin by considering a class of integral transforms that includes the Laplace transform. Suppose that f and k are functions defined on $[0, \infty)$. We define the transform $\mathcal{K}f$ by

$$(1.1) \quad \mathcal{K}f(x) = \int_0^\infty k(xy)f(y)dy.$$

Here the kernel function $k(t)$ is typically a smooth and rapidly decaying function.

*Received by the editors April 4, 2006; accepted for publication (in revised form) February 5, 2007; published electronically August 4, 2008.

<http://www.siam.org/journals/sirev/50-3/65727.html>

[†]Department of Mathematics, University of Pennsylvania, Philadelphia, PA (cle@math.upenn.edu). The research of this author was partially supported by NSF grants DMS02-03705 and DMS06-03973 and by the Francis J. Carey term chair.

[‡]Department of Bioengineering, University of Pennsylvania, Philadelphia, PA (schotland@seas.upenn.edu). The research of this author was partially supported by NSF grant DMS05-54100.

The Laplace transform is defined by $k(t) = e^{-t}$. A comprehensive exposition of the classical theory of the Laplace transform is given in [1]. Fast algorithms for the forward transform are given in [11] and [8]. In this note we present a method for the rapid computation of both the forward and inverse transforms for linear operators of this type. Our approach is essentially that of a “twisted” eigenfunction expansion. The underlying unitary transformation from $L^2([0, \infty); dx)$ to $L^2((-\infty, \infty); \frac{ds}{2\pi})$ is defined by

$$(1.2) \quad \tilde{f}(s) = \int_0^\infty f(x) x^{-\frac{1}{2}-is} dx.$$

This transform is of course nothing but a slight reparametrization of the Mellin transform. The application of this transform to study operators of the type in (1.1) appears in [7].

It is easy to see the connection between the transform, $f \mapsto \tilde{f}$, and operators whose kernels are functions of xy . An elementary change of variables shows that, for $\operatorname{Re} \alpha > -1$,

$$(1.3) \quad \mathcal{K}x^\alpha = \bar{k}(\alpha)x^{-1-\alpha},$$

where, provisionally, we set

$$(1.4) \quad \bar{k}(\alpha) = \int_0^\infty k(t)t^\alpha dt.$$

Hence, for any real s the subspace generated by $\{x^{-\frac{1}{2}+is}, x^{-\frac{1}{2}-is}\}$ is an invariant subspace under the action of $\mathcal{K}$. The mapping properties of $\mathcal{K}$ and $\mathcal{K}^{-1}$ are determined by the behavior of the function $\bar{k}(-\frac{1}{2}-is) = \tilde{k}(s)$ for $s \in \mathbb{R}$. For the classical Laplace transform, $k(t) = e^{-t}$ and

$$(1.5) \quad \tilde{k}(s) = \Gamma\left(\frac{1}{2} - is\right).$$

The well-known difficulties of inverting the Laplace transform stem from the fact that

$$(1.6) \quad \left| \Gamma\left(\frac{1}{2} - is\right) \right| = \sqrt{\frac{\pi}{\cosh \pi s}}.$$

We study the harmonic analysis and sampling theory relevant to the transform $f \mapsto \tilde{f}$ and then explain how to use it to approximate $\mathcal{K}$ and $\mathcal{K}^{-1}$. We then consider how the regularized inverse affects noisy measurements, and we show that translational stationarity of the noise process interacts in a nasty way with the multiplicative group structure underlying the inversion process. We also investigate how the choices of low-pass filter and cut-off function, used in signal acquisition, affect the reconstructed signal; an unfortunate choice here can induce a surprisingly severe distortion of the reconstructed signal. The paper concludes with numerical experiments.

2. Harmonic Analysis. The basic observation is that the transform $f \mapsto \tilde{f}$ is simply related to the Fourier transform. This gives the Parseval theorem and inversion formulae for smooth data with compact support. For such data the map $f \mapsto \tilde{f}$ and its inverse are defined by absolutely convergent integrals.

THEOREM 2.1. *If f is a smooth function with compact support in $(0, \infty)$, then*

$$(2.1) \quad \int_0^\infty |f(x)|^2 dx = \frac{1}{2\pi} \int_{-\infty}^\infty |\tilde{f}(s)|^2 ds$$

and

$$(2.2) \quad f(x) = \frac{1}{2\pi} \int_{-\infty}^\infty \tilde{f}(s) x^{-\frac{1}{2}+is} ds.$$

Proof. We use a simple change of dependent and independent variables, setting $f(x) = \frac{g(x)}{\sqrt{x}}$ and $x = e^t$. With these changes of variables we see that

$$(2.3) \quad \tilde{f}(s) = \int_{-\infty}^\infty g(e^t) e^{-its} dt.$$

Both statements in the theorem now follow from the change-of-variables formula, the standard Parseval theorem, and the Fourier inversion formula applied to $G(t) = g(e^t)$. $\square$

There is a natural notion of convolution connected to this transform: we define $f \star g(x)$ by

$$(2.4) \quad f \star g(x) = \int_0^\infty f\left(\frac{x}{y}\right) g(y) \frac{dy}{y}.$$

Note that this is *not* the usual convolution associated to the Laplace transform, which is defined by

$$(2.5) \quad f * g(x) = \int_0^x f(y) g(x-y) dy$$

and satisfies $\mathcal{L}(f * g) = \mathcal{L}f \cdot \mathcal{L}g$. There is a formula for $\widetilde{f \star g}$ in terms of $\tilde{f}$ and $\tilde{g}$, but it is complicated and requires analytic continuation. On the other hand, a simple calculation proves the following proposition.

