5-22-2024

Computational Tools for Exploring Eigenvector Localization

Robyn Ashley Markee Reid
Portland State University

Follow this and additional works at: https://pdxscholar.library.pdx.edu/open_access_etds

Part of the Mathematics Commons

Let us know how access to this document benefits you.

Computational Tools for Exploring Eigenvector Localization

by

Robyn Ashley Markee Reid

A dissertation submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy
in
Mathematical Sciences

Dissertation Committee:
Jeffrey Ovall, Chair
Jay Gopalakrishnan
Stefan Steinerberger
Panayot Vassilevski

Portland State University
2024

Abstract

We develop computational tools for exploring eigenvector localization for a class of selfadjoint, elliptic eigenvalue problems regardless of the cause for localization. The user inputs a desired region R (not necessarily connected), a tolerance for the amount localization in R, and the desired energy range $[a, b]$. The tool outputs eigenvectors concentrated within the tolerance inside R and within $[a, b]$. We develop ample theory that justifies our algorithm, which involves a complex, compact perturbation of the operator $\mathcal{L}$, $\mathcal{L}_s = \mathcal{L} + \mathrm{i} s \chi_R$, for some (small) $s > 0$. Our central idea can be summarized as follows: if (λ, ψ) is an eigenpair of $\mathcal{L}$ with ψ highly localized in R, then there will be an eigenpair (μ, ϕ) of the shifted operator $\mathcal{L}_s$ such that μ is near $\lambda + \mathrm{i}\,s$ and ϕ is near ψ i.e. the eigenpair (μ, ϕ) of $\mathcal{L}_s$ is close to the eigenpair $(\lambda + \mathrm{i}s, \psi)$ of $\mathcal{L} + \mathrm{i}s$. The algorithm finds eigenpairs (μ, ϕ) of $\mathcal{L}_s$ with $\Im(\mu)$ near s and $\Re(\mu)$ in or near $[a, b]$. In practice, the algorithm proves robust, immediately eliminating *over* 90% of unwanted eigenmodes, and many examples are provided. A post processing feature is included, which consists of (a few) inverse iterations that further eliminate unwanted eigenvectors and even identify computationally difficult cases.

Although most of the theory is developed with one kind of operator in mind, we show that it applies directly to the magnetic Laplacian, $\hat{H}(A) := (-\mathrm{i}\nabla - A(x))^2$ as well. We additionally provide a method for *a priori* prediction of where eigen-

functions may localize for the magnetic Laplacian. Essentially, if an eigenvector achieves a global maximum (in modulus) at $x_0 \in \Omega$, then A behaves similarly to a conservative vector field in a neighborhood around x_0 that depends on the eigenvalue. We provide numerical examples to show the eigenmodes localize where $|\text{curl}A|$ is small, at least for low energies, as expected based on our theory.

Dedication

To my mom, who loved me even when I didn't love myself, and from whom I inherited my resilience and work ethic. I have been so lucky to have been loved by you. To my dad. It is often like looking in a mirror, and I would not have it any other way.

Acknowledgments

Thank you, Dr. Ovall. You have been an excellent mentor. For all of your patience and all of your effort. For critiquing me, even when I didn't (yet) know how to handle it, and for holding me to a standard. Thank you for turning so many of my mistakes into a teaching moment. Thank you for constantly encouraging me, and for being humble when I felt I could not compare. This was *such* a long, and often taxing, journey that you guided me through. Thank you.

Also, thank you Jay, Panayot, and Stefan for taking the time and effort to support me through the end of this long process.

Table of Contents

List of Tables

List of Figures

1 Introduction

1.1 Motivating Problem

Eigenvector localization is the phenomena where the eigenvector "mass" is concentrated in some relatively small portion of the domain. It can be induced by complex geometry, properties of the operator such as an inhomogeneous or discontinuous potential, and/or the boundary conditions,. Examples of highly localized eigenmodes together with their "localization measures", which indicates how concentrated an eigenmode is in a given region and will be introduced formally later, are pictured in Figure 1.1.

This work develops new computational tools for exploring localization in self-adjoint, elliptic eigenvalue problems of the form:

$$\mathcal{L}\psi \doteq -\nabla \cdot (A\nabla\psi) + V\psi = \lambda\psi \text{ in } \Omega \quad , \quad \psi = 0 \text{ on } \partial\Omega \, , \psi \not\equiv 0 \text{ in } \Omega, \qquad (1.1)$$

where $\Omega \subset \mathbb{R}^d$ is a bounded, connected, open set. For convenience, the potential $V \in L^\infty(\Omega)$ is non-negative. The diffusion matrix $A : \Omega \to \mathbb{R}^{d\times d}$ is symmetric and uniformly elliptic, i.e. there are constants $c, C > 0$ such that

$$c\mathbf{v}^t\mathbf{v} \leq \mathbf{v}^t A(x)\mathbf{v} \leq C\mathbf{v}^t\mathbf{v} \text{ for all } \mathbf{v} \in \mathbb{R}^d \text{ and a.e. } x \in \Omega \ .$$

For $d = 2$, $A \in [L^\infty(\Omega)]^{d\times d}$. Otherwise, A is uniformly Lipschitz in each of its

Potential   Geometry

Figure 1.1: Left: the eigenvector is concentrated in a small region due to a higher energy potential barrier outside this region. The eigenvector remains localized until its corresponding eigenvalue surpasses the potential barrier. Right: the eigenvector is geometrically localized from thin channels that "squeeze" the eigenmode into one bulb of the Threebulb domain. Both eigenmodes have localization measures $\tau = 0.99$ and $\delta = 0.14$ in their respective regions.

components.

The unbounded operator $\mathcal{L} : \text{Dom}(\mathcal{L}) = \{v \in H_0^1(\Omega) : \mathcal{L}v \in L^2(\Omega)\} \to L^2(\Omega)$ has a real spectrum that consists of the sequence $\inf_\Omega V < \lambda_1 < \lambda_2 \leq \lambda_3 \leq \cdots$, and has no finite accumulation points; its eigenspace $E(\lambda, \mathcal{L}) = \{v \in \text{Dom}(\mathcal{L}) : \mathcal{L}v = \lambda v\}$ is finite dimensional for each $\lambda \in \text{Spec}(\mathcal{L})$. These conditions on $\mathcal{L}$ guarantee the *unique continuation property*, which states that if $v \in H^1(\Omega)$, $\mathcal{L}v = 0$ in Ω, and $v = 0$ on some non-empty open subset of Ω, then v must be identically zero on the whole domain Ω (cf. [1, 15, 20]). Consequently, no eigenvector can be identically zero on any open subset of Ω, but it may be "nearly zero" outside of a small subdomain; a localized eigenmode is thus almost zero everywhere except in some portion of the domain. We now formally define the measures of localisation in a nonempty, open, proper (not necessarily connected) subset R, which we will refer to as the subdomain or region of interest. For any function $v \in L^2(\Omega)$ and

subdomain R, we define $\delta(v, R), \tau(v, R) \in (0, 1)$:

$$\delta(v, R) = \frac{||v||_{L^2(\Omega \setminus R)}}{||v||_{L^2(\Omega)}} \quad , \quad \tau(v, R) = \frac{||v||_{L^2(R)}}{||v||_{L^2(\Omega)}} \quad . \tag{1.2}$$

Observe that τ is a measure of how concentrated or localized v is within R, while $\delta(v, R)$ measures how concentrated v is outside of R. Thus, they are complementary quantities, meaning that $\delta^2(v, R) + \tau^2(v, R) = 1$. Given some reasonably small δ^*, which may depend on the size of R relative to the size of Ω, we say v is localized in R within the tolerance δ^* if $\delta^* \geq \delta(v, R)$. In practice, we have (often) used $\delta^* \leq 0.4$, meaning $\tau(v, R) \geq 0.165$, but the issue of selecting a suitable tolerance will be considered later. For ease, we will just say that v is localized to mean that v is localized in some given R within some small tolerance, δ^*.

A natural generalization of (1.2) would be to replace L^2 with L^p for some $p \geq 1$. In this case, we would say that v is L^p-localized. There are multiple methods to quantify localization. D. J. Thouless [36] introduces the confinement of an eigenmode by the ratio of the "existence surface" with the surface area of the domain in $\mathbb{R}^2$. For $p < q$, [18] generalizes the existence surface as

$$S_{p,q}(v) = \left(\frac{||v||_{L_p(\Omega)}}{||v||_{L_q(\Omega))}} \right)^{\frac{1}{\frac{1}{p} - \frac{1}{q}}},$$

where a function v is considered localized when $S_{p,q}(v) \ll |\Omega|$. In practice, Thouless [36] and Felix et al [12] have used $p = 2$ and $q = 4$, which is supported by the fact that a higher power norm diminishes any function that is small in some region R

faster than a lower power norm. The choice of $p = 2$ and $q = 4$ is sometimes called the "Inverse Participation Ratio". Notice that this definition does not specify a region in which the eigenmode is localized, simply that it is localized somewhere, whereas the concept of being L_p localized is region dependent. If the user wanted to quantify localization in a given R, it would be necessary to do so with the L_p-localized measure.

While the scope of this work primarily focuses on geometric localization and eigenvector localization caused by a scalar potential, we highlight that our tool applies regardless of the cause for the localization, and even provide examples where the localization occurs due to a vector potential. Furthermore, the tool is agnostic to relative energy of the eigenvalues, which is often one key factor for predicting if and where localization may occur. Higher energy eigenvalues, or those that are far from the first eigenvalue, may not localize as frequently, or even may localize in unexpected places. Figure 1.2 provides the first 40 eigenmodes of the Three Bulb domain, where any localization is due to the geometry of the domain. When the eigenvalues are relatively small, more localized eigenmodes tend to occur, both in how concentrated each eigenvector is and in the number of eigenmodes exhibiting localization. Then as one moves up in the spectrum, the functions tend to spread more evenly throughout the domain, which decreases the amount of localization.

The rest of this dissertation is outlined as follows. For the remainder of this chapter, we cover existing methods and supporting theory while discussing their limitations and compare them with our approach. We discuss recent work that

Figure 1.2: The first 40 eigenmodes of the Three Bulb domain. Localization becomes less frequent as one goes higher into the spectrum.

provides insight into the mechanisms driving localization, describe a numerical method that can provide candidate subdomains for where localization may occur, and also estimate of the smallest eigenvalue whose eigenvector is concentrated in a such a subdomain [1, 3, 4, 6, 13, 35]. A few applications of localization are also provided. In Chapter 2, we present new computational tools that allow the user to predict and model localization for (1.1) regardless of the reason for the localization –this is our algorithm, and the primary focus of this dissertation. Rigorous mathematical theory supporting the algorithm is also presented here. Background Information, like an introduction to eigenvalue solvers such as Pythonic FEAST [22] or an introduction to finite element methods, is provided in Chapter 3. We use the finite element software NGSolve [24], which interfaces through Python [38],

to run Pythonic FEAST in order to solve the eigenvalue problems in this work. Extensive numerical experiments that provide a proof of concept for our algorithm are given in Chapter 4; examples using the magnetic Laplacian operator are additionally provided. Chapter 5 builds on the foundations of Chapter 3 and the theory in Chapter 2 to provide guidance for how to set some of the parameters in the software, focusing primarily on how one should set the shift and tolerance values. This chapter focuses on the practical considerations one should keep in mind when using the computational tools, and is additionally supported by theory and numerical experiments. Chapter 6 and chapter 7 focus on what the driving mechanisms for geometric and potential localization are, respectively. Guidance on how to more efficiently and purposefully manipulate both types of localization is explored, along with insight into how one may "optimize" geometric or potential localization. Finally, Chapter 8 covers our new theory regarding a priori prediction for where low-energy eigenmodes will localize for eigenvalue problems using the magnetic Laplacian operator. Supporting numerical experiments furnish this chapter as well. The appendices contain additional experiments and supporting material that may be of interest, but which is not for understanding the core information.

1.2   Existing Methods and Theory

We begin by introducing the Landscape Function u of $\mathcal{L}$. For second-order differential operators, the landscape function u is the solution to the problem $\mathcal{L}u = 1$

in Ω with $u|_{\delta\Omega} = 0$. In 2012, Mayboroda and Filoche used the landscape function, where $\mathcal{L}$ was the Schrödinger type operator $\mathcal{L} = -\Delta + V$, to predict the location, shape, and size of the region where **ground state** eigenfunctions may localize [13]. In this context, ground state eigenvectors are the eigenvectors with the smallest eigenvalues that are localized to a given region. Essentially, they're the first eigenmodes localized within any given subdomain. The general result is that every eigenmode satisfies $|\psi(x)| \leq \lambda||\psi||_\infty u(x)$ for all $x \in \overline{\Omega}$.

Theorem 1.2.1 (Pointwise Stability Result). *Suppose that $\mathcal{L}u = 1$ in Ω, $u = 0$ on $\partial\Omega$. If v is sufficiently regular, then $|v(x)| \leq ||v||_{L^\infty(\partial\Omega)} + ||\mathcal{L}v||_{L^\infty(\Omega)}u(x)$.*

Proof (via Maximum Principle). Let $w = \pm v - ||\mathcal{L}v||_{L^\infty(\Omega)}u$. Then $\mathcal{L}w = \pm\mathcal{L}v - ||\mathcal{L}v||_{L^\infty(\Omega)} \leq 0$, so

$$w(x) \leq \max_{y\in\overline{\Omega}} w(y) \leq \max_{y\in\partial\Omega}\max\{w(y),0\} = \max_{y\in\partial\Omega}\max\{\pm v(y),0\} \leq ||v||_{L^\infty(\partial\Omega)}$$

Therefore, $\pm v(x) \leq ||v||_{L^\infty(\partial\Omega)} + ||\mathcal{L}v||_{L^\infty(\Omega)}u(x)$. $\qquad\square$

This general result immediately implies that u provides pointwise control of a properly normalized eigenmode, $\psi/(\lambda||\psi||_{L^\infty(\Omega)})$.

Theorem 1.2.2 (Filoche/Mayboroda, PNAS 2012). *Suppose that $\mathcal{L}u = 1$ in Ω, $u = 0$ on $\partial\Omega$, and that $\mathcal{L}\psi = \lambda\psi$ in Ω and $\psi = 0$ on $\partial\Omega$. Then $|\psi(x)| \leq \lambda||\psi||_{L^\infty(\Omega)}u(x)$.*

Consequently, since u is composed of a series of peaks and valleys, one can map out where the eigenmodes are likely to be concentrated, at least lower in

Figure 1.3: The landscape function and its negative given in grey in each picture, and the first six eigenvectors are superimposed in black. Regions where the magnitude of the landscape function are the highest are the first places localization occurs. These regions also indicate where the largest number of eigenvectors will concentrate early in the spectrum, where the two largest peaks already have localized non-ground states well before several of the other regions have ground states.

the spectrum. Applying the Watershed Principle to $\frac{1}{u}$ reveals the regions of the domain where localization early in the spectrum may occur by separating the domain along the lines where the eigenfunctions must be small, and therefore identifying subdomains where the eigenfunctions can concentrate. In practice, this method provides a very useful model for where to hunt for localized **low energy or ground state eigenvectors**. Some utility is gained by examining the magnitude of the landscape function; Figure 1.3 illustrates that regions with higher magnitude are more likely to be the first regions with localized eigenmodes. In this case, the first three eigenmodes are shown inside the three largest peaks of the landscape function. Furthermore, these are the regions where the most eigenfunctions will concentrate early on.

While the landscape function may still have pointwise control over higher en-

ergy eigenmodes, not much information is gained with higher energies. Potential regions of localization are identified for the ground states, however it is notable that ground states can occur higher in the spectrum. Figure 1.4 exemplifies how the utility of the pointwise control from the landscape function can deteriorate with increasing energy. Here, the unit interval is sliced into 4 equally measured intervals with potential $V = 80^2 \chi_{(\frac{1}{4}, \frac{1}{2})} + 400^2 \chi_{(\frac{3}{4}, 1)}$. The pictures show the normalised eigenmode drawn in black and the landscape function in grey. In each case, the landscape function dominates the eigenmodes as expected, and correctly indicates that one may find localized eigenmodes in the first and third quarters of the domain. This pointwise control, however, quickly becomes trivial since the landscape function does not decipher in which region, or union of regions, any specific eigenmode will be concentrated in or what the shape or behavior of the eigenmodes will be.

The landscape function provides a very useful starting point to suggest where one may find localized, low energy eigenvectors. It does not provide any localization information for higher energies, where strong localization can "unexpectedly" occur. This means that the potential regions suggested by the landscape function are not exhaustive, and this localization was observed when hunting for higher energy eigenpairs using our algorithm. Furthermore, even for low energies, one must actively go hunting to pinpoint the behavior, location, or the degree of localization for the eigenfunctions. For example, Figure 1.3 illustrates how looking at the landscape function by itself, does not show what the eigenmodes actually look like. Some eigenmodes even change sign, but there is no indication provided by
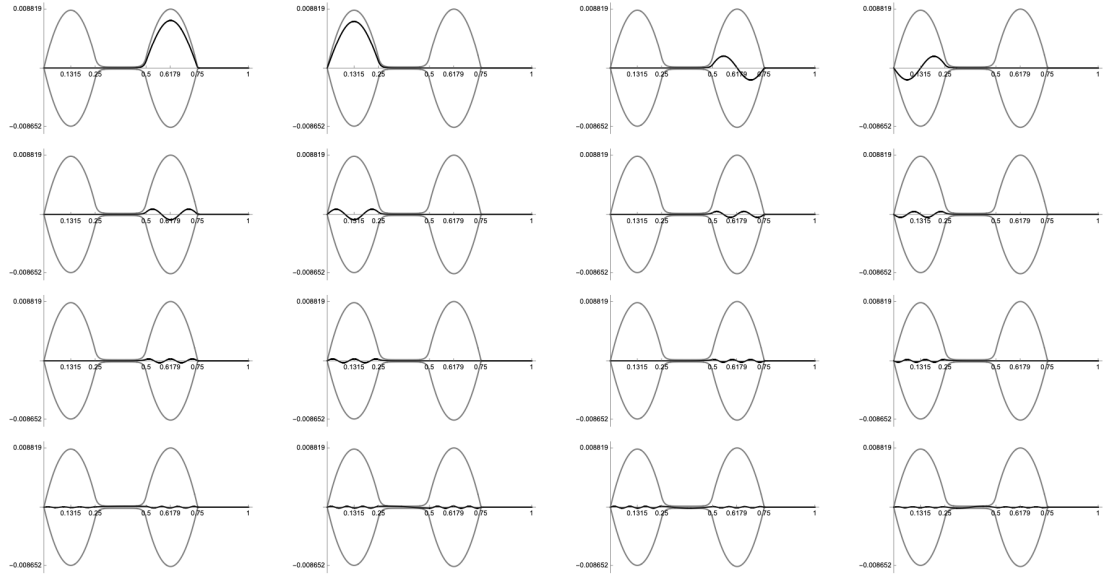
Figure 1.4: The landscape function is pictured in grey together with the first 16 normalised eigenfunctions in black. After the first two eigenfunctions, the pointwise control of the landscape function over the eigenmodes are no longer meaningful. No further regions of potential localization are suggested.

the landscape function that this would happen. Also, none of the eigenfunctions are concentrated to the all of the potential regions suggested by the landscape function.

Arnold et al [6] showed that replacing the potential V with the *effective potential* $\frac{1}{u}$ captures the effects of both the potential and kinetic energy, as opposed to just the potential energy. As a result, eigenfunctions with eigenvalues λ have most of their mass concentrated in the spatially confined region $E(\hat{\lambda}+\gamma) = \{x \in \Omega : \frac{1}{u(x)} \leq \lambda + \gamma\}$ for some suitable, small $\gamma > 0$. They developed an approach to provide estimates of several eigenvalues together with potential *regions* of localization for eigenvectors whose eigenvalues are near the estimates. This approach is outlined below:

1. Compute the landscape function, u.

2. Determine local minima of the associated *effective potential* $W = 1/u$, $W_k = W(x_k)$ for $1 \leq k \leq N$, with $W_k \leq W_{k+1}$.

3. Estimate N eigenvalues as $\hat{\lambda}_k = (1 + d/4)W_k$ where $\Omega \subset \mathbb{R}^d$. The factor $(1 + d/4)$ is supported empirically and heuristically.

4. Choose the set R_k to be the connected component of $\{x \in \Omega : W(x) \leq C\}$ containing x_k, where $C \geq \hat{\lambda}_k$ is a parameter to be set by the user. It is expected that $\delta(\psi, R_k)$ is small, where ψ is an eigenvector associated with the eigenvalue of $\mathcal{L}$ estimated by $\hat{\lambda}_k$.

The authors mention the option of solving for the smallest eigenpair of $\mathcal{L}$ with homogeneous Dirichlet boundary conditions on R_k, which would limit the user to

only one eigenvector estimate per region. Alternatively, it is suggested in [8] that a ground state with peak at x_0 and eigenvalue λ can be well approximated by $e^{-\rho_\lambda(x,x_0)}$, where ρ_λ is an Agmon distance between x and x_0, depending on the effective potential W. This approximation is further supported in Theorem 2.1 and several experiments by the authors in [7]. Furthermore, this method only approximates ground state eigenvectors in each region, and not necessarily the first N eigenvectors since some of these may not be ground state eigenvectors. In Figure 1.5, the spatially confined region $E(\hat{\lambda} + \gamma)$ indicates that there may be localization in each of the square bulbs, correctly identifying that the first, second, and fifth eigenmodes (these are ground state eigenmodes) are concentrated in the left, right, and middle bulb, respectively. Notice that these regions are given in columns 1, 2, and 4, where R_k is the left, right, and middle bulb, respectively. This method does not identify where the third or fourth eigenvectors are. In terms of locating any additional eigenmodes, this method only indicates that they may be in one of the square bulbs or any union thereof. No further utility is gained by computing the spatially confined region for non-ground state eigenfunctions, as seen by Figure 1.5. In the first column, the region E is drawn for the first to third, fifth, and fourteenth eigenmodes together with their estimated eigenvalues $\hat{\lambda}$, or actual eigenvalues λ, when no estimate is available. The actual eigenfunctions are shown in the right column together with their actual eigenvalues. The last row represents "higher" energy eigenmodes; observe that it is not clear from the spatially confined region if the eigenvectors are even localized at all.

Steinerberger [35] aimed to extract the localization information from the land-

$E(\hat\lambda_1 + \gamma)$: $\hat\lambda_1 = 1.26655$

Ψ_1 where $\lambda_1 = 1.22975$

$E(\hat\lambda_2 + \gamma)$: $\hat\lambda_2 = 2.23167$

Ψ_2 where $\lambda_2 = 2.1715$

$E(\lambda_3 + \gamma)$

Ψ_3 where $\lambda_3 = 3.05277$

$E(\hat\lambda_5 + \gamma)$: $\hat\lambda_5 = 4.57127$

Ψ_5 where $\lambda_5 = 4.50287$

$E(\lambda_{14} + \gamma)$

Ψ_{14} where $\lambda_{14} = 9.38303$
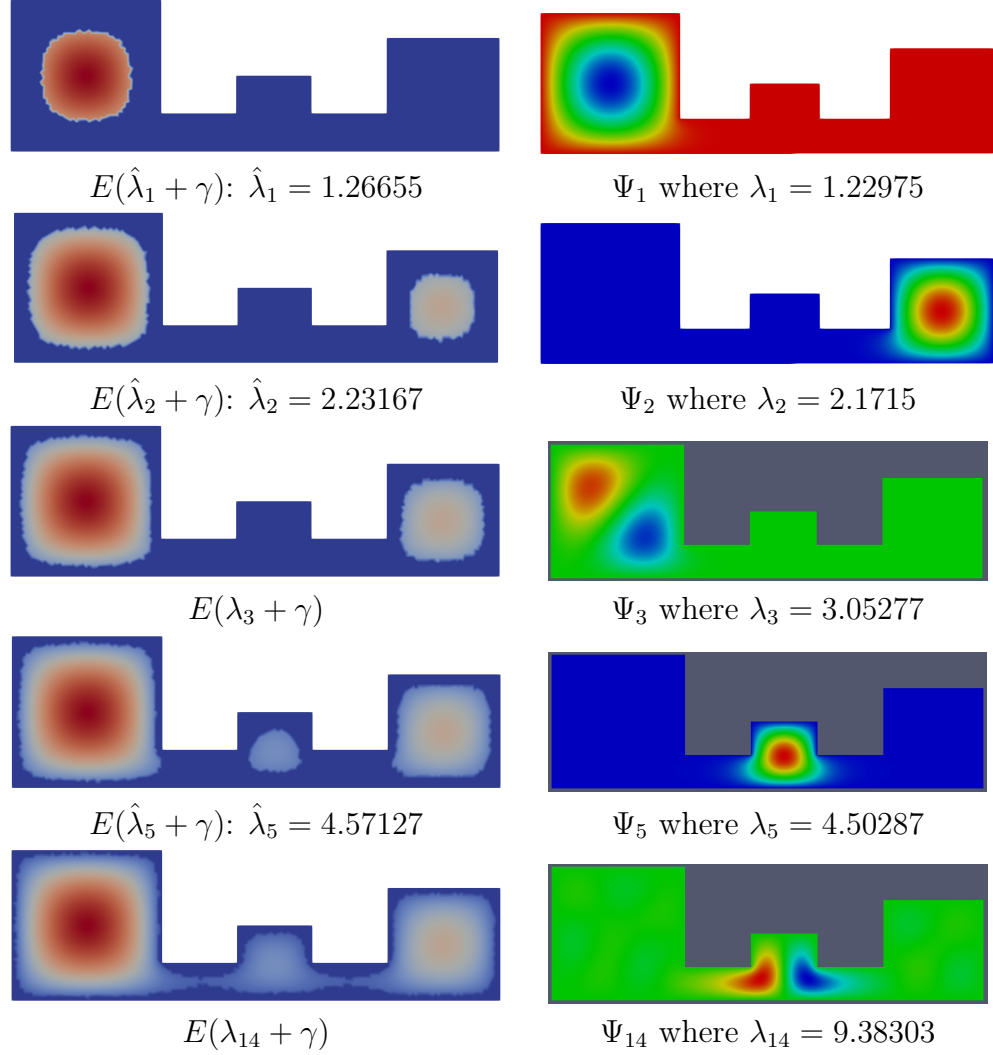
Figure 1.5: The spatially confined regions $E(\hat\lambda + \gamma)$ are pictured in the left column together with their estimated eigenvalues, $\hat\lambda$, or actual eigenvalues λ when no estimate could be made. Here $\gamma = 0.02$. The actual eigenmode is pictured in the right column together with the corresponding eigenvalue, λ.

scape function. He started by interpreting the localized eigenstates as critical points of the energy potential for $\Omega \subset \mathbb{R}^d$:

$$J(\psi) = \int_\Omega |\nabla \psi(x)|^2 dx + \int_\Omega V(x)\psi(x)^2 dx$$

Lu and team [28] then pointed out that the relevant quantity was the second integral, the integral average over V, and then replaced the potential by a slightly smoothed version:

$$-\Delta \psi(x) + (V * k_t)(x)\psi(x) = \lambda \psi(x) + error(x, t)$$

where $k_t : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is a Gaussian type kernel that is smooth, radial, with probability density centered at the origin and having most of its mass at a ball of radius $\sqrt{t}$. When V is large in a region, low frequency eigenfunctions are small, thereby isolating subregions as indicators for localization. Experimental evidence shows that the convolution $k_t * V$ agrees well and has as comparable predictive power as the effective potential $W = \frac{1}{u}$, where u is the landscape function. In fact, for small t, the error is also relatively small. Steinerberger's approach uses a Fast Fourier Transform to lower the computational cost and can be extended to other operators like the fractional or bi-Laplacian $\Delta^\alpha, \alpha \in (0, 1) \cup \{2\}$. While this method is a useful tool to indicate where localization may be found, even for a wider range of operators, one still has to go hunting through the spectrum to actually identify these eigenmodes.

A two part iterative approach from [3, 4] computes approximations of multiple

localized eigenpairs within a given domain in the presence of a highly disordered, structured potential, V. In the first phase, a coarse mesh $\mathcal{T}^H$ is used to provide potential regions of localization along with a basis for a rough approximation of the space that will be used to approximate eigenpairs in the second phase. The basis is obtained by using finite element hat functions that are associated with the vertices of $\mathcal{T}^H$ to perform a few approximate inverse iterations; one preconditioned conjugate gradient (PCG) step per iteration, involving Rayleigh quotients and a parameter $\eta \in (0, 1)$ used for decreasing the set of functions after each iteration, is performed. Then a fine mesh $\mathcal{T}^\varepsilon$ is generated as a refinement of the coarse mesh $\mathcal{T}^H$ that is deemed suitable for resolving (via a finite element method) the lowermost part of the spectrum of $\mathcal{L}$, where the user determines the number of N eigenpairs desired. In phase two, a few steps of approximate inverse iterations on $\mathcal{T}^\varepsilon$ are used, starting with the functions obtained from phase one; here, three PCG steps are used per iteration, as opposed to one. Each inverse iteration and a Rayleigh-Ritz procedure on the remaining set of functions results in approximate eigenpairs. If needed, a similar mechanism is used to decrease the number of functions for the next iteration. Ultimately, this results in a set of at least $2N$ approximate eigenpairs being obtained, where the first N are kept. This algorithm provides approximations of eigenvalues *and* eigenvectors, where the quality of these approximations can be controlled by the parameters in the discretization. Furthermore, in contrast to Arnold et al [6], more than one eigenvector localized in a given region can be found. The authors assume that the first N eigenvectors are localized, which is a reasonable assumption in the presence of such V, but

is limiting when there are other causes for localization (e.g. geometry, boundary conditions, etc.) or for the eigenpairs that are higher in the spectrum.

Existing approaches either only suggest potential regions where localization may occur or are primarily restricted to the low-energy or ground state regime when computing approximate eigenpairs. Methods that are not restricted to the ground state have stringent conditions that must be met and are limited to localization that occurs due to the potential. These approaches are best suited for problems where the targeted part of the spectrum contains many eigenvectors. Moreover, none of these methods allow for a priori control of how localized an eigenvector must be to be considered localized enough. In other words, for what value of $\delta(v, R)$ is the eigenvector considered no longer localized, or similarly to what tolerance δ^*? Our approach provides localized eigenpairs in a user determined region of interest, within the user determined tolerance, and for eigenpairs anywhere from low to high energy. As previously stated, localization can occur in unexpected places that are not presently predicted by any method. Example 1.6 provides an example of what this looks like in practice, which highlights a key benefit of being able to identify localization at high energies.

Example 1.2.3. For $\Omega = (0, 1)$, consider the operator

$$\mathcal{L} = -\frac{d^2}{dx^2} + \sum_{k=1}^{4} V_k \chi_{R_k} \quad , \quad R_k = \frac{1}{4}(k - 1, k) ,$$

with homogeneous Dirichlet boundary conditions and with

$V = 80^2 \chi_{(\frac{1}{4}, \frac{1}{2})} + 400^2 \chi_{(\frac{3}{4}, 1)}$ When $\lambda > 400^2$, the eigenfunction, pictured in Fig-

Figure 1.6: The landscape function (left) predicts that localization will occur primarily in the first and third subintervals, with a small amount in the second subinterval and none in the fourth subinterval. The eigenvector corresponding to the first eigenvalue that is larger than 400^2 is shown on the right. This eigenfunction is highly localized to $\tau = 0.9983301$ in the fourth subinterval, for Example 1.6.

ure 1.6 is highly concentrated in the fourth subinterval, with $\tau = 0.998330$. However, the landscape function indicates that the eigenmodes will localize only in the regions of the zero potential, the first and third subintervals. Figure 1.6 illustrates how localization can occur in unexpected locations that other methods would not have predicted, specifically for high energies.

1.3 Real World Applications

Surely the most famous example of localization is Anderson localization. Early models for predicting electrical conductivity outlined free moving electrons encompassed in metal or crystal lattice structures composed of positive ions. Electrical conductivity was believed to be directly proportional to the mean free path, which measured the mean distance an electron would travel before colliding with an ion,

but this length was experimentally determined in metals to be almost two orders of magnitude larger than the lattice constant, the distance between unit cells. Surely the electrons were not traveling for a distance one hundred times greater than the length of the encompassing lattice before colliding into that lattice! Measuring electrical resistance, however, supported this claim. Resistance results when electrons deviate from their original trajectory, or scatter, from colliding with an ion. With quantum mechanics came the discovery that wave-like electrons diffract from regular lattice sides, indicating that no resistance occurs unless the electron meets an imperfection in the lattice. Naturally, the more imperfections the more resistance. Moreover, a critical amount of these randomly placed imperfections can stop all electrical conductivity by physically trapping electrons in some subsection of the lattice. This is the phenomena is known as Anderson Localisation, a groundbreaking discovery that won the 1977 Physics Nobel Prize and which one would have to "resort to the indignity of numerical simulations to settle even the simplest questions about it" [5, 23, 26, 29]. Anderson Localisation explains the metal-insulator transition observed in semiconductors, paving the way for the Anderson insulator, which has no conductivity along partial bands of trapped electrons where the density of electronic states is also nonzero. This proved that quantized electrical transport of charge and energy can be enhanced by disorder.

A more widely accessible example of localization is the "Noise Abatement Wall" or "Fractal Wall"^{TM} [34]. It essentially traps sound waves, at frequencies typically experienced in practice, inside carefully placed, cleverly shaped holes located where the eigenfunctions of the wave equation localize. Normally, rooms or walls are

covered with a material that dampens acoustics; this can be as natural as curtains or pictures, or as intentional as foam or profiled diffusers. Irregular geometry in the room or along the wall will induce localization of the eigenfunctions of the corresponding wave equation. That localization enhances damping and thereby absorption. This is why an empty room often has an echo that is no longer observed when decorated and filled with furniture. Presumably, the reader is convinced that optimizing this absorption would prove incredibly useful, and hence the "Noise Abatement Wall": providing an effective sound barrier (98% on average of the incident acoustic energy in the audio spectrum) and minimizing traffic noise along busy roads in the neighborhoods of France! One might even be lucky enough to observe the Fractal Wall in action, but until then Figure 1.7 provides an illustration that will just have to do.

It should come as no surprise that waves are everywhere. It's worth saying twice: waves are everywhere! And not just from sound, light, or mechanical vibrations. Quantum Mechanics has disclosed that all matter has wavelike character at the atomic level. Recently, even classical gravitational waves have been detected. Since waves are everywhere (worth saying thrice), then obviously understanding their movements and patterns would prove to be extremely beneficial. Consequently, the phenomenon of localization is a rapidly growing field. Figure 1.8 graphs the number of citations of the original Anderson paper each year following the year it was discovered in 1958, which serves as a proxy of the level of research activity in this area of. Though slow to start, research in localization has grown (almost) exponentially, even gaining the attention of 2022 Nobel Prize winner,

Figure 1.7: The Noise Abatement Wall [34] effectively confines noise for housing along the busy road. It contains holes ranging from 30cm to submillimeter pores because the irregular geometry makes it a bad resonator. Absorption of acoustics are strongly linked to the spectral properties of the Laplacian operator. Image retrieved from IPAM Research Articles on April 15, 2024.

