

The Laplacian: An Exploration and Historical Survey Tailored for Translation Surfaces

Joshua Southerland

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

Master of Science

University of Washington
2019

Committee:
Jayadev Athreya
Bobby Wilson

Program Authorized to Offer Degree:
Mathematics

©Copyright 2019

Joshua Southerland

University of Washington

Abstract

The Laplacian: An Exploration and Historical Survey Tailored for Translation Surfaces

Joshua Southerland

Chair of the Supervisory Committee:

Jayadev Athreya

Department of Mathematics

This thesis is a historical survey of the Laplacian as an operator on L^2 -functions specifically geared towards building the understanding necessary to define a Laplacian on a translation surface. The author explores the role the Laplacian has played historically in analysis and geometry, with a particular interest in the connections between the Laplacian and the geodesics. The primary thread the author follows develops a representation-theoretic perspective of the Laplacian, which proves advantageous when working on symmetric spaces. The other appeals to a functional-analytic perspective in more abstract settings. In the final section, the author proposes a starting point for defining a Laplacian on a translation surface.

THE LAPLACIAN: AN EXPLORATION AND HISTORICAL SURVEY TAILORED FOR TRANSLATION SURFACES

JOSH SOUTHERLAND

CONTENTS

1.	Introduction	5
2.	Definition and First Properties	5
3.	The Laplacian on Graphs	11
3.1.	A discrete version of Harmonic Functions	11
3.2.	A discrete version of the Laplacian	14
3.3.	The Laplacian as an Operator	15
4.	Fourier Transform on \mathbb{R}^n	16
5.	Fourier Series on \mathbb{T} and a Decomposition of $L^2(\mathbb{T})$	21
5.1.	Fourier Series and Poisson Summation	21
5.2.	A Representation-Theoretic Viewpoint	22
5.3.	Translation Invariant Subspaces of $L^2(\mathbb{T})$	22
5.4.	Eigenfunctions of the Laplacian	23
5.5.	Homomorphisms of the Convolution Algebra on $L^2(\mathbb{T})$	23
5.6.	Remark on the finite abelian case	23
6.	Key Example: \mathbb{T}^2	25
6.1.	A Representation-Theoretic Viewpoint	25
6.2.	Translation Invariant Subspaces of $L^2(\mathbb{T}^2)$	25
6.3.	Eigenfunctions of the Laplacian	25
6.4.	Homomorphisms of the Convolution Algebra on $L^2(\mathbb{T}^2)$	25
7.	Issues with the non-abelian case	26
8.	Some Representation Theory: Compact Groups and the Peter-Weyl Theorem	26
8.1.	Definitions and Schur's Lemma	27
8.2.	Schur's Orthogonality Relations	30
8.3.	The Peter-Weyl Theorem	30
8.4.	Remark on Locally Compact Groups	31
9.	Symmetric Spaces	31
9.1.	Eigenfunctions of the Laplacian on a Compact Riemannian Manifold	31
9.2.	The flat plane, \mathbb{R}^2	38
9.3.	The sphere, S^2	38
9.4.	The Hyperbolic Plane, \mathbb{H}	38
10.	Connections to Geodesics	38
10.1.	Eigenvalues and Isospectrality	38
10.2.	Huber's Theorem	38
10.3.	Sunada's Theorem	39
10.4.	Can you hear the shape of a Translation Surface?	39
10.5.	Eigenfunctions and Quantum Ergodicity	39

11. Terminus a Quo: Laplacian on a Translation Surface	40
11.1. Translation Surface as a Polygon	40
11.2. Translation Surface as a Riemannian Manifold with Cone Points	40
12. Conclusion and Next Steps	40
References	41

1. INTRODUCTION

During the Autumn quarter of 2018, Professor Jayadev Athreya posed the following question: *How would one go about defining a Laplacian for a Translation Surface? Can this be defined in a way such that one can recover information about geodesics on the surface?* Given that translation surfaces are “flat surfaces” with a flat metric, one might believe that the definition of the Laplacian is obvious, and in fact, already well-understood. However, translation surfaces contain singularities, which are an obstruction to the usual definition. Each translation surface is topologically a compact, orientable surface (see [35]), identifiable by a genus [23]. By applying the Gauss-Bonnet Theorem [21], we find they will have negative curvature for genus 2 or greater. Hence, even though the translation surface has a flat metric, for genus 2 or higher, we are guaranteed singularities where the negative curvature accumulates. It is these points that are of interest.

We could begin answering Professor Athreya’s question by naively defining the Laplacian as the operator which has an averaging property, giving us a mean value theorem when we integrate loops around singularities. However, there are many subtleties that would need to be investigated. First, we would need to notice that the loops surrounding a singularity have “angle” greater than 2π . Second, we would need to notice that a loop could contain more than one singularity. Does the existence of multiple singularities effect the averaging? Should it?

One might think a better approach would be to define the Laplacian more abstractly, where the mean value theorem is a consequence. We might consider, for instance, using the definition of the Laplacian on a manifold. However, one would have to be wary: these translation surfaces, while (almost) smooth manifolds, do not have a smooth structure at the singularities. These are manifolds with singularities, so one would need to consider a definition of the Laplacian on cone manifolds. And in order to effectively evaluate a definition, one would need to understand what results to expect from the definition.

While both of these techniques would yield insight, we must first understand the purpose of defining the Laplacian. What value comes from having a Laplacian on a surface? What roles has the Laplacian played? We start instead with an historical survey, reviewing some of the elementary roles the Laplacian has played in analysis, geared towards understanding connections to geodesics. Our journey will inevitably lead us towards a representation-theoretic perspective of the Laplacian, and this will become the primary thread we will follow. We will not leave the spectral-theoretic perspective behind though. In fact, the further we progress, the more abstract spaces become (and less obvious it becomes where symmetries are hidden), the more spectral methods will play an important role in establishing properties of the Laplacian. Along our journey, we will build an understanding of what to expect from a Laplacian on a translation surface, and look for hints that may guide us to the most fruitful generalization of the Laplacian. In the final section, we propose a starting point for the definition, following the work of Cheeger [5] [6].

2. DEFINITION AND FIRST PROPERTIES

The first definition of a Laplacian that one normally sees is in a Multivariable Calculus course, where the Laplacian of a twice differentiable real-valued function is defined as the divergence of the gradient of

the function [28]. We will begin with the same definition, but with one modification. We will follow the geometer's convention of taking the negative [25].