PROPOSITION 2.2. *With $f \star g$ defined in (2.4), we have*

$$(2.6) \quad \widetilde{f \star g}(s) = \tilde{f}(s) \tilde{g}(s).$$

Using a standard density argument we obtain an extension of $f \mapsto \tilde{f}$ as a unitary operator from $L^2([0, \infty); dx)$ to $L^2((-\infty, \infty); \frac{ds}{2\pi})$, and the inversion formula holds in the “limit-in-the-mean” sense. The analogue of the Nyquist sampling theorem follows easily, given the intimate connection with the Fourier transform.

PROPOSITION 2.3. *If $f(x)$ is supported in the interval $[L^{-1}, L]$ for an $L > 1$, then f is determined by any set of samples of $\tilde{f}$ of the form*

$$(2.7) \quad \{\tilde{f}(s_0 + j\Delta s) : j \in \mathbb{Z}\},$$

provided that $\Delta s \leq \frac{\pi}{\log L}$. On the other hand, if $\tilde{f}$ is supported in an interval $[-B, B]$, then f is determined by the samples

$$(2.8) \quad \{f(\lambda_0 \lambda^j) : j \in \mathbb{Z}\}$$

for any $\lambda_0 \in (0, \infty)$, provided $\lambda < e^{\frac{\pi}{B}}$.

Recalling that $g(t) = G(e^t)$, these statements follow easily from the fact that $\tilde{f}(s) = \widehat{G}(s)$; see (2.3). We also have natural discrete and finite analogues of this

transform. From the identity

$$(2.9) \quad \int_0^\infty f(x)x^{-\frac{1}{2}-is}dx = \int_{-\infty}^\infty f(e^t)e^{\frac{t}{2}}e^{-its}dt,$$

we see that the correct finite sum approximation, given the evenly spaced data in (2.8), is

$$(2.10) \quad \tilde{f}(s) \approx \lambda_0^{\frac{1}{2}-is} \sum_{j=-\infty}^\infty \lambda^{\frac{j}{2}} f(\lambda_0 \lambda^j) \lambda^{-isj} \log \lambda.$$

With the samples in (2.7) we get an approximate inversion formula

$$(2.11) \quad f(x) \approx \frac{x^{-\frac{1}{2}+s_0}}{2\pi} \sum_{j=-\infty}^\infty \tilde{f}(s_0 + j\Delta s) x^{ij\Delta s} \Delta s.$$

The right-hand side in (2.11) is a log-periodic function, with log-period $\frac{2\pi}{\Delta s}$. If N is a power of two, then the finite versions of these transforms, using N -samples, can be computed using order of $N \log_2 N$ operations. Indeed the difference between the finite Fourier transform and the finite versions of the transforms in (2.10)–(2.11) involves multiplication by diagonal matrices.

3. Analysis of $\mathcal{K}$. To analyze the operator $\mathcal{K}$ we use the inversion formula (2.2). Applying $\mathcal{K}$ to both sides gives

$$(3.1) \quad \begin{aligned} \mathcal{K}f(x) &= \frac{1}{2\pi} \int_{-\infty}^\infty \tilde{f}(s) \mathcal{K}(x^{-\frac{1}{2}+is}) ds \\ &= \frac{1}{2\pi} \int_{-\infty}^\infty \tilde{f}(s) \tilde{k}(-s) x^{-\frac{1}{2}-is} ds. \end{aligned}$$

This proves the following formula:

$$(3.2) \quad \widetilde{\mathcal{K}f}(s) = \tilde{k}(s) \tilde{f}(-s).$$

Using the finite version of $f \mapsto \tilde{f}$ and its inverse, we can use (3.2) to obtain a fast ($O(N \log_2 N)$) algorithm for approximately computing $\mathcal{K}f(x)$, provided samples of f are collected on a log-uniformly spaced sample set, as in (2.8). Both Rokhlin and Strain have given $O(N)$ algorithms for the forward Laplace transform with samples on essentially arbitrary sets; see [8, 11]. The $O(N \log_2 N)$ bound assumes that the necessary values of $\{\tilde{k}(s_j)\}$ have been computed and stored.

The Parseval formula implies that the transform $f \mapsto \mathcal{K}f$ is bounded as a map from $L^2([0, \infty))$ to itself if and only if $\|\tilde{k}(s)\|_{L^\infty(\mathbb{R})} < \infty$. This formalism extends to tempered kernels k for which $\tilde{k}(s)$ is defined distributionally. If $\tilde{k}(-s)$ is a bounded measurable function, then $\mathcal{K}$ is boundedly invertible if the essential infimum of $|\tilde{k}(-s)|$ is positive. In this case we have the following formula for the inverse:

$$(3.3) \quad \mathcal{K}^{-1}g(x) = \frac{1}{2\pi} \int_{-\infty}^\infty \frac{\tilde{g}(-s)}{\tilde{k}(-s)} x^{-\frac{1}{2}+is} ds.$$

More generally we obtain a regularized inverse by choosing a measurable cut-off function $\tilde{\psi}$ such that $\tilde{\psi}(t) = 1$ for t sufficiently large and $\tilde{\psi}(t) = 0$ in a neighborhood of $t = 0$. The regularized inverse defined by $\tilde{\psi}$ is given by

$$(3.4) \quad \mathcal{K}_\psi^{-1}g(x) \stackrel{d}{=} \frac{1}{2\pi} \int_{-\infty}^\infty \frac{\tilde{g}(-s) \tilde{\psi}(|\tilde{k}(-s)|)}{\tilde{k}(-s)} x^{-\frac{1}{2}+is} ds.$$

The cut-off function can be smooth or sharp. Depending upon the data, a function that approaches zero sufficiently rapidly as $t \rightarrow 0$, and one as $t \rightarrow \infty$, for example, $e^{-\frac{1}{t^2}}$, could also be used.