Figure 1.8: Citation and publication report for 1958 P. Anderson paper "Absence of diffusion in certain random lattices," showing impressive growth in research related to localization. Report generated by Web of Science, April 15, 2024.

Alain Aspect, and 2022 Fields Medal winner, Hugo Duminil-Copin, both leading projects for the prestigious Simons Collaboration on Localization of Waves.

2   The Algorithm

The theory in this chapter, with the exception of Lemma 2.1.6 and Remark 2.1.3, was published in [33]:

J. S. Ovall and R. Reid.   An algorithm for identifying eigenvectors exhibiting strong spatial localization. *Math. Comp.*, 92(341):1005–1031, 2023. https://doi.org/10.1090/mcom/3734

Jeffrey .S Ovall and Robyn Reid

Author contributions to the article cited above.

J.S. Ovall: analysis, simulation, writing, and editing

R. Reid: analysis, simulation, writing, and editing

2.1 Supporting Mathematical Theory

Consider a complex-shifted version of the operator, $\mathcal{L}$, defined by $\mathcal{L}_s = \mathcal{L} + \mathrm{i}s\,\chi_R$ on subdomain R where $s > 0$ is sufficiently small. Notice $\mathcal{L}$ is normal and $\mathrm{Dom}(\mathcal{L}_s) = \mathrm{Dom}(\mathcal{L})$. The idea for $\mathcal{L}_s$ comes from Marletta's work [2,30] combating the effects of spectral pollution when computing eigenvalues for operators having essential spectrum. Though our shifted operator does not have an essential spectrum, we may still apply the concept of dissipative perturbation to shift the eigenvalues of interest up and out of "pollution". Here, "pollution" refers to eigenvalues whose corresponding eigenfunctions are not localized in the subdomain R.

Our central idea can be summarized as follows: if (λ, ψ) is an eigenpair of $\mathcal{L}$ with ψ highly localized in R, then there should be an eigenpair (μ, ϕ) of $\mathcal{L}_s$ such that μ is near $\lambda + \mathrm{i}\,s$ and ϕ is near ψ i.e. the eigenpair (μ, ϕ) of $\mathcal{L}_s$ is close to the eigenpair $(\lambda + \mathrm{i}s, \psi)$ of $\mathcal{L} + \mathrm{i}s$. To establish this, as well as other supporting theory, first suppose that the eigenpair (μ, ϕ) is an eigenpair of $\mathcal{L}_s$, then:

$$\mu = \frac{(\mathcal{L}\phi, \phi)}{||\phi||^2_{L^2(\Omega)}} + \mathrm{i}\,s[\tau(\phi, R)]^2 \; , \tag{2.1}$$

where $(\cdot, \cdot)$ is the complex inner-product on $L^2(\Omega)$.

Proof. By the Rayleigh Quotient, we have that:

$$\mu = \frac{(\mathcal{L}_s\phi, \phi)}{(\phi, \phi)} = \frac{((\mathcal{L} + \mathrm{i}s\chi_R)\phi, \phi)}{||\phi||^2_{L^2(\Omega)}} = \frac{(\mathcal{L}\phi, \phi)}{||\phi||^2_{L^2(\Omega)}} + \frac{(\mathrm{i}s\chi_R\phi, \phi)}{||\phi||^2_{L^2(\Omega)}}$$
$$= \frac{(\mathcal{L}\phi, \phi)}{||\phi||^2_{L^2(\Omega)}} + \mathrm{i}\,s[\tau(\phi, R)]^2.$$

$\square$

This identity, in combination with the unique continuation property and variational characterization of the eigenvalue of $\mathcal{L}$, can then be used to obtain the proposition below.

Proposition 2.1.1. *For any* $\mu \in \mathrm{Spec}(\mathcal{L}_s)$, $0 < \Im\mu < s$ *and* $\Re\mu > \lambda_1(\mathcal{L}) = \min(\mathrm{Spec}(\mathcal{L}))$.

Proof. Let (μ, ϕ) be an eigenpair of $\mathcal{L}_s$, with $||\phi||_{L^2(\Omega)} = 1$. For $\phi = \phi_1 + \mathfrak{i}\phi_2$ and $\mu = \mu_1 + \mathfrak{i}\mu_2$. We compute from $\mathcal{L}_s\phi = \mu\phi$:

$$(\mathcal{L} + \mathfrak{i}s\chi_R)(\phi_1 + \mathfrak{i}\phi_2) = (\mu_1 + \mathfrak{i}\mu_2)(\phi_1 + \mathfrak{i}\phi_2)$$

$$(\mathcal{L}\phi_1 - s\chi_R\phi_2) + \mathfrak{i}(\mathcal{L}\phi_2 + s\chi\phi_1) = \mu_1\phi_1 - \mu_2\phi_2 + \mathfrak{i}(\mu_2\phi_1 + \mu_1\phi_2)$$

Then by comparing the real and imaginary parts:

$$\mathcal{L}\phi_1 - s\chi_R\phi_2 = \mu_1\phi_1 - \mu_2\phi_2 \ , \tag{2.2}$$

$$\mathcal{L}\phi_2 + s\chi_R\phi_1 = \mu_2\phi_1 + \mu_1\phi_2 \ . \tag{2.3}$$

Suppose that $||\phi||_{L^2(R)} = 0$. Then $\chi_R\phi = 0$ above, implying that $\mu_2 = 0$; this gives $\mathcal{L}\phi_j = \mu_1\phi_j$ for $j = 1, 2$. Now suppose that $||\phi||_{L^2(R)} = 1$, then similarly $\mu_2 = s$ and $\chi_R\phi_j = \phi_j$; as before, we have that $\mathcal{L}\phi_j = \mu_1\phi_j$. In both cases $\phi_1 = \phi_2 = 0$ in Ω, which contradicts the claim that ϕ is an eigenvector. Since the weak unique continuation property guarantees that $0 < \tau(\phi, R) < 1$, we have that $\mu_2 = \Im\mu \in (0, s)$.

By the variational characterization of eigenvalues of $\mathcal{L}$, $\mu_1 \geq \lambda_1$. Recall that λ_1 is simple, so $\mu_1 = \lambda_1$ can only occur if $\phi = c\psi_1$ for some $c \in \mathbb{C}$ with $|c| = 1$, where ψ_1 is taken to be a real eigenvector of $\mathcal{L}$ for λ_1 and $||\psi_1||_{L^2(\Omega)} = 1$. Assuming that $\mu_1 = \lambda_1$, we have $\mathcal{L}_s\psi_1 = \mu\psi_1$ and thus (2.2) apply, giving $\mathrm{i}\,s\chi_R\psi_1 = \mathrm{i}\,s\mu_2\psi_1$. However, $s \neq 0$ and consequently ψ_1 is supported in R, which again implies that $\psi_1 = 0$ in Ω, contradicting that it is an eigenvector of $\mathcal{L}$. □

This proposition designates approximately where eigenvalues of the shifted operator should be if their corresponding eigenfunctions are sufficiently localized in R, and therefore hints at an approach for discerning between eigenpairs that do meet the localization criteria and those that do not. We have the following theorem:

Theorem 2.1.2. *Let (λ, ψ) be an eigenpair of $\mathcal{L}$ and s be sufficiently small. Then*

$$\mathrm{dist}(\lambda + \mathrm{i}\,s, \mathrm{Spec}(\mathcal{L}_s)) \leq s\,\delta(\psi, R) \ . \tag{2.4}$$

Consequently, if (λ, ψ) is an eigenpair of $\mathcal{L}$ with $\lambda \in [a, b]$ and $\delta(\psi, R) \leq \delta^*$, then there is an eigenpair (μ, ϕ) of $\mathcal{L}_s$ in the region

$$U = U(a, b, s, \delta^*) = \{z \in \mathbb{C} : \mathrm{dist}(z, L) \leq s\delta^*, \ \Im z < s\}, \tag{2.5}$$

where $L = [a, b] + \mathrm{i}\,s$, pictured in Figure 2.1.

Proof. As pointed out on page 68 by Kato [25], $\mathcal{L}_s$ is a holomorphic perturbation of the selfadjoint operator $\mathcal{L}$, when s is sufficiently small. As a result, there is an

Figure 2.1: The region $U = U(a, b, s, \delta^*)$ from Theorem 2.1.2 supplies an area where eigenfunctions that meet the localization criteria must be found. Increasing s and δ^* lengthen this search region past $[a, b]$.

eigenpair $(\mu(s), \phi(s))$ of $\mathcal{L}_s$, depending holomorphically on s, such that $\mu(s) \to \lambda$ and $\phi(s) \to \psi$ in norm, where (λ, ϕ) is an eigenpair of $\mathcal{L}$. Denote the L^2 inner product by $(.,.)$ and use that $\mathcal{L}$ is selfadjoint to compute:

$$(\mu(s)\phi(s), \psi) = (\mathcal{L}_s\phi(s), \psi) = ((\mathcal{L} + is\chi_R)\phi(s), \psi) = (\mathcal{L}\phi(s), \psi) + is(\chi_R\phi(s), \psi)$$

$$= (\phi(s), \mathcal{L}\psi) + is(\chi_R\phi(s), \psi) = \lambda(\phi(s), \psi) + is(\chi_R\phi(s), \psi)$$

$$= \lambda(\phi(s), \psi) + is(\phi(s), \psi) - is(\chi_{\Omega \backslash R}\phi(s), \psi)$$

Rearranging the above equality and utilizing that $\chi_{\Omega \backslash R}$ is an orthogonal L^2 projection gives:

$$\frac{\lambda + is - \mu(s)}{is} = \frac{(\chi_{\Omega \backslash R}\phi(s), \chi_{\Omega \backslash R}\psi)}{(\phi(s), \psi)} .$$

Thus,

$$\frac{|\lambda + is - \mu(s)|}{s} \to \frac{||\psi||^2_{L^2(\Omega \backslash R)}}{||\psi||^2_{L^2(\Omega)}} = [\delta(\psi, R)]^2 \quad \text{as} \quad s \to 0 .$$

Since $\delta(\psi, R) < 1$ this implies that:

$$|\lambda + \mathrm{i}s - \mu(s)| \le s\delta(\psi, R) ,$$

when $s > 0$ is sufficiently small. This completes the proof. $\square$

The first part of Theorem 2.1.2 implies that for corresponding eigenfunctions, the eigenvalues of $\mathcal{L}$ and the real parts of the eigenvalues of $\mathcal{L}_s$ should be close; the same is said for s and the imaginary parts of the eigenvalues of $\mathcal{L}_s$. Notice that a large shift s that makes $s\delta > 1$ renders this inequality essentially meaningless, but setting s too small can negatively affect the algorithm, which we later consider. In the same vein, too large a δ can have the same effect, but this is less of a concern since $\delta \in (0, 1]$, whereas there are no theoretical bounds on s for this Theorem. Naturally, one might ask in practice how small must s be to be small enough to maintain the utility of Theorem 2.1.2, in which we have observed a trade-off between the size of s and the size of $\delta(v, R)$. This topic will be discussed more later.

Observe that (2.4) does not make any statement about how localized the eigenfunctions are. This result is definitely meaningful when the corresponding eigenvector ϕ is localized in R because then $\delta(\phi, R)$ is small, and we assume s is sufficiently small. However, when the eigenvector is very localized in the complement of R, $\delta(\phi, R) \approx 1$ and (2.4) is still meaningful for $s \le 1$. Ultimately since $\delta(\phi, R) \in [0, 1]$, then how useful (2.4) is relies mostly on the size of s, where s has an increased range when the eigenfunction is localized in R.

The second part of Theorem 2.1.2 indicates that if there is a localized eigenpair of $\mathcal{L}$, then there will be a corresponding eigenpair of $\mathcal{L}_s$ in the bottom half of the "Bunimovich stadium", pictured in Figure 2.1. This result is particularly useful because it supplies for any eigenpair (λ, ψ) of $\mathcal{L}$, the search region where a corresponding (μ, ϕ) will be, if it exists. Whether μ is in the bottom half of the Bunimovich stadium or not can then be used as a minimal requirement for where to hunt for eigenpairs of the shifted operator.

The Bunimovich stadium first appeared in 1979 where Leonid Bunimovich illustrated how a collection of particles, initially moving in the same direction, will become distributed uniformly throughout the domain in time, which is quite the opposite phenomenon from localization [11]. We point out that our use of the Bunimovich stadium is merely a coincidence.

While we briefly cover the basics here, we point out that section 3.2 contains a detailed discussion about how the Bunimovich Stadium is implemented within the context of FEAST. Denote by γ the boundary of $\tilde{U}$, and we will call it the Bunimovich curve. We build the Bunimovich curve with a counterclockwise, unit speed parameterization $z(t) = x(t) + \mathfrak{i}y(t)$, that begins at the point $b + \mathfrak{i}(s - r)$,

where:

$$(x(t), y(t)) = \begin{cases} r\left(\sin(\frac{t}{r}), -cos(\frac{t}{r})\right) + (b, s) & t \in [0, t_1] \\ (b + t_1 - t, s + r) & t \in [t_1, t_2] \\ r\left(\sin(\frac{t+a-b}{r}), \cos(\frac{t+a-b}{r})\right) + (a, s) & t \in [t_2, t_3] \\ (a - t_3 + t, s - r) & t \in [t_3, P] \end{cases}$$

and

$$r = s\delta^*, \quad t_1 = \pi r, \quad t_2 = t_1 + b - a, \quad t_3 = t_2 + \pi r, \quad P = 2\pi r + 2(b - a) . \quad (2.6)$$

This parameterization is made $P-$periodic by setting $z(t + P) = z(t)$. The trapezoid rule is applied to the Cauchy Integral Formula, $\dfrac{1}{2\pi i} \oint_\gamma (\xi - z)^{-1} d\xi$, to obtain the associated n-point rational filter function for γ. Note that the trapezoid rule converges rapidly for periodic functions and is extremely accurate when integrating periodic functions over their intervals. For $h = \frac{P}{n}$, the filter function, quadrature points, and quadrature weights are given by:

$$f(z) = \sum_{k=0}^{n-1} w_k(z_k - z)^{-1}, \quad z_k = z(kh + \frac{\pi r}{2}), \quad w_k = \frac{hz'(kh + \frac{\pi r}{2})}{2\pi i} . \quad (2.7)$$

The quadrature points, z_k, and weights, w_k, are offset by $\pi r/2$ to give a more symmetric distribution of the points. Figure 2.2 illustrates how the points sit along the Bunimovich curve.

More details about how this filter function is obtained and its purpose is detailed

Figure 2.2: Bunimovich curve, γ, with the quadrature points of $f(z)$ overlaid on it.

Figure 2.3: Contour plot of $|f(z)|$ for $\Re z \in [a - 2r, b + 2r]$ and $\Im z \in [s - 2r, s]$. The lower half of γ is given by the red curve.

in Chapter 3. Important to note now is that the filter function is used to clearly distinguish between points inside the contour, γ, and points outside; an ideal filter would map points inside to one and outside to zero. Figure 2.3 provides a contour plot of $|f(z)|$ to show how well it differentiates between points inside and outside the contour. The black curves drawn are the contours $|f(z)| = 2^j$ for $j \in [-8, 2]$. Observe that the curve further from γ is the $|f(z)| = 2^{-8}$ contour, and the the mapped values, which are color coded, rapidly decay in this direction; this shows that the $|f(z)|$ rapidly decay away from γ as desired. The red curve indicates the lower half of γ where $\Im z \leq s$, which is notable because this region is where the imaginary parts of the eigenvalues of $\mathcal{L}_s$ will be, as stated in Theorem 2.1.2.

Remark 2.1.3. We refer the reader to Kato's [25] for more details of the claims made in this remark. The perturbed operator $\mathcal{L}_s$ is holomorphic of type A. A holomorphic

operator $T(\xi)$ that is defined for perturbation ξ on domain $Dom(\xi) = D_0 \in \mathbb{C}$, belongs to the family of holomorphic operators of type A if it is a close-able operator $T : X \to Y$ between Banach spaces, and it meets the following two conditions. One: the domain of $T(\xi)$, denoted by D is independent of the perturbation ξ. Two: $T(x)$ is holomorphic for all $\xi \in D_0$ and $u \in D$. As discussed, $\mathcal{L}_s : H_0^1 \to L^2$ is a closeable, holomorphic operator between Hilbert Spaces. Clearly, any shift does not change the $Dom(\mathcal{L}_s)$, nor does it make the operator not analytic for any $u \in H_0^1$, so $\mathcal{L}_s$ is holomorphic of type A. Additionally, Theorem 2.6 provides a criterion to determine when an operator is holomorphic of type A, of which $\mathcal{L}_s$ is exemplified as a particular, easy case in Remark 2.7.

Thus, we may apply the key result from Kato, which states that if the $\mathrm{Spec}(\mathcal{L})$ is separated into two parts by some closed contour Γ, then $\mathrm{Spec}(\mathcal{L}_s)$ is also separated into two parts by the same Γ for sufficiently small s. Further, "sufficiently small" can be estimated, where separation of the spectrum occurs *at least if*:

$$s < r_0 := \min_{\gamma \in \Gamma}(||R(\gamma, \mathcal{L})||_{L^2})^{-1} , \tag{2.8}$$

where r_0 is called the radius of perturbation. We note that separation of the spectra can occur for larger s, and we suspect that it often does, however this $s < r_0$ *guarantees* it. For a contour Γ centered at the target eigenvalue λ, the radius should be less than $dist(\partial\Gamma, \mathrm{Spec}(\mathcal{L}))$ to guarantee that the $\mathrm{Spec}(\mathcal{L}_s)$ is separated by that contour. Observe that Γ can contain eigenvalues outside and inside it– as in the contour can surround multiple eigenvalues, but only one eigenvalue is

considered the target with respect to the radius of perturbation; this can pair down to setting r_0 equal to one half of the distance between the nearest eigenvalue and the target eigenvalue. In the cases where there are multiple eigenvalues that are very close to each other, this distance can be very small, which in turn requires an even smaller s. Though not the topic here, a very small s can negatively impact the algorithm or requires a lot more computational resources.

The following is the complement to Theorem 2.1.2:

Theorem 2.1.4. *Let (μ, ϕ) be an eigenpair of $\mathcal{L}_s$. Then*

$$s\,[\delta(\phi, R)]^2 \leq \mathrm{dist}(\mu, \mathrm{Spec}(\mathcal{L} + \mathrm{i}\,s)) \leq s\,\delta(\phi, R) \ , \tag{2.9}$$

$$\mathrm{dist}(\Re\mu, \mathrm{Spec}(\mathcal{L})) \leq s\,\delta(\phi, R)\,\tau(\phi, R) \ . \tag{2.10}$$

Let $\lambda = \arg\min\{|\sigma - \Re\mu| : \sigma \in \mathrm{Spec}(\mathcal{L})\}$. If $\Lambda \subset \mathrm{Spec}(\mathcal{L})$ contains λ, then

$$\inf_{v \in E(\Lambda, \mathcal{L})} \frac{||\phi - v||_{L^2(\Omega)}}{||\phi||_{L^2(\Omega)}} \leq \frac{s\,\delta(\phi, R)\,\tau(\phi, R)}{\mathrm{dist}(\Re\mu, \mathrm{Spec}(\mathcal{L}) \setminus (\Lambda \cup \{\Re\mu\}))} \ , \tag{2.11}$$

where $E(\Lambda, \mathcal{L})$ is the corresponding invariant subspace.

Proof. For simplicity, denote $\delta = \delta(\phi, R)$, $\tau = \tau(\phi, R)$, and $\mu_1 = \Re\mu$. Since $\mu \notin \mathrm{Spec}(\mathcal{L} + \mathrm{i}\,s)$, we directly compute $(\mathcal{L} + \mathrm{i}\,s - \mu)\phi = (\mathcal{L} + \mathrm{i}\,s)\phi - (\mathcal{L} + \mathrm{i}\,s\chi_R\phi) = \mathrm{i}\,s\,\chi_{\Omega \setminus R}\phi$. Following the same argument as in (2.4), we obtain the upper bound for (2.4). To obtain the lower bound, notice: $\mu = \mu_1 + \mathrm{i}\,s\tau^2$, then $|d + \mathrm{i}\,s - \mu| = |d - \mu_1 + \mathrm{i}\,s\delta^2| \geq s\delta^2$ for any $d \in \mathbb{R}$.

Now, using $\phi = \phi_1 + i\phi_2$ and assuming $\mu_1 \notin \mathrm{Spec}(\mathcal{L})$ (else 2.10 is trivial):

$$(\mathcal{L} - \mu_1)\phi = \mathcal{L}\phi_1 + i\mathcal{L}\phi_2 - \mu_1\phi_1 - i\mu_1\phi_2$$

$$= is\tau^2\chi_{\Omega\backslash R}\phi_1 - s\tau^2\chi_{\Omega\backslash R}\phi_2 - \delta^2 is\chi_R\phi_1 + \delta^2 s\chi_R\phi_2$$

Comparing the real (and imaginary) parts gives $(\mathcal{L} - \mu_1)\phi_1 = -s(\tau^2\chi_{\Omega\backslash R}\phi_2 - \delta^2\chi_R\phi_2)$. Taking the inverse of $(\mathcal{L} - \mu_1)$, applying the norm, and then squaring both sides:

$$||\phi||^2_{L^2(\Omega)} \leq ||(\mathcal{L} - \mu_1)^{-1}||^2|| - is\delta^2\chi_R\phi + is\tau^2\chi_{\Omega/R}\phi||^2 \qquad (2.12)$$

$$= ||(\mathcal{L} - \mu_1)^{-1}||^2(s^2\delta^4||\phi||^2_R + s^2\tau^4||\phi||^2_{\Omega\backslash R}). \qquad (2.13)$$

Where we used that $\chi_{\Omega\backslash R}\phi\chi_R\phi = 0$. Rearranging again and using the fact that $\delta^2 + \tau^2 = 1$ gives

$$||(\mathcal{L} - \mu_1)^{-1}||^{-2} \leq s^2\delta^4\tau^2 + s^2\tau^4\delta^2 = s^2\delta^2\tau^2 \qquad (2.14)$$

Thus, (2.10) is proved.

Denote by $P = \frac{1}{2\pi i}\int_\gamma (z - \mathcal{L})^{-1}\,dz$ as the orthogonal, spectral projector for $E(\Lambda, \mathcal{L})$, where γ is a simple closed contour that encloses $\Lambda \cup \{\mu_1\}$ and excludes $\mathrm{Spec}(\mathcal{L}) \backslash \Lambda$. This implies that P and $\mathcal{L}$ commute, thus $(I - P)\mathcal{L} = \mathcal{L}(I - P)$. Further using the fact that $(I - P)^2 = (I - P)$ gives:

$$I - P = ((\mathcal{L} + is)(I - P) - \mu)^{-1}(I - P)(\mathcal{L} + is - \mu), \qquad (2.15)$$

This in combination with (2.14) results in:

$$\inf_{v \in E(\Lambda, \mathcal{L})} \frac{||\phi - v||_{L^2(\Omega)}}{||\phi||_{L^2(\Omega)}} = \frac{||(I - P)\phi||_{L^2(\Omega}}{||\phi||_{L^2(\Omega)}}$$

$$\leq ||(\mathcal{L} + \mathfrak{i}\, s)(I - P) - \mu)^{-1}|| s\delta(\phi, R)\tau(\phi, R) \ ,$$

where (2.11) is attained from recognizing that

$$||(\mathcal{L} + \mathfrak{i}\, s)(I - P) - \mu)^{-1}|| = \frac{1}{\text{dist}(\mu_1, \text{Spec}(\mathcal{L})) \setminus (\Lambda \cup \mu_1)} \ .$$

$\square$

For the bounds in (2.9) to be meaningful, the right hand side must be less than one. Notice when there are two eigenvectors whose eigenvalues are close and are included in the contour Λ, then the $\text{dist}\,(\Re\mu, \text{Spec}(\mathcal{L})) \setminus (\Lambda \cup \{\Re\mu\})$ in the denominator blows up, which means that s may have to be very tiny for the bound to be meaningful. In these cases, an eigenvector of $\mathcal{L}_s$ may not need to be as "faithful" a representation of any one eigenvector of $\mathcal{L}$; this can be observed when the eigenvector of the perturbed operator appears as a linear combination of the eigenvectors with the close eigenvalues. When two eigenvectors are identical except in sign, then that linear combination may be an eigenvector that (falsely) appears localized in one region of the domain.

In the context of Theorem 2.1.4, if some scaling, $c\phi$, of ϕ is close to a (real) eigenvector ψ of $\mathcal{L}$, which happens when the upper bound of (2.11) is small, then

the imaginary part of $c\phi$ should be small. This motivates normalizing ϕ,

$$\phi \longleftarrow c\phi \text{ where } c = \arg\min\{||\Im(d\phi)||_{L^2(\Omega)} : |d| = 1\} . \tag{2.16}$$

Given a nonzero function $\phi = \phi_1 + i\phi_2$ such that $||\phi||_{L^2(\Omega)} = 1$, the minimization problem $\alpha = \min\{||\Im(d\phi)||_{L^2(\Omega)} : |d| = 1\}$ can be cast as a 2×2 eigenvalue problem whose smallest eigenvalue is α^2 and with semidefinite matrix:

$$\begin{pmatrix} ||\phi_2||^2_{L^2(\Omega)} & \int_\Omega \phi_1\phi_2 \, dx \\ \int_\Omega \phi_1\phi_2 \, dx & ||\phi_1||^2_{L^2(\Omega)} \end{pmatrix} . \tag{2.17}$$

The eigenvectors, $\mathbf{c} = (c_1, c_2)$, are real and unit length, and are related to the optimal scalar by $c = c_1 + ic_2$. However, a better, more effective minimization procedure for our purposes arises from Theorem 2.1.4 and is described below.

Proposition 2.1.5. *Let (μ, ϕ) be an eigenpair of $\mathcal{L}_s$. It holds that:*

$$||(\mathcal{L} - \Re\mu)\Re\phi||^2_{L^2(\Omega)} = s^2(\tau^4||\Im\phi||^2_{L^2(\Omega\setminus R)} + \delta^4||\Im\phi||^2_{L^2(R)}) \tag{2.18}$$

Proof. In the proof of Theorem 2.1.4, we obtained the identity: $(\mathcal{L} - \mu_1)\phi = is\tau^2\chi_{\Omega\setminus R}\phi_1 - s\tau^2\chi_{\Omega\setminus R}\phi_2 - \delta^2is\chi_R\phi_1 + \delta^2s\chi_R\phi_2$, then compared the real and imaginary parts to obtain

$$(\mathcal{L} - \mu_1)\phi_1 = -s(\tau^2\chi_{\Omega\setminus R}\phi_2 - \delta^2\chi_R\phi_2).$$

In the same manner as before, we obtain the identity 2.18. □

Notice that ϕ can be scaled by any $c = e^{i\theta}$ while remaining an eigenvector of unit length, and that τ and δ are scale invariant. Proposition 2.1.5 suggests that the residual $(\tau^4||\Im\phi||^2_{L^2(\Omega\backslash R)} + \delta^4||\Im\phi||^2_{L^2(R)})$, not $|\Im\phi|$, should be minimized. In some cases, $|\Im\phi|$ is not reduced. We begin by normalizing ϕ:

$$\phi \longleftarrow c\phi \text{ where } c = \arg\min\{\tau^4||\Im(b\phi)||^2_{L^2(\Omega\backslash R)} + \delta^4||\Im(b\phi)||^2_{L^2(R)} : |b| = 1\} .$$

(2.19)

We can recast this minimization problem as a 2×2 Hermitian eigenvalue problem using the matrix:

$$\delta^4 \begin{pmatrix} ||\phi_2||^2_{L^2(R)} & \int_R \phi_1\phi_2 \, dx \\ \int_R \phi_1\phi_2 \, dx & ||\phi_1||^2_{L^2(R)} \end{pmatrix} + \tau^4 \begin{pmatrix} ||\phi_2||^2_{L^2(\Omega\backslash R)} & \int_{\Omega\backslash R} \phi_1\phi_2 \, dx \\ \int_{\Omega\backslash R} \phi_1\phi_2 \, dx & ||\phi_1||^2_{L^2(\Omega\backslash R)} \end{pmatrix} . \quad (2.20)$$

Figure 2.4 exemplifies how normalizing ϕ in this manner returns eigenvectors of $\mathcal{L}_s$ that more closely resemble their corresponding eigenvectors of $\mathcal{L}$, and is therefore the superior minimization procedure.

The renormalization procedure is not dependent on the size of s, despite its appearance in Proposition 2.1.5. However, as we have noted, the quality of the returned eigenvectors of $\mathcal{L}_s$ often rely (at least partially) on the size of s. Therefore it is best to consider what s is appropriate before a renormalization is even relevant. The proceeding lemma provides a sufficiently small s in theory, however note that this is not something that will be computed, and furthermore it is often (much)

Original ϕ	Renormalized ϕ	ψ
561.36244 + 99.56455i	561.36244 + 99.56455i	561.35749
42.141598	0.0059877762	
7153.5089 + 12.1049i	7153.5089 + 12.1049i	7154.5778
15.114742	5.0622032	
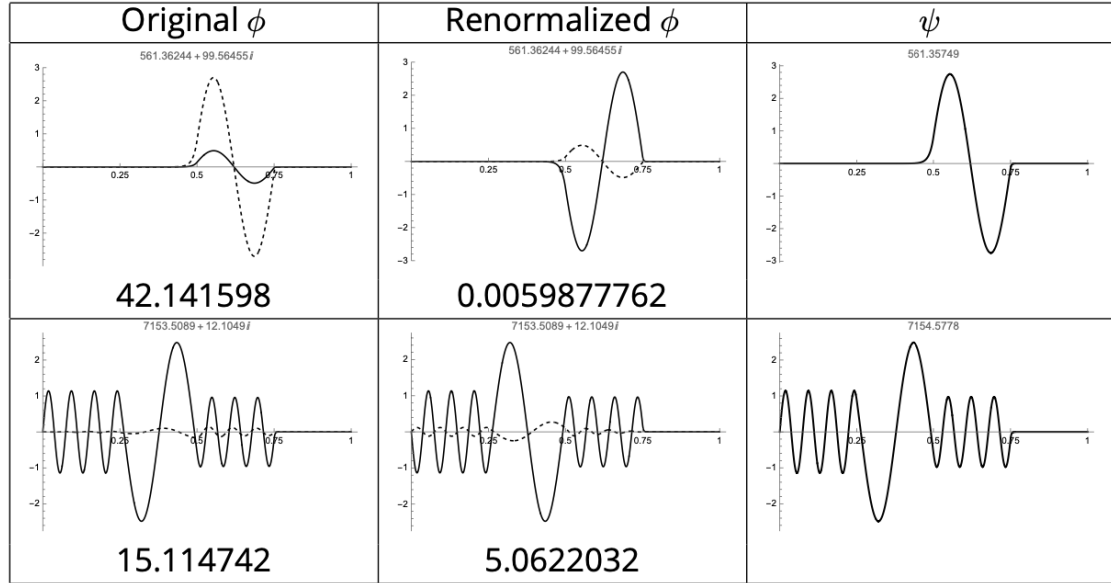
Figure 2.4: Each row corresponds to an eigenmode. The first column represents the unchanged eigenmode ϕ of $\mathcal{L}_s$, the second illustrates the renormalized (as in Proposition 2.1.5) ϕ, and the last column shows the corresponding ψ of $\mathcal{L}$. The real and imaginary parts of the eigenmode are shown in solid and dotted lines, respectively, while the value of the residual is included in the first two columns. Notice the renormalized eigenmode in column 2 much better resembles the eigenmode in column 3.

smaller than what is practical or necessary. Practical sizes for s will be discussed in Chapter 5.

Lemma 2.1.6. *Let (λ, ψ) be an eigenpair of the selfadjoint operator $\mathcal{L}$, $(\mu(s), \phi(s))$ be an eigenpair of $\mathcal{L}_s$ that depends on s, and assume the λ are simple. The size of s is such that:*

$$s \leq \left(\frac{1 - \delta(\psi, R)}{1 + \delta(\psi, R)}\right)\frac{r}{2} , \tag{2.21}$$

where $r = r_0$ is the radius of perturbation described by (2.8) and $\delta(v, R) = \dfrac{||v||_{L^2(\Omega \setminus R)}}{||v||_{L^2(\Omega)}}$ for any $v \in L^2(\Omega)$ and subdomain R.

Proof. From the proof in Theorem 2.1.2, we have $\dfrac{\lambda + \mathrm{i}s - \mu(s)}{\mathrm{i}s} = \dfrac{(\chi_{\Omega \setminus R}\phi(s), \psi)_{L^2(\Omega)}}{(\phi(s), \psi)_{L^2(\Omega)}}$. Taking the modulus, using the Cauchy-Schwarz Inequality, using that $\chi_{\Omega \setminus R}^2 = \chi_{\Omega \setminus R}$, gives:

$$\frac{|\lambda + \mathrm{i}s - \mu(s)|}{s} \leq \frac{||\phi(s)||_{L^2(\Omega \setminus R)}||\psi||_{L^2(\Omega \setminus R)}}{|\,||\psi||_{L^2(\Omega)}^2 - (\psi - \phi(s), \psi)_{L^2(\Omega)}|}$$

$$\leq \frac{||\phi(s)||_{L^2(\Omega \setminus R)}||\psi||_{L^2(\Omega \setminus R)}}{||\psi||_{L^2(\Omega)}^2 - ||\psi - \phi(s)||_{L^2(\Omega)}||\psi||_{L^2(\Omega)}} = \delta(\psi, R)\frac{||\phi(s)||_{L^2(\Omega \setminus R)}}{||\psi||_{L^2(\Omega)} - ||\psi - \phi(s)||_{L^2(\Omega)}}$$

We want to show that the latter term $\dfrac{||\phi(s)||_{L^2(\Omega \setminus R)}}{||\psi||_{L^2(\Omega)} - ||\psi - \phi(s)||_{L^2(\Omega)}} \leq 1$, since (by definition) $\delta(\psi, R) \leq 1$. This makes it so $\dfrac{|\lambda + \mathrm{i}s - \mu(s)|}{s}$ is meaningful. Assume that all of the λ are simple and denote by P the orthogonal spectral projector onto

the span(ψ). From the proof of Theorem 2.1.4, we have the identity:

$$I - P = ((\mathcal{L} + \mathrm{i}s)(I - P) - \mu(s))^{-1} (1 - P)(\mathcal{L} + \mathrm{i}s - \mu(s))$$

Notice $(I - P)\phi(s) = \phi(s) - c(s)\psi$ for some c since P is the orthogonal spectral projector onto span(ψ). Applying $\phi(s)$ to both sides of the above identity, applying the Cauchy-Schwarz Inequality, and using that $(\mathcal{L} + \mathrm{i}s - \mu(s))\phi(s) = \mathrm{i}s\chi_{\Omega\backslash R}\phi(s)$ gives:

$$||\phi(s) - c(s)\psi|| \leq || ((\mathcal{L} + \mathrm{i}s)(I - P) - \mu(s))^{-1} || \; ||(I - P)|| \; ||\phi(s)||_{L^2(\Omega\backslash R)}s$$

$$\leq || ((\mathcal{L} + \mathrm{i}s)(I - P) - \mu(s))^{-1} || \; ||\phi(s)||_{L^2(\Omega\backslash R)}s \; ,$$

where the last inequality is because $||(I - P)|| \leq 1$. Since $\mathcal{L} + \mathrm{i}s$ is a normal operator, $\mathcal{L}$ is selfadjoint, and $(I - P)$ is idempotent, then $(\mathcal{L} + \mathrm{i}s)(I - P)$ is also normal. As before, recognize that

$$|| ((\mathcal{L} + \mathrm{i}s)(I - P) - \mu(s))^{-1} || = \frac{1}{\mathrm{dist}\,(\mu(s), \mathrm{Spec}((\mathcal{L} + \mathrm{i}s)(I - P)))}$$

Notice that $\mathrm{Spec}(\mathcal{L} + \mathrm{i}s)$ is $\lambda + \mathrm{i}s$, for specific ϕ, $(I - P)(\lambda + \mathrm{i}s)\phi = 0$, and that applying $(I - P)$ to any other ϕ returns the same ϕ. Therefore, $\gamma :=$ dist $(\mu(s), \mathrm{Spec}((\mathcal{L} + \mathrm{i}s)(I - P))) = \mathrm{dist}(\mu(s), \lambda_b + is)$, where the target eigenvalue λ is distinct from its closest and second closest eigenvalues λ_b and λ_a, respectively. Figure 2.5 illustrates this for some $\mu(s)$ corresponding to the target λ.