Definition 2.1 (Laplacian on \mathbb{R}^n). For a twice differentiable, real-valued function on \mathbb{R}^n we define the Laplacian, Δ , as

$$\Delta f = -\operatorname{div}(\operatorname{grad}(f)). \quad (1)$$

For twice differentiable functions on \mathbb{R}^n , we can compute the following.

$$\Delta f = -\operatorname{div}(\operatorname{grad}(f)) = -\sum_{i=1}^n \frac{\partial^2 f}{\partial x_i^2} \quad (2)$$

For some of us, this is the end of the story. Just another way a Calculus instructor can ask us to take two derivatives. For others, the Laplacian reappears in a first course on Partial Differential Equations, where one is introduced to Laplace's equation

$$\Delta u = 0.$$

Solutions to Laplace's equations are called harmonic functions [30]. The astute student will recognize that solutions to Laplace's equation, these harmonic functions, have an interesting geometric property. At every point, when you sum second derivatives in each coordinate direction summed vanishes: this is probably easiest visualized in two dimensions, where the concavity of the function in the x-direction must cancel with the concavity of the function in the y-direction.




FIGURE 1. Saddle Point

In other words, every point on the surface is a saddle point, and as calculus students know, saddle points cannot be local extrema, so it is not a stretch to postulate that if the maximum or minimum values of such functions exist, they must lie on the boundary of the domain. This is indeed correct, and called the *Maximum Principle*. It turns out however, the easiest way to prove this statement is to recognize another key characteristic of harmonic functions. The astute student may recognize that the increase and decrease of a function away from a point must be somehow balanced since second derivatives cancel, so one could postulate that the value at a point is precisely the average of the function along a curve surrounding that point. This, too, is true, and called the *Mean Value Property*. The mean value property and corresponding maximum principle enable us to characterize the Laplacian in a different way, namely as an "averaging operator" on functions. The Laplacian is measuring how much a function deviates from satisfying the mean value property.

Since we have defined harmonic functions on \mathbb{R}^2 above, we give a more thorough statement below, and prove them using the machinery taught in an undergraduate complex analysis course [27]. As Hadamard's quote [15] goes, "It has been written that the shortest and best way between two truths of the real domain often passes through the imaginary one." The reader should note that while these proofs work for \mathbb{R}^2 , the following theorems hold for \mathbb{R}^n , but the proofs necessarily take a different form.

Theorem 2.1 (Mean Value Property). Let $u : \mathbb{C} \rightarrow \mathbb{R}$ be a harmonic function on an open domain $\Omega \subset \mathbb{C}$. For any $z_0 \in \Omega$, u satisfies the mean-value property:

$$u(z_0) = \int_0^1 u(z_0 + re^{2\pi it}) \text{ for any } r < \text{dist}(z_0, \partial\Omega) \quad (3)$$

Proof. Fix $z_0 \in \Omega$, let $r < \text{dist}(z_0, \partial\Omega)$. Let $B_R(z_0) \subset \Omega$ be an open ball with radius $r < R < \text{dist}(z_0, \partial\Omega)$. $B_R(z_0)$ is simply-connected, so u has a harmonic conjugate in $B_R(z_0)$. Let $f = u + iv$, where v is the harmonic conjugate, and we have that f is a holomorphic function. Let γ be the boundary of a ball of radius r about z_0 traversed counterclockwise. Parametrize this by the curve $z_0 + re^{2\pi it}$ for $0 \leq t \leq 1$. We can then use Cauchy's integral formula:

$$f(z_0) = \frac{1}{2\pi i} \int_{\gamma} \frac{f(z)}{z - z_0} dz = \int_0^1 f(z_0 + re^{2\pi it}) dt.$$

Take the real part of $f(z_0)$ and this completes the proof.

$$u(z_0) = \Re(f(z_0)) = \Re\left(\int_0^1 f(z_0 + re^{2\pi it}) dt\right) = \int_0^1 u(z_0 + re^{2\pi it}) dt. \quad \blacksquare$$

Corollary 2.1 (Maximum Principle). For u as above, if u attains a local minimum or maximum on the interior of Ω , then u is constant. In particular, if Ω is an open, bounded and $u \in C(\overline{\Omega})$, then

$$\min_{\tilde{z} \in \partial\Omega} u(\tilde{z}) \leq u(z) \leq \max_{\tilde{z} \in \partial\Omega} u(\tilde{z}) \text{ for any } z \in \Omega,$$

where we achieve equality if and only if u is constant.

Proof. For the first statement, assume otherwise. First assume that u attains maximum on the interior of Ω . Let z_0 be a point where u achieves its maximum. Then, it follows from the mean value theorem that u must be constant on a sufficiently small neighborhood of z_0 . Otherwise, if there is a point in the neighborhood \tilde{z} such that $u(\tilde{z}) < u(z_0)$, then let $r = \text{dist}(z_0, \tilde{z})$ and notice

$$u(z_0) = \int_0^1 u(z_0 + re^{2\pi it}) dt = \int_{\partial B_r(z_0)} u(z) dz.$$

Now notice that $\tilde{z} \in \partial B_r(z_0)$. Since u is harmonic, u is continuous, there exists a connected neighborhood U of \tilde{z} such that that for any $z \in U$, $u(z) < u(z_0)$. $U \cap \partial B_r(z_0)$ is nonempty, open and connected in the subspace topology of $\partial B_r(z_0)$. Thus, we can split the integral into two parts and compute.

$$u(z_0) = \int_{\partial B_r(z_0) \setminus (U \cap \partial B_r(z_0))} u(z) dz + \int_{U \cap \partial B_r(z_0)} u(z) dz < u(z_0).$$

Thus, we see a contradiction unless u is constant on any disk of radius r centered at the maximum in Ω . Since any two points in the connected domain Ω have a path, we can fatten the path to show any two points are contained in a simply-connected domain, and use the argument above (with minimal modifications) to show u must be constant on a disk in this domain. Since the domain is simply connected, there exists a harmonic conjugate to u , and thus a holomorphic function f whose real part is u . Apply the uniqueness theorem to the function f , and we see that u must be constant on the entire simply-connected domain. Hence, we can conclude u must be globally constant. If we assume u has a minimum, then $-u(z)$ is a harmonic function with a maximum. Apply the argument above, and we see that the u cannot contain a minimum unless u is constant, as desired. \blacksquare

Next, for students who took that course in Partial Differential Equations, they saw integration by parts at work. In particular, they may have been asked to show the following relationship. Let f and g be functions with continuous second derivatives, and the first integral below be integrable. Then,

$$\int_{\mathbb{R}^2} g(x,y)\Delta f(x,y)dx dy = - \int_{\mathbb{R}^2} \nabla g(x,y) \cdot \nabla f(x,y)dx dy = \int_{\mathbb{R}^2} \Delta g(x,y)f(x,y)dx dy. \quad (4)$$

The proof of the equation above works in any dimension: use integration by parts, and recall that for any function f to be integrable on \mathbb{R}^n , we must have that $f(x) \rightarrow 0$ as $x \rightarrow \infty$.