The adjoint operator $\mathcal{K}^*$ to $\mathcal{K}$ is defined by the kernel function $\bar{k}(t)$. It is a simple calculation to show that the multiplier defined by the self-adjoint operator $\mathcal{K}^*\mathcal{K}$ is $|k(-s)|^2$. This shows that the generalized singular values of $\mathcal{K}$ are simply the values of $|k(-s)|$ for $s \in \mathbb{R}$. Thus the regularized inverse in (3.4) is very close in spirit to that given by truncating a singular value decomposition, but without the necessity of finding the exact singular vectors. A Tikhonov-type regularized inverse is given by

$$(3.5) \quad \mathcal{K}_\lambda^{-1}g(x) \stackrel{d}{=} \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{\tilde{g}(-s)\overline{\tilde{k}(-s)}}{|\tilde{k}(-s)|^2 + \lambda^2} x^{-\frac{1}{2}+is} ds.$$

As the spectral formulae for $\mathcal{K}_\psi^{-1}$ and $\mathcal{K}_\lambda^{-1}$ involve the transforms $f \leftrightarrow \tilde{f}$ and simple multiplication operators for $N = 2^n$, these operators can be implemented in $O(N \log N)$ operations using the fast discrete versions of these transforms. As before, this assumes that the needed samples of the multiplier $\{\tilde{k}(-s_j)\}$ have been computed in advanced and stored.

REMARK 1 (some history). *Much of the analysis presented in this section, including the inversion formula (3.3), appears in a 1978 paper of McWhirter and Pike [7]. In this paper, ideas from information theory are applied to quantify the information content in the Laplace transform of a function. A detailed spectral analysis of the Laplace transform, building on the work of McWhirter and Pike, appears in a 1998 paper of Boumenir and Al-Shuaibi [3]. The work of McWhirter and Pike was, in turn, preceded by a 1959 paper of Gardner, Gardner, and Meinke [5], which contains a similar, though not identical, formula for the case of the Laplace transform. A comparative review of numerical techniques for inverting the Laplace transform is given in the 1979 paper of Davies and Martin [4]. In [2], Bertero and Grunbaum consider similar questions for the Laplace transform of compactly supported data. It seems that the horror of the inverse Laplace transform is something each generation must discover for itself.*

REMARK 2 (a little operator algebra). *Operators with kernels of the form $k(xy)$ are not quite an algebra. If we compose two such kernels, then we get a kernel of "convolution type," $h(\frac{x}{y})y^{-1}$:*

$$(3.6) \quad \begin{aligned} h\left(\frac{x}{y}\right) \frac{1}{y} &= \int_0^\infty k_1(xz)k_2(zy)dz \\ &= \frac{1}{y} \int_0^\infty k_1\left(t\frac{x}{y}\right)k_2(t)dt. \end{aligned}$$

Composing a kernel of convolution type with a function of xy again leads to a kernel which is a function of xy . Composing two kernels of convolution type gives a kernel of convolution type. If we let $\mathcal{A}_0$ be the operators of convolution type and $\mathcal{A}_1$ those with kernels depending on xy , then we see that $\mathcal{A}_0 \oplus \mathcal{A}_1$ is a $\mathbb{Z}_2$ -graded algebra. The subalgebra $\mathcal{A}_0$ is commutative, though $\mathcal{A}_1$ is not.

4. The Classical Laplace Transform. The approach in section 3 applies to the analysis of the classical Laplace transform. In this case $k(t) = e^{-t}$, so that

$$(4.1) \quad \tilde{k}(-s) = \Gamma\left(\frac{1}{2} + is\right).$$

The Laplace transform is a self-adjoint operator and the invariant subspaces

$$\{x^{-\frac{1}{2}+is}, x^{-\frac{1}{2}-is}\}$$

can be further split into generalized eigenspaces. In fact, there is a smooth, real-valued function $\phi(s)$ so that the eigenspaces are spanned by

$$(4.2) \quad \varphi_+(s; x) = \operatorname{Re}(e^{i\phi(s)} x^{-\frac{1}{2}+is}), \quad \varphi_-(s; x) = \operatorname{Im}(e^{i\phi(s)} x^{-\frac{1}{2}+is}).$$

The corresponding eigenvalues are given by

$$(4.3) \quad \lambda_{\pm}(s) = \pm \sqrt{\frac{\pi}{\cosh \pi s}}.$$

The completeness of the generalized eigenbasis follows easily from Theorem 2.1 and implies the following result.

PROPOSITION 4.1. *The Laplace transform*

$$(4.4) \quad \mathcal{L}f(x) = \int_0^\infty e^{-xy} f(y) dy$$

is a bounded self-adjoint operator on $L^2([0, \infty))$, with purely absolutely continuous spectrum, of multiplicity one, lying in the interval $[-\sqrt{\pi}, \sqrt{\pi}]$.

This result appears in [3]. From (4.3) it follows that the information in the Laplace transform decays very rapidly with increasing oscillation. It gives a quantitative explanation of the notorious difficulty of retaining significant detail when inverting the Laplace transform. If $\tilde{\psi}$ is a function tending to zero sufficiently rapidly, as $|s| \rightarrow \infty$, then a regularized inverse for the Laplace transform is given by

$$(4.5) \quad \mathcal{L}_\psi^{-1}g(x) \stackrel{d}{=} \frac{1}{2\pi} \int_{-\infty}^\infty \frac{\tilde{g}(-s)\tilde{\psi}(s)}{\Gamma(\frac{1}{2}+is)} x^{-\frac{1}{2}+is} ds.$$

Note that we use a slightly simplified notation for this special case.