Figure 2.5: For the proof of Lemma 2.1.6. Expressing the distance between $\mu(s)$ and $\mathrm{Spec}((\mathcal{L} + is)(I - P))$ is the next closest eigenvalue to λ, of which μ approximates, with the shift is added.

By definition, $\mathrm{dist}(\mu(s), \lambda_b + is) \geq \frac{r}{2}$, thus:

$$||\phi(s) - c(s)\psi|| \leq \frac{s}{(r/2)} ||\phi(s)||_{L^2(\Omega \backslash R)} \ . \tag{2.22}$$

Now consider the following where $c(s)\psi$ has been added and subtracted, and the triangle inequality has been applied:

$$||\phi(s)||_{L^2(\Omega \backslash R)} \leq ||c(s)\psi||_{L^2(\Omega \backslash R)} + ||\phi(s) - c(s)\psi||_{L^2(\Omega \backslash R)}$$

$$\leq ||c(s)\psi||_{L^2(\Omega \backslash R)} + \frac{s}{(r/2)} ||\phi(s)||_{L^2(\Omega \backslash R)} \ ,$$

where (2.22) was used for the last inequality. Rearranging:

$$||\phi(s)||_{L^2(\Omega \backslash R)} \leq \frac{1}{1 - (\frac{s}{(r/2)})} ||c(s)\psi||_{L^2(\Omega \backslash R)} \ . \tag{2.23}$$

Going back to the first identity to apply (2.22) and (2.23):

$$\frac{\lambda + is - \mu(s)}{s} \le \delta(\psi, R) \frac{(\frac{1}{1-s/(r/2)}) \|c(s)\psi\|_{L^2(\Omega \setminus R)}}{\|c(s)\psi\|_{L^2(\Omega)} - \frac{s}{(r/2)} \|\phi(s)\|_{L^2(\Omega \setminus R)}}$$

$$\le \delta(\psi, R) \frac{(\frac{1}{1-s/(r/2)}) \|c(s)\psi\|_{L^2(\Omega \setminus R)}}{\|c(s)\psi\|_{L^2(\Omega)} - \frac{s}{(r/2)} \frac{1}{(1-\frac{s}{(r/2)})} \|c(s)\psi\|_{L^2(\Omega \setminus R)}}$$

$$= \delta(\psi, R) \frac{(\frac{1}{1-s/(r/2)}) \delta(\psi, R)}{1 - \frac{s}{(r/2)} \frac{1}{(1-\frac{s}{(r/2)})} \delta(\psi, R)} = \delta(\psi, R) \frac{\delta(\psi, R)}{1 - \frac{s}{(r/2)} - \frac{s}{(r/2)} \delta(\psi, R)}$$

$$= \delta(\psi, R) \frac{\delta(\psi, R)}{1 - (1 + \delta(\psi, R)) \frac{s}{(r/2)}} \ .$$

As before, since $\delta(\psi, R) \le 1$ by definition, we want $\dfrac{\delta(\psi, R)}{1 - (1 + \delta(\psi, R)) \frac{s}{(r/2)}} \le 1$. So, it must be that:

$$\delta(\psi, R) \le \frac{\delta(\psi, R)}{1 - (1 + \delta(\psi, R)) \frac{s}{(r/2)}}$$

$$(1 + \delta(\psi, R)) \frac{s}{(r/2)} \le 1 - \delta(\psi, R) \ .$$

And finally,

$$s \le \left(\frac{1 - \delta(\psi, R)}{1 + \delta(\psi, R)} \right) \frac{r}{2} \ .$$

$\square$

The limit on s in Lemma 2.1.6 provides an upper bound that depends on how localized the target eigenvectors are and how far apart the eigenvalues of interest are from their nearest neighbors. A very localized eigenmode has a wider range of

"sufficiently small" than one that is less localized, or not localized at all, in R, and the same goes for eigenvalues that are further apart. It is well known that when computing eigenmodes via subspace iteration, the eigenvectors may mix when the dimension of the subspace is greater than 1 and the eigenvalues are "too" close to each other. It is important to note that this mixing may produce vectors that are falsely localized in some portion(s) of the domain, and we will call these cases "false positives". When the eigenvalues are very close together, then one may have to compensate with an even smaller s to avoid this, which is accounted for in Lemma 2.1.6. We reiterate that this s is *guaranteed* to be sufficiently small, but it may not be necessary or realistic in many practical applications, which will be seen in later chapters.

2.2   Unpacking our Algorithm

Our key task is to find localized eigenpairs of (1.1):

> Given an interval $[a, b]$ and a (small) tolerance $\delta^* > 0$, find all eigenpairs (λ, ψ) for which $\delta(\psi, R) \leq \delta^*$ and $\lambda \in [a, b]$, or determine that there are not any, for some subdomain R, which may be given or found by our algorithm.   (T)

Our intuition, which is explained in detail in the previous section, is that if (λ, ψ) is an eigenpair of $\mathcal{L}$ with ψ highly localized in R, then there will be an eigenpair (μ, ϕ) of $\mathcal{L}_s$ such that μ is near $\lambda + \mathrm{i}\, s$ and ϕ is near ψ. Our Algorithm 2.1 to

achieve our key task (T) is detailed below. The "*" symbol indicates that this line does not apply to the magnetic Schrödinger case because both of these lines assume the eigenvectors of $\mathcal{L}$ are real, which they are not when $\mathcal{L}$ is the magnetic Schrödinger. In the case of the magnetic Schrödinger, line 7 must be dropped and in line 8, we use $(\Re\mu, \phi)$ instead of $(\Re\mu, \Re\phi)$ since the eigenvalues of the magnetic Schrödinger are real.

Algorithm 2.1 Template

1: **procedure** LOCALIZE(a, b, s, δ^*, R)
2:      determine all eigenpairs (μ, ϕ) of $\mathcal{L}_s$ with $\mu \in U(a, b, s, \delta^*)$ ▷ Theorem 2.1.2
3:      **if** no eigenvalues are found in 2 **then**
4:         exit   ▷ There are no eigenpairs (λ, ψ) of $\mathcal{L}$ with $\lambda \in [a, b]$ and $\delta(\psi, R) \leq \delta^*$
5:      **else**
6:         **for** each eigenpair (μ, ϕ) found in 2 **do**
7:            normalize ϕ, $\phi \longleftarrow c\phi$ ▷ *, as in Proposition 2.1.5 in (2.20)
8:            post-process $(\Re\mu, \Re\phi)$ to obtain (approximate) eigenpair $(\tilde{\lambda}, \tilde{\psi})$ of $\mathcal{L}$
                                                              ▷ *
9:            **if** $\delta(\tilde{\psi}, R) \leq \delta^*$ **then**
10:              accept $(\tilde{\lambda}, \tilde{\psi})$
11:            **else**
12:              reject $(\tilde{\lambda}, \tilde{\psi})$
13:            **end if**
14:         **end for**
15:      **end if**
16:      **return** accepted (approximate) eigenpairs $(\tilde{\lambda}, \tilde{\psi})$
17: **end procedure**

The post-processing step given in Algorithm 2.1 involves performing (a few) inverse operations given in Algorithm 2.2. The purpose of the post-processing is to find the (approximate) desired eigenpair of the original operator $\mathcal{L}$ from the eigenpair of the shifted operator that was found by our algorithm. Theorem 2.1.4

and Proposition 2.1.5 indicate that a good starting point for the post-processing is $(\mu_1, \phi_1) = (\Re\mu, \Re\phi)$ since it is close to the desired eigenpair. Further, Proposition 2.1.5 provides an efficient way to compute the initial residual, which may already be smaller than the prescribed tolerance, and therefore no post-processing is needed.

Algorithm 2.2 Approximate Eigenpair Post-Processing

1: **procedure** POST-PROCESS($\mu_1, \phi_1, \text{tol}$)
2: $\quad \tilde{\lambda} \longleftarrow \mu_1$
3: $\quad \tilde{\psi} \longleftarrow \phi_1/||\phi_1||_{L^2(\Omega)}$
4: $\quad$ **while** $||\mathcal{L}\tilde{\psi} - \tilde{\lambda}\tilde{\psi}||_{L^2(\Omega)} > \text{tol}$ **do**
5: $\qquad \tilde{\psi} \longleftarrow (\mu_1 - \mathcal{L})^{-1}\tilde{\psi}$
6: $\qquad \tilde{\psi} \longleftarrow \tilde{\psi}/||\tilde{\psi}||_{L^2(\Omega)}$
7: $\qquad \tilde{\lambda} \longleftarrow (\mathcal{L}\tilde{\psi}, \tilde{\psi})$
8: $\quad$ **end while**
9: $\quad$ **return** post-processed (approximate) eigenpair $(\tilde{\lambda}, \tilde{\psi})$
10: **end procedure**

It is expected that the post-processing will converge in no more than a few iterations; no convergence after a few iterations may indicate that (μ, ϕ) is an approximation of more than one eigenpair of $\mathcal{L}$. These cases occur when there are multiple eigenmodes with very close eigenvalues. As discussed, this could produce false positives. In practice, we have often seen that the inverse iterations converge to a more faithful representation of one eigenfunction of $\mathcal{L}$ that is not truly localized. And this has the (fortunate) effect of alerting the user to the potential false positive.

3   Background Information and Methodology

This chapter dives into some of the basic background information such as eigen-value solvers and descriptions of a few different methods that are commonly used. Filtered Subspace iteration, as it relates to FEAST, and the FEAST algorithm [22] are explained as well. Pythonic FEAST [16, 17] used in conjunction with the finite element tool NGSolve [24], interfaced with Python [38], are the tools used to perform all of the computations herein.

3.1   Eigenvalue solvers

An eigenvalue solver aims to compute eigenvalues $\lambda \in \mathbb{C}$ and their corresponding eigenvectors $\vec{x} \in \mathbb{C}^n$ satisfying:

$$A\vec{x} = \lambda\vec{x} \text{ or equivalently } A\vec{x} = \lambda M\vec{x} \ , \tag{3.1}$$

where the stiffness matrix A and the mass matrix M are large and sparse, as is typical when A and M arise from the discretization of a differential operator. A common example, and one used throughout this work, is:

$$\begin{cases} (-\Delta + V)\psi = \lambda\psi \ \text{ in } \ \Omega \\ \\ \psi|_{\delta\Omega} = 0 \end{cases}$$

where $\Omega \subset \mathbb{R}^d$ is a bounded, connected, open set, the potential $V \in L^\infty(\Omega)$ is non-negative, and the eigenfunction $\psi \in W := H_0^1 = \{\psi \in H^1 : \psi|_{\delta\Omega} = 0\}$. The strong form of the problem (above) is converted into its weak form by multiplying by a test function v and integrating by parts. In this case, the variational formulation is to find $\lambda \in \mathbb{C}$ and $\psi \in H_0^1$ satisfying

$$\int_\Omega \nabla\psi\nabla v dx + \int_\Omega V\psi v dx = \lambda \int_\Omega \psi v dx \ \forall v \in W \ .$$

Here, we used the zero boundary conditions to eliminate the boundary term that arose from using the Divergence Theorem. To make the problem suitable for solving via finite element methods, we choose a finite dimensional subspace $W_h \subset W$ that consists of continuous piecewise polynomial functions on a triangulation of Ω. The triangulation is called quasi-uniform if all of the triangles have the same diameter. If the local polynomial degree is fixed and the triangulation is quasi-uniform, then the diameter of the triangles $h = 1/N$ is proportional to $\dim(W_h) = N$. The problem becomes:

$$\int_\Omega \nabla\psi_h\nabla v_h dx + \int_\Omega V\psi_h v_h dx = \lambda_h \int_\Omega \psi_h v_h dx \ \forall v \in W_h \ .$$

Expanding ψ_h via the finite element basis functions $\{\phi_j\}_{j=1}^N$ of W_h, $\psi_h = \sum_{j=1}^N x_j\phi_j$, and substituting ϕ_i for each v_h results in:

$$\forall \ i \in [0, N], \ \sum_{j=1}^N \left(\int_\Omega \nabla\phi_j\nabla\phi_i + V\phi_j\phi_i dx\right)\psi_j = \lambda_h \sum_{j=1}^N \left(\int_\Omega \phi_j\phi_i dx\right)\psi_j \ .$$

Thus, the discretized eigenvalue problem is complete after assembling:

$$A = \left(\int_\Omega \nabla\phi_j \nabla\phi_i + V\phi_j\phi_i dx \right) \text{ and } M = \left(\int_\Omega \phi_j\phi_i dx \right)$$

A suitable eigenvalue solver can be used from here. There are many such solvers, and a few will be summarized here such as the Power Method and Subspace Iteration. For simplicity, we describe the methods for the standard eigenvalue problem $A\vec{x} = \lambda\vec{x}$.

3.1.1 Power Method and Shifted Power Methods

The power method begins with some normalized initial guess vector $\vec{x}_0$, which is either an approximation of the dominant eigenvector or chosen at random. The matrix A is then applied to $\vec{x}_0$, propagating another vector that is immediately normalized to produce $\vec{x}_1$. This is an iterative procedure that produces a sequence of vectors $\vec{x}_n$ that converges to the dominant eigenvector, given that some basic assumptions are satisfied. Eigenvalues are then recovered via Rayleigh quotients. The necessary assumptions are that there is an eigenvalue λ_1 of A that is strictly greater than the other eigenvalues in magnitude (with a larger disparity preferred) and that A is diagonalizable in the sense $A = V\Lambda V^{-1}$, where Λ is the diagonal matrix of eigenvalues. The diagonalizable condition on A guarantees that it has n linearly independent eigenvectors that form a basis on the underlying space, and therefore there is some $\vec{x}_0$ guaranteed to have a nonzero component in the direction of an eigenvector associated with λ_1. However, this condition is already

(almost surely) guaranteed: in the unlikely chance that $\vec{x}_0$ was chosen to have a zero component in the direction of the associated eigenvector, floating point arithmetic would generally perturb this to have a nonzero coefficient.

Denote by λ_2 the second largest eigenvalue. The power method convergence rate is geometric with ratio $\frac{|\lambda_2|}{|\lambda_1|}$, making it clear why λ_1 must be simple and dominate, and why a larger disparity between λ_1 and λ_2 is desired. Convergence may not even happen if the dominant eigenvalue is too close to the next eigenvalue. Multiplication of A by $\vec{x}_k$ is the computationally most expensive step with $\mathcal{O}(n^2)$ when A is dense and $\mathcal{O}(n)$ when A is sparse, so the Power Method is most effective for very large sparse matrices. Even though the convergence rate can sometimes be slow and only one eigenpair results from the power method, it is still a decent starting point for more sophisticated methods because it is simple to implement and only requires a "black box" that multiplies a matrix by a vector, which is a great advantage for problems where the matrix is mostly zeros or has some special structure. [14, 31]

Utilizing that the eigenvalues of A^{-1} are the reciprocals of the eigenvalues of A, one can instead find the smallest eigenvalue of A by implementing the power method with A^{-1} substituted for A; this is known as the Inverse Power Method, and it has the same basic assumptions and results as the Power Method above.

If one instead wishes to find other eigenpairs as opposed to just the largest or the smallest, the Shifted Power Method can be used. This method incorporates a shift of the operator, $A - \sigma I$, by some user specified parameter σ. This new matrix has precisely the same eigenvectors as A, and $\text{Spec}(A - \sigma I) = \lambda - \sigma : \lambda \in \text{Spec}(A)$.

Applying the power method to $A - \sigma I$ yields an eigenvector of A whose eigenvalue is furthest from σ. As before, the corresponding eigenvalue of A is obtained from a Rayleigh quotient of A. Similarly, applying the inverse power method to $A - \sigma I$ yields an eigenvector of A whose eigenvalue is closest to σ.

Even for a reasonable shift, the Shifted Power Method converges linearly (at best!). To speed up the algorithm, shift dynamically! A prime example is used in the Rayleigh Quotient Iterations Method, where the shift is chosen dynamically by applying the Rayleigh quotient at each step:

$$\sigma_{k+1} = \frac{(v^{(k)}, Av^{(k)})}{||v^{(k)}||^2} \ .$$

This method is locally quadratic for asymmetric matrices and locally cubic for symmetric matrices. However, in the case of clustered eigenvalues, these methods may not converge or are greatly slowed; even a carefully selected shift cannot be chosen to increase the needed gap between successive eigenvalues. In these instances, one can turn to Subspace Iteration. [31]

3.1.2 Subspace Iteration and Filtered Subspace Iteration

Instead of focusing on finding an eigenvector associated with the dominant eigenvalue, Subspace Iteration instead finds the invariant subspace associated with the p dominant eigenvalues, $|\lambda_1| \geq ... \geq |\lambda_p| > |\lambda_{p+1}| \geq ... \geq |\lambda_n|$. One computes $\mathcal{V}_k = A^{(k)} \mathcal{V}_0$ for some initial subspace $\mathcal{V}_0$ often chosen as a p-dimensional subspace of random initial vectors.

Computing the $\mathcal{V}_k$ requires computing a basis for the space first, which can be achieved via QR decomposition. Essentially, any matrix $A \in \mathbb{C}^{m \times p}$ with linearly independent columns and $m \geq p$ can be factored into the product of an $m \times p$ matrix X with orthonormal columns and a $p \times p$ upper triangular matrix, R. Given some gap between $|\lambda_p|$ and $|\lambda_{p+1}|$, these iterations converge to an orthonormal basis for the invariant subspace spanned by the first p eigenvectors of A. Particularly, the first n columns of X_k converge to a basis for the dominant n-dimensional invariant subspace.

Algorithm 1.3 Subspace Iteration

1: **procedure** FIND EIGENSPACE(A, X, tol)
2:     compute eigenvectors corresponding to the p largest eigenvalues of A
3:     Set $X^{(0)} := X$, $k = 0$ $\triangleright A \in \mathbb{C}^{n \times n}$ $\triangleright X \in \mathbb{C}^{n \times p}$
4:     **repeat**
5:         $k = k + 1$
6:         $Z^{(k)} := A X^{(k-1)}$
7:         $X^{(k)} R^{(k)} := Z^{(k)}$ $\triangleright$ QR factorization of $Z^{(k)}$ to prevent the columns of $X^{(k)}$ from converging to the eigenvector with the largest modulus
8:     **until** $||(I - X^{(k)} X^{(k)*}) X^{(k-1)}|| > tol$
9:     **Converged** $A = X^{(k)} R^{(k)} X^{-(k)}$ is an approximation of the partial Schur decomposition
10: **end procedure**

The convergence rate is like that of the power iteration, from which it is based on, and is often realized in practice. One iteration converges linearly and the rate of convergence is related to the ratio of the moduli of the eigenvalues. When the eigenvalues are close, convergence is slow. Like the power iteration, it is better if the eigenvalues are far apart. Recall that subspace iteration focuses on finding the entire invariant subspace associated w p eigenvalues. The value $m > p$ denotes the

number of initial vectors at the beginning of the iterations, so when m is small, a relatively large number of iterations are required to converge. However when m is large, each iteration requires more computation though there are less of them. In an effort to optimize, it is standard to set m to the number of expected vectors p or *just* above it to be safe, but not choose m so much larger that each iteration is unnecessarily slowed. [9]

Filtered subspace iteration differs from basic subspace iteration by first transforming the eigenspace of the cluster Λ to the dominant eigenspace of the filter, another bounded operator. Subspace iteration is then applied. Computable finite rank approximations of the resolvent at a few points are required to apply this method; these approximations are obtained via some discretization process that are explored in much more detail in [16]. The FEAST algorithm, which is the primary eigenvalue solver used throughout this work, performs filtered subspace iteration.

3.2 FEAST

FEAST [16, 17, 22] is subspace iteration with a "filtered" version of the operator/matrix A, given by $B = f(A)$. The filter function is chosen to be a rational approximation of the characteristic function for some search region $\tilde{U}$, so that the eigenvalues of A that are in $\tilde{U}$ get mapped to the dominant eigenvalues of B: $\lambda \in \mathrm{Spec}(A) \to f(\lambda) \in \mathrm{Spec}(B)$. This f greatly enlarges $|f(\lambda)|$ for $\lambda \in \mathrm{Spec}(A) \cap \widetilde{U}$ relative to $\lambda \in \mathrm{Spec}(A) \cap \widetilde{U}^c$. An ideal filter maps the desired eigenvalues to one

and all others to zero. Assuming A is some unbounded operator on $L^2(\Omega)$, whose domain $\text{Dom}(A) \subset L^2(\Omega)$, then the filter function $f = f(z)$ is rational, bounded, and continuous on $\text{Spec}(A)$, so $B = f(A)$ is bounded and $\text{Dom}(f(A)) = L^2(\Omega)$. It's important to note that even though f changes the eigenvalues λ of A to $f(\lambda)$, the eigenfunctions are unchanged; the eigenfunctions of $f(A)$ are precisely the same as the eigenfunctions of A. The idea for the filter function comes from Cauchy's Integral Formula:

$$F(\xi) = \chi_{\tilde{U}}(\xi) = \frac{1}{2\pi i} \oint_{\partial \tilde{U}} \frac{1}{z - \xi} dz = \begin{cases} 0, & \xi \in \tilde{U}^c \setminus \partial \tilde{U} \\ 1, & \xi \in \tilde{U} \end{cases} \tag{3.2}$$

We assume $\partial \tilde{U}$ is some simple, closed loop, such as a circle of nonzero radius, and that ξ is a point not in $\partial \tilde{U}$. In this case, f takes the form of the rational quadrature approximation of (3.2) given by:

$$f(\xi) = \sum_{k=0}^{n-1} \frac{w_k}{z_k - \xi} \approx \frac{1}{2\pi i} \oint_{\partial \tilde{U}} \frac{1}{z - \xi} dz = F(\xi) \ ,$$

where w_k are the quadrature weights and the z_k are the quadrature nodes. While there are other ways of obtaining such f that are considered in [19, 37], they will not be explored here.

This motivates the definition of the main computational tool for the FEAST algorithm, the orthogonal spectral projector $\mathcal{S}$ of A associated with Λ. $\mathcal{S} : H \to H$ is a bounded, linear operator obtained from the Dunford-Taylor Integral with the

form:

$$S = \frac{1}{2\pi i} \int_{\partial \tilde{U}} (zM - A)^{-1} dz \ .$$

We then approximate S in the sense of:

$$S \approx S_k = \sum_{k=0}^{n-1} (z_k M - A)^{-1} \ .$$

S_k is iteratively applied to an initial, random, finite dimensional subspace, $E_{h,0} \subset$ Dom(A) that satisfies $S_k E_{h,0} = E_h$, where $E_h := Range(S_k) = E(\Lambda, A)$ is the target invariant subspace and h is the standard meshing parameter. Thus FEAST iterations generate a collection of subspaces having the form $\{E_{h,k+1} = f(A)E_{h,k}\}$ that converge to E_h with respect to the subspace gap. The subspace gap measures the distance between two linear subspaces V, W of E_h, and is defined:

$$gap_{E_h}(V, W) = \max \left(\sup_{v \in \overline{B}_V} \text{dist}_E(v, W), \sup_{w \in \overline{B}_w} \text{dist}_E(w, V) \right) \ ,$$

where $\overline{B}_w$ is the closed, unit ball of W. The rate of convergence of $E_{h,k}$ to E_h is the dist$(E_h, E_{h,k}) = O(\kappa^k)$. This rate is controlled by the ratio:

$$\kappa := \frac{\sup_{\lambda \in \text{Spec}(A) \backslash \Lambda} |f(\xi)|}{\inf_{\lambda \in \Lambda} |f(\xi)|} < 1 \ . \tag{3.3}$$

Recall that the filter function (ideally) maps eigenvalues in the contour to 1 and those outside to 0. Having a filter function with clear distinction between the

eigenvalues of interest and the remaining eigenvalues makes κ small, as desired.

In order to recover the target eigenvalues Λ, a Rayleigh-Ritz procedure is used on a finite rank operator $A_k : E_k \mapsto E_k$ to obtain approximations $\Lambda_{h,k}$ of Λ that converge to Λ_h with respect to the Hausdorff metric. We define the Hausdorff metric between two sets of numbers $X, Y \in \mathbb{C}$ as

$$\text{dist}(X, Y) = \max \left(\sup_{x \in X} \inf_{y \in Y} |x - y|, \sup_{y \in Y} \inf_{x \in X} |x - y| \right) .$$

This rate of convergence is $\text{dist}(\Lambda_h, \Lambda_{h,k}) = O(\kappa^{2k})$. Again, it can be seen that a "good" filter function can greatly affect the speed of the algorithm. [16, 17, 33]

3.2.1 Computational Complexity of FEAST

It may come as no surprise that the major cost of FEAST is solving each system:

$$X_{j+1} = \sum_{k=1}^{N} w_k (z_k M - A)^{-1} X_j ,$$

where M and A are $N \times N$ matrices, X is an $N \times p$ matrix, and these should be sparse matrices. The N number of quadrature points used determines the number of matrix solves per linear system per iteration. In practice, this discretized version of the finite rank approximation of the resolvent term $(z - A)^{-1}$ depends on h. This is because the quadrature points, z_k, and weights, w_k, both relate to the mesh parameter, h. This relationship is clearly defined for our Bunimovich curve in 2.6 in section 2.1.

With resolvent discretization, subspace iteration can be applied as described before for each N quadrature point. The m_0 number of initial guesses in FEAST, which quickly becomes the m number of eigenvectors in the contour, determines the number of right hand side vectors b to the linear problem $Ax = b$. For example, for $m = 10$ vectors given in a single iteration, there are 10 systems of linear equations. A sparse direct solver involving LU factorization, known as UMFPACK, is used for the linear solves. The computational complexity for one solve is on the order of $\mathcal{O}(p \cdot N^{3/2})$ [27]. However, each quadrature point has an associated matrix where a sparse solve is performed and for each of these associate matrices there are m number of linear systems to solve. Therefore, the computational complexity becomes $\mathcal{O}(p \cdot m \cdot N^{3/2})$. One could choose a different solver, such as MUMPS, and the computational complexity would become $\mathcal{O}(p \cdot m \cdot N \log(N))$. We see that these matrix solves can be slow, especially when m is chosen large.

With respect to the question about how long a typical FEAST simulation takes, there are several considerations beyond choosing a suitable number of initial vectors m. As discussed, choosing a larger number of quadrature points N, a larger search interval $[a, b]$ (or just having an interval dense with eigenmodes), a finer mesh, or a higher polynomial degree will all increase computational time. The choice of contour also matters in practice; for the Bunimovich curve, the relationship between N, $[a, b]$, and some other parameters is outlined as the aspect ratio in 5.1. Having a poor aspect ratio makes it difficult for the algorithm to determine which eigenmodes to filter, and therefore more iterations are needed to resolve this issue; alternatively, some eigenmodes would not be filtered, and the more eigenvectors

found per iteration, the more time and expense.

Still, the general takeaway should be that the most costly step is building the rational filter, which increases in expense for larger N. The linear solves are typically faster, however have more parameters that can affect time and expense. For actual times and memory consumption of a typical simulation, we refer the reader to Chapter 5 to find Tables 5.1, 5.2, and 5.3.

4   The Algorithm

4.1   Proof of Concept

This section provides examples of the algorithm in action coupled with the post processing. These examples range between one and two dimensional domains and are shown for localization that derives from the domain geometry, potential, or properties of the operator. There are 265 eigenmodes of the original operators computed in this section; 66 of those *strictly* meet whatever criteria are given and 80 are returned before any post processing. Right off the bat, the algorithm eliminated *over* 90% of the unwanted eigenvectors, where some of the few unwanted eigenmodes were *very* close to meeting the criteria. With the short post-processing step, over 99% were clearly eliminated. Details of the post-processing are covered in the post-processing section.

Example 4.1.1. Figures 4.1 and 4.2 summarize the algorithm performance on the four dumbbell variations featured in the "Geometry" chapter. In total, there are 48 eigenpairs of $\mathcal{L}$ with 14 that are localized within the criteria given in each case, all of which were returned. There were a total of 2 undesired vectors found, which were quickly tossed after post-processing. The algorithm immediately filtered $\approx 94\%$ of the undesired eigenpairs and then 100% in the post processing. Note that smaller s, such as $s = 0.01$ could have been used so that 100% of the unwanted eigenmodes were immediately filtered, but such s would require many iterations. The aspect

ratio in all cases remained below 18, and only a few iterations were needed for each case.

Example 4.1.2. Figure 4.3 provides the first 20 eigenpairs (though 23 total were computed) of $\mathcal{L}$ for the "Delitsyan" domain. It is coupled with Figure 4.4, which pictures the moduli of the 5 returned for the shifted operator $\mathcal{L}_s$ after the post-processing, where the right half of the domain is R. This region is far less likely to exhibit localization, with only two modes localized to $\tau = 0.916$. Prior to post-processing, 5 eigenfunctions of $\mathcal{L}_s$ were found, all of which contained significant eigenfunction mass in R, but may have bled too much into the much larger left half of the domain; the three that did not meet the tolerance had tau values $\tau = 0.867, 0.770,$ and 0.795, respectively. Thus, 85.7% of the unwanted eigenvectors were filtered immediately, and then 100% in the post processing, and *all* of the desired eigevectors were found. Note that a sufficient aspect ratio was maintained for these experiments, where the search region size did not exceed 5, and so only 3 iterations of the FEAST algorithm were implemented.

Example 4.1.3. Much like the Delitsyan example, the Threebulb domain provides another case where purely geometric localization occurs. In Chapter 5 Figure 5.3, the first 64 eigenmodes of the Threebulb were combed for eigenvectors concentrated to $\delta^* = 0.25$ (meaning $\tau \geq 0.9682$) in the left, right, and middle bulbs. There are 27 eigenvectors that strictly meet this tolerance with 5 that are *very* close, in the sense of having $\tau \in (0.956, 0.965)$. There are 3 that are pretty close in the sense of $\tau \in (0.942, 0.95)$. When $s = 1$ for search regions having length 10, the aspect ratio was 41 and there were 36 eigenmodes returned, which means

$\mu = 2.11001 + 0.4984i \qquad 5.47257 + 0.49996i \qquad 5.09189 + 0.4859i \qquad 8.72859 + 0.4998i$

Moduli of the eigenmodes of $\mathcal{L}_s$ for $s = 0.5$, $\delta^* = 0.25$, $R =$ right bulb, and aspect ratio of 17. This corresponds to the first twelve eigenfunctions of $\mathcal{L}$ featured in Figure 6.4. All four concentrated eigenfunctions were found without any false positives.

$\mu = 2.11003 + 0.4984i \qquad 5.09118 + 0.4886i$

Moduli of the eigenmodes of $\mathcal{L}_s$ for $s = 0.5$, $\delta^* = 0.2$, $R =$ right bulb, and aspect ratio of 16. This corresponds to the first twelve eigenfunctions of $\mathcal{L}$ featured in Figure 6.6. Both localized eigenfunctions were found without any false positives.

$\mu = 2.04385 + 0.4985i \qquad 5.29481 + 0.49996i \qquad 4.93914 + 0.4874i$

$8.44572 + 0.4998i \qquad 10.60746 + 0.49999i$

Moduli of the eigenmodes of $\mathcal{L}_s$ for $s = 0.5$, $\delta^* = 0.35$, $R =$ left bulb, and aspect ratio of ≈ 12.5. This corresponds to the first twelve eigenfunctions of $\mathcal{L}$ featured in Figure 6.5. All five concentrated eigenvectors were found without any false positives.