The importance of such a formula may be lost on students when they first encounter it (it was lost on me). It certainly says "under nice circumstances, you can move the Laplacian from one function to another." But it says much more. If you realize the Laplacian is a differential operator and equip this set of functions with continuous second derivatives in $L^2(\mathbb{R}^n)$ with the usual L^2 norm, you have just shown that the Laplacian is self-adjoint:

$$\langle g, \Delta f \rangle_{L^2} = \langle \Delta g, f \rangle_{L^2}. \quad (5)$$

This fact turns out to have important consequences in the spectral theory of the operator. We see this in Section 9 below, but it is hidden by the fact that we are using the fact that the heat kernel is self-adjoint (the heat kernel and Laplacian share a very close relationship).

Next, we solve a variant of Laplace's equation:

$$\Delta f = \lambda f. \quad (6)$$

This is the eigenvalue problem for the Laplacian. We think of f as a vector in $L^2(\mathbb{R})$, and Δ as a linear operator. The solution to this equation only requires elementary theory from Ordinary Differential Equations (but do not forget the negative in the definition of the Laplacian!). With a bit of normalization, the reason for which will be apparent in a moment, we see that $f(x) = e^{-2\pi i \gamma x}$ for any $\gamma \in \mathbb{R}$, and the corresponding eigenvalue is $\lambda_\gamma = 4\pi^2 \gamma^2$. We call the function $e^{-2\pi i \gamma x}$ an eigenfunction with eigenvalue λ . One should notice that the functions $e^{-2\pi i \gamma x}$ and $e^{2\pi i \gamma x}$ both have the same eigenvalues: they are both eigenfunctions for the same eigenvalues. In fact, these functions are reflections of each other. If $f(x) = e^{-2\pi i \gamma x}$, then $f(-x) = e^{2\pi i \gamma x}$.

Next, we ask if we can solve the same equation, only for functions on $[0, 1]$, where we impose $f(0) = f(1)$. In other words, we are looking for functions supported on the circle, or torus, $\mathbb{T} = \mathbb{R}/\mathbb{Z}$. These types of functions can be extended continuously to all of \mathbb{R} where we recognize them as periodic functions of period 1. It should not come as a surprise that we can look at the solutions above and restrict to those that are supported on \mathbb{T} , namely the eigenfunctions that are periodic with period 1. Keeping with our normalization above, the solutions are of the form $f(x) = e^{-2\pi i n x}$, where $n \in \mathbb{Z}$. The corresponding eigenvalues are $\lambda_n = 4\pi^2 n^2$. As above, we should notice that $e^{-2\pi i n x}$ and $e^{2\pi i n x}$ are both eigenfunctions with the same eigenvalue.

We now begin to stray a bit beyond the usual undergraduate math education. Here we will give a definition of the Laplace-Beltrami operator [21] (actually the negative of the Laplace-Beltrami operator, keeping with our convention). This is a generalization of the Laplacian defined above, but on smooth manifolds with a Riemannian metric g of any dimension. For the advanced reader, you may recall that the usual Laplace-Beltrami operator (also called the connection Laplacian) has an expression in terms of the Levi-Cevita connection on a Riemannian manifold, $\Delta = \text{tr}(\nabla^2)$. The Bochner-Laplacian on a vector bundle with a fiber metric and compatible connection can be expressed as $\Delta = -\text{tr}(\nabla^2)$, so the following definition may be more appropriately called a special case of the Bochner-Laplacian.

Definition 2.2 ((Negative of the) Laplace-Beltrami Operator). For a real-valued function with continuous second derivatives on a smooth Riemannian manifold (M, g) , we define the *Laplace-Beltrami* operator as follows.

$$\Delta f = -\operatorname{div}(\operatorname{grad} f)$$

where $\operatorname{div}(X)$ for a vector field X and dV_g volume form associated with the metric g is

$$\operatorname{div}(X)dV_g = d(X \lrcorner dV_g)$$

and where $\operatorname{grad}(f)$ is

$$\operatorname{grad}(f) = df^\sharp.$$

If the metric is given by $g_{ij}dx^i dx^j$ in smooth coordinates (x^i) .

$$\operatorname{grad}(f) = g^{ij} \frac{\partial f}{\partial x_i} \frac{\partial}{\partial x_j}.$$

One can show that on a Riemannian manifold (M, g) , we have the following coordinate representation for the Laplacian:

Proposition 2.1 (Coordinate Representation of the Laplace-Beltrami Operator). Let (M, g) be a Riemannian Manifold with or without boundary, and let (x^i) be any smooth local coordinates on an open set $U \subset M$. The coordinate representation of the Laplacian is as follows:

$$\Delta f = -\frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} \left(g^{ij} \sqrt{\det g} \frac{\partial f}{\partial x^j} \right) \quad (7)$$

where $\det g = \det(g_{kl})$ is the determinant of the component matrix of g in these coordinates. On \mathbb{R}^n , with the Euclidean metric and standard coordinates, this reduces to:

$$\Delta f = \sum_{i=1}^n \frac{\partial^2 f}{(\partial x^i)^2}. \quad (8)$$

Proof. We will use the definition of the divergence given above. In smooth local coordinates, we may assume $dV_g = \sqrt{\det g} dx^1 \wedge \cdots \wedge dx^n$. Computing, we find

$$\begin{aligned} \sum_{i=1}^n X^i \frac{\partial}{\partial x^i} \lrcorner dV_g &= \sum_{i=1}^n X^i \frac{\partial}{\partial x^i} \lrcorner \sqrt{\det g} dx^1 \wedge \cdots \wedge dx^n \\ &= \sum_{j=1}^n (-1)^{j-1} dx^j \left(\sum_{i=1}^n X^i \frac{\partial}{\partial x^i} \right) \sqrt{\det g} dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^n \\ &= \sum_{j=1}^n (-1)^{j-1} X^j \sqrt{\det g} dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^n \end{aligned}$$

Then,

$$\begin{aligned}
d\left(\sum_{i=1}^n X^i \frac{\partial}{\partial x^i} \lrcorner dV_g\right) &= d\left(\sum_{j=1}^n X^j \sqrt{\det g} dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^n\right) \\
&= \sum_{j=1}^n (-1)^{j-1} d\left(X^j \sqrt{\det g} dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^n\right) \\
&= \sum_{j=1}^n (-1)^{j-1} \frac{\partial}{\partial x^j} \left(X^j \sqrt{\det g}\right) dx^j \wedge dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^n \\
&= \sum_{j=1}^n \frac{\partial}{\partial x^j} \left(X^j \sqrt{\det g}\right) dx^1 \wedge \cdots \wedge dx^j \wedge \cdots \wedge dx^n \\
&= \sum_{j=1}^n \frac{\partial}{\partial x^j} \left(X^j \sqrt{\det g}\right) \sqrt{\det g} \frac{1}{\sqrt{\det g}} dx^1 \wedge \cdots \wedge dx^j \wedge \cdots \wedge dx^n \\
&= \frac{1}{\sqrt{\det g}} \sum_{j=1}^n \frac{\partial}{\partial x^j} \left(X^j \sqrt{\det g}\right) dV_g.
\end{aligned}$$

