Math 6620: Analysis of Numerical Methods, II
Approximation with Fourier Series

See Hesthaven, S. Gottlieb, and D. Gottlieb 2007, Chapters 1-2,
Canuto et al. 2011, Chapters 2.1, 5.1,

Shen, Tang, and Wang 2011, Chapter 2

Akil Narayan[1]

[1] Department of Mathematics, and Scientific Computing and Imaging (SCI) Institute
University of Utah

High order approximations

We are now very familiar with our rather standard approximation to u_{xx} on an equidistant grid:

$$D_+ D_- \; D_0 u_j^n \approx u_{xx} + \mathcal{O}(h^2)$$

Note that the h^2 truncation error is a direct result of our choice of 3-point stencil.

Using more points in the stencil allows us to attain higher order truncation errors.

$$\frac{1}{12h^2}\left[-u_{j-2}^n + 16u_{j-1}^n - 30u_j^n + 16u_{j+1}^n - u_{j+2}^n\right] \approx u_{xx} + \mathcal{O}h^2. \quad \mathcal{O}(h^4)$$

High order approximations

We are now very familiar with our rather standard approximation to u_{xx} on an equidistant grid:

$$D_0 u_j^n \approx u_{xx} + \mathcal{O}(h^2)$$

Note that the h^2 truncation error is a direct result of our choice of 3-point stencil.

Using more points in the stencil allows us to attain higher order truncation errors.

$$\frac{1}{12h^2}\left[-u_{j-2}^n + 16u_{j-1}^n - 30u_j^n + 16u_{j+1}^n - u_{j+2}^n\right] \approx u_{xx} + \mathcal{O}h^2.$$

In general, using $2p+1$ points allows us to achieve $\mathcal{O}(h^{2k})$ LTE.

Why stop here? Why not take p as large as possible?

This requires a stencil spreading over the whole domain, globally coupling all degrees of freedom.

Is it worth it?

Before solving differential equations, let's answer some basic approximation theory questions first.

The simplest example of an approximation scheme that globally couples all degrees of freedom is a Fourier Series.

Before solving differential equations, let's answer some basic approximation theory questions first.

The simplest example of an approximation scheme that globally couples all degrees of freedom is a Fourier Series.

Consider a given $u : [0, 2\pi] \to \mathbb{C}$, which we represent as a sum of complex exponentials,

$$u(x) \approx \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x), \qquad \phi_k(x) = \frac{1}{\sqrt{2\pi}} e^{ikx}.$$

The most straightforward strategy to identify $\widehat{u}_k$ is to choose them to minimize a loss,

$$\widehat{u}_k = \underset{\widehat{u}_k, k \in \mathbb{Z}}{\arg\min} \left\| u(x) - \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x) \right\|_2^2,$$

where we have introduced the norm and a corresponding inner product,

$$\langle f, g \rangle := \int_0^{2\pi} f(x) \overline{g}(x) \mathrm{d}x, \qquad \|f\|_2^2 := \langle f, f \rangle,$$

where $\overline{z}$ is the complex conjugate of z. [1]

[1] We are mostly interested in real-valued functions, so the introduction of complex arithmetic is somewhat artificial here. We could write the basis as real-valued $\sin kx$ and $\cos kx$ functions with real coefficients. This achieves the same results but uses somewhat more technical formulas.

We have conveniently chosen the basis ϕ_k so that,

$$\langle \phi_k, \phi_\ell \rangle = \begin{cases} 1, & k = \ell \\ 0, & k \neq \ell \end{cases}$$

Such basis functions are orthonormal.

There is a unique solution for the $\widehat{u}_k$ that minimizes the loss, and using basis orthonormality the solution has a fairly simple expression,

$$\widehat{u}_k = \langle u, \phi_k \rangle = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} u(x) e^{-ikx} \mathrm{d}x.$$

We have conveniently chosen the basis ϕ_k so that,

$$\langle \phi_k, \phi_\ell \rangle = \begin{cases} 1, & k = \ell \\ 0, & k \neq \ell \end{cases}$$

Such basis functions are orthonormal.