Figure 4.1: The algorithm applied to 36 eigenmodes of $\mathcal{L}$ over three similar domains from the "Geometric Localization" chapter. All desired eigenfunctions were returned without any that did not meet the criteria.
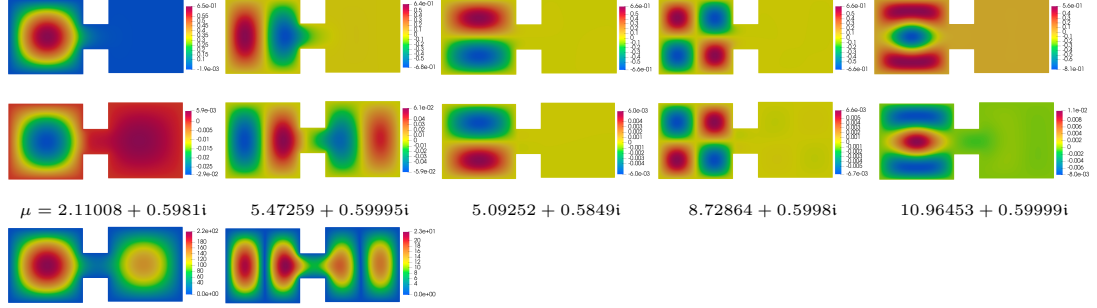
Figure 4.2: Real and imaginary parts of the eigenvectors of $\mathcal{L}_s$ and associated eigenvalues are featured in the top rows, respectively. Here $\delta^* = 0.2$, $s = 0.6$, R = left bulb, and the aspect ratio is ≈ 17.67. Corresponding eigenpairs of $\mathcal{L}$ are pictured in Figure 6.7. There are only 3 eigenfunctions that are sufficiently concentrated in the left bulb, but five were returned. Post processing, pictured in the third row, indicated that the first 2 eigenmodes were not sufficiently localized and were thus discarded. A smaller s, such as $s = 0.01$, could have instead been used to avoid returning any false positives.

immediately the algorithm discarded 28 eigenvectors. The desired 27 were all returned, plus the close 8 and the additional was a linear combination of two eigenvectors that had $\tau = 0.946$ and $\tau = 0.875$ in the right bulb. In terms of per bulb performance, there was 1 sufficiently localized eigemode in the middle bulb, which was identified without any false positives. For the right bulb, 5 additional were returned. This means the algorithm, before post processing, returned 100% of the desired eigenfunctions and only returned eigenfunctions that did not quite meet the tolerance 5.6% of the time, but were often very close. We highlight that improving the aspect ratio improves this performance, however this must be done by choosing a smaller search interval, which comes at the expense of computational resources and time. In this experiment, only three iterations of the algorithm were needed to sweep through the entire 64 eigenvectors, and these all converged
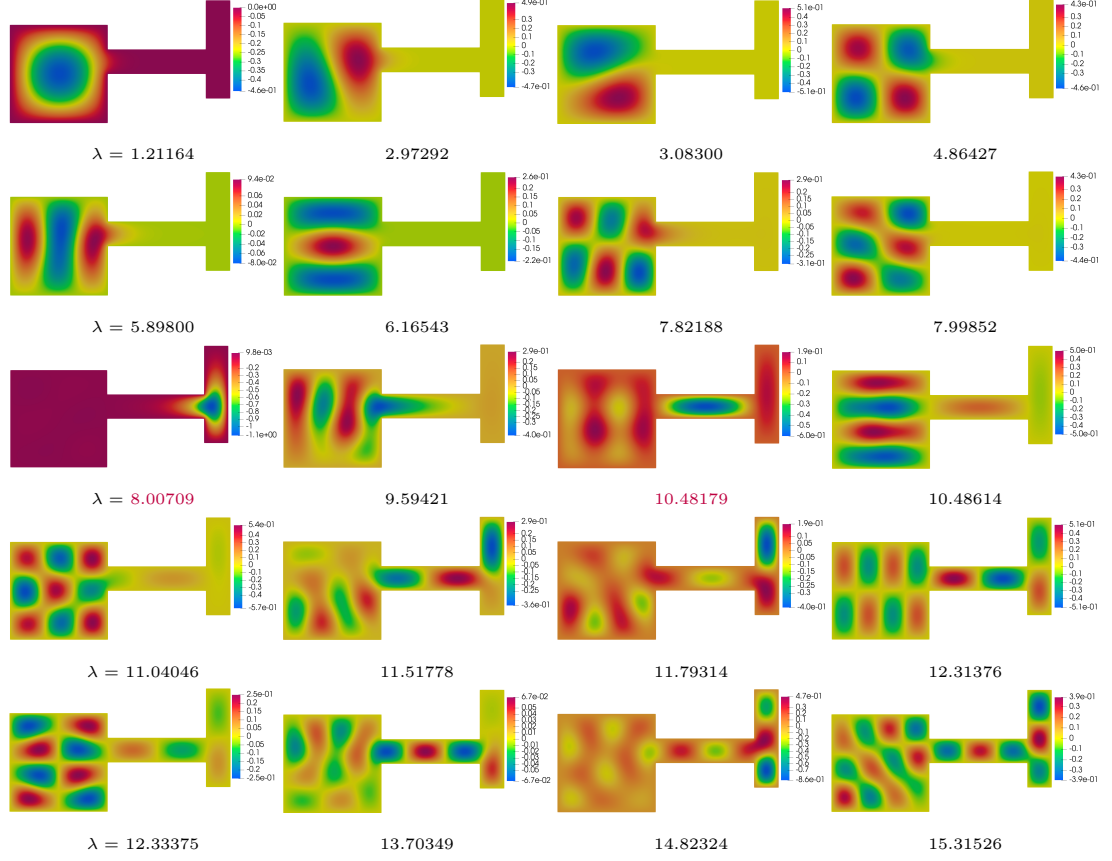
Figure 4.3: The first 20 eigenpairs of the "Delitsyan domain". The two purple eigenvalues correspond to the only eigenmodes that meet the tolerance $\delta^* = 0.4$ in the right half of the domain. The returned eigenmodes of $\mathcal{L}_s$ are pictured in Figure 4.4.

$\mu = 8.0071 + 0.999i$ $10.4815 + 0.919i$

$\mu = 12.0899 + 0.751i$ $13.7644 + 0.592i$ $15.2647 + 0.632i$

11.8584   13.7035   15.3153

Figure 4.4: The 5 returned (real parts of the) eigenmodes after the post processing for $s = 1$, and corresponding to Figure 4.3 are shown in the top two rows. The top row met the tolerance of $\delta^* = 0.4$. The bottom row contains the modulus of the post processed eigenmodes that did not meet the tolerance and were returned eigenvectors of $\mathcal{L}_s$. Post-processing, easier seen via the modulus, clearly identifies these to be non sufficiently localized and were then discarded. Of 23 total eigenmodes (not all pictured) only two met the localization criteria and were returned

within 3 FEAST iterations. One could also pick s sufficiently small to ensure there were no false positives, but again this can be cumbersome since the search interval must also decrease accordingly to maintain the accuracy and resolution. For a typical use case, quickly identifying all of the desired eigenvectors and then having the additional (if present) thrown out in the post processing is a much more useful, practical approach than trying to optimize the parameters to ensure only the precisely wanted eigenmodes are returned, which can be time consuming. The *very* close eigenvectors may even prove unexpectedly desired.

Example 4.1.4. Testing the algorithm on cases where an interesting potential induces localization proves fruitful as well. A one dimensional case on the unit interval chopped into 4 equally sliced regions and having potential $V = [0, 80^2, 0, 160^2]$ is given in Figure 4.5 and Figure 4.6 with a very strict $\delta = 0.2$. After some post processing, only the 3 eigenmodes that met the tolerance were returned, even though there were 2 *very* close ones, exemplifying how rigorous the post-processing can be. The real and imaginary parts of the 3 accepted eigenfunctions are shown; notice the imaginary parts are five orders of magnitude smaller than the real parts, which really drives home the idea that the eigenvectors of $\mathcal{L}_s$ closely resemble the (real) eigenvectors of $\mathcal{L}$, as is consistent with (2.16).

Example 4.1.5. A simple two dimensional example where the localization arises from the potential is illustrated with the square domain, where no eigenmode of $\mathcal{L}$ can concentrate in any subdomain without it being the result of some potential. Figure 4.7 provides the first 20 eigenpairs on the unit square that is chopped into 16 total sub-squares, each having a random, constant potential in $[0, 1500]$. A total

$\lambda = 136.59654$ 143.18099 545.74931 572.08290 1225.36986

$\mu =$ 143.18099 + 0.99892i 572.08290 + 0.99555i

$\lambda = 1284.57374$ 2171.22973 2276.20569 3374.97283 3537.72619

$\mu = 1284.57374 + 0.98937i$
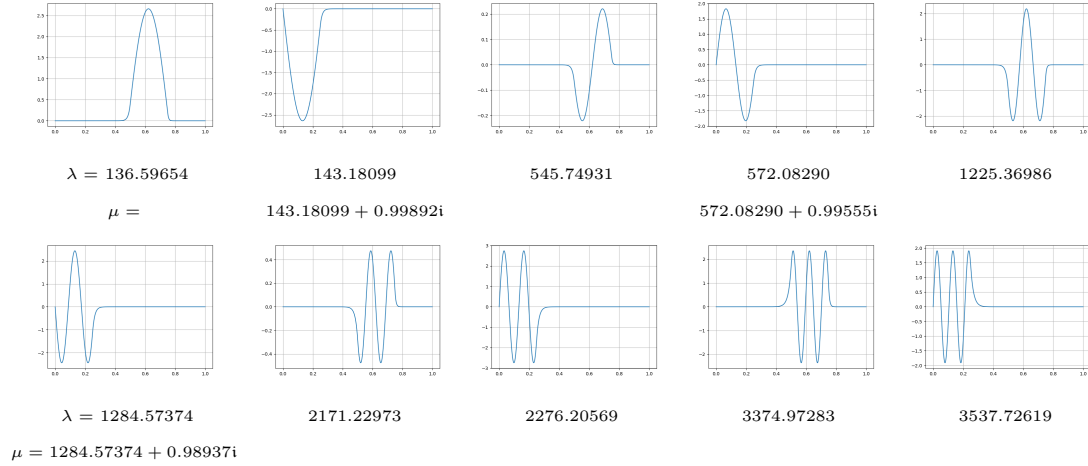
Figure 4.5: Unit interval chopped into 4 equally sliced regions having potential $V = [0, 80^2, 0, 160^2]$. The first ten eigenpairs (λ, ψ) of $\mathcal{L}$ are pictured along with the eigenvalues, μ, of $\mathcal{L}_s$ if they were returned for $s = 1$ and $\delta^* = 0.2$ for $R = [0, 1/4)$ and successfully post processed. This is a strict tolerance, which even the eighth eigenfunction does not *quite* meet. This example illustrates how rigorous the algorithm can be.

Figure 4.6: The real (top) and imaginary (bottom) parts of the eigenvectors of $\mathcal{L}_s$ returned. Notice, as is consistent with 2.1.2, the imaginary parts are small, since the eigenvectors ϕ correspond to are real. These eigenmodes are normalized in the sense of (2.16).

of 30 eigenpairs were found, though only 20 are pictured, where 7 strictly met the criterion $\tau \geq 0.89$ in R. Prior to post processing, the algorithm returned the desired 7 plus 2 additional that are close with $\tau = 0.864$ and 0.853. The real and imaginary parts of these eigenvectors are pictured, plus the potential, landscape, and shift functions in Figure 4.8. Notice as expected that the imaginary parts are two orders of magnitude smaller than the real parts.

Example 4.1.6. A more complicated example of localization that arises from a potential is given in Figure 4.9, which contains the first 30 eigenpairs on a unit square that is chopped into 256 total sub-squares, each having a random, constant potential in $[0, 10, 000]$. There are only four sufficiently localized eigenmodes in R, four that are within 5% of the localization tolerance, and one within 10%– these nine were returned, and their real and imaginary parts are seen in Figure 4.10. Since the size of a domain is inversely related to the size of the eigenvalues, tiny regions containing concentrated eigenfunctions are more likely to bleed than larger regions; as a result, setting strict tolerances for relatively tiny regions can be counterproductive. The eigenmodes that don't quite make the cut in Figure 4.10 are still very concentrated in one small subdomain together with a small surrounding area, and this (slightly) larger region still makes up an insubstantial part of the domain that approximates R well. Thus, it may be wise to either provide a larger tolerance or extend R to include a small buffer around the desired region of interest.

Figure 4.11 provides an example where R is a small, connected region of Ω. The first 30 eigenfunctions of $\mathcal{L}$ are shown in Figure 4.9, however a total of 42 were
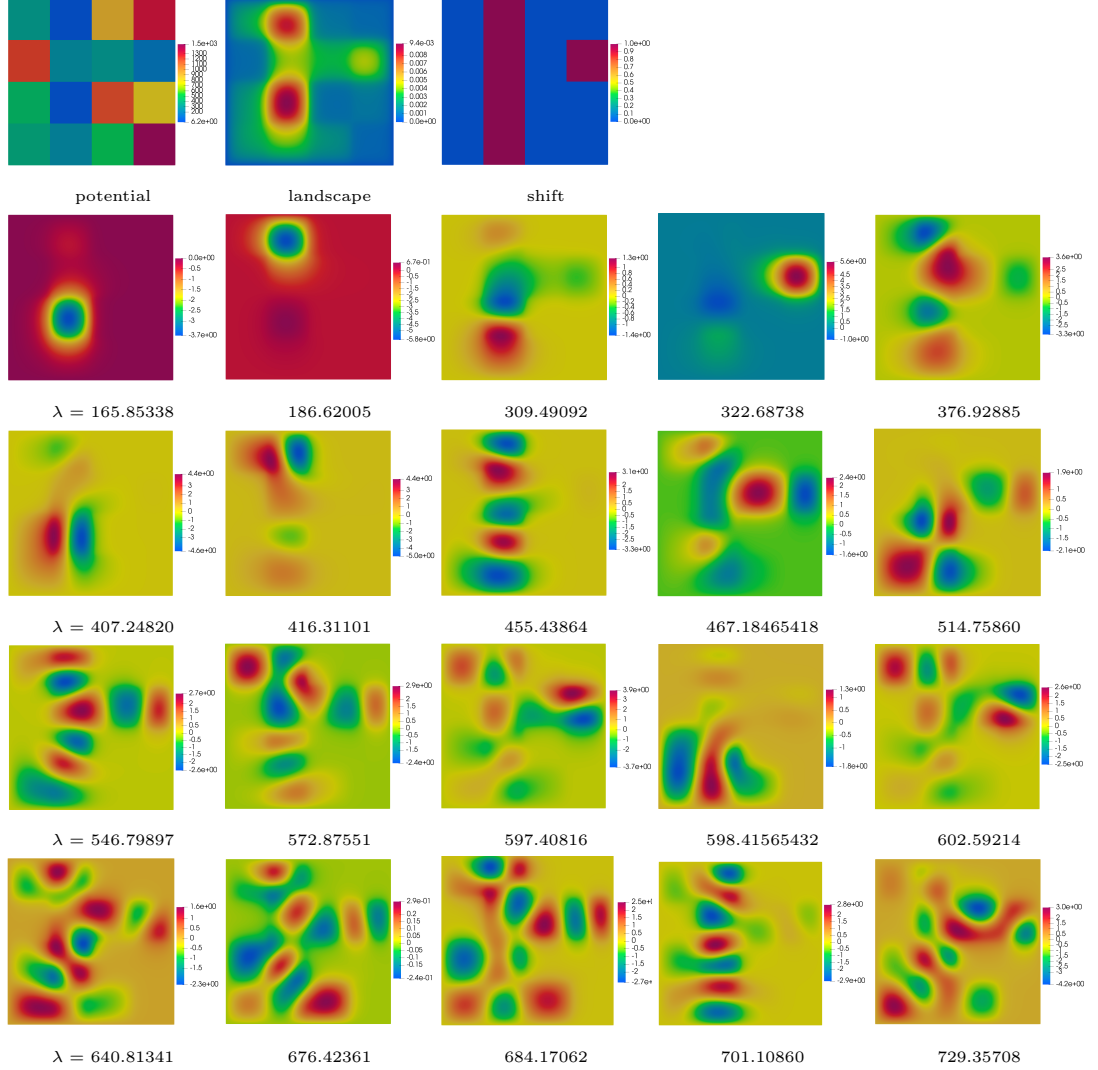
Figure 4.7: Unit square chopped into 4×4 sub-squares each with constant potential generated randomly from 0 - 1500. The top row provides the potential, landscape, and shift function, where the shift function is automatically created from the $n = 5$ most likely regions to contain localized eigenmodes based off the landscape function. The n number of regions are user selected. The remaining rows provide the first 20 eigenpairs of $\mathcal{L}$. The modes of $\mathcal{L}_s$ are in Figure 4.8.

Figure 4.8: Each pair of rows contains the real (top) and imaginary (bottom) part of the returned eigenmodes of $\mathcal{L}_s$ for $s = 1$ and $\delta^* = 0.45$ paired with the eigenvalues. The first 20 eigenpairs of $\mathcal{L}$ are pictured in Figure 4.7 before any post-processing. There are nine eigenpairs of $\mathcal{L}_s$ returned of a total of 30 (not all pictured), where all either meet the tolerance of $\tau = 0.89$ or are very close (2). This experiments illustrates how well the algorithm distinguishes localized eigenmodes before post processing, even, and how well the automatic region selection identifies potential regions of localization.

$\lambda = 1737.18570$ 1796.09081   1999.89455   2122.98736   2200.67329   2213.02292

$\lambda = 2395.31057$ 2590.76880   2717.67666   2729.69019   2732.12664   2761.57635

$\lambda = 2779.55276$ 2781.72949   2813.70955   2905.60294   3097.48696   3118.04829

$\lambda = 3137.94650$ 3156.48816   3181.35500   3242.66122   3330.41920   3343.35926

$\lambda = 3354.39846$ 3390.88154   427.85186   3435.37347   3439.72593   3495.24804
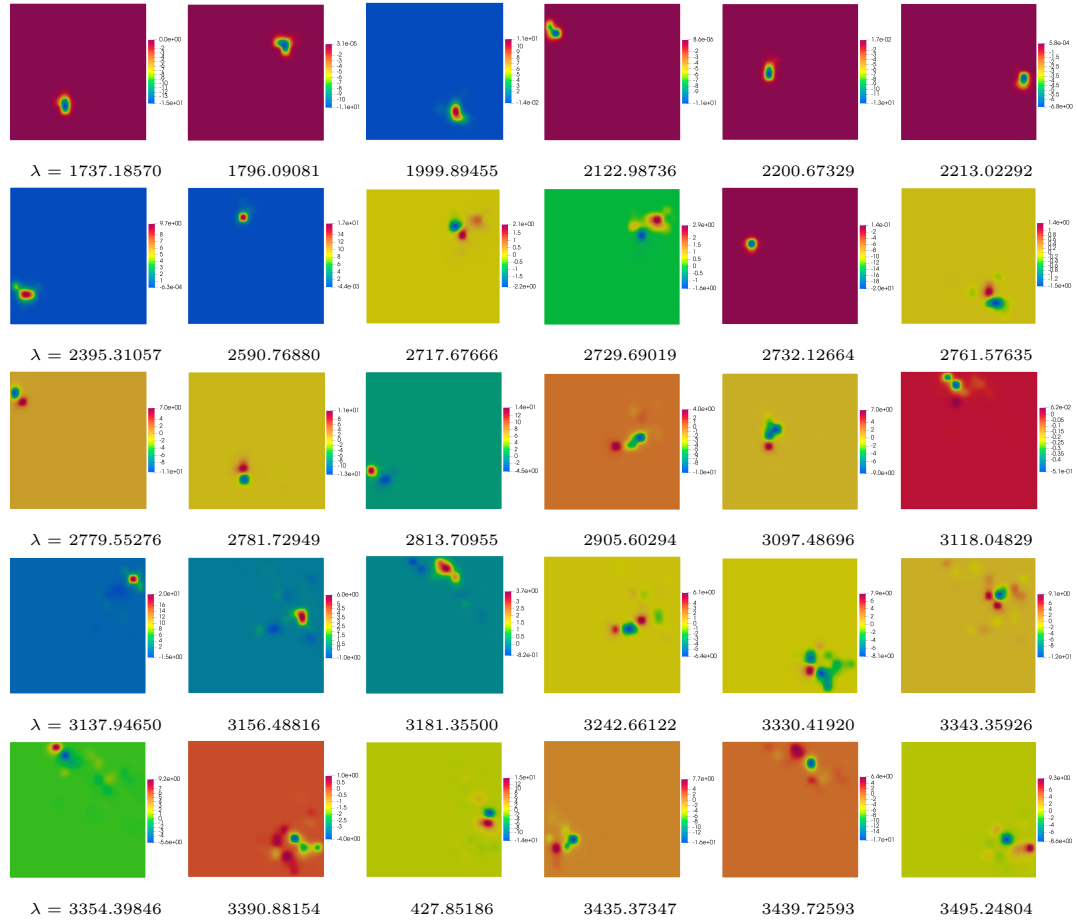
Figure 4.9: The first 30 eigenpairs of a unit square domain chopped into 16×16 sub-squares with constant, random potential in each ranging from 0 to 10,000.
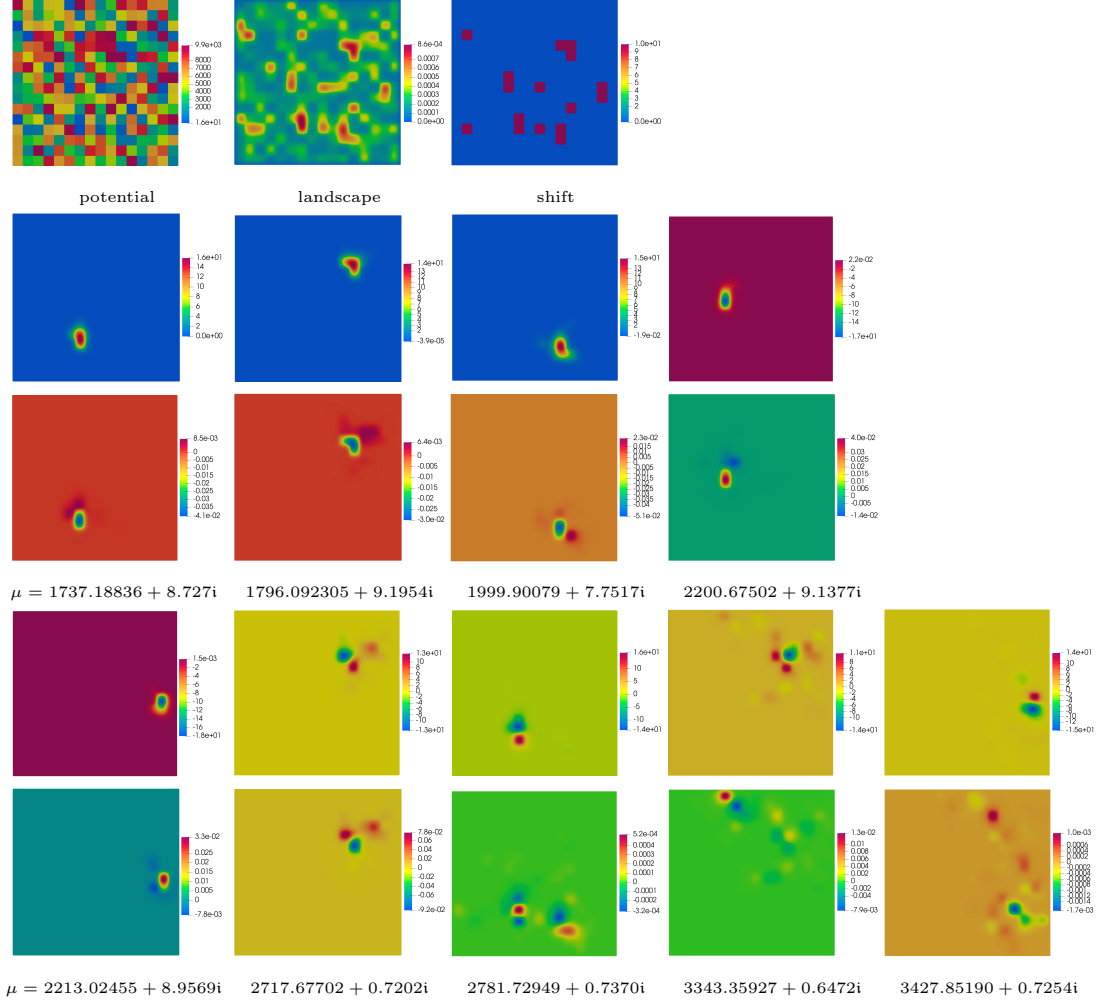
Figure 4.10: The potential function, landscape function, and shift functions are pictured in the top row. Next, each pair of rows pictures the real then imaginary parts of the eigenmodes of $\mathcal{L}_s$ for $s = 1$ and $\delta^* = 0.4$ related to the experiment in Figure 4.9. There are 4 eigenmodes (1, 2, 4, 5) of 30 that meet the tolerance, while all but one of the remaining are within 5%. The 8th is within 10%.

shift $\lambda = 1976.13493 + 0.9051\mathfrak{i}$ $2784.94972 + 0.9231$

Figure 4.11: The first 30 eigenpairs of the original operator are pictured in Figure 4.9, however 42 were computed in total. For the complex shift pictured on the left, $s = 1$ and $\delta^* = 0.4$, there were only two sufficiently localized eigenvectors of $\mathcal{L}_s$. Both were returned.

computed. Out of these 42, only two eigenfunctions met the tolerance of $\delta^* = 0.4$. The algorithm returned both eigenmodes without any false positives and no post processing was needed. This example illustrates how rigorous the algorithm is, even for extremely choppy and huge potentials.

4.2 Post-Processing

The post processing consists of (a few) shifted inverse iterations using the real part of the eigenvalue μ_1 as the initial shift and the real part of the corresponding eigenvector ϕ_1 as the initial test vector, though we note that the landscape function could be used instead. The method of shifted inverse iteration is explained in the "Background Information and Methodology" chapter in the "Power Method and Shifted Power Methods" section. This section provides examples of the post-processing effectively eliminating false positives, even when the experiment is pathological.

Figure 4.12 tests the post processing on the pathological dumbbell case with $s = 0.5$ in the right bulb. As expected, any practical s returns eigenmodes of

$\mathcal{L}_s$ that are localized in the right bulb when they should not be; these real parts and their corresponding (real parts of) eigenvalues μ_1 are pictured below the λ pairs, and notice $\mu_1 \approx \dfrac{\lambda_1 + \lambda_2}{2}$. The next two row duos are the eigenpairs of the three shifted inverse iterations when the initial test vector is either the real part of the eigenvector of $\mathcal{L}_s$ or the landscape function, respectively. When using ϕ_1, the vectors of the inverse iteration indicate there is no localization and are eliminated; in some cases, these eigenvalues differ significantly from the μ_1, which serves as a red flag that the returned ϕ_1 do not represent the ψ well. However, λ_1 and λ_2 are *very* this can be unreliable, while otherwise the inverse iterations are faithful representations of the eigenfunctions of $\mathcal{L}$. Similar can be said when using the landscape function as the initial test function, with the exception of the fourth column. Here the eigenmode has a tiny amount of mass in the left bulb, but it is very faint and could easily be missed. Though the presence of mass in the left bulb when it was not present for the shifted eigenmode should indicate that the ϕ and ψ could be significantly different; this should then lead to further investigation, as with the other examples.

4.3   Magnetic Laplacian

Since the theory presented in this manuscript has been developed via algebraic manipulation and relies only on the assumption that the eigenvalues are real, it applies to other selfadjoint operators with real spectrum. One such example is the magnetic Laplacian, where a couple of examples applying our algorithm to this
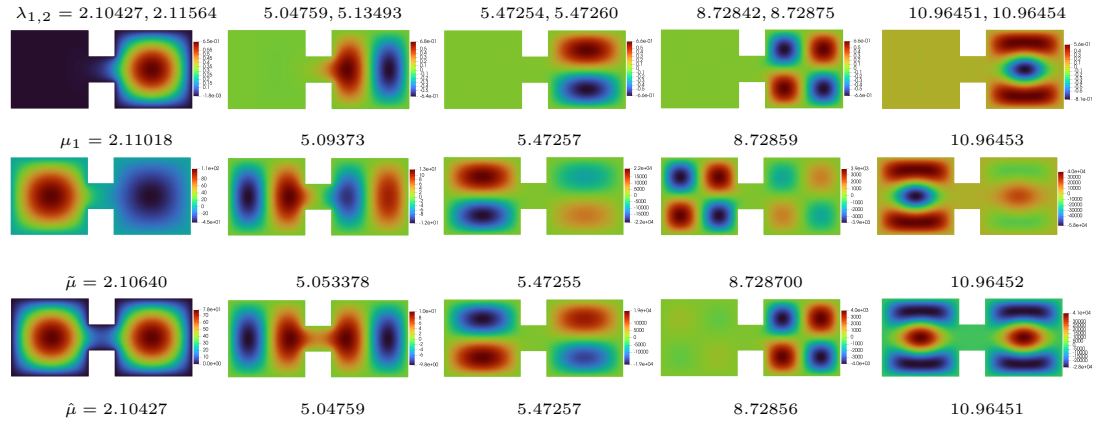
Figure 4.12: The top row of numbers lists the first five pairs of eigenvalues whose eigenvectors of $\mathcal{L}$ are not localized in either bulb, but whose eigenvectors of $\mathcal{L}_s$ are localized in whatever bulb hosts any practical s. The top row of pictures has the real parts of these eigenvectors of $\mathcal{L}_s$, ϕ_1, when the shift is in the right bulb. The next two rows contain the eigenpairs when the initial test vector used for post-processing is either ϕ_1 or the landscape function, respectively. It is not consistent which initial test vector highlights that the eigenmode is not actually localized, but overall the shifted inverse iterations do a good job of identifying that localization doesn't occur in either bulb.

operator are provided in this section. Recall, the magnetic Laplacian operator $\hat{H}$ applied to some differentiable vector field $A : \Omega \to \mathbf{R}^d$ is defined by:

$$\hat{H}(A) := \left(-\mathsf{i}\nabla - A(x)\right)^2 \ .$$

For these examples, the eigenpairs of the original operator are pictured in Figure 8.3 outlined in Example 8.2.2. In summary, the magnetic field $A = -a\left(\cos(f(x,y)), \sin(f(x,y))\right)$ where $f(x,y) = \pi\sin(\pi x)\cos(\pi y)$ and $a = 50$. The domain is the square $\Omega = [-1,1] \times [-1,1]$. Figure 4.13 gives the results when the R is the middle $\frac{1}{3} \times \frac{1}{3}$ sub-square of the domain, $s = 1$, and $\delta^* = 0.4$. There are nine eigenfunctions in total, where two are not at all localized in R and two are centered in R, but have too much mass outside of R. The algorithm returned all five desired eigenvectors.

Another, more selective example with the same domain and operator is provided in Figure 4.14. Here $s = 1$, $\delta^* = 0.3$, and R is the top right corner. There was one eigenmode, of nine total, that met the requirement; this one was returned, exemplifying that the algorithm is robust, even for other operators. The mesh was adaptively refined to the problem $-\Delta u = 1$ with zero boundary conditions, using to the NGSolve adaptive refinement procedure, which utilizes the Zienkiewicz-Zhu type error estimator, to a tolerance of 400,000.

Figure 4.13: For $s = 1$, and $\delta^* = 0.4$, all eigenpairs meeting the criteria were found; four eigenpairs were filtered. See Figure 8.3 for the 9 eigenpairs of the original operator. The domain is uniformly meshed with piecewise, cubic polynomials and with edge length $h = 0.01$
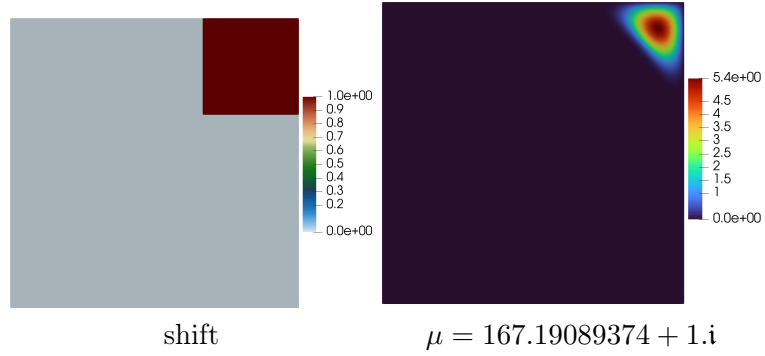
Figure 4.14: For $s = 1$, and $\delta^* = 0.3$, the one eigenpair, of nine, that met the requirement was returned. This example illustrates the robustness of the algorithm even for other operators and for computationally difficult problems. See Figure 8.3 for the 9 eigenpairs of the original operator.

5   The size of s and δ^*

Naturally, the question arises about what the correct size is for the complex shift s. The crux of our theory, Theorem 2.1.2 and Theorem 2.1.4, indicate that there must be a balancing act between s and the chosen δ^*. In order for the eigenvector results in Theorem 2.1.4 to be meaningful, it should be that $s\delta^* \leq 1$. Moreover, δ^* has a limited practical range for what it means to be localized. So the size of s should only be considered with respect to the practical range of δ^*. When $\delta^* = 0.1$, only eigenvectors that have $\tau \geq 0.9949$ can meet the requirement, which is already a *very* strict tolerance. Though δ^* can theoretically approach 0, values smaller than 0.1 are likely unreasonable. This immediately provides a reasonable upper limit for s as 10 to match the lower bound of $\delta^* = 0.1$.

The (practical) upper limit for δ^* is more up to the discretion of the user than its lower limit. Figure 5.1 provides a few 1D examples of eigenmodes that are localized for $\tau \in (0.839, 0.89)$, where the upper bound corresponds to $\delta^* \approx 0.45$. This figure exhibits how even τ values as high as 0.89 have a nontrivial amount of eigenvector mass outside of R; one may not so readily call these eigenfunctions localized. In a similar fashion, Figure 5.2 provides a few 2D examples of eigenmodes where $\tau \in (0.89, 0.943)$. It may be a stretch to claim these eigenmodes are sufficiently concentrated. One significant consideration when deciding the appropriate size of δ^* is how the size of R compares to the size of the domain Ω. In the cases that were

Figure 5.1: 1D Eigenmodes pictured with their corresponding τ values. In each case $\tau > 0.8$. The eigenfunction mass outside R is nontrivial, which may be too much to be considered localized.

Figure 5.2: 2D Eigenmodes along with their corresponding τ values. The amount of the eigenvector mass outside of R is nontrivial, even though τ seems high

.

just presented, R is not dramatically smaller than Ω. As seen by the definition of τ, when R is smaller an eigenvector must be even more concentrated to boost τ. An extreme example is given in Figure 4.10 where R is as little as $1/256$ times the area of Ω; even the eigenmodes that are not quite concentrated to the $\delta^* \leq 0.4$ are very clearly localized to a domain that approximates R well. In these cases, it is (probably) better to increase R a small amount in the surrounding area than to increase the tolerance since the prior ensures the eigenfunction is still concentrated in the same general region, while the latter allows it to occupy *anywhere* as long as it remains within the tolerance. The last example is fairly extreme, while the other two are more typical use cases. With this line of reasoning, we consider the practical range for the tolerance to be $\delta^* \in [0.1, 0.45]$.

An upper limit for δ^* does not necessitate a corresponding lower limit for s because Theorem 2.1.4 could have $s \to 0$ for whatever δ^* to meet $s\delta^* \leq 1$. Since a localized eigenpair of $\mathcal{L}_s$ has $\Im\mu \gtrsim s$ and (in this sense) is distinct from other, non-localized eigenfunctions, there must be a practical lower limit on s that keeps this distinction clear, especially when the computations are subject to errors in floating point arithmetic. For example, it is obvious that $\Im\mu = 0.99$ and $s = 1$ are close, but $\Im\mu = 0.00001$ and $s = 0.00002$ are not close, even though they are closer. Here is where the theoretical and practical size limitations on s can differ, and this disparity is exacerbated by computational limitations. Plus, note that the operators $\mathcal{L}$ and $\mathcal{L}_s$ presented in the theory are no longer the same operators once they are discretized. Our theory provides no bounds on s sufficiently small, but practical bounds on s do exist. We will focus this section on what size of s the user should set when implementing the software, for which there are several considerations.

First, define the aspect ratio, which affects the quality of the rational filter $f(z)$ by:

$$\text{Aspect Ratio} = \frac{(b - a) + s\delta^*}{s\delta^*}. \tag{5.1}$$

Recall from section 3.2 that, for the Bunimovich contour U, the contrast between $\mu \in \operatorname{Spec}(\mathcal{L}_s) \setminus U$ and $\mu \in U$ determines the number of FEAST iterations required. If the contrast is greater, which happens when the aspect ratio is smaller, fewer iterations are needed. Making s too small without compensating with an

equally smaller search interval $[a, b]$ can blow up/destroy the filter, thereby slowing convergence and affecting accuracy. Examples of this are given throughout this section.

Observe that for the Bunimovich contour pictured in Figure 2.1, increasing s or δ^* extends the search interval by $s\delta^*$. A user looking for localized eigenmodes with the corresponding $\mu_1 \in [a, b]$, has increased chance of finding eigenmodes having μ_1 outside of this region when s is larger. This will always be an unintended consequence of increasing s. However, some extension is necessary because the eigenvectors of $\mathcal{L}_s$ have eigenvalues that converge down to those of $\mathcal{L}$, albeit Theorem 2.1.2 indicates this should be small. It would be unreasonable to consider eigenvectors of $\mathcal{L}_s$ with μ_1 differing from λ by integer values as the same. Even μ_1 that differ from λ on the order of tenths should not be considered corresponding eigenvalues! This means that the length of the extension, given by $s\delta^*$, is already big enough when it is on the order of tenths. So given $\delta^* \in [0.1, 0.4]$, perhaps s does not need to be bigger than 1.