Using summation notation,

$$\operatorname{div}\left(X^i \frac{\partial}{\partial x^i}\right) = \frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} \left(X^i \sqrt{\det g}\right)$$

Notice, if g is the Euclidean metric on \mathbb{R}^n where $g_{ij} = \delta_i^j$ in coordinates, then the equation reduces to

$$\operatorname{div}\left(X^i \frac{\partial}{\partial x^i}\right) = \sum_{i=1}^n \frac{\partial X^i}{\partial x^i}.$$

Now, we will use this coordinate expression to derive the expression for the Laplacian. Below, we use the Einstein summation convention.

$$\begin{aligned}
\Delta f &= -\operatorname{div}(\operatorname{grad} f) \\
&= -\operatorname{div}((df)^\#) \\
&= -\operatorname{div}\left(g^{ij} \frac{\partial f}{\partial x_i} \frac{\partial}{\partial x_j}\right) \\
&= -\frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} \left(g^{ij} \sqrt{\det g} \frac{\partial f}{\partial x^j}\right)
\end{aligned}$$

Furthermore, if g is the Euclidean metric on \mathbb{R}^n , we see that the expression for the Laplacian reduces to

$$\Delta f = \frac{\partial}{\partial x^i} \left(\delta_i^j \frac{\partial f}{\partial x^j}\right) = \frac{\partial^2 f}{(\partial x^i)^2}.$$

Removing the Einstein summation convention, we can clearly see that the coordinate expression reduces to the usual expression of the Laplacian on \mathbb{R}^n .

$$\Delta f = \sum_{i=1}^n \frac{\partial^2 f}{(\partial x^i)^2}$$

■

In the course of the proof above, one should notice that we also gave a proof of the following proposition.

Proposition 2.2 (Coordinate Representation of the Divergence). With assumptions as above, the coordinate representation of the divergence is as follows.

$$\operatorname{div}\left(X^i \frac{\partial}{\partial x^i}\right) = \frac{1}{\sqrt{\det g}} \frac{\partial}{\partial x^i} (X^i \sqrt{\det g}) \quad (9)$$

On \mathbb{R}^n with the Euclidean metric and standard coordinates, this reduces to:

$$\operatorname{div}\left(X^i \frac{\partial}{\partial x^i}\right) = \sum_{i=1}^n \frac{\partial X^i}{\partial x^i}. \quad (10)$$

On a Riemannian manifold, the gradient has the same interpretation as it has in a multivariable calculus class: its points in the direction in which the function f increases the fastest. This has a very important consequence, especially since the Laplacian is defined with the gradient. *The Laplacian depends on the metric.* Even though there was no reference to a metric on \mathbb{R} above, we have been implicitly using the fact that \mathbb{R} carries with it the usual Euclidean metric.

The Laplace-Beltrami operator was generalized to the Laplace-de Rham operator [25], also known as the Hodge Laplacian, which not only acts on smooth functions on the manifold, but also smooth forms. In fact, both the Laplace-Beltrami Operator and the Laplace-de Rham operator act on L^2 of some space, either functions on the manifolds or forms on the manifold.

To do this, one must realize that the Laplace-Beltrami operator can be thought of as $\Delta f = \delta df$, where d is the exterior derivative, and δ its formal adjoint. (The existence of this formal adjoint requires a metric!) To extend beyond functions (0-forms), the definition becomes the following:

Definition 2.3 (Laplace-de Rham Operator or Hodge Operator). The Laplacian for k -forms on a Riemannian manifold (M, g) is given by

$$\Delta^k = \delta^{k+1} d^k + d^{k-1} \delta^k \quad (11)$$

where d^k denotes exterior differentiation on k -forms and δ^{k+1} its formal adjoint.

The interested reader should consult Rosenberg's text for a more thorough introduction [25]. There are other variants of the Laplacian that arise in Differential Geometry, such as the conformal Laplacian [27] [25].

3. THE LAPLACIAN ON GRAPHS

We will now leave the domain of existing definitions, and construct a playground on which we can gain a bit of intuition, and get a gist of the fundamental properties underlying the Laplacian and harmonic functions. In addition, we will introduce some elementary notions from Graph Theory. As such, this section has a two-fold purpose: building intuition about the Laplacian while being introduced concrete discrete models. In fact, this introduction will be valuable for the reader who goes on to study Sunada's topological construction of isospectral Riemann surfaces, which uses the Cayley Graph of a group [2].

3.1. A discrete version of Harmonic Functions. We begin with a discretized version of a harmonic function. Recall from Section 2 that harmonic functions, $u : \Omega \mapsto \mathbb{R}$, where $\Omega \subset \mathbb{R}^m$, satisfy the mean value property:

$$u(z_0) = \int_{\partial D(z_0, r)} u(z) dz$$

where $\partial D(z_0, r)$ is contained in the Ω . We will take inspiration from this property and define a discretized version of harmonic functions as follows.

Definition 3.1 (Harmonic Function on a Graph). Let $G = (V, E)$ be a connected graph, where V is the set of vertices and E the set of edges. Let $u : V \mapsto \mathbb{R}$ be a function on the vertices of G . We say u is harmonic if

$$u(v) = \frac{1}{d(v)} \sum_{w \in V: (v,w) \in E} u(w) \quad (12)$$

for all $v \in V$, where $d(v)$ is the number of vertices w adjacent to v , i.e. the cardinality of $\{w \in V : (v, w) \in E\}$.

Now we have a version of a harmonic function which satisfies a "mean value property". Notice that in this setting, we have decided to weigh each vertex equally. This was a choice we made: we could have weighed each vertex differently. This choice amounts to choosing a measure on the vertices. So while we have not explicitly defined a metric on the space, the harmonic functions are dependent on something like "length" coming from the choice of measure.

We will now give a few examples of a convenient subset of graphs, regular graphs, and try to identify harmonic functions on these graphs.

Definition 3.2 (Regular). We say a graph $G(V, E)$ is a regular graph if the number of vertices adjacent to any vertex is the same. We then write $d(v) = d$.

Example 3.1 (Connected Graph on S^1 with No Boundary). Consider a regular degree 2 graph $G = (V, E)$ on the circle S^1 that has no boundary. In the case of the circle, this means every vertex on the graph has two corresponding edges. The case for $n = 7$ is shown in the image below.




FIGURE 2. Finite Graph on a Circle, C_7

What can we say about harmonic functions on this graph? They are all constant!