From (4.5) we can derive a bound on the resolution available in $\mathcal{L}_\psi^{-1}F$, given the accuracy of our measurements of $F = \mathcal{L}f$ and the desired accuracy in the reconstruction. The singular values of $\mathcal{L}$ corresponding to the frequencies $\pm s$ have magnitude about $\sqrt{2\pi}e^{-\frac{\pi|s|}{2}}$. If the uncertainty in our measurements is ϵ and we are willing to tolerate an uncertainty of $\eta > \epsilon$ in our output, then this approximation for $|\Gamma(\frac{1}{2}+is)|$ shows that the essential support of $\tilde{\psi}$ should lie in the interval $[-s_{\max}, s_{\max}]$, where

$$(4.6) \quad s_{\max} \leq \frac{2}{\pi} \log \left[\sqrt{2\pi} \frac{\eta}{\epsilon} \right].$$

Measured on a *logarithmic scale*, the maximum spatial resolution of $\mathcal{L}_\psi^{-1}F$ is therefore roughly $\pi/s_{\max}$. As an example, suppose that $\epsilon = 10^{-10}$ and $\eta = 10^{-1}$; then $s_{\max} \approx 24$. Thus, close to $x = 1$, we get a spatial resolution of about .1 and an accuracy of about 10^{-1} , provided the data is measured with 10 significant digits! On the bright side, not many terms are required to do the inversion. This estimate is consistent with equation (19) in [10].

The generalized eigenbasis $\{x^{-\frac{1}{2}+is}\}$ can also be profitably used to analyze another ill-conditioned operator of general interest: the continuous Hilbert matrix, $\mathcal{H}$. This follows because the kernel function of $\mathcal{L}^2$ is given by

$$(4.7) \quad \frac{1}{x+y} = \int_0^\infty e^{-xz} e^{-zy} dz.$$

Indeed, it is elementary to see that

$$(4.8) \quad \mathcal{H}(x^{-\frac{1}{2}+is}) = \frac{\pi}{\cosh \pi s} x^{-\frac{1}{2}+is} \text{ for } s \in \mathbb{R}.$$

This gives a different proof that the spectrum of $\mathcal{H}$ is precisely $[0, \pi]$ with multiplicity two. Moreover, at least for log-uniformly spaced samples, we can rapidly compute $\mathcal{H}f$ and a regularized inverse for $\mathcal{H}$. A regularized inverse is given by

$$(4.9) \quad \mathcal{H}_\psi^{-1}g(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{\psi}(s) \frac{\cosh \pi s}{\pi} \tilde{g}(s) x^{-\frac{1}{2}+is} ds.$$

5. Noise and Filtration Analysis. In many different experimental contexts one measures samples of the Laplace transform of a function f , which are inevitably filtered and contaminated by noise. In this section we examine how regularized inverses, of the type given in (4.5), affect the noise variance. We consider the case that the measurement process operates in continuous time; the noise is modeled as a white noise process $n(t)$ with mean zero and covariance

$$(5.1) \quad \langle n(t)n(s) \rangle = \sigma^2 \delta(t-s).$$

We then turn to the interaction of the regularized inverse with a variety of standard filtering operations.

Our model for the measured data is

$$(5.2) \quad \begin{aligned} M(t) &= \int_0^\infty f(x) e^{-xt} dx + n(t) \\ &= m(t) + n(t). \end{aligned}$$

Suppose that $\tilde{\psi}(s)$ is a cut-off function, and set

$$(5.3) \quad g_\psi(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{\tilde{g}(-s)\tilde{\psi}(s)}{\Gamma(\frac{1}{2}+is)} x^{-\frac{1}{2}+is} ds.$$

The function reconstructed from the measurements using (5.3) is $m_\psi + n_\psi$, where, by virtue of (2.6), we can express the terms as

$$(5.4) \quad m_\psi(x) = f \star \psi(x)$$

and

$$(5.5) \quad \begin{aligned} n_\psi(x) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{\Psi}(s) \int_0^\infty n(y)(xy)^{-\frac{1}{2}+is} dy ds \\ &= \int_0^\infty n(y) \Psi(xy) dy. \end{aligned}$$

Here Ψ is the inverse transform of

$$(5.6) \quad \tilde{\Psi}(s) = \frac{\tilde{\psi}(s)}{\Gamma(\frac{1}{2}+is)}.$$

We assume that Ψ is a square integrable function. As

$$(5.7) \quad \left| \Gamma\left(\frac{1}{2}+is\right) \right| = \sqrt{\frac{\pi}{\cosh \pi s}},$$

this means that $\tilde{\psi}$ is rapidly decreasing and therefore that ψ must be smooth.

As $n_\psi(x)$ is a linear combination of mean zero random variables, it follows that $\langle n_\psi(x) \rangle = 0$ for all $x > 0$. Using (5.1), we now compute the covariance:

$$(5.8) \quad \begin{aligned} \langle n_\psi(x) \overline{n_\psi(y)} \rangle &= \left\langle \int_0^\infty \int_0^{\infty} \Psi(xu) n(u) \overline{\Psi(yv) n(v)} du dv \right\rangle \\ &= \sigma^2 \int_0^\infty \Psi(xu) \overline{\Psi(yu)} du. \end{aligned}$$

Letting $\tau = xu$ in the last integral, we obtain

$$(5.9) \quad \langle n_\psi(x) \overline{n_\psi(y)} \rangle = \frac{\sigma^2}{x} \int_0^\infty \Psi(\tau) \overline{\Psi\left(\frac{y}{x}\tau\right)} d\tau.$$

Using the Parseval relation, this becomes

$$(5.10) \quad \langle n_\psi(x) \overline{n_\psi(y)} \rangle = \frac{\sigma^2}{\sqrt{xy}} \int_{-\infty}^\infty |\tilde{\Psi}(s)|^2 \left(\frac{y}{x}\right)^{-is} ds.$$

In the natural complete metric, $\frac{dx}{x}$, of $\mathbb{R}_+^\times$, the distance from x to y is

$$d^\times(x, y) = \left| \log \frac{y}{x} \right|.$$

Hence, if $\tilde{\Psi}$ is smooth, then the correlations are rapidly decreasing as $d^\times(x, y)$ diverges. On the other hand, we see that the covariance diverges as xy tends to zero. This would seem to be a result of the interaction between the translational symmetry of the noise process and the multiplicative symmetry of the Laplace transform. The noise process is, in other words, “unaware” of the underlying group structure implicit in the measured signal and the inversion process. Evaluating at $x = y$, we get

$$(5.11) \quad \langle |n_\psi(x)|^2 \rangle = \frac{\sigma^2}{x} \int_{-\infty}^\infty |\tilde{\Psi}(s)|^2 ds.$$

Thus without filtration the white noise causes the variance in $M_\psi(x)$ to diverge as $x \rightarrow 0$. In the next section, we see this prediction strikingly confirmed.