Figure 5.3 plots $|\mu_1 - \lambda|$ versus $|s - \mu_2|$ for the returned eigenpairs of $\mathcal{L}_s$ when $\delta^* = 0.25$ and $s = 1$ for R as the left, middle, or right bulb of the Threebulb domain. Of 64 possible eigenvectors, 36 are returned, of which 27 are strictly localized, 4 are *very* close to meeting the tolerance, and 5 do not meet the requirement; in Figure 5.3, these are plotted as grey scale, orange, and red dots, respectively. Notice the clear distinction between the localized and non-localized vectors. The red dots are clearly separate, having either $|\mu_1 - \lambda| > 0.1$ and/or $|s - \mu_2| > 0.1$. The orange dots that are *almost* localized enough, though closer in, follow the same

(a) All 36 returned eigenmodes.

(b) Zoomed in on the clustered 27 of the 36 eigenmodes from the left Figure. The only orange dot represents an eigenmode with $\tau = 0.946$ in the right bulb
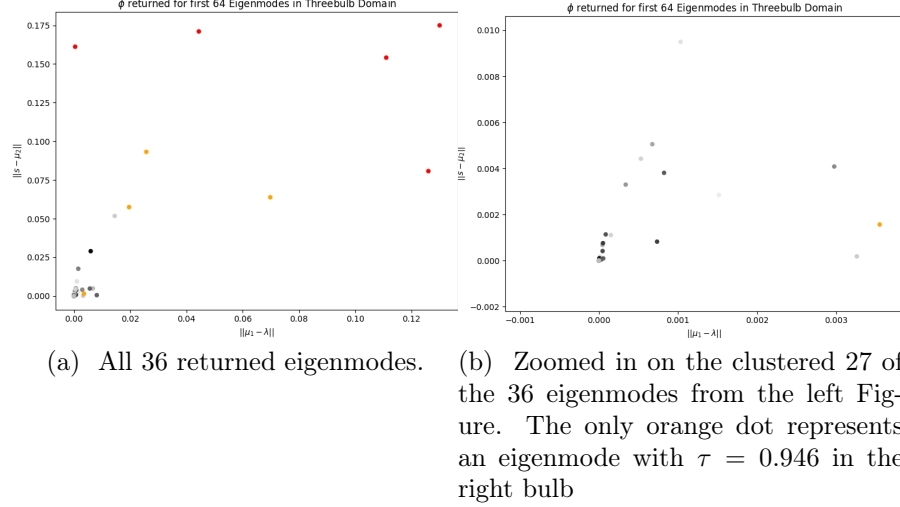
Figure 5.3: The $|\lambda - \mu_1|$ against $|s - \mu_2|$ of the 36 returned eigenmodes when $\delta^* = 0.25$, $s = 1$ and R is the left, middle, and right bulbs of the Threebulb domain. Red modes do not meet the tolerance, while orange ones *almost* do. The false positive in the right figure comes from a linear combination of the 37th and 38th mode that have close eigenvalues. These results are consistent with Theorem 2.1.2, where $\Re\mu$ is close to the corresponding eigenvalue λ and $\Im\mu$ is close to the complex shift s

trend with the exception of one false positive. On the right, Figure 5.3 provides a zoomed in view where 27 of the 36 points are located, including this false positive, which appears to be a linear combination of the 37th and 38th eigenmodes, which is included in the Appendix. We note that we do have robust post processing to "weed out" such false positives, which was detailed in section 4.2. In general, we see that the greyscale dots representing localized eigenmodes are (almost completely) distinct, and in all cases are such that $|\mu_1 - \lambda| < 0.02$ and $|s - \mu_2| \lesssim 0.052$.

For contrast, Figure 5.4 plots the $|\lambda - \mu|$ versus $|s - \mu_2|$ for the first 64 eigenmodes of the Threebulb domain, except with $s = 0.1$, for $\delta^* = 0.25$ and R as the left,

middle, or right bulb. There were 36 returned eigenmodes, with results similar to the case when $s = 1$. Since s was ten times smaller, the search interval $[a, b]$ would have to also be ten times smaller to maintain the aspect ratio of 41; thus there were ten times as many quadrature points computed, (at least) ten times as many iterations of FEAST, and therefore more time and computational expense. Even extending to $b - a = 2$ was too much, and sufficiently localized eigenmodes were sometimes not returned. For example, for $[a, b] = [8, 10]$, only one of three sufficiently localized eigenmodes are returned, which occurs due to the poor aspect ratio destroying the quality of the filter. Another example on the same domain can be seen in Figure 5.2 where $\delta^* = 0.2$, $[a, b] = [24, 30]$, and when $s \leq 0.1$; there are eigenmodes returned, but none of them are remotely localized! Again, partitioning the search region into six or more subintervals corrects this, but at six times the cost and with no gain from the $s = 1$ case.

To highlight the distinction between real and imaginary parts of localized eigenmodes versus non localized eigenmodes, the averages for the localized eigenfunctions are $|\lambda - \mu| \approx 10^{-5}$ and $|s - \mu_2| \approx 10^{-4}$ versus non-localized eigenvectors with $|\lambda - \mu| \approx 10^{-3}$ and $|s - \mu_2| \approx 10^{-2}$. In all cases, localized eigenmodes are such that $|\mu_1 - \lambda| < 5.42\text{e-}4$ and $|s - \mu_2| \lessapprox 5.8\text{e-}3$ versus nonlocalized eigenmodes having $|\mu_1 - \lambda| > 3.7\text{e-}3$ and $|s - \mu_2| > 1.07\text{e-}3$.

As just exemplified, decreasing s significantly requires smaller search regions in order to maintain an acceptable aspect ratio, which may otherwise negatively affect the algorithm. Equivalently, increasing s may require extending $[a, b]$, but too far can increase the computational cost and time needed to achieve the desired

(a) All 37 returned eigenmodes.

(b) Zoomed in on the clustered 25 of the 37 eigenmodes from the left Figure.
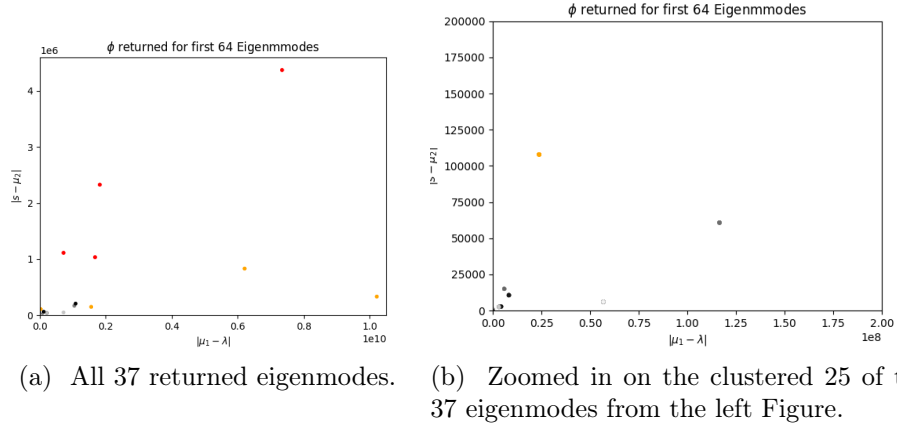
Figure 5.4: The $|\lambda - \mu_1|$ against $|s - \mu_2|$ of the 37 returned eigenmodes when $\delta^* = 0.25$, $s = 1$ and R is the left, middle, and right bulbs of the Threebulb domain. The 5 red modes do not meet the tolerance, while the 5 orange ones *almost* do. These results are consistent with Theorem 2.1.2 that proposes that μ for localized eigenvectors will be close to s.

task. As expected, there is a direct relationship between the number of eigenvectors found outside of the desired $[a, b]$ and the size of $[a - s\delta^*, b + s\delta^*]$; consequently, the "hunting" phase of the algorithm, described by Algorithm 2.1, will consume more memory and time for bigger $s\delta^*$. The same trend follows in the post-processing phase described by Algorithm 2.2 since all of the additional eigenvectors of $\mathcal{L}_s$ that don't correspond to eigenvectors of $\mathcal{L}$ within $[a, b]$ must be post-processed and discarded. On the contrary, setting the $s\delta^*$ too small slows the algorithm as well since either the size of the search region must comparatively decrease else the algorithm may fail to identify localized eigenmodes. Table 5.1 illustrates this for low to high energy cases with varying degrees of complexity in the localization behavior. Highlighted with red text are cases where there were insufficiently many

eigenmodes returned or they were insufficiently concentrated; the "worst perfor-mance" for each category is also red, such as the slowest wall-clock time or biggest memory consumption. The entries in each column reflect the difference relative to the "best performance" case. In summary, too small and too large s both could fail to return the correct set of eigenmodes, given the same search region, while also increasing computational time. Relatively "safe" values for s in all cases ranged from $[0.5, 1]$ given reasonable δ^*. We note that since the search region was constant during these experiments, the expanding aspect ratio as s decreases is the cause for the sub par results. Observe that $b - a = 1$ for the experiments presented in Table 5.1, and for $s \leq 1$ the algorithm returned the correct number of eigenvec-tors, but for larger search regions given in Tables 5.2 and 5.3, it did not. These experiments also serve to illustrate how a poor aspect ratio affects the algorithm and provide simulation run times and memory consumption data.

Moreover, s can be so large that it induces "false positives", which are localized eigenmodes of $\mathcal{L}_s$ that do not correspond to localized eigenmodes of $\mathcal{L}$. Exam-ple 5.0.1 demonstrates a case of this on the unit interval.

Example 5.0.1. Let $\mathcal{L} = -\frac{d^2}{dx^2}$ on $(0, 1)$ with zero boundary conditions, which has eigenpairs given by $(\lambda_n, \psi_n) = ((n\pi)^2, \sin(n\pi x))$. Take $R = (0, 1/4)$ and using

Three bulb where R = left bulb for search region $[4.5, 5.5]$ and $\delta^* = 0.25$				
s	# of Eigenvectors Returned	Peak Memory (GB)	Wall-Clock Time (s)	CPU Time (m:s:ms)
Base	1	7.03368	9.727	2:33:59
0.01	1	+0.096984	+0.252	0:00:48
0.1	1	0	0	0
0.5	1	+0.096204	0.006	0:05:82
1	1	+0.096984	0.019	0:07:49
5	3	+0.096980	1.33	0:22:59
10	5	+0.96988	5.05	1:19:21
100	28	+0.114004	492.71	132:49:53

Table 5.1: Varying sizes of s for small search region of length 1. Considerable time, without benefit, is added when shifting s by greater than one since the contour is significantly enlarged.

that for integer n:

$$\frac{1}{4} - \frac{1}{2\pi} \le \tau^2(\psi_n, R) = \frac{\pi n - 2\sin(\frac{\pi n}{2})}{4\pi n - 2\sin(2\pi n)} \le \frac{1}{4}\left(\frac{1 - 2(\frac{\pi n}{x})}{\pi n}\right) \le \frac{1}{4}\left(1 + \frac{2}{3\pi}\right) ,$$

thereby squeezing $\tau(\psi_n, R)^2$ to $\frac{1}{4}$ (the limit as $n \to \infty$ is $\frac{1}{4}$, even). This shows there are no localized eigenvectors in R (there are highly localized eigenvectors in the complement, however). Now considering $\mathcal{L}_s\phi = \mu\phi$ for $s = 10^4$, we obtain $\mu = 149.02494 + 9991.7736i$ and $\delta(\phi, R) = 0.02868$, so ϕ is highly localized in R.

Because $\mu_2 = s\tau^2$, an eigenvector that is sufficiently concentrated in R for chosen $\delta^* \in (0, 0.45]$ must have $\mu_2 \in (0.89303s, s)$. To the point of Example 5.0.1, choosing s too big can shift some μ_2 close enough to s so that it isn't filtered during the FEAST iterations. Figure 5.5 illustrates this; the real parts of the eigenvalues

Three bulb where R = left bulb for search region $[24, 30]$, $\delta = 0.2$				
s	# of Eigenvectors Returned	Peak Memory (GB)	Wall-Clock Time (s)	CPU Time (m:s:ms)
Base	3	7.331492	19.40	4:57:41
0.01	2**	+0.00018	87.48	23:13:26
0.1	1*	0	0	0
0.5	4	+0.00011	6.68	1:10:78
1	4	+0.00018	5.36	1:23:75
5	7	+0.01395	72.87	19:33:09
10	13	+0.04006	228.77	60:45:88
100	51	+0.219464	1014.60	208:46:24

Table 5.2: Varying sizes of s for small search region of length 6. There are 3 eigenmodes localized strictly within the tolerance, as well as a 4th eigenmode which is *very* close within 0.4%. As before, s too large induces false positives, while also finding acceptable localized eigenpairs that are outside the desired search region, essentially just adding to the needed computation time. However, s can also be too small for the given search region, which can lead to not finding all or any of the desired eigenpairs. When $s = 0.1$ or $s = 0.01$, no localized eigenmodes were found; not even the eigenmodes that were returned were localized.

Unit Square where $R = 16$ largest for search region $[1800, 2300]$, $\delta = 0.35$				
s	# of Eigenvectors Returned	Peak Memory (GB)	Wall-Clock Time (s)	CPU Time (m:s:ms)
Base	2	6.972268	14.40	3:07:13
0.01	0	+0.019972	122.412	33:08:33
0.1	2	0	9.894	3:26:72
1	2	0	11:459	3:47:65
2.857	3	+0.00038	7.719	2:47:36
10	3	0	13.372	4:16:59
100	5	0.00250	0	0

Table 5.3: A very high energy case where localization comes from a complicated 16×16 grid potential of random values in $[0, 10,000]$ The 16 cells with the largest L_2 norm, computed from the landscape function, were chosen for R. Since the eigenmodes are far apart, larger s is less likely to extend the contour to the point of finding excess eigenmodes outside the desired search region, however larger s still returns more eigenmodes outside of δ^*. $s = 0.01$ failed to converge.

are plotted along the x-axis, and they are mapped against their imaginary parts when $s \neq 0$ or against their τ^2 value when $s = 0$ (upper left). The behavior of the eigenvalues when $s = 0$ is the same as when $s = 1$ or 100, in the sense that eigenfunctions have s values proportional to their τ^2 value. This changed when $s = 10,000$, where the behavior of one eigenvalue, pointed out by the arrow, had its imaginary part pushed up near to s, creating a false positive.

One might simply take away from Example 5.0.1 and Figure 5.5 that $s = 10^4$ is the cut-off for "too large" with respect to inducing false positives, but this phenomenon is ultimately problem dependent. Even choosing s small does not ward off all possibilities of obtaining false positives. When there are eigenvectors whose eigenvalues are very close together, but distinct, a tiny perturbation can cause the
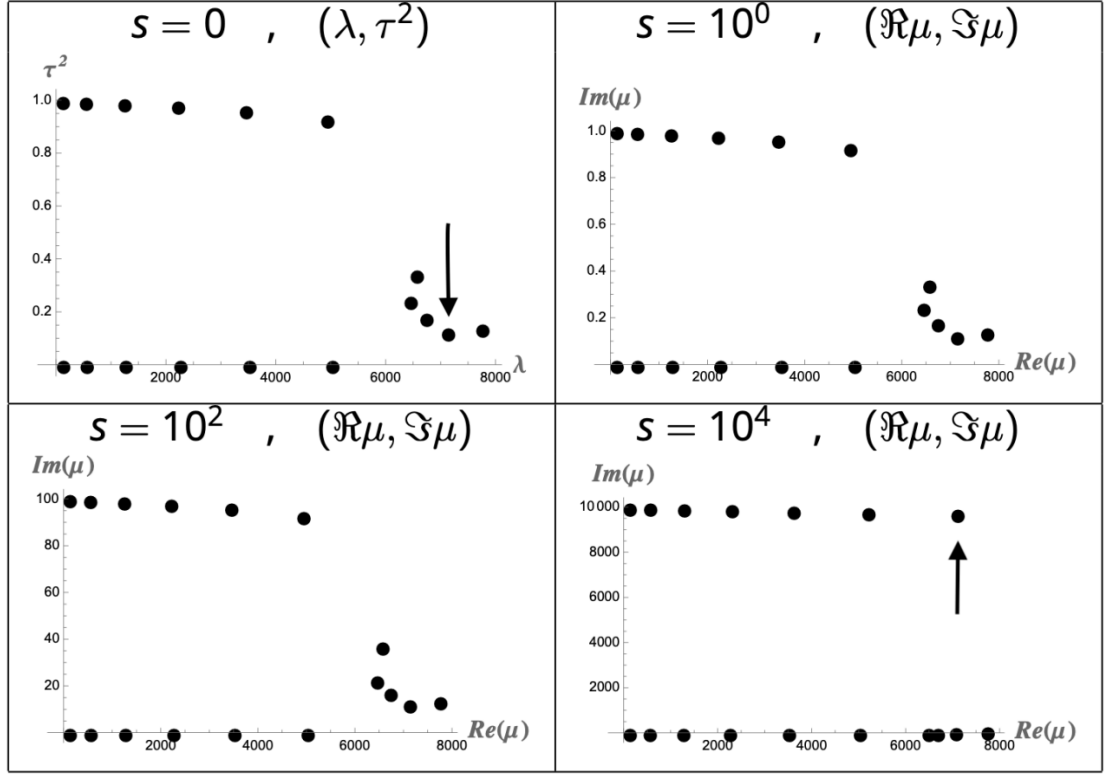
Figure 5.5: The top left plots the eigenvalues of $\mathcal{L}$ against their τ^2 value. The remaining three plots correspond to eigenvectors of $\mathcal{L}_s$ where the real parts of the eigenvalues are plotted against their imaginary parts. As indicated by the arrow, $s = 10^4$ is so large that it induces a false positive.

eigenmodes to mix; the resulting eigenvectors of $\mathcal{L}_s$ are linear combinations of these eigenvectors of $\mathcal{L}$ that may be falsely localised in one subdomain, which commonly occurs when the eigenvectors of $\mathcal{L}$ differ only by a sign in one region, but not the other. This can occur in symmetric domains with symmetric potential. Note that these are exceptional cases, and will be referred to as "pathological" cases. 1D and 2D pathological examples are given in Examples 5.0.3 and 5.0.2, respectively. Notice even a tiny shift such as $s = 10^{-5}$ is big enough to induce localization falsely in either region.

Example 5.0.2. Consider each of the four examples seen in Figure 5.6. The top row displays the two eigenpairs (λ_n, ψ_n) of $\mathcal{L} = -\Delta$ with Dirichlet boundary conditions, and the bottom row shows the eigenpair (μ, ψ) of $\mathcal{L}_s$ for $s = 0.1$ in the right bulb. There is no localization in either the left or right bulb of the dumbbell domain for the original operator $\mathcal{L}$. But the eigenfunctions of the shifted operator are linear combinations of the pair of ϕ with destructive interference in one bulb and constructive in the other that results in localized eigenvectors. Notice that even the μ are closer to being averages of the pair of λ than to either λ individually, which helps to illustrate that these ψ are just a lucky mixture of the pair of ϕ.

Example 5.0.3. A simple, one dimensional pathological case occurs on the unit interval when applying a symmetric, large potential about the center. These are Sturm-Liuoville problems where it can be shown that all eigenvalues are simple and the corresponding eigenvectors must have even or odd symmetry, which means that no eigenvector can be localized in the left or right subdomain. Even though the eigenvalues are distinct, pairs of them can be *very* close, becoming increasingly

(a) $\lambda = 5.47254$ (b) $\lambda = 5.47260$ (c) $\lambda = 8.72842$ (d) $\lambda = 8.72875$

(e) $\mu = 5.47257 + 0.009999i$ (f) $\mu = 8.72859 + 0.009994i$

(g) $\lambda = 10.96451$ (h) $\lambda = 10.96454$ (i) $\lambda = 21.931574$ (j) $\lambda = 21.931576$

(k) $\mu = 10.964527 + 0.009999i$ (l) $\mu = 21.931575 + 0.009999i$
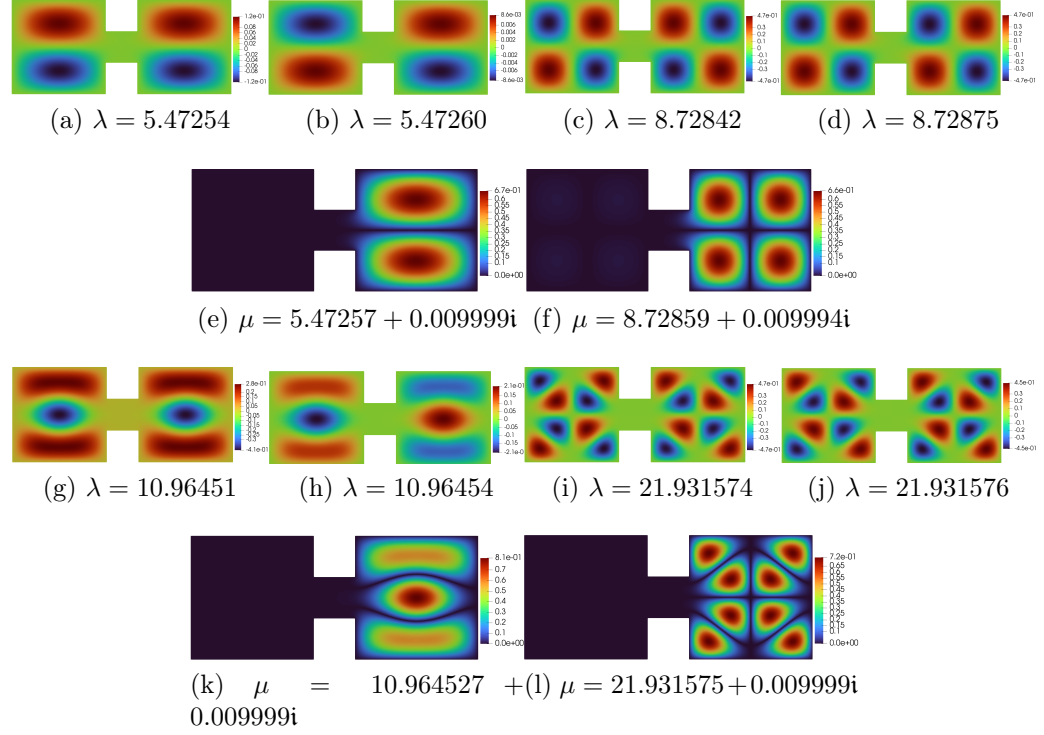
Figure 5.6: Four cases of false positives for the dumbbell domain when $s = 0.1$, $\delta = 0.25$, and R is the right bulb. Each block of three eigenmodes displays the eigenfunctions of $\mathcal{L}$, paired with their eigenvalues λ in the top row, and the corresponding (modulus of the) false positive paired with its eigenvalue μ. Each pair of λ_1 and λ_2 are close to each other while their respective eigenvectors ψ_1 and ψ_2 are identical up to a change in sign. When applying the small pertubation, $s = 0.1$, the resulting eigenfunction, ϕ of $\mathcal{L}_s$ (bottom row for each case), is a linear combination of ψ_1 and ψ_2. Consequently, half of ϕ is canceled out causing ϕ it to appear strongly localized in R. Notice even the eigenvalues μ of ϕ are close to an average of the respective λ_1 and λ_2.

close with increasing V. In fact, the eigenvalues are so close, that double precision is (easily) insufficient to capture their distinction. Consequently, the resulting eigenmodes are a linear combination of the pairs of eigenmodes with even and odd symmetry, and therefore they are falsely localized in the right or left subdomain.

Figure 5.7 contains the first 8 even/odd pairs of eigenpairs when $V = [0, 120^2, 120^2, 0]$ for a quartered, unit interval domain. Observe that no eigenmode is localized only in the left or right regions and the eigenvalues (mostly) match well beyond machine (double) precision, so typical computational tools using double precision are insufficient. In fact computing even these eigenfunctions, without any shift, is a difficult process that cannot be done via double precision! In order to obtain them the eigenmodes had to each be constructed, taking advantage of the even or odd symmetry, in Mathematica [39] and the eigenvalues were then computed to 200 digits of accuracy. Details of this process are in the Appendix section 8.2. Figure 5.8 exemplifies the sum of the even and odd eigenvectors, which are highly localized in the left region only. These are the kinds of eigenpairs returned when using standard eigenvalue solvers in double precision, even when $s = 0$. It should come as no surprise then, that essentially any shift induces these same false positives.

Example 5.0.4. In the case of the magnetic Laplacian $H(A) = (\hat{H} - V)(A)$, where $A = (-25 \times y, 25 \times x)$ on the unit disk centered at the origin, there are many tightly clustered eigenvalues. In some cases, the 16 eigenpairs of $H(A)$ in $[50, 51.3)$ have eigenvalues that agree to nine digits. Similar to Example 5.0.3, the eigenvalues are so close that only approximations of linear combinations of eigenvectors were

$\lambda_e = 147.8736094365089$ 591.2940076291312   1329.631565743891   2361.73184891566
82800336331012   82473892827510   99529994287306   773308506737495

$\lambda_o = 147.8736094365$ 591.29400762913   1329.63156574389199   2361.73184891566
08982800336331464   1282473892832050   529994292183   773308506820587

$|\lambda_e - \lambda_o| = 4.5266 \times 10^{-25}$ 4.5400×10^{-24} 4.87769×10^{-23} 8.3092×10^{-22}
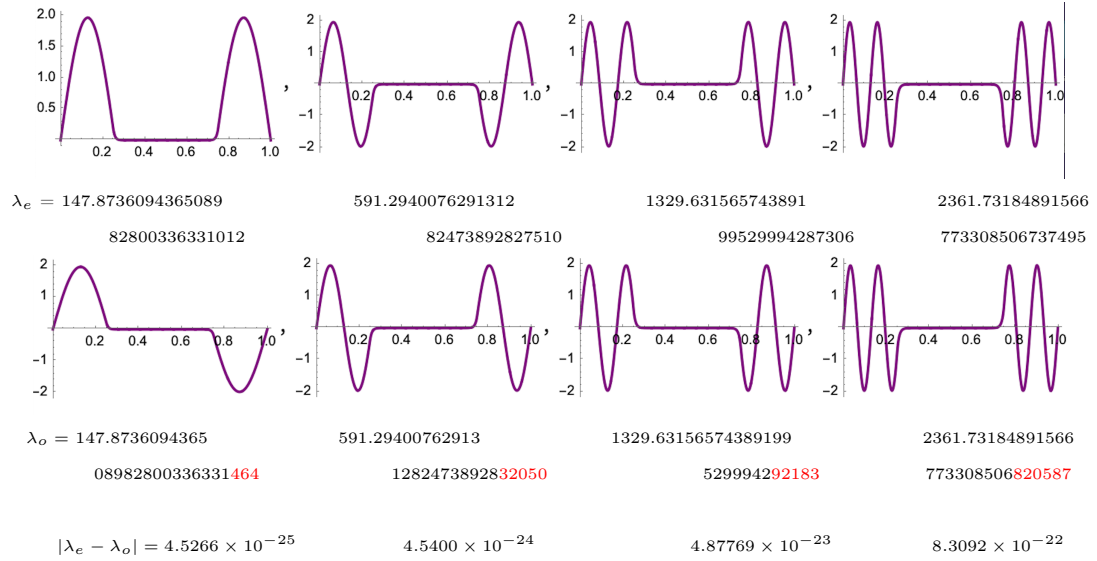
The first four pairs of eigenvectors together with their eigenvalues and the absolute difference between the corre-sponding even and odd eigenvalues.

$\lambda_e = 3685.7129532389$ 5298.55176702207   7195.169253765091   9366.0261786898989
5562789792592406   489122519380625   78002520731729   5929219084365

$\lambda_o = 3685.71295323895562$ 5298.551767022074   7195.1692537650   9366.026178689
789795314781   89122726627203   9178050483855788   899500728312005601

$|\lambda_e - \lambda_o| = 2.72237 \times 10^{-20}$ 2.07246×10^{-18} 4.7963×10^{-16} 5.4799×10^{-13}
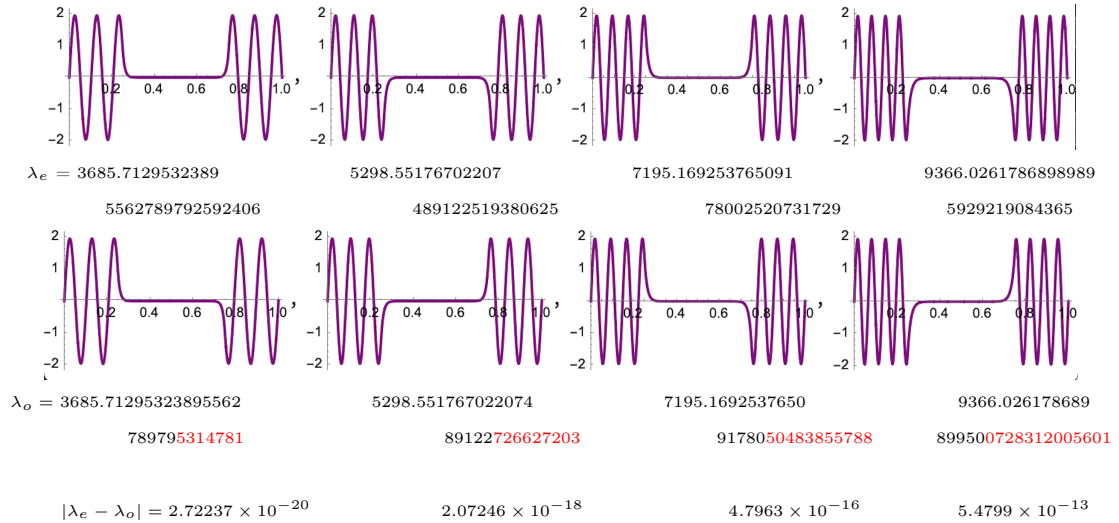
The fifth through eighth pairs of eigenvectors together with their eigenvalues and the absolute difference between the corresponding even and odd eigenvalues.

Figure 5.7: For the Sturm-Liouville problem $\mathcal{L}\psi := (\Delta + 120^2)\psi = \lambda\psi$ on $(0,1)$ with $\psi(0) = \psi(1) = 0$, the eigenvalues are simple and the eigenmodes are pairs of even and odd functions. The top row pictures the even function and its eigenvalue λ_e, while the second row pictures the odd functions and their eigenvalues λ_o; where the odd and even eigenvalue disagree is highlighted in red. These eigenvalues are very close, (mostly) well beyond double precision accuracy. The bottom provides the actual absolute difference between the even and odd eigenvalue pairs.

Figure 5.8: The sum of the even and odd eigenvector pairs are localized in only the left region. These are the kinds of linear combinations that are returned when using double precision, even for $s = 0$. Thus, any shift s is too large for such pathological cases.

returned. These (modulus of the) eigenpairs are given in Figure 5.9 along with the (modulus of the) eigenpairs of the shifted operator (in Row 4), $H(A)_s$ for $s = 1 \times 10^{-6}$, $\delta^* = 0.3$, and R is a circle of radius 0.35 centered at the origin. Notice that the eigenvectors of $H(A)_s$ are circular symmetry about the origin, as they should be. This indicates s is small enough to identify the sufficiently localized (actual) eigenvectors. This case serves as a $2D$ example to illustrate that having s smaller than half the gap between the desired eigenvalue and its nearest neighbor is unnecessary, even for such an exceptional case.

As we have pointed out, the size of s with respect to generating false positives is problem dependent. Table 5.4 tests for what s will the first 9 pairs of eigenpairs, (λ_a, ψ_a) and (λ_b, ψ_b) given by the rows, of the dumbbell domain return a false positive. The leftmost column gives the magnitude of agreement for each pair of eigenvalues, $|\lambda_a - \lambda_b|$; for example, the first two eigenvalues are 2.1042 and

$\lambda = 50.00000003$ 50.0000008   50.00000949   50.0000752

$\lambda = 50.00044565$ 50.00210558   50.00825585   50.02761343

$\lambda = 50.08038343$ 50.2068502   50.47665144   50.99497555

shift

$\mu = 50.00000003$ 50.0000008
$+i$ 9.817e-07 $+i$ 9.084e-07
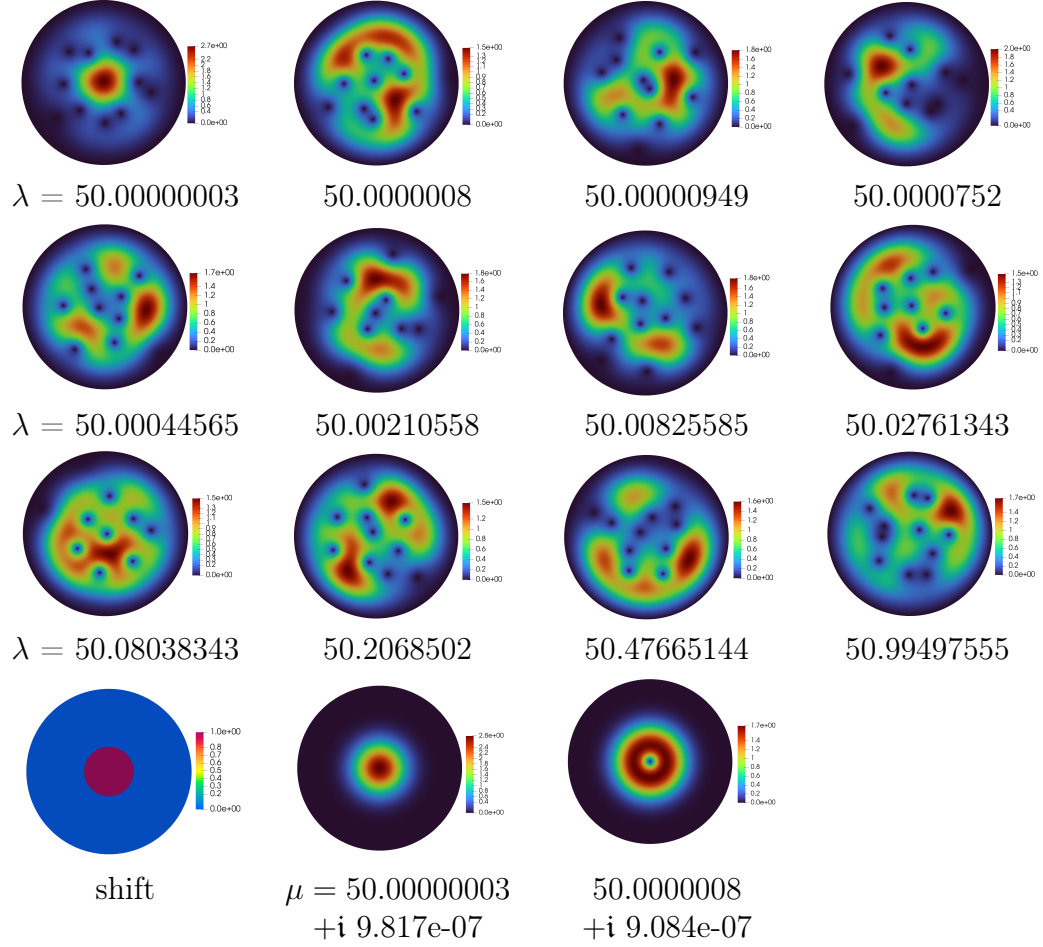
Figure 5.9: The (modulus of) eigenpairs of the magnetic Laplacian, $H(A) = (\hat{H} - V)(A)$, where $A = (-25 \times y, 25 \times x)$ on the unit disk centered at the origin, are pictured in the first 3 rows. In some cases, the eigenpairs of $H(A)$ in $[50, 51.3]$ have eigenvalues that agree to nine digits. Only 12 pairs were returned that were are all linear combinations of the actual eigenpairs, which should have circular symmetry about the origin. The (modulus of) eigenpairs for the complex shifted operator, $H(A)_s$, and the shift of $s = 1 \times 10^{-6}$ are pictured in Row 4, where $\delta^* = 0.3$. Such s was small enough to identify the sufficiently localized eigenpairs that are not approximations of linear combinations of $H(A)$. Here $h = 0.01$ and $p = 3$.