There is a unique solution for the $\widehat{u}_k$ that minimizes the loss, and using basis orthonormality the solution has a fairly simple expression,

$$\widehat{u}_k = \langle u, \phi_k \rangle = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} u(x) e^{-ikx} \mathrm{d}x.$$

This gives us a first taste of some functional analysis: Define,

$$L^2 = L^2\left([0, 2\pi]; \mathbb{C}\right) = \left\{ f : [0, 2\pi] \to \mathbb{C} \;\middle|\; \|f\|_2^2 < \infty \right\}.$$

Then Fourier Series representations are complete in L^2:

$$u \in L^2 \quad \Longrightarrow \quad \lim_{N \to \infty} \left\| u(x) - \sum_{k=-N}^{N} \widehat{u}_k \phi_k(x) \right\|_2 = 0,$$

and orthonormality of the basis results in Parseval's identity,

$$u \in L^2 \quad \Longrightarrow \quad \|u\|_2^2 = \sum_{k \in \mathbb{Z}} \left| \widehat{u}_k \right|^2.$$

$$u(x) \overset{L^2}{=} \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x), \qquad \widehat{u}_k = \langle u, \phi_k \rangle$$

This is all well and good, but how does this serve us *computationally*?

With finite storage, we have to truncate the infinite series,

$$u(x) \approx u_N(x) := \sum_{|k| \leqslant N} \widehat{u}_k \phi_k(x)$$

How well does u_N approximate u?

Comp "cost": $2N+1$

Comp "cost" of FD: M

$M \sim 2N+1$

$$u(x) \overset{L^2}{=} \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x), \qquad \widehat{u}_k = \langle u, \phi_k \rangle$$

This is all well and good, but how does this serve us *computationally*?

With finite storage, we have to truncate the infinite series,

$$u(x) \approx u_N(x) := \sum_{|k| \leqslant N} \widehat{u}_k \phi_k(x)$$

How well does u_N approximate u?

There is, of course, the other pesky issue that in practice we cannot actually compute $\widehat{u}_k$ exactly, and so must resort to additional approximations.
But let's focus on one sin at a time....

$$u(x) \stackrel{L^2}{=} \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x), \qquad \widehat{u}_k = \langle u, \phi_k \rangle$$

This is all well and good, but how does this serve us *computationally*?

With finite storage, we have to truncate the infinite series,

$$u(x) \approx u_N(x) := \sum_{|k| \leqslant N} \widehat{u}_k \phi_k(x)$$

How well does u_N approximate u?

There is, of course, the other pesky issue that in practice we cannot actually compute $\widehat{u}_k$ exactly, and so must resort to additional approximations.
But let's focus on one sin at a time....

So our question regards how *compressible* the infinite series is with respect to the truncation N:

$$\|u - u_N\|_2 \stackrel{?}{\lesssim} h(N),$$

for some function $h(N)$.

- h decays quickly with $N \rightarrow u$ is very compressible
- h decays slowly with $N \rightarrow u$ is not very compressible

Projections

Before investigating Fourier approximation results, it's worthwhile to introduce additional concepts: Projections.

Given an operator $P : L^2 \to V$, where $V \subset L^2$ is some subspace of L^2, then P is a projection operator if

$$P^2 = P.$$

The action $u \mapsto Pu$ projects u onto V.
The action $u \mapsto (I - P)u$ projects u onto some subspace W such that $V \oplus W = L^2$.