A realistic measurement process involves low-pass filtering, which we model as (ordinary) convolution with a compactly supported function φ , i.e.,

$$(5.12) \quad \begin{aligned} M_\varphi(t) &= \int_{-\infty}^\infty \left[\int_0^\infty f(x) e^{-x\tau} dx + n(\tau) \right] \varphi(t - \tau) d\tau \\ &= m_\varphi(t) + n_\varphi(t). \end{aligned}$$

As the Laplace transform is only defined for $t > 0$, one might want to restrict φ to be supported in $(-\infty, 0]$. A straightforward calculation shows that the formula for the covariance is replaced by

$$(5.13) \quad \langle n_{\varphi\psi}(x) \overline{n_{\varphi\psi}(y)} \rangle = \sigma^2 \int_0^\infty \int_0^\infty \Psi(xu) \overline{\Psi(yv)} \Phi(u - v) du dv,$$

where

$$(5.14) \quad \Phi(z) = \int_{-\infty}^\infty \varphi(t) \varphi(z + t) dt.$$

This formula makes it clear that the final answer results from a combination of the additive and multiplicative group structures on $\mathbb{R}$ and $\mathbb{R}_+$, respectively.

Even if Φ is smooth and with compact support, the variance $\langle |n_{\varphi\psi}(x)|^2 \rangle$ typically diverges as x tends to zero. To see this we observe that if we change variables, letting $a = xu$ and $b = xv$, then

$$(5.15) \quad \langle |n_{\varphi\psi}(x)|^2 \rangle = \sigma^2 \int_0^\infty \int_0^\infty \Phi\left(\frac{a-b}{x}\right) \frac{\Psi(a)\overline{\Psi(b)}dad b}{x^2}.$$

If Ψ and Φ are nonnegative and Φ is bounded from below in the interval $[-\alpha, \alpha]$, then the variance is bounded below by a constant multiple of

$$(5.16) \quad \frac{1}{x^2} \iint_{|a-b| < \alpha x} |\Psi(a)\Psi(b)|dad b.$$

It is not difficult to see that the integral in (5.16) behaves like $O(x)$ as x tends to zero. Hence, the covariance still diverges like x^{-1} . That the covariance should still diverge as $x \rightarrow 0$ is not that surprising, as $n_\varphi(t)$ does not decay as $t \rightarrow \infty$ and so it cannot be the Laplace transform of a bounded function. If noise is present in the measurements (or even numerical noise caused by round-off error in an approximate inversion formula), then these formulae indicate the difficulty of using measurements of $\mathcal{L}f$ to reliably determine values of $f(x)$ for small values of x .

It is interesting to further investigate the signal part m_φ of M_φ in (5.12). A simple change of variables shows that, ignoring the noise, we actually measure the following Laplace transform:

$$(5.17) \quad \int_0^\infty f(x)\varphi_{\mathcal{L}}(x)e^{-xt}dx,$$

where

$$(5.18) \quad \varphi_{\mathcal{L}}(x) = \int_{-\infty}^\infty e^{xw}\varphi(w)dw.$$

The regularized inverse then gives

$$(5.19) \quad f_{\varphi\psi}(x) = \int_0^\infty \psi\left(\frac{x}{y}\right) f(y)\varphi_{\mathcal{L}}(y)\frac{dy}{y}.$$

Hence, we obtain a smeared version of $f\varphi_{\mathcal{L}}$.

If, for example, $\varphi(x) = (\epsilon)^{-1}\chi_{[0,\epsilon]}(x)$, then

$$(5.20) \quad \varphi_{\mathcal{L}}(x) = \frac{e^{\epsilon x} - 1}{\epsilon x}.$$

This shows that the signal filtration process may have the effect of *exponentially* amplifying the values of f as $x \rightarrow \infty$. If we instead use $\varphi(x) = \epsilon^{-1}\chi_{[-\epsilon,0]}(x)$, then

$$(5.21) \quad \varphi_{\mathcal{L}}(x) = \frac{1 - e^{-\epsilon x}}{\epsilon x},$$

which does not exponentially amplify values of f , but rather linearly attenuates them, as $x \rightarrow \infty$. As noted above, this suggests that one should use an anticausal filter to

avoid exponential growth of the reconstructed signal. This can be thought of as essentially a labeling question: Suppose that the j th sample is a weighted average of values of $\mathcal{L}f$ over an interval $[t_j, t_{j+1}]$. For the purposes of inversion, this should be interpreted as the measured value of $\mathcal{L}f$ at the time t_j .

Windowing is one final type of filtering often done in acquisition of signals. When the measured signal has an interpretation as a Laplace transform, then it is usually a rapidly decaying function of time. The first few samples may be unreliable, and so we replace the measured data (see (5.12)) with $X(t)M(t)$; here X is a function vanishing at zero and rapidly rising to 1. If $X = \mathcal{L}(\chi)$, then it follows from the remarks in section 2 that

$$(5.22) \quad \mathcal{L}^{-1}(XM) = \chi * \mathcal{L}^{-1}(M).$$

In light of the extraordinary instability of $\mathcal{L}^{-1}$, it's difficult to say definitively whether or not such windowing will have a dramatic effect on the result of applying $\mathcal{L}_\psi^{-1}$. In numerical experiments this sort of windowing does not have important effects for some choices of window function and can dramatically degrade the reconstruction for other choices.