The size of s and pathological cases							
Order	$s = 1$	1e-1	1e-2	1e-3	1e-4	1e-5	1e-6
1e-1	✓	✓	✗	✗	✗	✗	✗
1	✓	✗	✗	✗	✗	✗	✗
1e-4	✓	✓	✓	✓	✓	✗	✗
1e-3	✓	✓	✓	✓	✗	✗	✗
1	✓	✗	✗	✗	✗	✗	✗
1e-4	✓	✓	✓	✓	✓	✗	✗
1e-2	✓	✓	✓	✗	✗	✗	✗
1e-1	✓	✓	✗	✗	✗	✗	✗
1e-5	✓	✓	✓	✓	✓	✓	✗

Table 5.4: For given $\delta^* = 0.35$ and R as the right bulb of the dumbbell domain, what s induces a pathological false positive? The left column gives the magnitude of agreement for each pair of eigenvalues, $|\lambda_a - \lambda_b|$. Other columns are given a checkmark to indicate a false positive was returned, while a "✗" symbol indicates that no false positive was returned for the given s in the top row. Observe that in each case, s as big as the order of agreement or larger induced the false positives.

2.1156, so they agree to the order of 10^{-1}. All of the remaining columns represent experiments with the given shift s in the right bulb with $\delta^* = 0.35$. A checkmark indicates a returned false positive, so the experiment resulted in an eigenmode localized in R. A "✗" symbol indicates that no false positive was returned. Observe that in each case, s is as big as the order of agreement or larger induced the false positives.

For s near the "small enough" boundary, however, false positives are not cut and dry– one can actually observe the "onset" with increasing s. Figure 5.10 gives an example of this on the ninth pair of eigenpairs for the dumbbell domain, displaying (μ, ψ) of $\mathcal{L}_s$ for varying s approaching the "sufficiently small" cutoff. As seen from Table 5.4, the order of agreement for the ninth pair is 1e-5, so it is

$\mu = 21.931575 + \mathrm{i}\ 4.7197\text{e-}6 \qquad 21.931575 + \mathrm{i}\ 8.85051\text{e-}6 \qquad 21.931575 + \mathrm{i}\ 4.9973\text{e-}5$

Figure 5.10: From left to right $s = 5 \times 10^{-5}$, $\quad 9 \times 10^{-6}$, and 5×10^{-6}. As the shift approaches sufficiently small, the (modulus of the) eigenvector approaches the eigenvector of the unshifted operator. This is consistent with Theorem 2.1.4, where $\psi \to \phi$ as $s \to 0$.

expected using that as s generates a false positive which it does, but this is not the actual limit. The shift can be smaller still, and one can see the eigenmode become closer and closer to its actual, non-concentrated counterpart, ϕ. This is consistent with Theorem 2.1.4, where $\psi \to \phi$ as $s \to 0$.

As given by Kato, a safe s depends on the radius of perturbation r_0 given in (2.8), which is essentially the distance between the target eigenvalue and the next nearest eigenvalue. This bound is quantified in Lemma 2.1.6, where $\frac{r_0}{2}$ is a worst case when the eigenfunctions are completely localized in R. Figure 5.11 pictures four dumbbell eigenpairs together with the smallest shifts (within 10^{x-1}, where x is the current order of magnitude) in the right bulb that returned a false positive; these eigenmodes are clearly not localized in R, so a user would easily determine these were false positives. In all cases, $s < r_0$ was sufficiently small.

Even in the cases where $\mu_1 \approx \lambda$ and $\mu_2 \approx s$, false positives can still occur, though non-pathological cases are uncommon. One example was referenced in Figures 5.3 and 5.4, which was represented by an orange dot that was not distinct from the grey-scale dots like the other orange or red dots did. Here $0.08 = s > r_0 = 0.033$ was the smallest s (up to 10^{x-1}) that induced this false positive, giving

Figure 5.11: The modulus of eigenmodes of the dumbbell domain together with their eigenvalues and the smallest shifts (within 10^{x-1}, where x is the current order of magnitude) in the right bulb that returned a false positive. While $\mu_1 \approx \lambda$ in each case, notice μ_2 is often far from s. A μ_2 significantly different from s indicates that the eigenmode is not as localized in R, and perhaps not as much as desired resulting in a false positive.

$\mu = 18.61759 + 0.783\mathrm{i}$. However, in order to have an acceptable aspect ratio with this s and δ^*, then $[a, b] = [18.6, 18.7]$, but this is only realistically viable because the eigenvalue was known beforehand; in many cases where false positives occur, choosing the "safe" s sufficiently small so as to guarantee separation of $\mathrm{Spec}(\mathcal{L}_s)$ is not economical. Recall that smaller s requires smaller search regions in order to maintain an acceptable aspect ratio. For the examples here, maintaining an acceptable aspect ratio require $b - a$ be on the order of 10^{-5} to 10^{-2}, which is extremely inefficient, even if the user only wanted to sweep over a tiny, total energy range.

In summary, for a user with little to no prior knowledge about the eigenpairs, aiming for $s \lessapprox 1$ is a good starting point that balances time and computational expense while also running a low risk of generating false positives. A practical

$\delta^* \in [0.1, 0.45]$, where the right choice is dependent on the need and the size of the region relative to the whole domain. For a user with prior knowledge that the eigenvalues of interest will be close, choose a smaller s, perhaps even on the order of $s < r_0$.

6   Geometric Localization

6.1   Overview

As one would expect, geometric localization occurs from the geometry of the domain, whether by "trapping" or "squeezing" eigenmodes or by irregularities in the geometry. Delitsyan et al [3,18] summarizes how geometric localization is induced when narrow channels or branches of the wave-guide trap the eigenmode in a subdomain; this usually happens when the spatial variations of the eigenmode are larger than the size c of the channel, $\pi\lambda^{-\frac{1}{2}} \geq c$. The bridges between the bulbs of the Threebulb domain serve as an example of this type geometric localization. Narrowing the channel, thereby decreasing c, increases localization. Comparing the eigenfunctions of the Threebulb Figure 6.8 versus the eigenfunctions when the bridges have been thinned by half, Figure 6.10, exemplifies this. Continuing to thin these channels traps higher energy eigenmodes and allows for more and more localized eigenpairs.

Creating irregularities in the geometry of the domain can also generate localization. Figure 6.1 illustrates that these irregularities need not be extreme; simply cutting the corner of the domain or creating a slightly rough edge can be sufficient.

Bridge location, size, and symmetry also play a key role. For example, an eigenmode that has its minimum along the boundary of the channel is less likely to "bleed out" into the rest of the domain than an eigenmode maximizing along

Figure 6.1: First three Dirichlet eigenmodes with domains (a) 25×1 rectangle, (b) right trapezoid with height 25, base 1, and top 0.9 right trapezoid, (c) right triangle with height 25 and base 1 (half of (a)). [18]

$$\lambda_{23} = 12.3170, \tau = 0.794 \qquad \lambda_{24} = 12.3350, \tau = 0.957$$

Figure 6.2: Eigenpairs 23 (left) and 24 (right) together with their eigenvalues and degrees of concentration of the Threebulb domain. Even though they are rotations of each other, one is significantly more concentrated in the left bulb because it has low energy along the bridge.

the channel. Two eigenmodes can have close eigenvalues, but if they are shaped differently so that one has a line of low energy along the channel and the other has high energy, they will have significantly different τ values. Figure 6.2 highlights how the 23rd and 24th eigenmodes of the Threebulb domain have very different concentrations in the left bulb, even though the eigenfunctions are just rotations of each other.

As previously discussed, the dumbbell domain does not have localization in either square, despite the narrow channel connection. This is due to the symmetry of the domain, so breaking the symmetry, like offsetting the bridge, corrects this. Enlarging one square, which would also disrupt the symmetry, would have the added benefit of forcing localization with preference in the larger square. One might think inducing geometric localisation is then as simple as introducing asymmetry or narrowing a bridge, and while this *can* introduce the desired phenomena, it can also raise the energy of these wanted eigenpairs. How the geometry is manipulated may also introduce more localized eigenmodes in R than another geometric manipulation. Some clever combination of introducing bridges, location, relative

Figure 6.3: The first twelve eigenvectors of the symmetric dumbbell domain occur in pairs with no localization concentrated primarily in the left or right bulb.

size, and asymmetry may optimize geometric localization in a desired region.

6.2 Exploration in Symmetry via the Dumbbell Domain

We have mentioned that disrupting the symmetry of the dumbbell domain can induce localization. The method to achieve this, however, can provide the user some control for where, how, and for what relative energy the localization occurs. The experiments in this section illustrate this point by manipulating the dumbbell domain.

In all of the experiments, breaking the symmetry of the dumbbell caused the eigenmode mass to favor one side of the domain over the other; in many cases, this
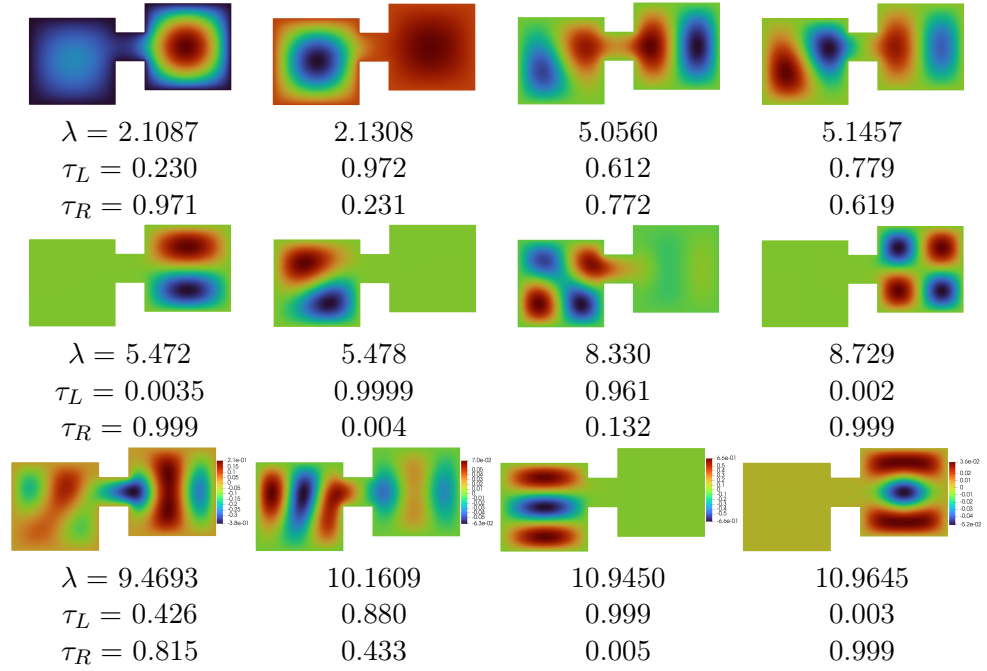
Figure 6.4: The first twelve eigenvectors of the dumbbell where the bridge is shifted up by 0.5. In general, the eigenmodes alternate between the left and right bulbs, where amount of concentration is dependent on the shape of the eigenmode itself. For each complementary pair, the distance between the corresponding eigenvalues is larger than those of the unaltered dumbbell domain, and in most cases the eigenvalues are also larger.

$\lambda = 2.0434$	2.1106	4.9287	5.1026
$\tau_L = 0.995$	0.082	0.959	0.251
$\tau_R = 0.082$	0.995	0.242	0.960
$\lambda = 5.2948$	5.4726	8.4457	8.7289
$\tau_L = 0.9999$	0.0002	0.9998	0.0016
$\tau_R = 0.0004$	0.9999	0.0006	0.9998
$\lambda = 9.1032$	9.7842	10.6075	10.9645
$\tau_L = 0.735$	0.576	0.999	0.002
$\tau_R = 0.509$	0.790	0.003	0.999

Figure 6.5: The first twelve eigenvectors of the dumbbell where the left bulb is larger than the right bulb by 0.05; the bulbs remain centered on the bridge. Eigenmodes alternate between the left and right bulbs, favoring the larger bulb, whose eigenvalues are smaller than the corresponding "pair eigenvalues" of the right bulb. Notice that in most cases, the eigenvalues of even the right bulb are smaller than the correlating eigenvalues of the dumbbell domain.

$\lambda = 0.7792$	1.9292	1.9734	2.1101
$\tau_L = 0.9998$	0.9989	0.9999	0.0234
$\tau_R = 0.0017$	0.0023	5e-6	0.9982
$\lambda = 3.1562$	3.8204	3.9478	5.0712
$\tau_L = 0.9999$	0.9968	0.9999	0.0764
$\tau_R = $ 1e-5	0.0182	5e-5	0.6347
$\lambda = 5.1039$	5.1273	5.4726	6.4254
$\tau_L = 0.645$	0.999	6e-5	0.989
$\tau_R = 0.758$	6e-5	0.999	0.049
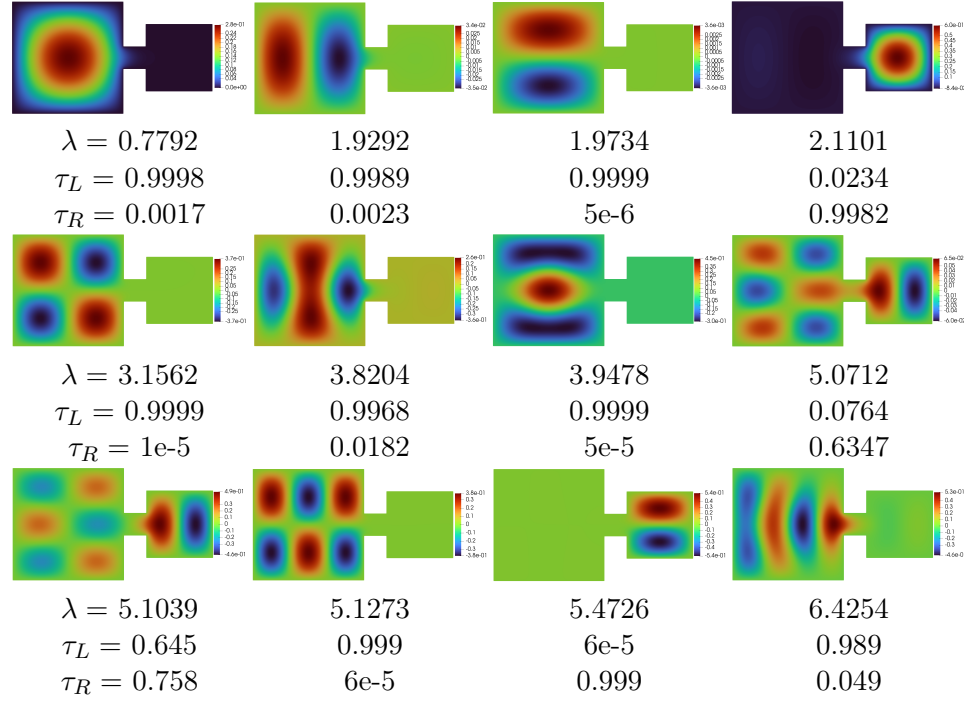
Figure 6.6: The first twelve eigenmodes of the dumbbell where the left bulb is much larger, a 5×5 versus a 3×3 square. The larger left bulb is greatly favored for localization, having many concentrated eigenfunctions with much lower energy. The right bulb, having not changed in size, still contains eigenvector mass with about the same energies as before. Though now this mass is not evenly concentrated, giving way to localized eigenmodes in both bulbs.

$$\lambda = 2.1065 \qquad 2.1185 \qquad 5.0595 \qquad 5.1465$$
$$\tau_L = 0.839 \qquad 0.541 \qquad 0.784 \qquad 0.604$$
$$\tau_R = 0.540 \qquad 0.840 \qquad 0.597 \qquad 0.790$$

$$\lambda = 5.4723 \qquad 5.4739 \qquad 8.7286 \qquad 8.7297$$
$$\tau_L = 0.999 \qquad 0.020 \qquad 0.992 \qquad 0.131$$
$$\tau_R = 0.020 \qquad 0.999 \qquad 0.127 \qquad 0.991$$

$$\lambda = 9.2735 \qquad 9.9362 \qquad 10.9645 \qquad 10.9647$$
$$\tau_L = 0.665 \qquad 0.656 \qquad 0.997 \qquad 0.072$$
$$\tau_R = 0.592 \qquad 0.728 \qquad 0.072 \qquad 0.997$$

Figure 6.7: The first twelve eigenvectors of the dumbbell where the right bulb is shifted up by 0.05; the right bulb is no longer centered on the bridge. The slight break in symmetry is sufficient to force eigenvector localization, often resulting in very localized eigenmodes occurring in each bulb. The eigenvalues are largely consistent with those of the dumbbell domain, however the distance between pairs is greater.

resulted in highly localized eigenvectors. Furthermore, the distance between the eigenvalue "pairs" increased, but it was not consistent how the energies compared in size just by disrupting the symmetry. Of course, the less the domain changed, the less the eigenvalues did; to induce localization more consistently in one bulb, increasing the size of that bulb, even minutely, does so with the added benefit of generally decreasing those eigenvalues as well. When not modifying the size, *how much* localization occurs depends on the shape of the eigenmode. Mode seven in Figure 6.4 is more concentrated than eight because the bridge is centered on the right bulb where the eigemode is zero along that line, which therefore decreased bleeding. Similarly, this occurs in modes nine and ten, where the value of the eigenmode is small near the top, so the high bridge placement for the left bulb preserves the concentration more than in the right bulb where the bridge is centered adjacent to where this eigenmode is larger. When examining Figure 6.7, notice that the left bulb is favored in each pair; in this instance, the perfectly centered bridge is slightly more conducive to localized eigenmodes, at least for low energy modes, where the wave behavior of the vectors are more likely to have a lower energy near the top, center, of bottom of the domain. It will not always be the case that a centered bridge gives rise to more concentrated vectors, as can be seen in the next section 6.3; instead the bridge should be place wherever the desired eigenmodes are most expected to have lower energy.

As a key takeaway, symmetry inhibits eigenvector localization; introducing localization can be as simple as introducing asymmetry, but how this is done affects where and when localization occurs. A larger region is more likely to host con-

centrated eigenmodes than a comparable, yet smaller, region. For domains with a bridge, the location of the bridge does affect the concentration, however optimal bridge placement depends on the shape and energy of the eigenmodes themselves; this is more deeply explored in the next section.

6.3   An Exploration via the ThreeBulb Domain

For brevity, only the first twelve eigenvectors are included for each experiment in this subsection. Eigenfunctions 13-40 are included for each experiment in the Appendix. For the reader who desires more examples or higher energy cases, consult the Appendix.

The ThreeBulb domain is a typical example where localization occurs due to the geometry; in particular, it exemplifies how the thin bridges placed between the bulbs can cut off eigenmodes with evenly distributed mass throughout the domain, into being concentrated in one bulb. Further, it nicely illustrates how size is the primary factor for where the eigenfunctions will concentrate first and most. Examining Figure 6.8, one can see that the large bulb hosts most of the concentrated eigenfunctions, having its ground state and all proceeding eigenfunctions before any other bulb; then the medium and large bulb, as expected. The effects of size, rather than location, can be seen in Figure 6.11 where the middle bulb is now the largest and the left bulb is the smallest. Despite the location of the large bulb, it still hosts just as many concentrated eigenmodes, and it obtains its series of eigenvectors before any other bulb.

Size of the bridge also matters; as it thickens, the domain approaches a rectangle, where no localization will occur, and the eigenvalues decrease. So, as the bridge thins, we would expect the opposite: more localized eigenfunctions with inflating eigenvalues. Figure 6.10 shows the first 12 eigenmodes with bridges that are thinned by half. When the bridge is thinned out, the bulbs become more isolated and host eigenmodes that behave increasingly independent of each other, are more concentrated since there is less bridge to leak into, and have larger eigenvalues.

Bridge location also plays a factor. Placing the bridge where one may expect the eigenmodes to be small can reduce bleeding, resulting in more concentration. We emphasize that moving the bridge is the same as moving the bulb on the bridge, and really it is just the relative location that matters. Perhaps unsurprisingly, the boundary conditions play into the effects of the bridge location. Zero boundary conditions force the (continuous) eigenvector to get smaller as it approaches the bottom (or top) of each bulb, in addition to actually being zero at these boundaries. Meanwhile, there is nothing preventing the eigenmodes from being large, or even maxing, in the center of each bulb. Therefore, we'd expect that centering the bridge would generally lead to more leaking. Figure 6.9 illustrates the first 12 eigenmodes of the ThreeBulb domain when the bridges are centered. In many cases, the functions rotated so as to be symmetric about the center of the bridge, resulting in less localization when the eigenmodes had either an odd number of oscillations or only oscillate in the x direction; when the number of oscillations is even, the eigenmode is zero along a line that is centered on the bridge, but when the number of oscillations is odd, the eigenmode has a high energy pocket centered on

the bridge opening, which then bleeds. For example, the ninth or eleventh modes in Figure 6.9 leak more easily because the eigenfunctions' high energy pockets are centered on the bridge. This phenomena is even worse for higher energies, where increased oscillation creates more, smaller pockets of high energy function that can more easily fit into the bridge opening. Note that the $13^{th} - 40^{th}$ eigenvectors for this domain, show highly localized eigenmodes in all three bulbs when there are an even number of oscillations while also showing that eigenfuctions are not localized at all when there are an odd number. Additionally when the bridge is centered, it allows the function to more easily bleed from above *and* below the bridge, as opposed to just the top *or* bottom, which hppend when the bridge is at the top or bottom of the domain. It is safe to assume that more opportunity to leak will lend to more leaking.

In contrast, the eigenmodes in Figure 6.10 rotated so as to center about the bridge openings. But with the zero boundary conditions and an even thinner opening where the eigenvectors don't have to be zero, there is less space for the (continuous) eigenmodes to achieve higher energy. Therefore, less likely to leak than the original Threebulb domain. As expected, these eigenvectors are (much) more localized. We mention that the bridge location and size may have less of an effect on the eigenfunctions if there were not zero boundary conditions.

One could also place the largest bulb in the center of the domain, but this exposes this bulb, which will host the most eigenfunctions, to more opportunities to leak. When compared to the original Threebulb domain, there are far more eigenfunctions in the center of the domain with (now) two openings to possibly

$\lambda = 1.2298$	2.1715	3.0528	3.0843
$\tau_L = 0.999$	2.5e-8	0.999	0.999
$\tau_M = 3.8e\text{-}5$	1.5e-4	3.4e-4	3.4e-6
$\tau_R = 6.6e\text{-}6$	0.999	1.3e-7	1.3e-9
$\lambda = 4.5028$	4.8706	5.3053	5.4827
$\tau_L = 0.003$	0.998	2.0e-5	2.8e-7
$\tau_M = 0.991$	0.004	0.003	4.7e-5
$\tau_R = 0.002$	2.0e-5	0.995	0.999
$\lambda = 6.0910$	6.16815	7.7167	8.0185
$\tau_L = 0.997$	0.999	0.985	0.9999
$\tau_M = 0.002$	4.7e-5	0.014	1.9e-4
$\tau_R = 1.7e\text{-}5$	3.4e-7	4.3e-4	1.3e-5
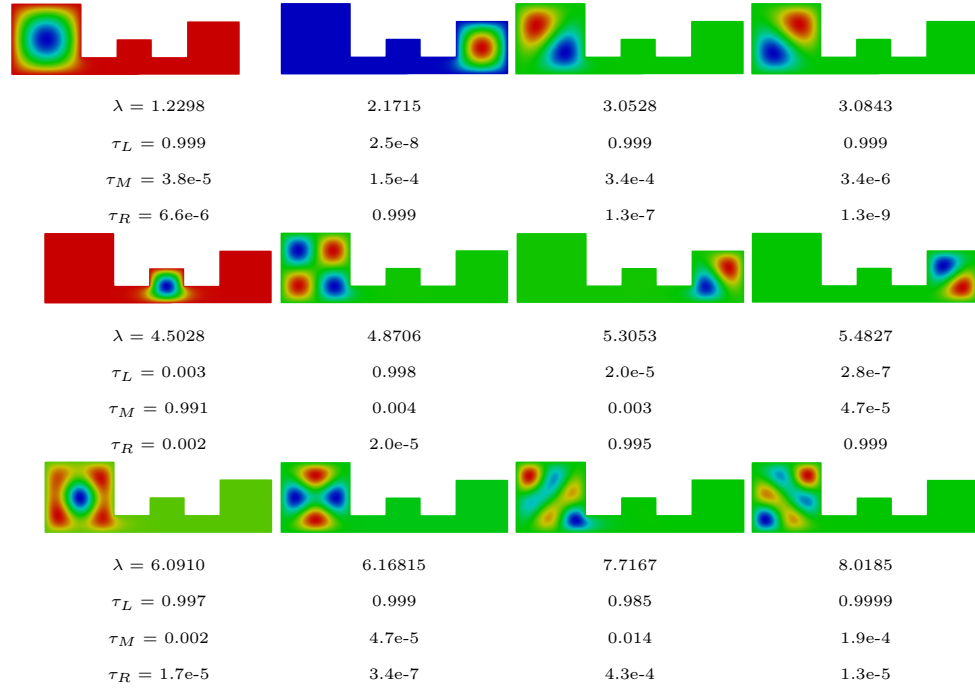
Figure 6.8: First 12 eigenpairs of the ThreeBulb domain.

$\lambda = 1.2080$	2.1101	2.9706	3.0823
$\tau_L = 0.999$	1.8e-5	0.998	0.999
$\tau_M = 1.3\text{e-}4$	4.1e-4	9.9e-4	2.7e-8
$\tau_R = 2.2\text{e-}8$	0.999	8.8e-7	0
$\lambda = 4.1680$	4.9268	5.0914	5.4726
$\tau_L = 0.003$	0.999	4.9e-5	6.5e-9
$\tau_M = 0.985$	1.1e-7	0.006	3e-7
$\tau_R = 0.004$	2.9e-9	0.990	0.999
$\lambda = 5.8196$	6.1683	7.8825	8.0000
$\tau_L = 0.989$	0.999	0.991	0.999
$\tau_M = 0.005$	2.4e-5	0.027	5.9e-7
$\tau_R = 7.9\text{e-}5$	3.2e-7	0.001	2e-8
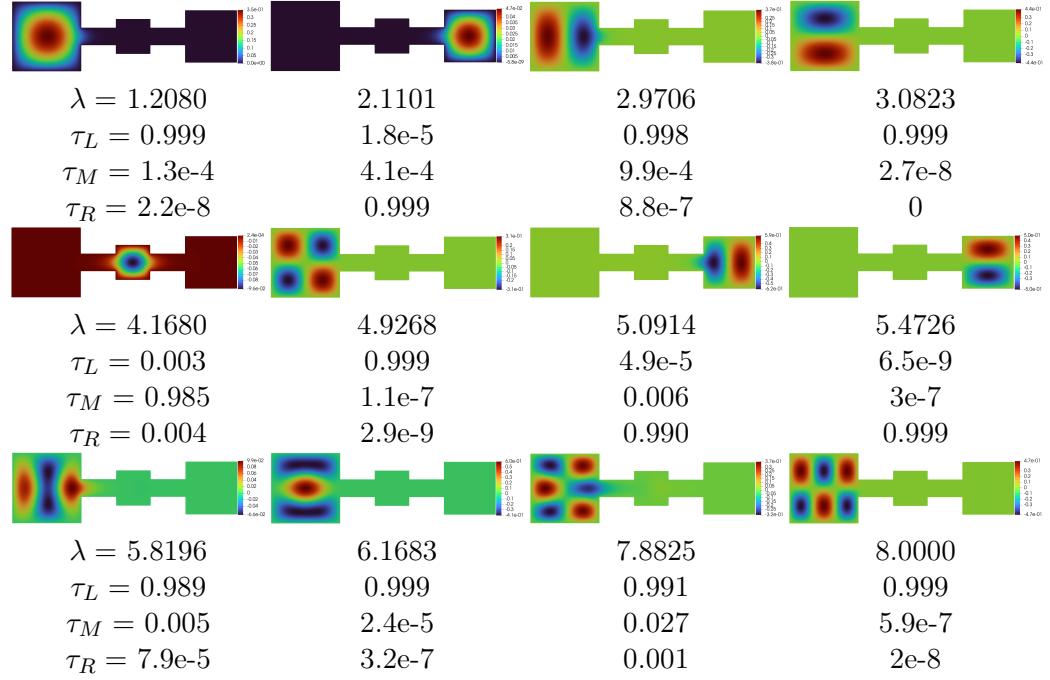
Figure 6.9: First 12 eigenvectors of the ThreeBulb where the bridges between each bulb are centered. These eigenmodes are generally less localized than those of the ThreeBulb domain, which is partially due to the zero boundary conditions and the centered bridge providing more opportunity for the eigenmode to bleed from both above *and* below. In most cases, the eigenvalues are also smaller from increased symmetry, as expected. This trend is even true for higher energies; notably, a ground state eigenfunction, and even another, actually formed in the bridges, though they were not localized.
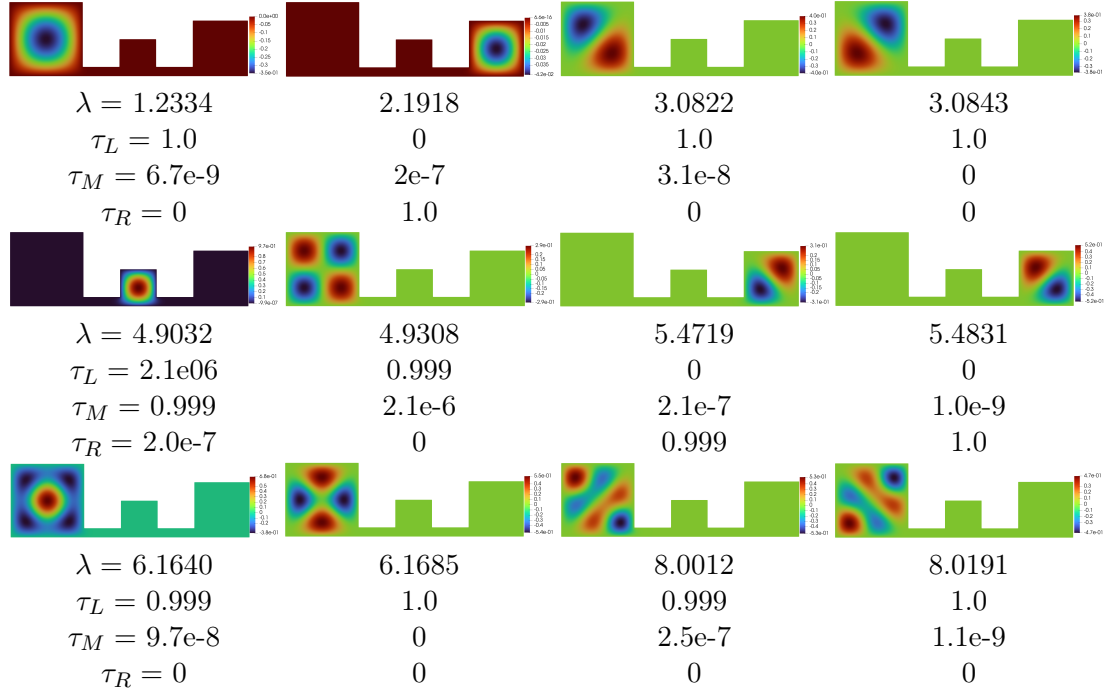
Figure 6.10:   The first 12 eigenvectors of the ThreeBulb where the bridges have been thinned by half. The bridges are $1 \times \frac{1}{2}$, as opposed to 1×1, positioned along the bottoms of the bulbs where the continuous eigenfunctions approach (and reach) zero, so the eigenmodes do not have as much space or energy to bleed. As expected, they are generally more localized. The eigenvalues are larger, and this disparity grows with increasing energy.

leak into, as opposed to only one. Again, it is safe to assume more opportunity to bleed will result in more bleeding. The eigenfunctions in the large, middle bulb were less concentrated. The eigenfunctions of the small, leftmost bulb now had less opportunity to bleed and were consequently more concentrated. This trend is even more dramatic for higher energies, particularly so when the large and small bulb are switch.

In summary, size is a strong indicator about where localization will occur most

	$\lambda = 1.2257$	2.1712	3.0474	3.0574
τ_L	3.7e-5	7e-9	2.5e-4	2.2e-4
τ_M	0.999	1.7e-4	0.999	0.999
τ_R	4.3e-5	0.999	2.3e-4	2.0e-4
	$\lambda = 4.6997$	4.8139	5.3037	5.4827
τ_L	0.995	0.011	1.6e-5	2.3e-7
τ_M	0.011	0.997	0.004	7e-4
τ_R	1.7e-5	0.002	0.995	0.999
	$\lambda = 6.0025$	6.1677	7.6670	7.7689
τ_L	0.003	5.1e-5	0.008	0.009
τ_M	0.994	0.999	0.984	0.985
τ_R	0.004	7.2e-5	0.017	0.200
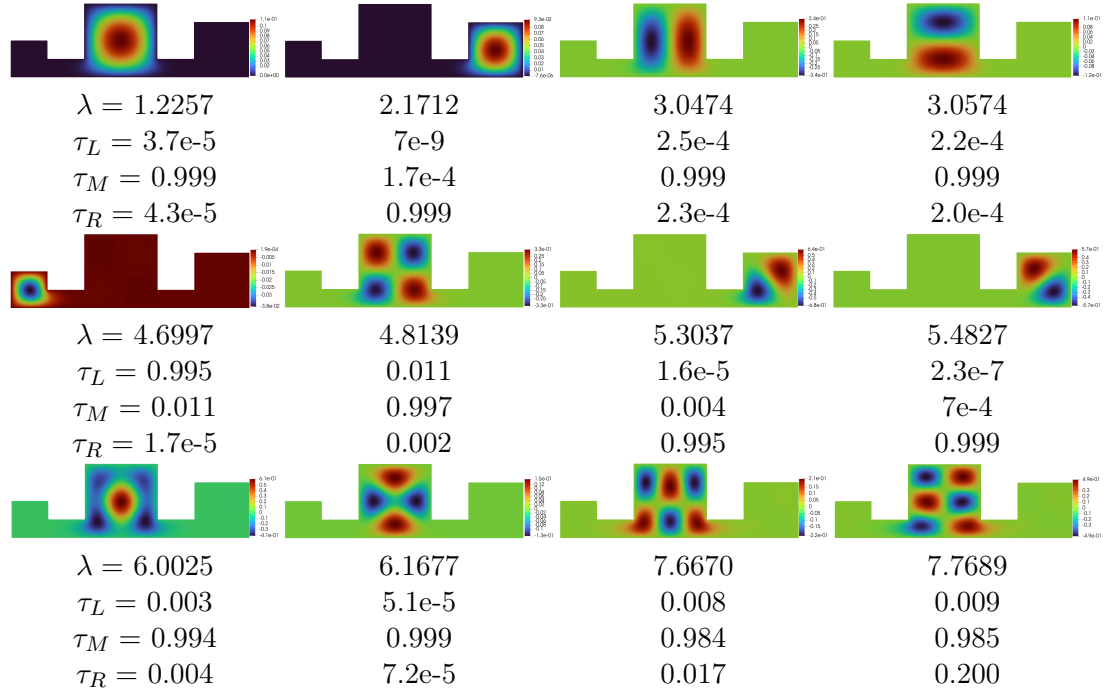
Figure 6.11: The first 12 eigenvectors of the ThreeBulb where the large and small bulb are switched. The large bulb, while still the primary location for localization, has its eigenmodes less localized by a small amount because the added bridge provides another space to leak into. Alternatively, the small bulb, now with only one bridge to leak into, has more concentrated eigenmodes. The eigenvalues are slightly larger.

and first. When placing bridges, note that eigenmodes may center on the bridges, which could greatly change how they concentrate. Boundary conditions can affect where an eigenmode is small, which in turn affects where an eigenmode is prone to bleeding. To increase localization, bridges should be placed where the magnitude of the eigenvector is small. Thinner bridges provide less space to eigenfunctions to leak into and continuing to thin them approaches having isolated regions; this results in subregions that behave more independently, therefore having greater localization at a (nominally) higher energy. More bridges means more potential for leakage, so place them sparingly, particularly for the largest regions of the domain.