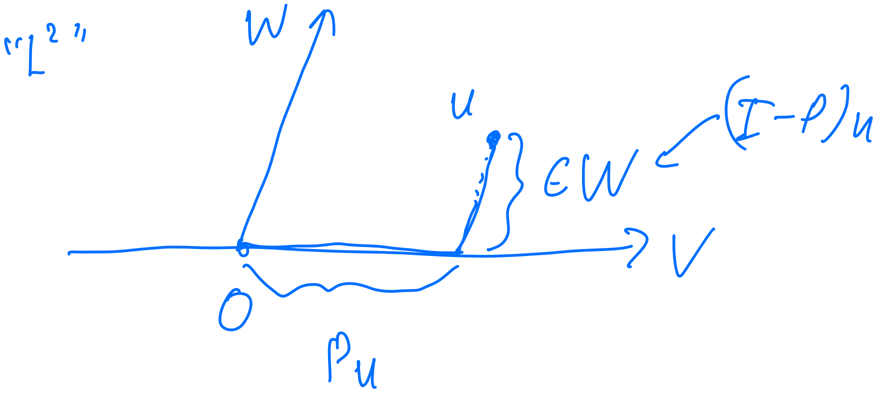

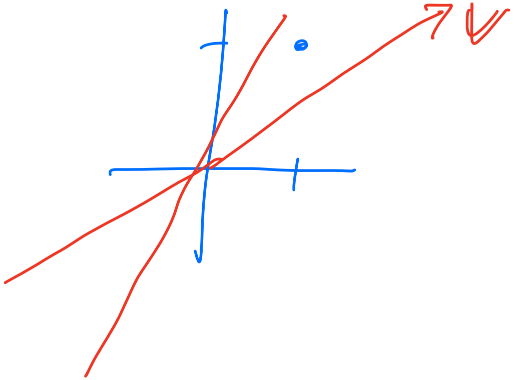

Projections

Before investigating Fourier approximation results, it's worthwhile to introduce additional concepts: Projections.

Given an operator $P : L^2 \to V$, where $V \subset L^2$ is some subspace of L^2, then P is a projection operator if

$$P^2 = P.$$

The action $u \mapsto Pu$ projects u onto V.
The action $u \mapsto (I - P)u$ projects u onto some subspace W such that $V \bigoplus W = L^2$.

For any projection operator P and any $u \in L^2$, we have,

$$(I - P)Pu = 0.$$

Before investigating Fourier approximation results, it's worthwhile to introduce additional concepts: Projections.

Given an operator $P : L^2 \to V$, where $V \subset L^2$ is some subspace of L^2, then P is a projection operator if

$$P^2 = P.$$

The action $u \mapsto Pu$ projects u onto V.
The action $u \mapsto (I - P)u$ projects u onto some subspace W such that $V \oplus W = L^2$.

For any projection operator P and any $u \in L^2$, we have,

$$(I - P)Pu = 0.$$

A projection operator P is orthogonal if $W \perp V$, equivalently if for every $u, v \in L^2$:

$$P = P^*, \qquad \langle P^* u, v \rangle := \langle u, Pv \rangle .$$

$\|Pu\| \leq \|u\|$ iff $P = P^*$
in general

We are considering the truncation,

$$\sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k(x) \stackrel{L^2}{=} u \approx u_N = \sum_{|k| \leqslant N} \widehat{u}_k \phi_k(x).$$

This truncation is an orthogonal projector.

Theorem

Define P_N as the operator,

$$P_N u = u_N = \sum_{|k| \leqslant N} \widehat{u}_k \phi_k(x), \qquad u \stackrel{L^2}{=} \sum_{k \in \mathbb{Z}} \widehat{u}_k \phi_k.$$

Then P_N is an orthogonal projection operator.

A basic approximation estimate, I

Can we bound $\|u - P_N u\|_2$? First note that,

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 \,.$$

A basic approximation estimate, I

Can we bound $\|u - P_N u\|_2$? First note that,

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 .$$

Integration by parts is our friend, and note that,

$$\begin{aligned} \widehat{u}_k &= \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} u(x) e^{-ikx} \mathrm{d}x \\ &= \frac{i}{k\sqrt{2\pi}} u(x) e^{-ikx} \Big|_0^{2\pi} - \frac{i}{k\sqrt{2\pi}} \int_0^{2\pi} u'(x) e^{-ikx} \mathrm{d}x. \end{aligned}$$

A basic approximation estimate, I

Can we bound $\|u - P_N u\|_2$? First note that,

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 .$$

Integration by parts is our friend, and note that,

$$\begin{aligned} \widehat{u}_k &= \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} u(x) e^{-ikx} \mathrm{d}x \\ &= \frac{i}{k\sqrt{2\pi}} u(x) e^{-ikx} \Big|_0^{2\pi} - \frac{i}{k\sqrt{2\pi}} \int_0^{2\pi} u'(x) e^{-ikx} \mathrm{d}x. \end{aligned}$$

Note that, conveniently, the first term vanishes if $u(0) = u(2\pi)$.
This is, of course, quite reasonable since we are approximating with periodic functions.