Proof. Assume otherwise. Then there exists a non-constant harmonic function u on the graph. There must be two adjacent vertices v_1 and v_2 on the graph such that $u(v_1) \neq u(v_2)$. Assume without loss of generality that $u(v_1) < u(v_2)$. Let v_3 be the other vertex connected to v_2 . Then since u is harmonic, $u(v_2) = \frac{1}{2}(u(v_1) + u(v_3))$, from which we can conclude $u(v_2) < u(v_3)$. Since G is a regular graph of degree two, we can repeat these steps for each new vertex until we arrive at the following expression: $u(v_1) < u(v_2) < \dots < u(v_7) < u(v_8)$. This is a contradiction! There are only 7 vertices on the graph, so the string of inequalities cannot hold. ■

All of the harmonic functions are constant, and one may think this is rather unenlightening. However, the proof is enlightening. It gives us an explicit way to interpret the mean value property in this setting.

Given the value of a harmonic function at a vertex, we can multiply this value by the degree of the graph, then distribute this value among the connected vertices. For instance, if the value of $u(v_2) = 10$ in the example above where the degree of the graph is 2, then we have 20 to distribute to $u(v_1)$ and $u(v_3)$. In the example, since the graph has only finitely many vertices, we see that the only possible distribution that works will be 10 to each.

Example 3.2 (Connected Graph on \mathbb{T}^2 with No Boundary). Now, consider a regular degree 4 graph $G = (V, E)$ on the torus (two-torus) \mathbb{T}^2 that has no boundary. In the case of the torus, this means every vertex on the graph has four corresponding edges. In the image below, we see a regular degree 4 graph with 16 vertices on \mathbb{T}^2 .




FIGURE 3. Finite Graph on \mathbb{T}^2

What can we say about harmonic functions on this graph? They are also all constant! In fact, a slight variation on the proof works as above since the graph has only finitely many vertices. The contradiction comes about since at some point in generating a sequence of vertices, some vertex must repeat, though not necessarily the vertex you started with.

What about the graphs above is causing this? There are two properties of these graphs causing this. First, there are a finite number of vertices. Second, if we are at an arbitrary vertex, we can get to any other vertex. We can name these properties and make them definitions.

Definition 3.3 (Compact Graph). We say a graph $G(V, E)$ is compact if the number of vertices $|V|$ is finite.

Definition 3.4 (Connected Graph). We say a graph $G(V, E)$ is connected if between any two vertices v and \tilde{v} on the graph there exists an edge set e_1, e_2, \dots, e_n for some $n \in \mathbb{N}$ and corresponding finite sequence of vertices such that $e_1 = (v, v_1)$, $e_j = (v_{j-1}, v_j)$ for $j \geq 2$ and $j < n$, and $e_n = (v_{n-1}, \tilde{v})$. For example, in Figure 3, the vertex 2 is connected to the vertex 11 with the series of edges $e_1 = (2, 6)$, $e_2 = (6, 10)$ and $e_3 = (10, 11)$. Notice that this "path" is not unique.

Now we can say that on *any* compact, connected graph, harmonic functions are precisely the constant functions. The proof runs exactly as before, but the contradiction comes about because eventually *some* vertex must repeat in the sequence we generate. The reader may object at this point since this seems to contradict intuition: should the harmonic functions be allowed to take a maximum on a boundary of the graph? Should we not have non-constant functions in this setting? Perhaps, and one could certainly explore other definitions, but with our toy model as constructed, we are led to this conclusion. However, one should notice that we have not defined what we mean by "boundary" of a graph. Below we will see that an entire compact graph is its own "interior". In other words, if we say the boundary is the graph without its interior, then all of the graphs we have seen so far have no "boundary".

Definition 3.5 (Interior). We say a vertex v is on the interior of a connected subset D of a graph $G(V, E)$ if every element in the edge set with v is an edge contained in D .

Now, we are in a position to see the fact that our harmonic functions are necessarily constant on compact, connected graphs as a direct consequence of another property: a maximum principle.

Theorem 3.1 (Maximum Principle for Harmonic Functions on Graphs). Let $G = (V, E)$ be a graph (not necessarily finite), and let $u : V \mapsto \mathbb{R}$ be a harmonic function. Given a connected subset $D \subset G(V, E)$, u does not have a maximum in the interior of D unless u is constant.

Proof. Assume otherwise. Let v_0 be a vertex on the interior of D where a maximum is achieved. Since G is connected, we may assume without loss of generality that v_0 is adjacent to a vertex \tilde{v} such that $u(v) > u(\tilde{v})$. Since v_0 is on the interior of D , every vertex w such that (v, w) is in the edge set of D is a vertex in D , including \tilde{v} . Since u is harmonic, we have the following.

$$u(v) = \frac{1}{d(v)} \sum_{w \in V: (v,w) \in E} u(w) < \frac{1}{d(v)} \left[\left(\sum_{i=1}^{d(v)-1} u(v) \right) + u(\tilde{v}) \right] < u(v).$$

We have a contradiction, hence u cannot attain a maximum in the interior of D unless u is constant. ■

3.2. A discrete version of the Laplacian. We can now take our intuition from Section 2 and define a Laplacian on the graph. The Laplacian should measure how far a function deviates from satisfying the mean value property, which is the same as measuring how far the function deviates from being harmonic. Various definitions of the Laplacian are given, depending on the desired normalization. See, for example, Chung's text [7].

Definition 3.6 (Laplacian on a Graph). We define the *Laplacian on a Graph* as follows.

$$\Delta f(v) = f(v) - \frac{1}{d(v)} \sum_{w \in V: (v,w) \in E} u(w) \quad (13)$$

Notice, if $\Delta f = 0$, then f is harmonic just as desired.

In the compact regular graph setting, if we define the diagonal matrix and adjacency matrix as follows, then the Laplacian has a very simple form. Assume $|V| = n$ and enumerate the vertices v_1, \dots, v_n . Let $d(v_i) = d$ for all $i \in \{1, \dots, n\}$.

Definition 3.7 (Diagonal Matrix, D , at a vertex v).

$$D = \begin{bmatrix} d & 0 & \cdots & 0 \\ 0 & d & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d \end{bmatrix} \text{ where } D \text{ is an } n \times n \text{ matrix.} \quad (14)$$

Definition 3.8 (Adjacency matrix, A , at a vertex v).

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} \text{ where } a_{ij} = 1 \text{ if } (v_i, v_j) \in E \text{ and } 0 \text{ otherwise.} \quad (15)$$

Now, if we consider the function $f : G(V, E) \rightarrow \mathbb{R}$ a column vector whose entries are $f(v_i)$, we can recognize the Laplacian as the following matrix: $\Delta = \frac{1}{d}(D - A)$. For the interested reader, there are generalizations to the non-regular case [7]. One may notice that we did not actually need to define the

diagonal or adjacency matrix, since the Laplacian reduces to $\Delta = (I_n - \frac{1}{d}A)$. However, the diagonal matrix and adjacency matrix are natural objects that arise in graph theory, so we will follow convention.