6.4 A Note on Higher Energies and Geometric Localization

Quite often, eigenmodes have oscillatory behavior that increases with increasing energy. More oscillatory behavior means smaller, higher energy "pieces" of the the eigenmode that more easily leak through bridges, and are therefore more difficult to confine by the domain alone. This happens, for example, for higher energies in the Threebulb domain, and can be seen in the Appendix. Because of this, geometric localization becomes increasingly difficult to generate as one moves up in the spectrum. A geometry that is sufficient in one part of the spectrum may not be at another, often higher, part of the spectrum. Thus, the desired energy range should be considered before creating or altering a domain for geometric localization, and in general, one should shoot for lower energy eigenmodes for more success with

geometric localization.

7   Potential Localization

We acknowledge that the user is likely confined to some geometric constraints, making it unrealistic to change the domain significantly. Thus, a more typical approach to inducing localization involves manipulating the potential V by placing a potential barrier that is larger than the energy of the eigenpair outside of the region of interest. Localization from the potential often provides a more reliable method for controlling the shape, and to some extent even the energy, of the eigenmodes than geometric localization. But, it can be computationally more expensive. To help minimize expense, one should pay attention to *how* the potential barrier is placed since it can make a significant difference in the eigenpairs. A relatively larger potential magnitude outside of the region of interest generally increases both the number of localized eigenmodes and how concentrated those eigenfunctions are. Just continuing to increase the potential magnitude can significantly increase the eigenvalues though, which means a user would have to sweep through a (often much) wider energy range to locate the (same) desired eigenvectors. In this sense, just blindly throwing an arbitrarily large barrier can be counterproductive. Furthermore, the (spatial) size and arrangement of the potential barrier affects the eigenpairs as well. A larger potential barrier, or the more high potential material in the domain present, increases the eigenvalues without always being beneficial to inducing more localization. Due to the inverse relationship between the size of the

domain and the eigenvalues, smaller domains can be more sensitive to the effects of high potential material, as well. Therefore, it is often best to use the potential barrier sparingly, with sort of a "less is more" approach. Examples are displayed in this section to support these claims; more examples of localization occurring purely from the potential are provided in section 4.1.

Figure 7.1 illustrates how a larger potential magnitude surrounding the region of interest can induce more localized eigenmdoes and with increased concentration, but at the cost of having significantly larger eigenvalues. If only the ground state eigenvector was desired, than perhaps the second of even the first rows, where the potential is 200 and 100, respectively, is sufficient, depending on how tightly confined one may want that eigenfunction. Observe that ground state eigenvalues increase by over 40% from row 1 to row 2. If only the first four eigenmodes were desired, than pushing the potential above 1000, as seen in row 4, would be inefficient. Though a potential of 10,000 seems ridiculous, row 5 shows how an arbitrarily large potential may not provide any more benefit of a smaller, yet sufficiently large, potential, and can even slow the computations. Also notice, as exemplified best in the first row, that once the eigenvalues surpass the potential, the eigenpairs proceed to occupy the whole domain, indicating that the magnitude of the potential really does determines the localization.

Instead of just having one uniform barrier throughout the entire domain, one could introduce a higher potential, local barrier immediately along the region of interest in order to increase concentration. Though potentially to a lesser extent, presence of higher energy (local) material does increase the eigenvalues though,

Figure 7.1: Unit square with zero potential inside a small sub-square and larger constant potential outside. The constant potential for each row from top to bottom is 100, 200, 500, 1000, and 10,000. The first columns display the landscape functions, while the remaining columns correspond to the first four eigenfunctions. As the potential increases relative to the region of interest, more eigenmodes become localized and each is also more concentrated in this region. Notice that this comes at the cost of greatly increasing the eigenvalues as well. One would have to sweep a much larger range in order to find even the ground state eigenmode when the potential is higher.

but this trade-off between increased localization and eigenvalues may prove more beneficial than just increasing the uniform potential over the whole domain, which could increase the eigenvalues even more. This concept is exemplified in Figure 7.2.

The potential barrier does not need to be an impassable wall in order to induce concentration, exemplified in Figure 7.2 and Figure 7.3. Much like geometric localization, the eigenmodes can bleed out through channels of relatively low potential, but as before this largely depends on the shape and energy of the eigenmodes. When comparing Figure 7.2 and Figure 7.3, observe that the eigenfunctions of the domains with the fenced potential in Figure 7.2 leak more. The much smaller domain drives the eigenvalues up faster, which means an even bigger potential is needed to confine the eigenmodes, but then this bigger potential further increases the eigenvalues, which in turn lessens the likelihood of localization. In this sense, smaller domains can be more sensitive to the effects of the potential material. In Figure 7.3, the ThreeBulb domain is given a constant potential of 40 except in the region of interest R and the immediate surrounding area, where the potential is 0 and 100, respectively. In the first row, an impassable wall of high potential generates localized first and second state eigenmodes in R; each proceeding row has less high potential material with the same localized eigenmodes, but with lesser energy. The third and fourth rows even display a passable barrier where the eigenmodes could leak, but do not, and the "localized eigenvalues" are $\approx 10\%$ smaller! Notice also that even though the ground state eigenvalue in Row 4 is smaller than Row 3, the next localized eigenvalue is larger, indicating that the orientation of the potential barrier may have a better affect on eigenmodes of one shape versus
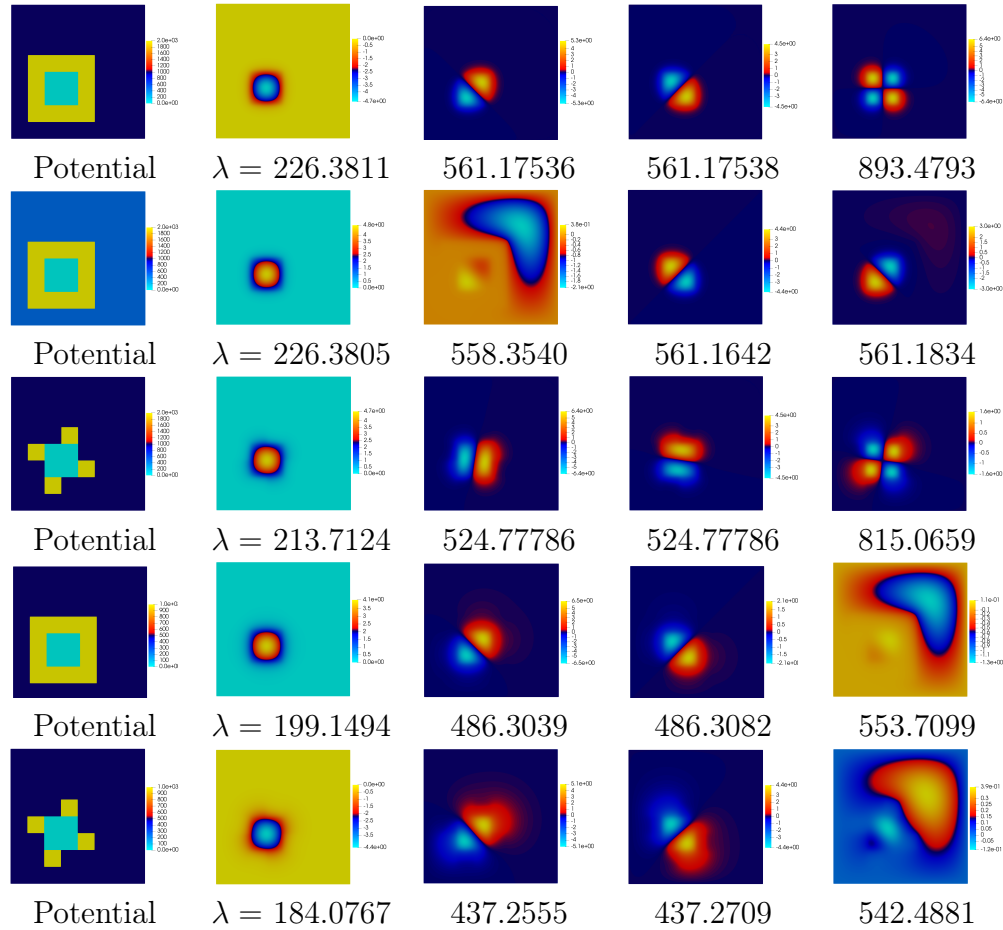
Figure 7.2: Unit square with zero potential inside a small sub-square surrounded immediately by a potential barrier of magnitude 2000 (top 3) or 1000 (bottom 2) and a smaller potential in the rest of the domain. The first column gives the potential function for each case. This example shows how a thin, higher potential barrier placed immediately around the region of interest, along with a smaller surrounding potential, can help to better concentrate eigenmodes. Row 3 has a surrounding potential of 1000 and an immediate potential of 2000; it has similar eigenvalue behavior to Row 4 and Row 5 of Figure 7.1, however its eigenvalues are larger than Row 4 and less than Row 5, while having eigenvectors more localized than Row 4 and less than Row 5. The last two rows have an outer potential of 500 and an immediate barrier of 1000. these are more localized than just having the surrounding potential of 500, but with bigger eigenvalues. Of particular interest is the final row, where the fence does not entirely surround the region of interest, but is sufficient to increase localization without increasing the eigenvalues as much as a solid barrier.
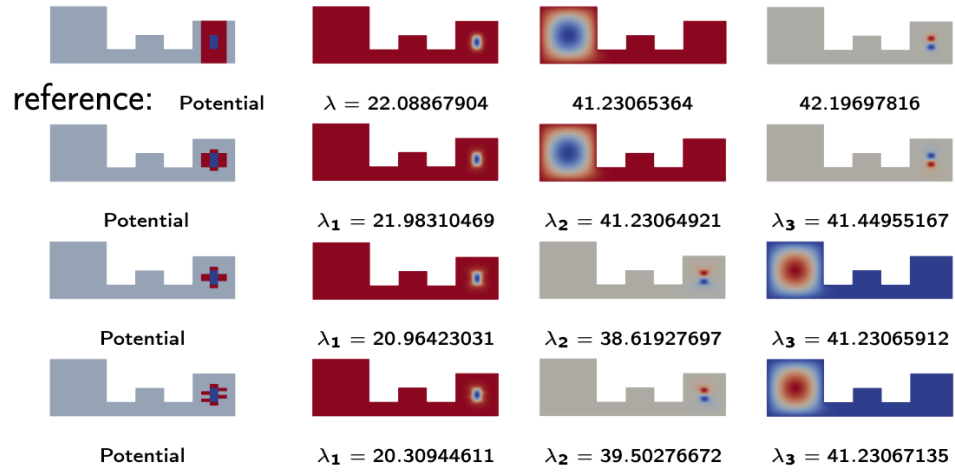
Figure 7.3: Domain potential of 40 with immediate potential barrier in right bulb of 100 around a 1.67×1 square R with 0 potential. The potential function followed by the first 3 eigenmodes are pictured in each row along with the corresponding eigenvalues.

another shape. There is just as much 100 potential material in Row 3 as Row 4, so the energy difference is due to the arrangement of the barrier.

8   Eigenfunctions of the Magnetic Laplacian and Localization

The theory in this chapter was published with citation: in [32]

J. S. Ovall, H. Quan, R. Reid, S. Steinerberger. On localization of eigenfunctions of the magnetic Laplacian. 2023.

Jeffrey S. Ovall, Hadrian Quan, Robyn Reid, and Stefan Steinerberger

Author contributions to the article cited above.

J.S. Ovall: theory, writing, editing, numerical experiment advising

H. Quan: theory, proofs, writing, editing

R. Reid: analysis, simulation, writing, and editing

S. Steinerberger: theory, proofs, writing, editing

All of the theory presented in this manuscript has been developed via algebraic manipulation and relies only on the assumption that the eigenvalues are real; therefore, in the case of the magnetic Laplacian operator, which is selfadjoint with real spectrum, all of this previous theory still holds directly. Examples of the algorithm applied to localization searches using the magnetic Laplacian are provided in the section 4.3. This chapter builds further to provide mechanisms for a priori prediction for where eigenfunctions may localize when the operator is the magnetic Laplacian operator. Results are supported by experimental evidence.

8.1 Theory

Suppose that $\Omega \subset \mathbf{R}^d$ is a bounded domain. Though unnecessary, assume Ω has a smooth boundary. The magnetic Schrödinger operator $\hat{H}$ applied to some differentiable vector field $A : \Omega \to \mathbf{R}^d$ is defined:

$$\hat{H}(A) := (-\mathrm{i}\nabla - A(x))^2 + V \ .$$

It is known from Filoche and Mayboroda, and even proven earlier, that for the eigenvalue problem with eigenpair (λ, ψ):

$$\frac{|\psi(x)|}{||\psi||_{L^\infty}} \le \lambda u \ , \tag{8.1}$$

where u is the landscape function. Hoskins, Quan and Steinerberger [21] proved that (8.1) remains true for the operator $\hat{H}$, indicating that the potential $\vec{0} \le V \in \mathbf{R}$

dominates where localization occurs, unless it is very small compared to A. We note that $V \geq 0$ is for convenience, and could be negative as long as it was bounded below. Naturally, the question arises about what dictates where localization occurs when the potential is small or not present.

Define the magnetic Laplacian H as the magnetic Schrödinger without the presence of the potential V: $H(A) = (\hat{H} - V)(A)$ and assume zero Dirichlet boundary conditions. We aim to understand how A and the eigenvalues dominate where and *if* eigenmodes of $H(A)$ will localize. Assuming $||\psi||_{L^2(\Omega)} = 1$, then using the Rayleigh Quotient gives:

$$\lambda = (H(A)\psi, \psi)_{L^2(\Omega)} = \left((-\mathbf{i}\nabla - A(x))^2 \psi, \psi\right) = \int_\Omega |(-\mathbf{i}\nabla - A(x))\psi(x)|^2$$

indicating that, at least for low energy eigenvalues, $\nabla\psi(x)$ tends to point in the direction of $-\mathbf{i}A(x)\psi(x)$. To minimize this expression, like one would do to obtain the first eigenvalue, it must be that $A(x)\psi(x) = -\mathbf{i}\nabla\psi$, which indicates that $A(x)\psi(x)$ behaves like the gradient of a function (at least) for low energies. This motivates the idea that the regions where A does behave like the gradient of a function are regions of special interest, and we will show that these are the exact regions where eigenvectors of $\lambda\psi = H(A)\psi$ will localize. Notice that this expression also stipulates that all of the eigenvalues are non-negative. By the Fundamental Theorem of Vector Calculus, any A can be expressed as the sum of a solenoidal

vector field F and an irrotational vector field $\nabla\phi$. So, we express:

$$A(x) = \nabla\phi + F(x) \qquad \phi \in (C)^1, \;\; \nabla \cdot F = 0 \;. \tag{8.2}$$

Performing the Helmholtz decomposition illustrates how the irrotational contribution of the vector field will not dictate where an eigenfunction will localize. This is because the irrotational component is removed after performing the Helmholtz decomposition. Instead, this component impacts the modulation of the eigenfunction in the sense $\psi \to e^{-i\phi}\psi$, as seen by the following Lemma:

Lemma 8.1.1. *Define the operator* $\tilde{H}g = e^{-i\phi}H(e^{i\phi}g)$. *We have* $\tilde{H} = (-i\nabla - F)^2$. *Then* $H\psi = \lambda\psi$ *if and only if* $\tilde{H}(e^{-i\phi}\psi) = \lambda e^{-i\phi}\psi$.

Proof. Assume $H\psi = \lambda\psi$, and compute:

$$\tilde{H}(e^{-i\phi}\psi) = e^{-i\phi}H(e^{i\phi}e^{-i\phi}\psi) = e^{-i\phi}H(\psi) = \lambda\psi e^{-i\phi} \;.$$

Conversely, assume $\tilde{H}(e^{-i\phi}\psi) = \lambda\psi e^{-i\phi}$:

$$\lambda\psi e^{-i\phi} = \tilde{H}(e^{-i\phi}\psi) = H\psi e^{-i\phi} \;,$$

where canceling $e^{-i\phi}$ gives the desired result. $\square$

We highlight that the key property of conservative vector fields is that their path integrals depend only on the endpoints of the path, not the route in between; this means that for path: $\gamma : [0, 1] \to \Omega, \quad \int_\gamma \nabla\phi \cdot dx = \phi(\gamma(1)) - \phi(\gamma(0))$. In this same
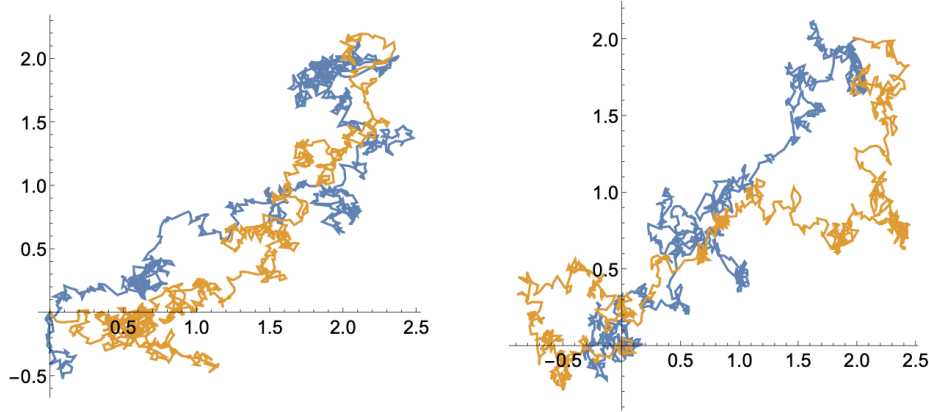
Figure 8.1:   Two sets of two equally likely Brownian motions conditioned upon $\omega(0) = (0,0)$ and $\omega(t) = (2,2)$. [32]

vein, we consider Brownian motion, $\omega(s) : [0,t] \to \Omega$ for time $t > 0$, conditioned only upon its starting point $\omega(0) = x \in \Omega$, its final point $\omega(t) = y \in \Omega$, and that the route remains in some domain Ω. Figure 8.1 illustrates two examples where $x = (0,0)$ and $y = (2,2)$; though the paths differ, they are equally likely and both allowable since they meet the necessary conditions.

For some vector field $A(x)$, then this Brownian motion can be considered in the context of some stochastic path integral:

$$\omega \to \int_0^t A(\omega(s)) \cdot ds \ ,$$

which can be considered as a real valued, random variable. If A is irrotational only, then the path integral is deterministic, so it is independent of ω and is almost surely equal to some point in the sense that it can be measured as a Dirac measure. If A contained some relatively large solenoidal component, then path

integral depends strongly on the path, being less concentrated in some fixed point. As one may expect, the more the solenoidal component decreases, the more the random variable becomes independent of the path, becoming closer and closer to that fixed point $\phi(x) - \phi(y)$. This paves the way for Theorem 8.1.2, which is valid for Neumann boundary conditions as well.

Theorem 8.1.2. *Let* (λ, ψ) *be an eigenpair of* $H(A)$, *with* $\psi|_{\delta\Omega} = 0$ *and the Helmholtz decomposition* $A = \nabla\phi + F$ *such that* $\nabla \cdot F = 0$. *Suppose that* $|\psi(x_0)| = ||\psi||_{L^\infty}$ *for some* $x_0 \in \Omega$. *Then* $\forall t > 0$:

$$\int_\Omega |\mathbf{E}_{\omega(0)=x_0, \omega(t)=y} e^{\mathrm{i}\int_0^t F \cdot d\omega(s)}| \frac{1}{(4\pi t)^{d/2}} e^{-\frac{||x_0-y||^2}{4t}} dy \geq e^{-\lambda t} \ . \qquad (8.3)$$

Proof. First, notice that

$$\lambda = (H(A)\psi, \psi)_{L^2(\Omega)} = \int_\Omega |(-\mathrm{i}\nabla - A(x))\psi(x)|^2 \geq 0 \ ,$$

indicating that all eigenvalues are nonnegative. Under the assumption that A is such that $A \cdot A$ and $\nabla A \cdot A$ are in the Kato class, there exists a Feynman-Katz formulation as expressed in Proposition 2.9 in Broderix et. al. [10]:

$$\left(e^{-tH(A)}\psi\right)(x) = \mathbf{E}\left(e^{-S_t(A|\omega)}\chi_\Omega(\omega, t)\psi(\omega(t))\right) \ , \qquad (8.4)$$

where the expectation is taken over all Brownian motion that starts at x, $\omega(0) = x$. The function $\chi_\Omega(\omega, t)$ acts as an indicator for when the Brownian motion is in the

domain:

$$
\chi_\Omega(\omega, t) =
\begin{cases}
1 & \text{if } \forall s \in [0, t], \omega(s) \in \Omega \\
\\
0 & \text{else}
\end{cases}
$$

And define:

$$
S_t(A|\omega) = \mathrm{i} \int_0^t A(\omega(s)) d\omega(s) + \frac{\mathrm{i}}{2} \int_0^t (\nabla \cdot A)(\omega(s)) ds \ . \tag{8.5}
$$

The time evolution operator for eigenfunctions can be represented:

$$
\left(e^{tH(A)} \psi \right)(x) = e^{-\lambda t} \psi(x) \tag{8.6}
$$

Carrying out the Helmholtz Decomposition on A and then utilizing Lemma 8.1.1 allows the substitution for the operator $\tilde{H}(F)g = (-\mathrm{i}\nabla - F)^2 g$, which has the same eigenvalues λ and the eigenvectors are just scaled, $g = e^{-\mathrm{i}\phi}\psi$. Since $e^{-\mathrm{i}\phi}$ is of complex modulus 1, then $|g(x)| = |\psi(x)|$, meaning both functions will localize in the same regions and be maximized at the same point $x_0 \in \Omega$. Therefore $|g(x_0)| = ||\psi||_{L^\infty}$. Immediately (8.4) can be rewritten in terms of using this substitution and (8.6) applied to x_0:

$$
\begin{aligned}
e^{-\lambda t} ||\psi||_{L^\infty} &= |\mathbf{E}\left(e^{-S_t(F|\omega)} \chi_\Omega(\omega, t) g(\omega(t)) \right)| \\
&= |\int_\Omega \left(\mathbf{E}_{\omega(0)=x,\omega(t)=y} e^{-S_t(F|\omega)} \right) g(y) \rho_t(x_0, y) dy| \\
&\leq ||g||_{L^\infty} \int_\Omega |\left(\mathbf{E}_{\omega(0)=x,\omega(t)=y} e^{-S_t(F|\omega)} \right)| \rho_t(x_0, y) dy \ ,
\end{aligned}
$$

where because the divergence free vector field F has now been substituted, (8.5) becomes the purely imaginary expression $S_t(F|\omega) = \mathfrak{i} \int_0^t F(\omega(s))d\omega(s)$. Here, the expectation is conditioned on $\omega(t) = y$, and since the probability distribution for particles that begin at x_0, go for t time, and remain inside the domain can be represented via the heat kernel, we have $y \to \rho_t(x_0, y)$. The last inequality is made by applying the triangle inequality and recognizing that $\rho_t(x_0, y) \geq 0$.

Finally, using that this $\rho_t(x_0, y)$ is bounded above by the heat kernel over $d-$dimensional Euclidean space provides the desired inequality.

$$
\begin{aligned}
e^{-\lambda t} &\leq \int_\Omega | \left(\mathbf{E}_{\omega(0)=x, \omega(t)=y} e^{-\mathfrak{i} \int_0^t F(\omega(s))d\omega(s)} \right) | \rho_t(x_0, y) dy \\
&\leq \int_\Omega | \left(\mathbf{E}_{\omega(0)=x, \omega(t)=y} e^{-\mathfrak{i} \int_0^t F(\omega(s))d\omega(s)} \right) | \frac{1}{(4\pi t)^{d/2}} e^{-\frac{||x_0-y||^2}{4t}} dy .
\end{aligned}
$$

$\square$

Note that instead of using the Gaussian heat kernel, defined on $\mathbf{R}^d$, the heat kernel on a given bounded domain could be substituted, which is a stronger result and could be advantageous. Essentially, Theorem 8.1.2 states that for an eigenpair (λ, ψ) of $H(A)$ where $\psi(x_0) = ||\psi||_{L^\infty}$, then there is an expectation for where most points y in some neighborhood of x_0, from the path integral x_0 to y, will be. The path integral must be relatively concentrated around some fixed value, which means that the vector field is either "close" to a conservative vector field or the integration path is short, and therefore λ is large. We coin the prior as "near-deterministic". We say that X is "near-deterministic" if it is likely to be

found close to some fixed value (mod 2π) in the sense:

$$\sup_{z \in \mathcal{T}} \mathcal{P}\left(|X \mod 2\pi - z| \leq \frac{1}{100}\right) \geq \frac{99}{100} \, ,$$

for $A = \nabla\phi + F$ such that $\nabla \cdot F = 0$, and where $\mathcal{T}$ represents the torus. For $y \in \Omega$ to be near deterministic, let $X = \int_0^t F d\omega(s)$ for any path integral that begins at x_0 and ends $\omega(t) = y$. Notice that this is one of the two ways that a random path integral can be highly concentrated, as needed for Theorem 8.1.2. Corollary 8.1.4 describes where this may occur, and Lemma 8.1.3 aids in its introduction.

Lemma 8.1.3. *For X as some real-valued, near deterministic, random variable, there exists some universal constant $c \in [0, 1)$ such that $|\mathbf{E}(e^{iX})| \leq c$.*

Proof. Notice that every point e^{iX} sits on the unit disk, and that the expectation $\mathbf{E}$ is a convex combination of these points. Since $\mathcal{S}^1$ is convex, then any convex combination of points in $\mathcal{S}^1$ must be inside the domain. $\square$

Corollary 8.1.4. *Let(λ, ψ) be an eigenpair of $H(A)$, with $\psi|_{\delta\Omega} = 0$ and the Helmholtz decomposition $A = \nabla\phi + F$ such that $\nabla \cdot F = 0$. Suppose that $|\psi(x_0)| = ||\psi||_{L^\infty}$ for some $x_0 \in \Omega$. For some universal constant $c \in (0, 1)$ that depends only on the spatial dimension d, let $t = \frac{c}{\lambda}$. Then a large fraction of points y in a $\sqrt{t}$-neighborhood of x_0 are near-deterministic:*

$$\{y \in B_{\sqrt{t}}(x_0) : y \text{ is near-deterministic}\} \geq \frac{9}{10}|B_{\sqrt{t}}(x_0)| \, .$$

Proof. By proof of contradiction, let $c \to 0$. Then for any $c > 0$ there exists an

eigenpair of $H(A)$, (λ, ψ), localized at x_0, such that for $t = \frac{\lambda}{c}$, there measure of points $t \in B_{\sqrt{t}}(x_0)$ that are near deterministic is $9/10^{th}$ of the measure. From the proof of Theorem 8.1.2, we have:

$$e^{-c} \leq \int_\Omega |\left(\mathbf{E}_{\omega(t)=y} e^{-S_t(F|\omega)}\right)| \rho_t(x, y) dy$$

Define the set:

$$\mathbf{A} = |y \in B_{\sqrt{t}}(x_0) : y \text{ is not near-deterministic}| .$$

Using that $\Omega = \mathbf{A} \cup \Omega \setminus \mathbf{A}$ and then applying Lemma 8.1.3 gives:

$$
\begin{aligned}
&\int_\Omega |\left(\mathbf{E}_{\omega(t)=y} e^{-S_t(F|\omega)}\right)| \rho_t(x, y) dy \\
&\leq \int_\mathbf{A} |\left(\mathbf{E}_{\omega(t)=y} e^{-S_t(F|\omega)}\right)| \rho_t(x, y) dy + \int_{\Omega \setminus \mathbf{A}} |\left(\mathbf{E}_{\omega(t)=y} e^{-S_t(F|\omega)}\right)| \rho_t(x, y) dy \\
&\leq k \int_\mathbf{A} \rho_t(x, y) dy + \int_{\Omega \setminus \mathbf{A}} \rho_t(x, y) dy ,
\end{aligned}
$$

for some $k \in [0, 1)$. Since $\mathbf{A}$ and the kernel $\rho_t(x, \cdot)$ each contain a positive proportion of the measure in the $\sqrt{t}$-neighborhood, we have:

$$\int_\mathbf{A} \rho_t(x, y) dy + \int_{\Omega \setminus \mathbf{A}} \rho_t(x, y) dy \leq 1 - c_2 ,$$

for some constant c_2. This is a contradiction when $c \to 0$ because then $t \to 0$. $\square$

Thus, far we have seen that the Helmholtz Decomposition (8.2) on the vector

field A separates A into the irrotational vector field $\nabla\phi$, which is does not impact where localization will occur, but instead simply modulates the eigenfuction, and the solenoidal component F that determines where the eigenmode will localize. Theorem 8.1.2 and Corollary 8.1.4 then demonstrated how the regions where F behaved similarly to a conservative vector field are the regions where localization should occur. These are the regions where $|\text{curl}F|$ is relatively small. To easily exemplify, assume $\Omega \in \mathbf{R}^2$, then for $F = (F_1, F_2)$, we have $|\text{curl}F| = \partial F_2/\partial x_1 - \partial F_1/\partial x_2$. Take $x_0 \in \Omega$ and define $A_{line}(x) = A(x_0) + J(x_0)(x - x_0)$ to be the linearization of A about x_0, where J is the Jacobian of A, then we can define the "nonlinearization" of A by $A_{nonline} = A - A_{line}$. Furthermore, A_{line} can be decomposed $A_{line} = A_1 + A_2$ such that for the 90^o counterclockwise rotation matrix $R = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$ we have:

$$A_1 = A(x_0) + \frac{1}{2}\left(J(x_0) + J^T(x_0)\right)(x - x_0) \tag{8.7}$$

$$A_2 = \frac{1}{2}\left(J(x_0) - J^T(x_0)\right)(x - x_0) = \frac{\text{curl}A(x_0)}{2}R(x - x_0) . \tag{8.8}$$

Noticing that $\frac{1}{2}\left(J(x_0) + J^T(x_0)\right)$ is the symmetric part of the Jacobian matrix of A helps to immediately conclude that A_1 is conservative. In fact, $A_1 = \nabla f$ for

$$f(x) = A(x_0) \cdot (x - x_0) + \frac{1}{2}(x - x_0)^t J(x_0)(x - x_0) . \tag{8.9}$$

Direct computation shows that A_2 is solenoidal. It becomes clear from this why

the curl$A$ is of special interest here, and is detailed in Corollary 8.1.5.

Corollary 8.1.5. *Let* (λ, ψ) *be an eigenpair of* $H(A)$, *with* $\psi|_{\delta\Omega} = 0$. *Suppose that* $|\psi(x_0)| = ||\psi||_{L^\infty}$ *for some* $x_0 \in \Omega$. *If* $A_{nonline}$ *is sufficiently small in a* $\frac{1}{\sqrt{\lambda}}$ *-neighborhood of* x_0, *then*

$$curlA(x_0) \leq c\lambda^2 \ ,$$

for some universal constant $c > 0$.