Note also that the remaining integral is the Fourier series coefficient for the derivative, $u'(x)$:

$$u'(x) = \sum_{|k| \in \mathbb{Z}} \widehat{u'}_k \phi_k(x), \qquad \widehat{u'}_k = \langle u', \phi_k \rangle .$$

A basic approximation estimate, I

Can we bound $\|u - P_N u\|_2$? First note that,

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 .$$

Integration by parts is our friend, and note that,

$$\begin{aligned}\widehat{u}_k &= \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} u(x) e^{-ikx} \mathrm{d}x \\ &= \frac{i}{k\sqrt{2\pi}} u(x) e^{-ikx}\Big|_0^{2\pi} - \frac{i}{k\sqrt{2\pi}} \int_0^{2\pi} u'(x) e^{-ikx} \mathrm{d}x.\end{aligned}$$

Note that, conveniently, the first term vanishes if $u(0) = u(2\pi)$.
This is, of course, quite reasonable since we are approximating with periodic functions.

Note also that the remaining integral is the Fourier series coefficient for the derivative, $u'(x)$:

$$u'(x) = \sum_{|k|\in\mathbb{Z}} \widehat{u'}_k \phi_k(x), \qquad \widehat{u'}_k = \langle u', \phi_k \rangle .$$

Thus, if u is periodic and $u' \in L^2$ (so that $\widehat{u'}_k$ is well-defined), then

$$\widehat{u}_k = -\frac{i}{k} \widehat{u'}_k .$$

A basic approximation estimate, II

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 ,$$

$$\widehat{u}_k = -\frac{i}{k}\widehat{u'}_k .$$

This very basic estimate for Fourier series coefficients implies:

$$\begin{aligned}\|u - P_N u\|_2^2 = \sum_{|k|>N} \frac{1}{|k|^2} \left|\widehat{u'}_k\right|^2 \leqslant \frac{1}{N^2} \sum_{|k|>N} \left|\widehat{u'}_k\right|^2 &\leqslant \frac{1}{N^2} \sum_{k\in\mathbb{Z}} \left|\widehat{u'}_k\right|^2 \\ &= \frac{1}{N^2} \|u'\|_2^2 ,\end{aligned}$$

where the last relation is Parseval's identity.

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} |\widehat{u}_k|^2 ,$$

$$\widehat{u}_k = -\frac{i}{k}\widehat{u'}_k .$$

This very basic estimate for Fourier series coefficients implies:

$$\|u - P_N u\|_2^2 = \sum_{|k|>N} \frac{1}{|k|^2} \left|\widehat{u'}_k\right|^2 \leqslant \frac{1}{N^2} \sum_{|k|>N} \left|\widehat{u'}_k\right|^2 \leqslant \frac{1}{N^2} \sum_{k\in\mathbb{Z}} \left|\widehat{u'}_k\right|^2$$

$$= \frac{1}{N^2} \|u'\|_2^2 ,$$

where the last relation is Parseval's identity.
We have just proven the following:

Theorem

Suppose $u, u' \in L^2$, *and that* $u(0) = u(2\pi)$. *Then,*

$$\|u - P_N u\|_2 \leqslant \frac{1}{N} \|u'\|_{L^2}$$

To generalize this result, some additional notation will be helpful.

Definition (Sobolev spaces)

Given $s \in \mathbb{N}_0 = \{0, 1, \ldots, \}$, the ($L^2$ periodic) Sobolev space of functions is given by,

$$H_p^s([0,2\pi];\mathbb{C}) := \Big\{ f : [0,2\pi] \to \mathbb{C} \;\Big|\; f^{(k)} \in L^2([0,2\pi];\mathbb{C}) \text{ for all } 0 \leqslant k \leqslant s, \; f^{(k)}(0) = f^{(k)}(2\pi) \text{ for all } 0 \leqslant k \leqslant s-1 \Big\}$$

The *norm* on H_p^s is defined as,

$$\|u\|^2_{H^s} := \sum_{k=0}^{s} \left\| u^{(k)} \right\|_2^2 .$$

Some specializations of interest:

- $s = 0 \Longrightarrow H^0 = L^2$
- $s > 0 \Longrightarrow$ continuous functions $\subset H_p^s$

To generalize this result, some additional notation will be helpful.