To see why this might be the case, we write $X(t) = 1 - \xi(t/\delta)$, where $\xi(0) = 1$ and ξ decays rapidly to zero as $t \rightarrow \infty$. Now suppose that there is a smooth function η so that $\xi = \mathcal{L}(\eta)$. The behavior of $\xi(t)$ as $t \rightarrow \infty$ is, in general, determined by the behavior of $\eta(x)$ near to $x = 0$. In particular, if $\xi(t) = O(t^{-(1+k)})$ at infinity for a $k > -1$, then $\eta(x) = O(x^k)$ near $x = 0$. A simple calculation shows that $\xi(t/\delta)$ is the Laplace transform of $\eta_\delta(x) = \delta\eta(\delta x)$. Thus, assuming that $\eta(x) = O(x^k)$, we see that the error incurred by multiplying the measurements by $X(t)$ can be expressed as

$$(5.23) \quad \begin{aligned} \eta_\delta * f_\psi \mathcal{K}(x) &= \delta \int_0^x \eta(y\delta) f_\psi(x-y) dy \\ &\propto \delta \int_0^x (y\delta)^k f_\psi(x-y) dy. \end{aligned}$$

Under the assumptions above, this evidently is a very modest source of error.

If the measured data is contaminated with noise, then it may be advisable to cut off, as $t \rightarrow \infty$, as well. A collection of functions, useful for this purpose, is provided by the functions

$$(5.24) \quad \chi^k(x) = \frac{x^k e^{-x}}{(k+1)!},$$

with Laplace transforms

$$(5.25) \quad \mathcal{L}(\chi^k) = \frac{1}{(1+t)^{k+1}}, \quad k \in \mathbb{N}.$$

The functions $\chi_\epsilon^k(x) = \epsilon^{-1} \chi^k(\epsilon^{-1}x)$ are easily seen to define an approximate identity for the $*$ -convolution, with

$$(5.26) \quad \mathcal{L}(\chi_\epsilon^k)(t) = \frac{1}{(1+\epsilon t)^{k+1}}.$$

Replacing the noise with $n_\varphi(t)(1+\epsilon t)^{-(k+1)}$ has the effect of controlling the divergence of the noise variance, as $x \rightarrow 0$, in the reconstructed function. It should also be

noted that $*$ -convolution with χ_ϵ^k has the effect of shifting peaks to the right by approximately ϵk . Of course, a slightly different choice of windowing function can produce dramatically different results. Indeed using the functions $(1+t^2)^{-k}$ instead of $(1+t)^{-2k}$ leads to much worse artifacts. It should be noted that $(1+t^2)^{-1} = \mathcal{L}(\sin x)$.

REMARK 3. *The Laplace transform can be applied to a large class of measures supported on $(0, \infty)$. Above we consider the case of absolutely continuous measures with reasonably smooth densities. In many applications one encounters atomic measures of the form*

$$(5.27) \quad f(x) = \sum_{j=1}^N a_j \delta(x - R_j), \text{ with } 0 < R_1 < R_2 < \cdots < R_N,$$

with the coefficients $\{a_j\}$ positive. This case arises frequently in NMR applications, where $R_j = T_{2j}^{-1}$ are spin-spin decay rates for different chemical species within a sample. The measurements are samples of

$$(5.28) \quad \mathcal{L}f(t) = \sum_{j=1}^N a_j e^{-tR_j},$$

usually taken at equally spaced points, and contaminated with noise. The methods introduced in the continuum case are not likely to work well in the present case, as

$$(5.29) \quad \widetilde{\mathcal{L}}f(s) = \Gamma\left(\frac{1}{2} - is\right) \sum_{j=1}^N a_j R_j^{-\frac{1}{2} + is},$$

and therefore $\widetilde{\mathcal{L}}f(-s)[\Gamma(\frac{1}{2} + is)]^{-1}$ does not decay as $|s| \rightarrow \infty$. Multiplying such data in the time domain by $(1 + \epsilon t)^{-(k+1)}$ considerably stabilizes the reconstruction process. This approach will be explored in a subsequent publication. Many algorithms have been introduced to handle this important case. Interesting examples are given in [5]. Suffice it to say that this inversion problem is also exponentially ill-posed. See Example 4.

6. Numerical Examples. We present examples to illustrate the behavior of the regularized inverse to the Laplace transform and the effects of noise on the reconstruction. The data we collect consists of log-uniformly equally spaced samples of the form

$$(6.1) \quad \{F(e^{kd}) : -(k_0 + N) \leq k \leq N - k_0 - 1\}.$$

In our experiments the choice of sample spacing d and offset k_0 has a very significant effect on the quality of the reconstruction. To implement the algorithm we use the fast Fourier transform (FFT) to compute approximations for the samples $\{\widetilde{F}(s_j)\}$. The frequencies are

$$(6.2) \quad \left\{ s_j = \frac{2\pi j}{2Nd} : j = -N, \dots, N-1 \right\}.$$

To employ the FFT we need to scale as indicated in Theorem 2.1, setting

$$(6.3) \quad g_k = e^{\frac{(k_0+k)d}{2}} F(e^{(k_0+k)d}).$$

We let $\{\hat{g}_k\}$ denote the $2N$ -point discrete Fourier transform of this sequence.