Example 3.3 (Connected Graph on \mathbb{T}^1 with No Boundary). Now, consider the following graph G on \mathbb{T}^1 . G is a regular graph of degree 2 with 5 vertices.




FIGURE 4. Finite Graph on a Circle, C_5

Using the formula above, we see that the Laplacian can be represented by the matrix:

$$L = \frac{1}{2} \left(\begin{bmatrix} 2 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 2 \end{bmatrix} - \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix} \right) = \frac{1}{2} \begin{bmatrix} 2 & -1 & 0 & 0 & -1 \\ -1 & 2 & -1 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \\ -1 & 0 & 0 & -1 & 2 \end{bmatrix}$$

Here, we can explicitly compute the eigenvalues of the matrix to find that the only eigenvalues are 0 with multiplicity 1, $\frac{5}{4} - \frac{\sqrt{5}}{4}$ with multiplicity 2, and $\frac{5}{4} + \frac{\sqrt{5}}{4}$ with multiplicity 2. Computations were done with Sage. Notice that all of the eigenvalues are positive. It will turn out that the eigenvalues will always be positive when we define the Laplacian as we did. Notice, this agrees with our original convention of using the negative in the definition of the Laplacian in Section 2. We will see later that this causes the eigenvalues to all be non-negative.


Example 3.4 (K_5). Now, consider a graph $G = K_5$, a regular graph of degree 4 with 5 vertices.

Using the formula above, we see that the Laplacian can be represented by the matrix:

$$L = \frac{1}{4} \left(\begin{bmatrix} 4 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 \\ 0 & 0 & 4 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 4 \end{bmatrix} - \begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 \end{bmatrix} \right) = \frac{1}{4} \begin{bmatrix} 4 & -1 & -1 & -1 & -1 \\ -1 & 4 & -1 & -1 & -1 \\ -1 & -1 & 4 & -1 & -1 \\ -1 & -1 & -1 & 4 & -1 \\ -1 & -1 & -1 & -1 & 4 \end{bmatrix}$$

Here, we can explicitly compute the eigenvalues of the matrix to find that the only eigenvalues are 0 with multiplicity 1 and $\frac{5}{4}$ with multiplicity 4.

3.3. The Laplacian as an Operator. It turns out that this Laplacian acts similarly to the Laplacian defined in Section 2, especially when we consider it an operator on $L^2(G(E, V))$. To see this, we need to define the gradient of a function on a graph. In this subsection, we follow McMullen's notes on the subject.

FIGURE 5. K_5

Definition 3.9 (Gradient on a Graph). The gradient $|\nabla f|$ on a graph is a function on the edges of the graph given by the following equation.

$$|\nabla f|(e) = |f(x) - f(y)| \text{ for } e = (x, y). \quad (16)$$

It is a short computation to conclude the following.

Theorem 3.2. For a regular graph $G(V, E)$ of degree d ,

$$\frac{1}{2d} \sum_E |\nabla f|^2 = \langle f, \Delta f \rangle = \sum_V f(x) \Delta f(x).$$

And we see that this is precisely what we wanted! The Laplacian is very closely related to the gradient, and we can now say, since our graphs are finite and multiplication commutative, that the Laplacian is a well-defined, self-adjoint operator on $L^2(V)$. The reader is encouraged to test out this new structure on the examples above.

4. FOURIER TRANSFORM ON \mathbb{R}^n

In this section, we will draw an important connection between Laplacian and the Fourier Transform. We will see that the Fourier Transform provides us with a type of spectral theorem for the Laplacian. The bulk of this section follows the presentation in Terras' text [31].

Definition 4.1 (Schwartz Space). The *Schwartz space*, \mathcal{S} is the space of all infinitely differentiable functions $f : \mathbb{R}^n \rightarrow \mathbb{C}$ such that $|x^a D^b f|$ is bounded for all $a, b \in \mathbb{Z}^n$, with $a_j \geq 0$ and $b_j \geq 0$. Functions in the Schwartz space are called *Schwartz functions*. Here we use standard multi-index notation:

$$\begin{aligned} a &:= (a_1, a_2, \dots, a_n) \\ x^a &:= x_1^{a_1} x_2^{a_2} \dots x_n^{a_n} \\ D^b &:= \frac{\partial^{|b|}}{\partial x_1^{b_1} \partial x_2^{b_2} \dots \partial x_n^{b_n}} \text{ where } |b| = b_1 + b_2 + \dots + b_n \end{aligned}$$

This definition may seem unwieldy at first glance, but it's main purpose is to capture a set of smooth functions that decay sufficiently fast (i.e. are $L^2(\mathbb{R}^n)$ integrable), and are dense in the set of all $L^2(\mathbb{R}^n)$ integrable functions [13]. There is another equivalent definition of the Schwartz space \mathcal{S} which does a

better job providing this intuition [13] (although for the proofs that follow, the definition above turns out to be the most useful).

Next, we will define a multiplication operation on functions (convolution). It turns out that convolution is the multiplication operation that plays nicely with the Fourier Transform. show that if we "multiply" two Schwarz functions in this way, we will get a Schwarz function.

Definition 4.2 (Convolution). We define the convolution operation $*$ on two functions $f, g \in L^1_{loc}(\mathbb{R}^m)$ where we assume one of the functions has bounded support

$$f * g(x) = \int_{\mathbb{R}^n} f(x-y)g(y)dy = \int_{\mathbb{R}^n} f(y)g(x-y)dy$$

Convolution is well-defined since the last two expressions are equivalent under a change of variables.

Next, we will define the Fourier Transform on Schwarz functions and follow this by a workhorse Theorem, which effectively gives us all of the properties necessary to prove the inversion formula.

Definition 4.3 (Fourier Transform of Schwartz functions). Let $f \in \mathcal{S}$. Then the Fourier Transform of f is

$$\hat{f}(y) = \int_{\mathbb{R}^n} f(x)e^{-2\pi i {}^t x y} dx \quad (17)$$

where $x \in \mathbb{R}^n$ is a column vector, and ${}^t x$ is the transpose. In other words, ${}^t x y$ is the dot product between x and y . This definition is well-defined since the integral above converges for functions in \mathcal{S} .

Theorem 4.1. The Fourier Transform of Schwartz Functions has the following properties.

- (1) If $f \in \mathcal{S}$, then $\hat{f} \in \mathcal{S}$.
- (2) $D^a(\hat{f}) = ((-2\pi i x)^a \hat{f})^\wedge$.
- (3) $(D^a f)^\wedge = (2\pi i x)^a \hat{f}$.
- (4) Convolution Theorem: $\widehat{(f * g)} = \hat{f} \cdot \hat{g}$.
- (5) Translation: Set $f_a(x) = f(x+a)$ for a, x in \mathbb{R}^m . Then

$$\hat{f}_a(x) = e^{2\pi i {}^t a x} \hat{f}(x).$$

- (6) Dilation: Let u be a positive real number and set ${}^u f(x) = f(ux)$ for x in \mathbb{R}^m . Then

$$\widehat{{}^u f}(x) = u^{-m} \hat{f}(u^{-1}x).$$

- (7) Let $f(x) = e^{-\pi \|x\|^2}$ for $x \in \mathbb{R}^n$. Then, $f = \hat{f}$.
- (8) $G_t = G_t^- = \hat{G}_t$, where we say $f^-(x) := f(-x)$.
- (9) $\int_{\mathbb{R}^m} \hat{f} g dx = \int_{\mathbb{R}^m} f \hat{g} dx$.