Definition (Sobolev spaces)

Given $s \in \mathbb{N}_0 = \{0, 1, \ldots, \}$, the ($L^2$ periodic) Sobolev space of functions is given by,

$$H_p^s([0, 2\pi]; \mathbb{C}) := \Big\{ f : [0, 2\pi] \to \mathbb{C} \;\Big|\; f^{(k)} \in L^2([0, 2\pi]; \mathbb{C}) \text{ for all } 0 \leqslant k \leqslant s, \\ f^{(k)}(0) = f^{(k)}(2\pi) \text{ for all } 0 \leqslant k \leqslant s-1 \Big\}$$

The *norm* on H_p^s is defined as,

$$\|u\|_{H^s}^2 := \sum_{k=0}^{s} \left\| u^{(k)} \right\|_2^2 .$$

Some specializations of interest:

- $s = 0 \implies H^0 = L^2$
- $s > 0 \implies$ continuous functions $\subset H_p^s$

The parameter s encodes the "amount" of smoothness that functions have, and the following inclusions hold:

$$H_p^r \subset H_p^s, \qquad r > s \geqslant 0.$$

The language of Sobolev spaces is the standard language in which to technically describe convergence rates of Fourier Series approximations.

Theorem

If $u \in H^s_p$*, then*

$$\|u - P_N u\|_{L^2} \leqslant N^{-s} \|u\|_{H^s_p}$$

Note that $s = 1$ is our previous result.

In terms of degrees of freedom, M, $\|u - P_N u\|_{L^2} \lesssim M^{-s} \|u\|_{H^s_p}$, which is *fantastic* for large s.

$M = 2N+1$

General approximation results

The language of Sobolev spaces is the standard language in which to technically describe convergence rates of Fourier Series approximations.

Theorem

If $u \in H^s$, then

$$\|u - P_N u\|_{L^2} \leqslant N^{-s} \|u\|_{H^s}$$

Note that $s = 1$ is our previous result.

In terms of degrees of freedom, M, $\|u - P_N u\|_{L^2} \lesssim M^{-s} \|u\|_{H_p^s}$, which is *fantastic* for large s.

Actually, something even stronger is true about Fourier approximation:

Theorem

If $u \in H^s$, then for every $0 \leqslant r < s$,

$$\|u - P_N u\|_{H_p^r} \leqslant N^{-(s-r)} \|u\|_{H_p^s}.$$

This result demonstrates tradeoff between smoothness of the function versus the strength of the norm under which convergence is sought.

Infinite regularity

The previous Fourier series results are essentially “good enough” to understand the basic point that smoothness of u translates into efficient approximations.
But to close the loop: what if u is infinitely differentiable?

Infinite regularity

The previous Fourier series results are essentially "good enough" to understand the basic point that smoothness of u translates into efficient approximations.
But to close the loop: what if u is infinitely differentiable?

Theorem

Let $u : [0, 2\pi) \to \mathbb{C}$ be the restriction of a function $f : \mathbb{C} \to \mathbb{C}$ to the unit circle. I.e., $u(x) := f(e^{ix})$.
Assume f is (complex) analytic in an annular neighborhood of the unit circle in $\mathbb{C}$. (This implies that u is infinitely differentiable.)

NB: $e^{-cN} <<<< N^{-s} \ \forall \ s.$

Then there exist constants $K, c > 0$ such that,

$$\|u - P_N u\|_{L^2} \leqslant K e^{-cN}.$$

Furthermore, for any $s \in \mathbb{N}_0$, there are constants $\widetilde{K}, \widetilde{c} > 0$ such that,

$$\|u - P_N u\|_{H_p^s} \leqslant \widetilde{K} e^{-\widetilde{c}N}.$$

The constants $K, c, \widetilde{K}, \widetilde{c}$ depend on the radii defining the annular region of analyticity.