The values $\{\Gamma(\frac{1}{2} + is_j)\}$ are computed and a cut-off function $\tilde{\psi}$ is selected. If $F = \mathcal{L}(f)$, then a computation shows that

$$(6.4) \quad f_{\psi}(e^{(m-k_0)d}) \approx \frac{e^{-\frac{(m-k_0)d}{2}}}{2N} \sum_{l=1-N}^N \frac{\hat{g}_{-l} \tilde{\psi}\left(\frac{2\pi ilm}{2N}\right) e^{\frac{2\pi ilm}{2N}}}{\Gamma\left(\frac{1}{2} + \frac{2\pi il}{2Nd}\right)}.$$

The change in the sample offset from $k + k_0$ to $m - k_0$ is a consequence of the fact that $\tilde{F}$ is evaluated at $-s$ in (4.5). The postmultiplication by exponential weights in (6.4) means that, with floating point arithmetic, the meaningful dynamical range of the computed values can vary dramatically from sample to sample. For our filter function we use

$$\tilde{\psi}(s) = \exp\left(-\frac{e^{\pi|s|}}{B}\right).$$

If $B = 10^k$, then the locus of points where

$$(6.5) \quad \frac{\tilde{\psi}(s)}{|\Gamma(\frac{1}{2} + is)|} = 10^{-15}$$

is given,

$$(6.6) \quad s \approx 0.733(k + 2.71);$$

hence, the theoretical logarithmic resolution grows in proportion to $\log B$.

For our numerical experiments we use two Laplace transform pairs:

$$(6.7) \quad \mathcal{L}(xe^{-x}) = \frac{1}{(1+t)^2} \quad \text{and} \quad \mathcal{L}(\sin x) = \frac{1}{1+t^2}.$$

We use the function $\frac{1}{(1+t)^2}$ as an example of a “good” Laplace transform, which is “easy” to invert, whereas the very similar function, $\frac{1}{1+t^2}$, is a “bad” Laplace transform, which is very hard to invert.

Example 1. For our first experiments we sample $\mathcal{L}(xe^{-x})$ and $\mathcal{L}(\sin x)$ at the points

$$(6.8) \quad \{e^{jd} : -2^{10} - 2^9 \leq j \leq 2^{10} - 2^9 - 1\},$$

where $d = 0.0488$. Note that the sample set is asymmetric about 1. The quality of the reconstruction depends sensitively on this choice of sample points. The choice of log-sample spacing, d , also dramatically affects the quality of the approximate inverse. In our cut-off function we set $B = 10^{20}$. These parameters have been chosen to give an “optimal” result for the function $1/(1+t)^2$. Figure 6.1(a) shows the computed approximation to $\mathcal{L}_{\psi}^{-1}(1/(1+t)^2)$, with a linear scale on the x -axis, and Figure 6.1(b) shows the same function with a $\log_2$ -scale on the x -axis. Figures 6.1(c) and (d) show the same results for $\mathcal{L}_{\psi}^{-1}(1/(1+t^2))$. The very large filter bandwidth used in these examples is needed to get a “reasonable” reconstruction of $\sin x$ over even the few cycles shown in Figure 6.1(c). One can get a good reconstruction of xe^{-x} with a much smaller bandwidth. In (c) we show the exact function as a dashed line. In (a) and (b) the reconstruction is indistinguishable from the exact result.

Example 2. Before considering the effects of noise, per se, we show how the reconstructions are altered if we change some of the parameters above. In Figures 6.2(a)

Fig. 6.1 Approximate inverse with well-chosen parameters and noise-free data. In (c) the exact result is shown as a dashed line.

and (b) we use samples of $\frac{1}{(1+t)^2}$ at the points in (6.8), with $d = 0.15$ or $d = 0.02$, and all other parameters the same as above. For Figure 6.2(c) we used the more symmetric set of sample points,

$$(6.9) \quad \{e^{jd} : -2^{10} - 2^5 \leq j \leq 2^{10} - 2^5 - 1\},$$

with all other parameters “optimal.” In Figure 6.2(d) we show the result of using the sample set in (6.9) and $d = 0.15$. The x -axes in these plots are shown in a $\log_2$ -scale. The errors are growing rapidly as x tends to zero, in good accord with the analysis in the previous section. In Figure 6.2(a) the maximum error is about 2.5×10^8 and in Figure 6.2(d) about 5×10^{23} ! Despite this, the reconstruction of xe^{-x} near to $x = 1$ (not shown) is quite accurate. This is possible because, as noted above, our reconstruction algorithm involves postmultiplication by exponential weights. Hence a reconstructed function can have a very large, x -dependent, but meaningful dynamic range. In these examples, the only sources of noise are round-off error and numerical

(a) The approximate Laplace inverse for the “good” function, with no noise and $d = 0.15$; $y\text{-max} = 2.5 \times 10^8$.

(b) The approximate Laplace inverse for the “good” function, with no noise and $d = 0.02$; $y\text{-max} = 200$.

(c) The approximate Laplace inverse for the “good” function, with no noise and a more symmetric sample set; $y\text{-max} = 8 \times 10^6$.

(d) The approximate Laplace inverse for the “good” function, with no noise, a more symmetric sample set, and $d = 0.15$; $y\text{-max} = 5 \times 10^{23}$.

Fig. 6.2 Details of the approximate inverse near $x = 0$, with poorly chosen parameters and noise-free data. The x -axes are shown with a $\log_2$ -scale; note the rapid divergence as $x \rightarrow 0$.

errors resulting from using a finite sum approximation to the Fourier transform. These produce the blow-up apparent in Figure 6.2 as $x \rightarrow 0$.