Proof. Using Lemma 8.1.1, we can conjugate $H(A)$ in order to remove the conservative part, resulting in div$A = 0$. Using the Taylor expansion of A gives $A = A_{line} + A_{nonline}$, and one can further decompose $A_{line} = A_1 + A_2$ as defined in (8.7), where the conservative part of A_1 is described by (8.9). Then because the conservative part of A has been removed and A_2 is solenoidal, we have:

$$0 = \nabla \cdot A(x_0) = \nabla \cdot A_{line}(x) = \nabla \cdot A_1(x) = \nabla \cdot (\nabla f(x)) = \Delta f(x)$$

$$curlA(x_0) = curlA_{line}(x) = curlA_2(x) \ .$$

Now applying Ito's lemma, using the assumption $\Delta f = \nabla \cdot A(x_0) = 0$, and noting that $\omega(t) = y$, we get:

$$\int_0^t A_1 \cdot d\omega(s) = f(\omega(t)) - \frac{1}{2} \int_0^t \Delta f(\omega(s))ds = f(\omega(t)) \ .$$

Since $A_{line} = A_1 + A_2$, substituting the above gives:

$$\int_0^t A_{line}(x, y) \cdot d\omega(s) = f(\omega(t)) + \frac{\text{curl}A(x_0)}{2} \int_0^t R(x - x_0) \cdot d(\omega(s)) \, ,$$

where $x = \omega(s)$. Since $f(\omega(t)) = f(y)$, then the first term is deterministic. Thus we must investigate the random variable, which is such because it is path dependent, $\int_0^t R(x - x_0) \cdot d(\omega(s))$. We have that $\omega(0) = x_0$ and without loss of generality assume that $x_0 = 0$. Notice that Brownian motion and our vector field are invariant under rotation, and therefore our random variable can only depend on t and $||\omega(t)|| = ||y||$. It is also a property for Brownian motion, that for any $\alpha > 0$, $\omega(\alpha t) \equiv \sqrt{\alpha} \cdot \omega(t)$, by which we mean both random processes are identical. Applying this property, we can say that the probability of a Brownian particle traveling along a fixed path from $\omega(0) = 0$ to $\omega(t) = 1$ is the same as the probability of the rescaled particle traveling from $\omega(0) = 0$ to $\omega(1) = \omega(\frac{1}{t}t) = \sqrt{1/t} \cdot \omega(t) = t^{-1/2}y$. Via a change of variables to change the limits of integration and substitute $s = t \cdot s$, we obtain:

$$\int_0^t R(x) \cdot d(\omega(s)) = t \int_0^1 R(x) \cdot d(\omega(s \cdot t)) = t^2 \int_0^1 R(x) \cdot d(\omega(s)) \, .$$

When $t = 1$ and $y \approx 1$, this random variable takes a deterministic mean value that depends only on the endpoint and with nonzero standard deviation that spread over the interval that is the length of $y \approx 1$. This interval then essentially spreads

over t^2 from which it is scaled by. Consequently,

$$\frac{\text{curl}A(x_0)}{2} \int_0^t R(x - x_0) \cdot d(\omega(s)) \text{ spreads over } t^2 \frac{|\text{curl}A(x_0)|}{2} .$$

For the eigenpair (λ, ψ) of $H(A)$ where ψ attains its maximum at x_0, we have the expectation conditioned upon $\omega(0) = x_0$ and $\omega(t) = y$:

$$\int_\Omega |\mathbf{E}e^{-S_t(A|\omega)}|\rho_t(x, y)dy \geq e^{-\lambda t} \geq 0.95 ,$$

where we grabbed the first inequality from the proof in Corollary 8.1.4 and the second from letting $t = 0.01/\lambda$. In order to satisfy this inequality, without considering the nonlinear part just yet, this path integral must be highly concentrated around most points. For some universal constant $c > 0$:

$$\frac{|\text{curl}A(x_0)|}{2} \left(\frac{0.01}{\lambda}\right)^2 \leq c$$

Then we can define some constant $c_1 = \frac{2c}{0.01^2}$ to simplify that statement:

$$|\text{curl}A(x_0)| \leq c_1\lambda^2 .$$

Finally considering the nonlinear part, we note that the higher order terms must be locally small, thereby:

$$|\int_0^t A_{nonline}(x, y) \cdot d\omega(s)| \ll |\text{curl}A(x_0)|\frac{1}{\lambda^2} .$$

□

While it is not yet clear if this type of decomposition is the best, it can be seen in the proceeding numerical examples that using sublevel sets of the $|\text{curl}A|$ leads to reasonable predictions of where eigenvectors may localize, at least early in the spectrum. Another a priori evaluation for where the eigenmodes may localize can be discovered by refining the landscape function, u, for $-\Delta u = 1$ in Ω with $u|_{\delta\Omega} = 0$. Rearranging the Landscape Inequality (8.1), we see that for any eigenpair (λ, ψ) of $H(A)$ satisfying Dirichlet boundary conditions:

$$|\psi(x)| \le \lambda ||\psi||_{L^\infty} u(x) \ .$$

Since the magnetic field A is not taken into account, it may come as no surprise that this inequality is not particularly informative; however, when A is irrotational, it accurately predicts where localization occurs in the ground state. For brevity, we denote:

$$\mathbf{E}_{x,t}(y) := \mathbf{E}_{\omega(0)=x,\omega(t)=y} e^{\mathrm{i} \int_0^t F(\omega(s))d\omega(s)} \ ,$$

where $\mathbf{E}_{x,t}(y) \in \{z \in \mathbb{C} : |z| = 1\}$. By integrating the result in Theorem 8.1.2, we will show when the landscape inequality is accurate for the ground state and present an improved landscape inequality in Corollary 8.1.6.

Corollary 8.1.6. *We have that:*

$$|\psi(x)| \leq \lambda ||\psi||_{l^\infty} \cdot \left(\int_0^\infty \int_\Omega |\mathbf{E}_{x,t}(y)|^2 \rho_t(x,y) dy dt \right)^{\frac{1}{2}} \sqrt{u(x)} \,, \qquad (8.10)$$

where we can bound the first integral:

$$\left(\int_0^\infty \int_\Omega |\mathbf{E}_{x,t}(y)|^2 \rho_t(x,y) dy dt \right)^{\frac{1}{2}} \leq \left(\int_0^\infty \int_\Omega \rho_t(x,y) dy dt \right)^{\frac{1}{2}} \leq \sqrt{u(x)} \,. \quad (8.11)$$

Proof. From the proof of Theorem 8.1.2, we have the inequality that holds for all x and for all t:

$$\int_\Omega |\mathbf{E}_{x,t}(y)| \rho_t(x,y) dy \geq e^{-\lambda t} \frac{|\psi(x)|}{||\psi||_{L^\infty}} \,.$$

Integrating both sides with respect to time and recognizing that $\rho_t(x,y) \geq 0$ gives:

$$\int_0^\infty \int_\Omega |\mathbf{E}_{x,t}(y)| \rho_t(x,y) dy dt \geq \frac{1}{\lambda} \frac{|\psi(x)|}{||\psi||_{L^\infty}} \,,$$

where the integration $\int_0^\infty e^{-\lambda t} dt = \frac{1}{\lambda}$. Applying the Cauchy-Schwarz Inequality to the left hand side where for $|\langle f, g \rangle| \leq ||f|| \, ||g||$, we define $f = |\mathbf{E}_{x,t}(y)| \sqrt{\rho_t(x,y)}$ and $g = \sqrt{\rho_t(x,y)}$. Then the inequality becomes:

$$\left(\int_0^\infty \int_\Omega |\mathbf{E}_{x,t}(y)|^2 \rho_t(x,y) dy dt \right)^{1/2} \left(\int_0^\infty \int_\Omega \rho_t(x,y) dy dt \right)^{1/2} \geq \frac{1}{\lambda} \frac{|\psi(x)|}{||\psi||_{L^\infty}} \,.$$

Since the eigenfunctions of the Laplacian on Ω, ϕ_k, form a Hilbert basis in L^2, any function in L^2 can be expressed in terms of the ϕ_k. With some manipulation,

which is illustrated in Remark 8.1.7, we apply this to our heat kernel and obtain the closed form expression:

$$\rho_t(x, y) = \sum_{k=1}^{\infty} e^{\lambda_k t} \phi_k(x) \phi_k(y) \ , \tag{8.12}$$

where the λ_k are the corresponding eigenvalues. Using this and switching the order of integration gives:

$$\int_{\Omega} \int_0^{\infty} \rho_t(x, y) dt dy = \sum_{k=1}^{\infty} \frac{1}{\lambda_k} \phi_k(x) \left(\int_{\Omega} \phi_k(y) dy \right) \ .$$

To focus further on this second integral, we expand the landscape function u in terms of the eigenfunctions ϕ_k.

$$\int_{\Omega} \phi_k dx = \langle 1, \phi_k \rangle = \langle -\Delta u, \phi_k \rangle = \langle u, -\Delta \phi_k \rangle = \lambda_k \langle u, \phi_k \rangle \ ,$$

where the first equality is simple since $\phi_k = 1 * \phi_k$ and the second is just substitution since $\Delta u = 1$ with zero boundary conditions. The third equality results from integrating by parts twice with zero boundary conditions, so each time the boundary term is vanished. And the last equality comes from substitution since $-\Delta \phi_k = \lambda_k \phi_k$. Finally, if we performed an eigenfunction expansion on u and compare it to the above equalities, we see:

$$\left(\int_0^{\infty} \int_{\Omega} \rho_t(x, y) dy dt \right) = u(x) \ .$$

Putting all of these pieces together results in the desired (8.10). For the first bound in (8.11), recall the first is because $|\mathbf{E}_{x,t}(y)| \leq 1|$ depending on how path dependent the integral is. $\qquad\square$

In the instance that the vector field is conservative, or close to it, then $|\mathbf{E}_{x,t}(y)| \approx 1|$, in which case (8.10) basically becomes the landscape inequality. But, the further the vector field is from being conservative, the smaller $|\mathbf{E}_{x,t}(y)|$ is, and therefore the more pessimistic the landscape inequality is; while the landscape inequality is still true, it is not necessarily informative when A is not (at least close to) conservative. Corollary 8.1.6 provides an improved bound where A is far from path independent, and that is at least as good as the landscape inequality.

Remark 8.1.7. This Remark serves to justify the statement made in (8.12). For the heat kernel $\rho_\Omega(t, x, y)$, we have:

$$\mathcal{H}(\rho) := \frac{\partial \rho}{\partial t} - \nabla \rho = \delta \ ,$$

such that δ is the dirac delta function. For fixed t and x, we may express $\rho_\Omega(t, x, y) = \sum_{k=1}^{\infty} c_k(t, x)\phi_k(y)$, where ϕ_k are the eigenfunctions of the Laplacian with zero boundary conditions. Applyinh $\mathcal{H}$ to this $\rho_\Omega(t, x, y)$ now results in:

$$\sum_{k=1}^{\infty} \left(\frac{\partial c_k(t, x)}{\partial t} + \lambda_k c_k(t, x) \right) \phi_k(y) = \delta_x(y) \ .$$

Integrating by both sides against $\phi_j(y)$ gives:

$$\frac{\partial c_{k,j}(t,x)}{\partial t} + \lambda_{k,j}c_{k,j}(t,x) = \int_\Omega \delta_x(y)\phi_j(y)dy \ .$$

Let $x \neq y$, then the right hand side is 0. We may write $c_{k,j}(t,x) = \sum_{k=1}^\infty f_{k,j}(t)\phi_k(x)$, and so we thus have:

$$\sum_{k=1}^\infty \left(f'_{k,m}(t) + \lambda_m f_{k,j}(t) \right) \phi_k(x) \ .$$

The eigenfunction is not identically zero for all k, and therefore by rearranging and solving the differential equation, it is revealed that $f_{k,j} = e^{\lambda_j t}$. Consequently, $c_k(t,x) = e^{\lambda_j t}\phi_j(x)$, which means $\rho_t(x,y) = \sum_{j=1}^\infty e^{\lambda_j t}\phi_j(x)\phi_j(y)$ as desired.

8.1.1 The Magnetic Schrödinger Equation

Though the focus has been the magnetic Laplacian operator, the (general) arguments made can be extended to the magnetic Schrödinger equation:

$$\hat{H}(A)\psi := (-\mathbf{i}\nabla - A(x))^2 \psi(x) + V(x)\psi(x)$$

for potential $V \in \mathbf{R}_{\geq 0}$. A similar Feynman-Kac representation is used:

$$\left(e^{-tH(A,V)}\psi\right)(x) = \mathbf{E}\left(e^{-S_t(A,V|\omega)}\chi_\Omega(\omega,t)\psi(\omega(t))\right) \ ,$$

where the expectation is taken over all the Brownian motion that started at x, as before, and the function $\chi_\Omega(\omega, t)$ is unchanged. The path integral now must take into account the added potential:

$$S_t(A, V|\omega) = \mathfrak{i} \int_0^t A(\omega(s))d\omega(s) + \frac{\mathfrak{i}}{2} \int_0^t (\nabla \cdot A)(\omega(s))ds + \int_0^t V(\omega(s))ds \ .$$

Using the same techniques as before, the $\text{div}(A)$ disappears with the Helmholtz decomposition and $V \equiv 0$, which proceeds in a very similar fashion as Theorem 8.1.2. However, the presence of V and it's interaction with A may significantly complicate the result as discussed before, and is not explored in this manuscript.

8.1.2   Neumann Boundary Conditions

Up until this point, our theory has considered only Dirichlet boundary conditions. Neumann boundary conditions are almost identical, since most of the theory does not directly utilize properties of the boundary conditions. In practice, *if* the eigenmodes being considered are localized inside of the domain and away from the boundary, by which we mean at least several wavelengths from the boundary, then there is very little difference when using Neumann boundary conditions. This is illustrated in the proof of the primary result, Theorem 8.1.2, where the Feynman-Kac is used in order to supply a reproducing identity of parabolic nature; for our purposes means that the modes have an exponential decay away from the boundary at the scale of a wavelength. Thus, whatever happens at the boundary is sufficiently small and therefore not a concern for the theory presented here.

Again, this assumes the eigenmodes are localized away from the boundary. The same cannot be said if they are localized near the boundary.

8.2   Numerical Examples

This section includes a few numerical examples that illustrate the utility of the primary theorem, which states that, for eigenvalues lower in the spectrum, eigenfunctions of $H(A)$ localized where curl$A$ is small. The examples are on the square $\Omega = [-1,1] \times [-1,1]$, uniformly meshed with piecewise, cubic polynomials and with edge length $h = 0.01$.

Example 8.2.1. First a simple example of a polynomial vector field. Let $A = -a(x^2 + y^2, x^2 - y^2)$ and $a = 1000$. Then curl$A = -2a(x - y)$. A stream plot of A along with a plot of $|\text{curl}A|$ and the real and imaginary parts of the first four eigenmodes are pictured in Figure 8.2. Though it isn't clear from A where the low energy eigenmodes will localize, the first four (pictured) modes are localized along the off diagonal where the curl is smallest. Furthermore, notice that the eigenvectors are oscillatory, even the ground state.

As suggested by Theorem 8.1.2, an eigenpair (λ, ψ) of $H(A)$ where $\psi(x_0) = ||\psi||_{L^\infty}$ has the expectation that most points y in some neighborhood of x_0, from the path integral x_0 to y, will be relatively concentrated around some fixed value when the vector field is either "close" to a conservative vector field or the integration path is short (when λ is large). This suggests that localization may occur where the magnitude of A is small compared to other regions, which can be observed in

Figure 8.2: Let $A = -a(x^2 + y^2, x^2 - y^2)$ and $a = 1000$ The plots of A and $|\mathrm{curl}A|$ in the top row. It is not clear from looking at A where the eigenmodes will localize, and the off diagonal is not not clearly an indicator where we see that the $\mathrm{curl}A$ is smallest. The second row contains the real (above) and imaginary (below) parts of the first four eigenfunctions together with their eigenvalues; they are localized only where $|\mathrm{curl}A|$ is small as expected. These eigenfunctions are oscillatory in their components, even the ground state.

practice. However, more complicated examples illustrate the paramount impact of the curl, and are considered in the next examples.

Example 8.2.2. Let $A = -a\left(\cos(f(x,y)), \sin(f(x,y))\right)$ where $f(x,y) = 5\pi \sin(x^2 + y^2)$ and $a = 50$. Figure 8.3 includes a stream plot of A and the $|\text{curl}A|$ together with the (modulus of the) first nine eigenmodes. The eigenmodes begin localizing in the center of the domain, where the $|\text{curl}A|$ is the smallest for the most area, and develop in this spiral, as is consistent with the $\text{curl}A$. Two eigenmodes even separately concentrate in the off diagonal corner, away from the center spiral, which is another set of regions where the curl is small. Since the curl is identical in both these corners, this is an example of two localized eigenfunctions corresponding to the same eigenvalue.

Example 8.2.3. Let $A = -a\left(\cos(f(x,y)), \sin(f(x,y))\right)$ where $f(x,y) = \pi \sin(\pi x) \cos(\pi y)$ and $a = 50$. Figure 8.4 contains a stream plot of A and the $|\text{curl}A|$ together with the (modulus of the) first nine eigenmodes. The regions where the $|\text{curl}A|$ are small, intricate curves throughout the domain, yet the eigenfunctions can clearly be seen localizing throughout these same regions. It is not clear from A where these early eigenvectors may concentrate, but it is from observing the $|\text{curl}A|$ plot, which further validates our theory.

Figure 8.3: Let $A = -a \left(\cos(f(x,y)), \sin(f(x,y)) \right)$ where $f(x,y) = 5\pi \sin(x^2 + y^2)$ and $a = 50$. A stream plot of A and the $|\mathrm{curl} A|$ together with the (modulus of the) first nine eigenmodes. The eigenfunctions concentrate in the spiral or the off-diagonal corners where the $|\mathrm{curl} A|$ is smallest.

Figure 8.4: Let $A = -a\left(\cos(f(x,y)), \sin(f(x,y))\right)$ where $f(x,y) = \pi \sin(\pi x)\cos(\pi y)$ and $a = 50$. The regions where the $|\mathrm{curl}A|$ are small, intricate curves throughout the domain, yet the eigenfunctions can clearly be seen localizing throughout these same regions.

Bibliography

[1] G. Alessandrini. On Courant's nodal domain theorem. *Forum Math.*, 10(5):521–532, 1998.

[2] S. Aljawi and M. Marletta. On the eigenvalues of spectral gaps of matrix-valued Schrödinger operators. *Numer. Algorithms*, 86(2):637–657, 2021.

[3] R. Altmann, P. Henning, and D. Peterseim. Quantitative Anderson localization of Schrödinger eigenstates under disorder potentials. *Math. Models Methods Appl. Sci.*, 30(5):917–955, 2020.

[4] R. Altmann and D. Peterseim. Localized computation of eigenstates of random Schrödinger operators. *SIAM J. Sci. Comput.*, 41(6):B1211–B1227, 2019.

[5] P. W. Anderson. Local moments and localized states. *Rev. Mod. Phys.*, 50:191–201, Apr 1978.

[6] D. N. Arnold, G. David, M. Filoche, D. Jerison, and S. Mayboroda. Computing spectra without solving eigenvalue problems. *SIAM J. Sci. Comput.*, 41(1):B69–B92, 2019.

[7] D. N. Arnold, G. David, M. Filoche, D. Jerison, and S. Mayboroda. Localization of eigenfunctions via an effective potential. *Comm. Partial Differential Equations*, 44(11):1186–1216, 2019.

[8] D. N. Arnold, G. David, D. Jerison, S. Mayboroda, and M. Filoche. Effective confining potential of quantum states in disordered media. *Phys. Rev. Lett.*, 116:056602, Feb 2016.

[9] K.-J. Bathe. The subspace iteration method – revisited. *Computers & Structures*, 126:177–183, 2013. Uncertainty Quantification in structural analysis and design: To commemorate Professor Gerhart I. Schueller for his life-time contribution in the area of computational stochastic mechanics.

[10] K. Broderix, D. Hundertmark, and H. Leschke. Continuity properties of Schrödinger semigroups with magnetic fields. *Rev. Math. Phys.*, 12(2):181–225, 2000.

[11] L. A. Bunimovich. On the ergodic properties of nowhere dispersing billiards. *Communications in Mathematical Physics*, 65(3):295 – 312, 1979.

[12] S. Felix, M. Asch, M. Filoche, and B. Sapoval. Localization and increased damping in irregular acoustic cavities. *JOURNAL OF SOUND AND VIBRATION*, 299(4-5):965–976, FEB 6 2007.

[13] M. Filoche and S. Mayboroda. Universal mechanism for Anderson and weak localization. *Proc. Natl. Acad. Sci. USA*, 109(37):14761–14766, 2012.

[14] W. Ford. *Numerical linear algebra with applications*. Elsevier/Academic Press, Amsterdam, 2015. Using MATLAB.

[15] N. Garofalo and F.-H. Lin. Unique continuation for elliptic operators: a geometric-variational approach. *Comm. Pure Appl. Math.*, 40(3):347–366, 1987.

[16] J. Gopalakrishnan, L. Grubišić, and J. Ovall. Spectral discretization errors in filtered subspace iteration. *Math. Comp.*, 89(321):203–228, 2020.

[17] J. Gopalakrishnan, L. Grubišić, J. Ovall, and B. Parker. Analysis of FEAST spectral approximations using the DPG discretization. *Comput. Methods Appl. Math.*, 19(2):251–266, 2019.

[18] D. S. Grebenkov and B.-T. Nguyen. Geometrical structure of Laplacian eigenfunctions. *SIAM Rev.*, 55(4):601–667, 2013.

[19] S. Güttel, E. Polizzi, P. T. P. Tang, and G. Viaud. Zolotarev quadrature rules and load balancing for the FEAST eigensolver. *SIAM J. Sci. Comput.*, 37(4):A2100–A2122, 2015.

[20] L. Hörmander. Uniqueness theorems for second order elliptic differential equations. *Comm. Partial Differential Equations*, 8(1):21–64, 1983.

[21] J. Hoskins, H. Quan, and S. Steinerberger. Magnetic Schrödinger operators and landscape functions. *arXiv Preprint*, 10 2022.

[22] Jay Gopalakrishnan and Benjamin Parker. Pythonic feast.

[23] T. Jiang, A. Fang, Z.-Q. Zhang, and C. T. Chan. Anomalous anderson localization behavior in gain-loss balanced non-hermitian systems. *Nanophotonics*, 10(1):443–452, 2021.

[24] Joachim Schöberl. Netgen/ngsolve.

[25] T. Kato. *Perturbation theory for linear operators*. Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1980 edition.

[26] A. Lagendijk, B. Tiggelen, and D. Wiersma. Fifty years of anderson localization. *Phys. Today*, 62:24–29, 08 2009.

[27] R. J. Lipton, D. J. Rose, and R. E. Tarjan. Generalized nested dissection. *SIAM J. Numer. Anal.*, 16(2):346–358, 1979.

[28] J. Lu, C. Murphey, and S. Steinerberger. Fast localization of eigenfunctions via smoothed potentials. *J. Sci. Comput.*, 90(1):Paper No. 38, 18, 2022.

[29] P. Markoš. Numerical analysis of the anderson localization. *Acta Physica Slovaca. Reviews and Tutorials*, 56(5), oct 2006.

[30] M. Marletta and R. Scheichl. Eigenvalues in spectral gaps of differential operators. *J. Spectr. Theory*, 2(3):293–320, 2012.

[31] J. H. Mathews and K. K. Fink. *Numerical Methods Using Matlab (4th Edition)*. Pearson, 4 edition, jan 2004.

[32] J. S. Ovall, H. Quan, R. Reid, and S. Steinerberger. On localization of eigenfunctions of the magnetic Laplacian. *arXiv Preprint*, 2023.

[33] J. S. Ovall and R. Reid. An algorithm for identifying eigenvectors exhibiting strong spatial localization. *Math. Comp.*, 92(341):1005–1031, 2023.

[34] B. Sapoval, M. Filoche, M. Chappat, and D. Peyrard. Noise abatement wall, 09 2003. Publication No. WO/2003/078740; International Application No. PCT/FR2003/000881; Application filed by Ecole Polytechnique and COLAS Corporation.

[35] S. Steinerberger. Regularized potentials of Schrödinger operators and a local landscape function. *Comm. Partial Differential Equations*, 46(7):1262–1279, 2021.

[36] D. Thouless. Electrons in disordered systems and the theory of localization. *Physics Reports*, 13(3):93–142, 1974.

[37] M. Van Barel. Designing rational filter functions for solving eigenvalue problems by contour integration. *Linear Algebra and its Applications*, 502:346–365, 2016. Structured Matrices: Theory and Applications.

[38] G. Van Rossum and F. L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

[39] Wolfram Research, Inc. Mathematica 14.0.

APPENDICES

Appendix A: Obtaining true eigenpairs for a 1D pathological case

We have:

$$\begin{cases} -\psi''(x) + V(x)\psi(x) = \lambda\psi(x) \in (0,1) \\[2mm] \psi(0) = \psi(1) = 0 \\[2mm] V|_{[1/4,3/4]} = 120^2 \\[2mm] V = 0, \ otherwise \end{cases}$$

We can construct the eigenfunctions in three parts $\psi = \psi_1 + \psi_2 + \psi_3$, on the intervals $[0, \frac{1}{4}), [\frac{1}{4}, \frac{3}{4}]$, and $(\frac{3}{4}, 1]$. The zero boundary conditions force ψ_1 and ψ_3 to be sine functions, and then assuming this ψ is even centered at $\frac{1}{2}$ forces ψ_2 to be a cosine function. Thus:

$$\begin{cases} \psi_1(x) = \sin(\sqrt{\lambda}x) \quad \in [0, 1/4) \\[2mm] \psi_2(x) = a\cos\left(\sqrt{\lambda - 120^2}(x - \frac{1}{2})\right) \quad \in [1/4, 3/4] \\[2mm] \psi_3(x) = \sin\left(\sqrt{\lambda}(1 - x)\right) \quad \in (3/4, 1] \end{cases}$$

where a is the constant to be found in order to account for the continuity of ψ at $\frac{1}{4}$ and $\frac{3}{4}$. Finding a is as simple as satisfying $\psi_1(1/4) = \psi_2(1/4)$ and $\psi_2(3/4) = \psi_3(3/4)$, which gives:

$$a = \frac{\sin(\sqrt{\lambda})/4}{\cos(\sqrt{\lambda - 120^2})/4}$$

We know $\psi \in H_0^1([0,1])$, so $\psi'(x)$ must also be continuous. This is as simple as satisfying $\psi_1'(1/4) = \psi_2'(1/4)$ and $\psi_2'(3/4) = \psi_3'(3/4)$ which, for example, is equivalent to determining when

$$f(\lambda) := 1 - \frac{\psi_2'}{\psi_1'} = 0 \ .$$

Notice the roots of $f(\lambda)$ are the eigenvalues, and they were roughly identified by plotting $f(\lambda)$. Then they were found to 200 digits of accuracy using Newton's method with the rough estimates as initial guesses. Lastly, the eigenfunctions were then computed directly.

If the eigenfunction is instead odd, then this process is the same with few notable exceptions. The middle ψ_2 must now be odd, represented by the sine function about $\frac{1}{2}$:

$$\psi_2(x) = b\frac{\sin\left(\sqrt{\lambda - 120^2}(x - 1/2)\right)}{\sqrt{\lambda - 120^2}} \ ,$$

where we divide by $\sqrt{\lambda - 120^2}$ to ensure the ψ_2 is real for $\lambda \leq 120^2$, since $\sin(\mathrm{i}x) = \mathrm{i}\sinh(x)$. The constant b serves the same purpose as a before:

$$b = -\sqrt{\lambda - 120^2}\frac{\sin\left(\frac{\sqrt{\lambda}}{4}\right)}{\sin\left(\frac{\sqrt{\lambda - 120^2}}{4}\right)}$$

The rest of the process is the same as before.

Appendix B: Geometric Localization Figures

Additional eigenpairs of $\mathcal{L}$ are included for the variations of the Threebulb domain that were discussed in Chapter 6. The reader can review the effects of these geometric changes for higher energies.

$\lambda = 8.3618$	9.3829	10.0287	10.2838	10.4845
$\tau_L = 1.8\text{e-}3$	0.134	0.836	0.305	0.999
$\tau_M = 0.035$	0.794	0.212	0.258	0.005
$\tau_R = 0.971$	0.118	0.173	0.795	0.008

$\lambda = 10.6460$	10.8446	10.9606	11.5621	11.7670
$\tau_L = 0.706$	0.725	0.015	0.741	0.321
$\tau_M = 0.402$	0.487	0.019	0.307	0.552
$\tau_R = 0.362$	0.293	0.999	0.261	0.530

$\lambda = 12.3170$	12.3350	13.2125	14.0434	14.2463
$\tau_L = 0.794$	0.957	0.067	0.750	0.006
$\tau_M = 0.385$	0.184	0.385	0.373	0.007
$\tau_R = 0.0146$	0.006	0.763	0.099	0.995

$\lambda = 15.2686$	15.4168	15.6731	16.0249	16.0409
$\tau_L = 0.188$	0.999	0.798	0.964	0.960
$\tau_M = 0.465$	0.009	0.390	0.180	0.192
$\tau_R = 0.605$	0.007	0.108	0.021	0.021

$\lambda = 17.3097$	17.5355	17.8693	18.5886	18.6220
$\tau_L = 0.765$	0.385	0.999	0.697	0.260
$\tau_M = 0.486$	0.456	0.018	0.404	0.139
$\tau_R = 0.240$	0.666	0.013	0.465	0.946

$\lambda = 18.9805$	19.7398	20.0948
$\tau_L = 0.314$	0.696	0.650
$\tau_M = 0.247$	0.348	0.296
$\tau_R = 0.874$	0.522	0.610

Figure 5: The 13th - 40th eigenfunctions of the ThreeBulb domain along with the eigenvalues and τ values for each bulb. As the energy is increased, there are less concentrated eigenmodes since the smaller, more frequent oscillations of the higher energy eigenmodes can more easily fit through the narrow channels of the bridge.

$\lambda = 8.5070$	8.7288	9.4620	9.5199	10.4797
$\tau_L = 0.077$	8.2e-6	0.351	0.846	0.999
$\tau_M = 0.852$	7e-5	0.087	0.143	7.5e-7
$\tau_R = 0.071$	0.999	0.811	0.351	3.1e-7

$\lambda = 10.9152$	10.9645	11.2396	11.8680	12.1255
$\tau_L = 0.946$	0.009	0.430	0.372	4.2e-5
$\tau_M = 0.091$	0.302	0.292	0.302	0.999
$\tau_R = 0.061$	0.999	0.282	0.361	0.001

$\lambda = 12.2729$	12.3370	13.3577	13.9365	14.1452
$\tau_L = 0.999$	1.0	0.501	0.470	4.4e-5
$\tau_M = 1.7e\text{-}5$	3.7e-7	0.353	0.204	2.1e-5
$\tau_R = 5.7e\text{-}6$	3.2e-7	0.361	1.9e-4	5.7e-6

$\lambda = 14.7980$	15.3554	15.6000	16.0364	16.0755
$\tau_L = 0.456$	0.999	0.947	0.997	0.561
$\tau_M = 0.341$	4.8e-6	0.109	0.031	0.362
$\tau_R = 3.2e\text{-}7$	6.6e-6	0.195	0.044	0.496

$\lambda = 17.4546$	17.8222	17.9061	18.6083	18.8036
$\tau_L = 0.580$	0.999	0.928	4.1e-5	6.6e-4
$\tau_M = 0.344$	1.6e-5	0.139	7.8e-5	0.995
$\tau_R = 0.482$	5.5e-7	0.216	0.999	7.1e-4

$\lambda = 19.0763$	19.5712	19.6143
$\tau_L = 0.519$	0.312	0.999
$\tau_M = 0.346$	0.170	5.6e-5
$\tau_R = 0.588$	0.888	1.6e-4

Figure 6: The 13th - 40th eigenpairs and τ values for each bulb of the ThreeBulb domain with centered bridges. The localization is decreased significantly, particularly for the eigenfunctions with an odd number of oscillations in the y-direction and as the energy increases, resulting in a higher number of oscillations that create small pockets of eigenfunction that can easily bleed into the bridge. Boundary conditions also play some role in the decreased localization.

$\lambda = 8.7506$	10.4781	10.4864	10.9408
$\tau_L = 0$	0.999	0.999	0
$\tau_M = 3.5\text{e-}7$	4.6e-7	5.7e-9	1.1e-6
$\tau_R = 0.999$	0	0	0.999
$\lambda = 10.9662$	11.0831	12.1903	12.2301
$\tau_L = 0$	0.999	0	0
$\tau_M = 1.3\text{e-}8$	1.1e-6	0.999	0.999
$\tau_R = 1$	0	1.9e-6	1.7e-6
$\lambda = 12.3054$	12.3370	14.1559	14.2561
$\tau_L = 0.999$	0.999	0	0
$\tau_M = 2\text{e-}4$	1.5e-7	3.2e-6	2.3e-8
$\tau_R = 0$	0	0.999	1.0

Figure 7: The 13th - 24th eigenfunctions paired with their eigenvalues and τ values of the Threebulb domain with thinned bridges. The eigenfunctions behave (fairly) independently in each bulb, being very concentrated, but at the cost of higher eigenvalues. The required energy to obtain localized bulbs increases at high energies when compared to the corresponding eigenfunction of the original Threebulb domain.

$\lambda = 8.3610$	9.8222	9.9863	10.2346	10.3521
$\tau_L = 1.5\text{e-}4$	0.457	0.341	0.501	0.122
$\tau_M = 0.045$	0.654	0.706	0.716	0.583
$\tau_R = 0.972$	0.089	0.302	0.312	0.744

$\lambda = 10.6043$	10.9600	11.0224	11.5700	12.1272
$\tau_L = 0.137$	0.013	0.217	0.229	0.408
$\tau_M = 0.950$	0.048	0.790	0.552	0.602
$\tau_R = 0.123$	0.998	0.330	0.481	0.281

$\lambda = 12.3159$	12.3262	12.9799	14.0291	14.2454
$\tau_L = 0.995$	0.068	0.121	0.367	0.010
$\tau_M = 0.083$	0.997	0.500	0.716	0.045
$\tau_R = 0.022$	0.024	0.694	0.232	0.998

$\lambda = 14.3851$	15.6049	15.8248	16.0232	16.0864
$\tau_L = 0.085$	0.246	0.347	0.037	0.180
$\tau_M = 0.751$	0.906	0.693	0.989	0.826
$\tau_R = 0.492$	0.130	0.357	0.098	0.369

$\lambda = 17.1447$	17.6100	18.1793	18.5952	18.6294
$\tau_L = 0.298$	0.284	0.254	0.342	0.232
$\tau_M = 0.808$	0.844	0.776	0.538	0.376
$\tau_R = 0.368$	0.337	0.472	0.702	0.869

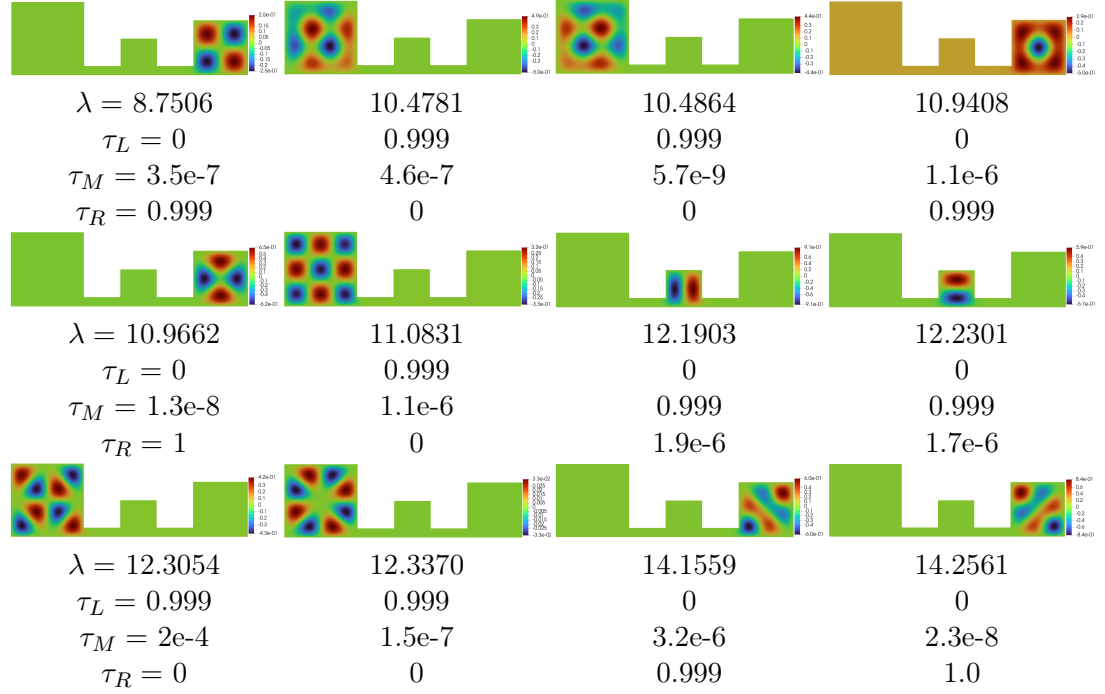| $\lambda = 19.0997$	19.7392	20.1988
$\tau_L = 0.186$	0.348	0.405
$\tau_M = 0.491$	0.696	0.652
$\tau_R = 0.810$	0.522	0.528

Figure 8: The 13th - 40th eigenpairs together with their τ values in each bulb of the Threebulb domain where the large and small bulbs are switched. The large bulb is still the primary host of the eigenmodes, however since it is now in the middle of two bridges, these eigenmodes are not as concentrated, esp with higher energy. In contrast, the small bulb now only has one channel for possible bleeding, so these eigenfuctions are more localized.