Before proving property four, we will require the following lemma showing that $\widehat{(f * g)}$ is a well-defined Fourier Transform. The reader is encouraged to consult Folland's text [13] for a proof.

Lemma 4.1. If f and g are in \mathcal{S} , then $f * g$ is in \mathcal{S}

With the lemma under our belt, we can begin a proof of our main theorem.

Proof. We begin by proving the second property. Note that we may pass the derivative since the function is integrable, as is its derivative. Additionally, note that here D^a is a differential operator with respect to y .

$$\begin{aligned}
D^a(\hat{f})(y) &= D^a\left(\int_{\mathbb{R}^m} f(x)e^{-2\pi i^t xy} dx\right) \\
&= \int_{\mathbb{R}^m} D^a(f(x)e^{-2\pi i^t xy}) dx \\
&= \int_{\mathbb{R}^m} f(x)D^a(e^{-2\pi i^t xy}) dx \\
&= \int_{\mathbb{R}^m} f(x)\frac{\partial^{|a|}}{\partial^{a_1}y_1 \cdots \partial^{a_m}y_m}\left(\prod_{j=1}^m e^{-2\pi i x_j y_j}\right) dx \\
&= \int_{\mathbb{R}^m} f(x)\left(\prod_{j=1}^m -2\pi i x_j\right)\left(\prod_{j=1}^m e^{-2\pi i x_j y_j}\right) dx \\
&= \int_{\mathbb{R}^m} f(x)(-2\pi i x)^a e^{-2\pi i^t x y} dx \\
&= ((-2\pi i x)^a f(x))^\wedge
\end{aligned}$$

Next, we show the third property. We first show this for $|a| = 1$, then show that the general case follows by induction. Note that here, D^a is a differential operator with respect to x .

Let $|a| = 1$:

$$\begin{aligned}
(D^a f)^\wedge &= \int_{\mathbb{R}^m} (D^a f(x))e^{-2\pi i^t xy} dx \\
&= \int_{\mathbb{R}^m} \frac{\partial f}{\partial x_i} e^{-2\pi i^t xy} dx, \text{ for some } i \\
&= \int_{\mathbb{R}^m} (2\pi i y_i) f(x) e^{-2\pi i^t xy} dx, \text{ by integration by parts} \\
&= \int_{\mathbb{R}^m} (2\pi i y)^a f(x) e^{-2\pi i^t xy} dx \\
&= (2\pi i y)^a \hat{f}(y)
\end{aligned}$$

Now, let D^a be any differential operator with $|a| > 1$, and define D^b such that $D^b \frac{\partial}{\partial x_i} = D^a$, for some i . Notice $|b| = |a| - 1$. Then, by induction, we can show the result follows.

$$\begin{aligned}
(D^a f)^\wedge &= \int_{\mathbb{R}^m} (D^a f(x))e^{-2\pi i^t xy} dx \\
&= \int_{\mathbb{R}^m} (D^b \frac{\partial f}{\partial x_i}(x))e^{-2\pi i^t xy} dx \\
&= (2\pi i y)^b \int_{\mathbb{R}^m} \frac{\partial f}{\partial x_i}(x) e^{-2\pi i^t xy} dx, \text{ by the inductive hypothesis} \\
&= (2\pi i y)^b (2\pi i y_i) \hat{f}(y), \text{ by replicating the base case for induction} \\
&= (2\pi i y)^a \hat{f}(y),
\end{aligned}$$

which completes the proof of the third property. To see the first property, we need only use these two properties. If $f \in \mathcal{S}$, then $(D^a f)$ is bounded for any a so we can conclude by the third property that $(2\pi i y)^a \hat{f}(y)$ is bounded for any a . Similarly, since $f \in \mathcal{S}$, we can conclude $(-2\pi i x)^a f(x)$ is bounded for any a . By property two, we can conclude $D^a(\hat{f})(y)$ is bounded for any a . Hence, we can conclude $\hat{f} \in \mathcal{S}$.

Now we prove property four.

$$\begin{aligned}
\widehat{(f * g)} &= \left(\int_{\mathbb{R}^m} f(x-z)g(z)dz \right)^\wedge \\
&= \int_{\mathbb{R}^m} \left(\int_{\mathbb{R}^m} f(x-z)g(z)dz \right) e^{-2\pi i xy} dy \\
&= \int_{\mathbb{R}^m} \left(\int_{\mathbb{R}^m} f(x-z)g(z) e^{-2\pi i xy} dy \right) dz, \text{ by Fubini, which we can apply because of the lemma.} \\
&= \int_{\mathbb{R}^m} g(z) \left(\int_{\mathbb{R}^m} f(x-z) e^{-2\pi i xy} dy \right) dz \\
&= \int_{\mathbb{R}^m} g(z) \left(\int_{\mathbb{R}^m} f(w) e^{-2\pi i wy} e^{2\pi i zy} dw \right) dz, \text{ letting } w = x - z \\
&= \left(\int_{\mathbb{R}^m} g(z) e^{2\pi i zy} dz \right) \left(\int_{\mathbb{R}^m} f(w) e^{-2\pi i wy} dw \right) \\
&= \hat{f} \cdot \hat{g}
\end{aligned}$$

We move on to property (5).

$$\begin{aligned}
\hat{f}_a(y) &= \int_{\mathbb{R}^m} f(x+a) e^{-2\pi i xy} dx \\
&= \int_{\mathbb{R}^m} f(w) e^{-2\pi i (w-a)y} dx, \text{ letting } w = x + a \\
&= \int_{\mathbb{R}^m} f(w) e^{-2\pi i wy} e^{2\pi i ay} dx \\
&= e^{2\pi i ay} \int_{\mathbb{R}^m} f(w) e^{-2\pi i wy} dx \\
&= e^{2\pi i ay} \hat{f}(y)
\end{aligned}$$

Next, we show property (6).

$$\begin{aligned}
\widehat{u}f(y) &= \int_{\mathbb{R}^m} f(ux) e^{-2\pi i xy} dx \\
&= \int_{\mathbb{R}^m} f(w) e^{-2\pi i (u^{-1}w)y} u^{-m} dx, \text{ where } u^{-m} \text{ is the Jacobian of } x = u^{-1}w \\
&= u^{-m} \int_{\mathbb{R}^m} f(w) e^{-2\pi i w(u^{-1}y)} dx \\
&= u^{-m} \hat{f}(u^{-1}y).
\end{aligned}$$

To prove property (7), compute $\frac{d\hat{f}}{dx}(x)$ using integration by parts and then solve the resulting differential equation. For property (8), let G_t be the Gauss kernel, and compute directly using property (6) and property (7). Property (9) is an application of Fubini with an appropriate change of variables. ■

We will now use our workhorse theorem to prove the Fourier Inversion formula for Schwartz functions.