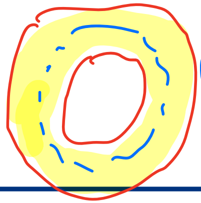
unit circle

The previous Fourier series results are essentially "good enough" to understand the basic point that smoothness of u translates into efficient approximations.
But to close the loop: what if u is infinitely differentiable?

Theorem

Let $u : [0, 2\pi) \to \mathbb{C}$ be the restriction of a function $f : \mathbb{C} \to \mathbb{C}$ to the unit circle. I.e., $u(x) := f(e^{ix})$.
Assume f is (complex) analytic in an annular neighborhood of the unit circle in $\mathbb{C}$. (This implies that u is infinitely differentiable.)

Then there exist constants $K, c > 0$ such that,

$$\|u - P_N u\|_{L^2} \leqslant K e^{-cN}.$$

Furthermore, for any $s \in \mathbb{N}_0$, there are constants $\widetilde{K}, \widetilde{c} > 0$ such that,

$$\|u - P_N u\|_{H_p^s} \leqslant \widetilde{K} e^{-\widetilde{c}N}.$$

The constants $K, c, \widetilde{K}, \widetilde{c}$ depend on the radii defining the annular region of analyticity.

Proof steps:

- $P_N u$ is a truncated Laurent series of f around the origin in $\mathbb{C}$.
- Convergence of the Laurent series in the region $r_1 \leqslant |z| \leqslant r_2$, where $r_1 < 1 < r_2$, can be used to estimate the truncated Laurent series coefficients.

The overall theme

The results we've described are *generic* lessons for nonperiodic global approximation as well:

- Such global methods have (rates of) accuracy that are limited only by functional regularity
  - Finite regularity $\Longrightarrow$ polynomial rates of error decay
  - Infinite regularity $\Longrightarrow$ superpolynomial (often exponential) rates of error decay
    (Note that real analyticity is not sufficient for complex analyticity; lack of complex analyticity generally downgrades pure exponential convergence to subexponential.)
- Global discretization methods are typically called *spectral methods*.

The results we've described are *generic* lessons for nonperiodic global approximation as well:

- Such global methods have (rates of) accuracy that are limited only by functional regularity
  - Finite regularity $\Longrightarrow$ polynomial rates of error decay
  - Infinite regularity $\Longrightarrow$ superpolynomial (often exponential) rates of error decay
    (Note that real analyticity is not sufficient for complex analyticity; lack of complex analyticity generally downgrades pure exponential convergence to subexponential.)
- Global discretization methods are typically called *spectral methods*.
- The very succinct punchline:

Smoothness $\Longrightarrow$ Compressibility

The results we've described are *generic* lessons for nonperiodic global approximation as well:

- Such global methods have (rates of) accuracy that are limited only by functional regularity
  - Finite regularity $\Longrightarrow$ polynomial rates of error decay
  - Infinite regularity $\Longrightarrow$ superpolynomial (often exponential) rates of error decay
    (Note that real analyticity is not sufficient for complex analyticity; lack of complex analyticity generally downgrades pure exponential convergence to subexponential.)
- Global discretization methods are typically called *spectral methods*.
- The very succinct punchline:

Smoothness $\Longrightarrow$ Compressibility

- All of this only applies to function approximation. For computing solutions to differential equations, this gives us tools to understand *consistency* of schemes.
- To really achieve convergence, we must understand *stability* as well.

References I

Canuto, Claudio et al. (2011). *Spectral Methods: Fundamentals in Single Domains*. 1st ed. 2006. Corr. 4th printing 2010 edition. Berlin ; New York: Springer. ISBN: 978-3-540-30725-9.

Hesthaven, Jan S., Sigal Gottlieb, and David Gottlieb (2007). *Spectral Methods for Time-Dependent Problems*. Cambridge University Press. ISBN: 0-521-79211-8.

Shen, Jie, Tao Tang, and Li-Lian Wang (2011). *Spectral Methods: Algorithms, Analysis and Applications*. Springer Science & Business Media. ISBN: 978-3-540-71041-7.