Example 3. We next examine what happens if the data is contaminated by noise. We consider only the good function $\frac{1}{(1+t)^2}$ and reduce the bandwidth B to 10^4 . We evaluate this function on the sample set in (6.8), with $d = 0.0488$. To these samples we add σ times a vector of random data, uniformly distributed in $[-1, 1]$. In Figure 6.3(a), we show the detail near 0, on the $\log_2$ -scale, with $\sigma = 10^{-4}$. While the reconstruction error is extremely large near $x = 0$, (1.5×10^7), the reconstruction of xe^{-x} still looks fine near $x = 1$. In Figure 6.3(b) we show the detail near 0, on the $\log_2$ -scale, with $\sigma = 10^{-2}$. The maximum error is about 1.5×10^9 ; Figure 6.3(c) shows the reconstruction of xe^{-x} with this data, which is clearly no longer usable. The bandwidth used in this example is very small (the cut-off function is less than 10^{-15} for $|s|$ a little larger than 4) and would probably not be usable for most realistic input data. For purposes of comparison, in Figure 6.3(d), we show the approximate inverse,

Fig. 6.3 The effects of noise in the approximate Laplace inverse. In (c) and (d) the exact result is shown as a dashed line.

with $B = 10^4$, applied to samples of $\frac{1}{1+t^2}$. In (c) and (d) we show the exact function as a dashed line, which now differs markedly from the reconstructions in both cases.

Example 4. For our last example we consider a type of function that arises in many practical contexts, e.g., in NMR:

$$(6.10) \quad F(t) = e^{-t} + e^{-2t} + e^{-4t} = \mathcal{L}(\delta(x-1) + \delta(x-2) + \delta(x-4)).$$

We use the sample set in (6.8) and $d = 0.0488$. As noted above, our algorithm is not especially well suited to this type of data; in order to obtain any information about the locations of the δ -functions on the right-hand side of (6.10) we need to use a very large filter bandwidth. For Figure 6.4(a) we use $B = 10^{30}$; for (b) we use $B = 10^{20}$; and for (c) we use $B = 10^{10}$. While the graphs in (a) and (b) have peaks in approximately the correct locations $x = 1, 2$, and 4 , they also have many artifactual peaks and assume negative values. The graph in (c) has no useful, accurate information. As shown in Figure 6.4(d), with the bandwidth $B = 10^{20}$, an extremely small level of

Fig. 6.4 The approximate Laplace inverse acting on a finite sum of decaying exponentials, $F(t) = e^{-t} + e^{-2t} + e^{-4t}$. The heights and exact locations of impulses are indicated by the dashed line.

noise ($\sigma = 10^{-10}$) completely destroys the useful information in Figure 6.4(b). The location and heights of the impulses are shown with a dashed line.

7. Conclusion. We have obtained fast, FFT-based, forward and inverse algorithms for Laplace-like transforms. For the case of the Laplace transform, we used the harmonic analysis of the transform $f \mapsto \tilde{f}$ to analyze the effects on the reconstructed function of noise in the measurements and various filtering operations. The noise variance is shown to diverge as x goes to zero, even under realistic assumptions about the nature of the measured data. In the discretely sampled case, there is a trade-off, determined by the cut-off filter bandwidth, between the point where this divergence becomes apparent in the reconstructed signal and the resolution available in the reconstruction. Whether a usable reconstruction can be obtained depends on

both the level of noise in the data and the location of the support of the signal. On one hand, our results make precise the difficulties entailed in inverting the Laplace transform, but on the other hand they also provide flexible tools for analyzing approximate inverses to operators in this class.

Acknowledgments. The first author would like to thank Prof. Leslie Greengard for providing the references [11] and [8] and helpful remarks, and Henry Ong for help with the numerical experiments. The authors would also like to thank the referees for many useful suggestions.

REFERENCES

- [1] R. E. BELLMAN AND R. S. ROTH, *The Laplace Transform*, World Scientific Press, Singapore, 1984.
- [2] M. BERTERO AND F. GRUNBAUM, *Commuting differential operators for the finite Laplace transform*, *Inverse Problems*, 1 (1985), pp. 181–192.
- [3] A. BOUMENIR AND A. AL-SHUAIBI, *The inverse Laplace transform and analytic pseudo-differential operators*, *J. Math. Anal. Appl.*, 228 (1998), pp. 16–36.
- [4] B. DAVIES AND B. MARTIN, *Numerical inversion of the Laplace transform: A survey and comparison of methods*, *J. Comput. Phys.*, 33 (1979), pp. 1–32.
- [5] D. C. GARDNER, J. C. GARDNER, AND W. W. MEINKE, *Methods for the analysis of multicomponent exponential decay curves*, *J. Chem. Phys.*, 31 (1959), pp. 978–986.
- [6] V. MARKEL AND J. SCHOTLAND, *Symmetries, inversion formulas, and image reconstruction in optical tomography*, *Phys. Rev. E*, 70 (2004), article 056616.
- [7] J. MCWHIRTER AND E. PIKE, *On the numerical inversion of the Laplace transform and similar Fredholm integral equations of the first kind*, *J. Phys. A Math. Gen.*, 11 (1978), pp. 1729–1745.
- [8] V. ROKHLIN, *A fast algorithm for discrete Laplace transformation*, *J. Complexity*, 4 (1988), pp. 12–32.
- [9] J. SCHOTLAND, *Continuous-wave diffusion imaging*, *J. Opt. Soc. Amer. A*, 14 (1997), pp. 275–279.
- [10] Y.-Q. SONG, L. VANKATARAMANAN, AND L. BURCAW, *Determining the resolution of the Laplace inversion spectrum*, *J. Chem. Phys.*, 122 (2005), article 104104.
- [11] J. STRAIN, *A fast Laplace transform based on Laguerre functions*, *Math. Comp.*, 58 (1992), pp. 275–283.
- [12] M. E. TAYLOR, *Estimates for approximate solutions to acoustic inverse scattering problems*, in *Inverse Problems in Wave Propagation*, G. Chavent, G. Papanicolaou, P. Sacks, and W. W. Symes, eds., IMA Vol. Math. Appl. 90, Springer-Verlag, New York, 1997, pp. 463–499.
- [13] Z.-M. WANG, G. Y. PANASYUK, V. A. MARKEL, AND J. C. SCHOTLAND, *Experimental demonstration of an analytic method for image reconstruction in optical diffusion tomography with large data sets*, *Optics Lett.*, 30 (2005), pp. 3338–3340.