Theorem 4.2 (Fourier Inversion Formula). For $f \in \mathcal{S}$, $\hat{\hat{f}} = f^{-}(x)$, which we can write as:

$$f(x) = \int_{\mathbb{R}^n} \hat{f}(y) e^{2\pi i yx} dx. \quad (18)$$

Proof. Use properties 8 and 9 from Theorem 9.1 with the fact that G_t is a positive Dirac sequence to conclude the inversion formula. ■

Now, we can compare the Fourier Transform to the Inverse Transform, and unpack the meaning of the statement at the beginning of this section: "The Fourier Transform gives us a 'spectral theorem' about the Laplacian". We will do this for \mathbb{R} first:

$$\begin{aligned}\hat{f}(y) &= \int_{\mathbb{R}^n} f(x)e^{-2\pi ixy} dx \\ f(x) &= \int_{\mathbb{R}^n} \hat{f}(y)e^{2\pi iyx} dx\end{aligned}$$

From Section 2 above, we computed the eigenfunctions of the Laplacian on \mathbb{R} . It turns out that the same statement holds in higher dimensions, where eigenfunctions are given by products of exponentials. Now, notice that the Fourier transform is integration against eigenfunctions of the Laplacian. The same holds for the inverse transform, but in addition, it also gives us a characterization of the function in terms of eigenfunctions of the Laplacian.

This suggests that one way to understand the Laplacian is to understand Fourier Integrals, and in particular, the algebraic structure of the set of functions on our space. However, the reader should keep in mind, all of the statements above apply to Schwartz functions, which may seem a bit restrictive. We will show how to extend these statement from the class of Schwartz functions to all of $L^2(\mathbb{R}^n)$. First, however, we will give a few consequences of the machinery we have built. The reader is encouraged to consult Folland's text for proofs [13].

Theorem 4.3. $f \mapsto \hat{f}$ is one-to-one, linear from \mathcal{S} onto \mathcal{S} .

In addition to property (4) above (Convolution Theorem), $(f * g)^\wedge = \hat{f} \cdot \hat{g}$, we also have the following.

Theorem 4.4. Let f and g be Schwartz functions. $\widehat{(fg)} = \hat{f} * \hat{g}$.

Next, given the inner product for f, g in \mathcal{S} by $(f, g) = \int f\bar{g}$ (notice that this the usual L^2 norm), we have the following.

Theorem 4.5 (Parseval's Identity). Let f and g be Schwartz functions. $(f, g) = (\hat{f}, \hat{g})$.

We also have the following Corollary (let $g = f$ in the above Theorem).

Corollary 4.1 (Plancharel Identity). For $f \in \mathcal{S}$, $\|f\|_{L^2} = \|\hat{f}\|_{L^2}$

Now, recall that Schwartz functions are dense in $L^2(\mathbb{R}^n)$. If we can successfully define the Fourier Transforms on $L^2(\mathbb{R}^n)$, then we can extend all of the above machinery to all functions in $L^2(\mathbb{R}^n)$. First, notice that the Fourier Transform as defined above will not work for arbitrary $L^2(\mathbb{R}^n)$ functions. However, we can define the transform as follows.

Definition 4.4 (Fourier Transform of functions in $L^2(\mathbb{R}^n)$).

$$\hat{f}(y) = \lim_{n \rightarrow \infty} \int_{-n}^n f(x)e^{-2\pi iyx} dx \quad (19)$$

where convergence is in $L^2(\mathbb{R}^n)$.

This definition makes sense; we can apply the Dominated Convergence Theorem to show that this Fourier Transform is well-defined. With a bit of bootstrapping, the properties above will now follow for the Fourier Integrals on $L^2(\mathbb{R}^n)$. In particular, we see that the Fourier Transform is an isometric isomorphism of $L^2(\mathbb{R}^n)$.

Thus, we can conclude that the inverse transform also gives us a decomposition of functions in $L^2(\mathbb{R}^n)$ in terms of the eigenfunctions of the Laplacian. This raises many questions, the most obvious of which is a question of whether the eigenfunctions are a basis for $L^2(\mathbb{R}^n)$. For $L^2(\mathbb{R}^n)$, this is true, and the argument is not hard (in fact, we have an orthonormal basis). It requires an application of the Stone-Weierstrauss theorem and a touch of bootstrapping. Does this hold for any space? We will give a generalization for Compact Riemann surfaces in Section 8.

Now, we can say that the Fourier transform is a change of basis to new coordinates that are especially suitable for understanding the Laplacian. We have effectively "diagonalized" the Laplacian. However, there is one quirky thing happening. The eigenfunctions of the Laplacian on \mathbb{R} are not actually in $L^2(\mathbb{R})$. For this reason, they are referred to as *generalized eigenfunctions* or *almost eigenfunctions*. When we reduce to the compact case (for example, $[0, 1]$), we will see that the eigenfunctions of the Laplacian are in L^2 , being smooth functions on a compact space.

5. FOURIER SERIES ON \mathbb{T} AND A DECOMPOSITION OF $L^2(\mathbb{T})$

5.1. Fourier Series and Poisson Summation. Since the Laplacian as it is normally defined on \mathbb{R}^n has such a nice connection to the Fourier Transform, we could ask if the same is true of the Laplacian on other spaces? What does a Fourier Transform look like on a different space? We will start to answer this question by recalling the Fourier Series, which turns out to be an example of a transform on the quotient of \mathbb{R} , $\mathbb{T} \cong \mathbb{R}/\mathbb{Z}$. Recall that functions $f \in L^2(\mathbb{T})$, f has a Fourier series as follows:

$$f(x) = \sum_{n \in \mathbb{Z}} \hat{f}(n) e^{2\pi i n x}, \text{ where} \quad (20)$$

$$\hat{f}(n) = \int_{\mathbb{R}} f(x) e^{-2\pi i n x} dx \quad (21)$$

where convergence is with respect to the usual L^2 norm. We should first notice that the formula for \hat{f} looks a lot like the formula for a Fourier Transform, but only defined for integers. The Fourier Series, on the other hand, is a decomposition of a function with respect to a set of complex exponentials indexed by integers. In fact, recall from Section 1 that eigenfunctions of the Laplacian applied to functions periodic of period 1 are precisely these exponentials. Similar themes begin to emerge: we have a transform from a function in $L^2(\mathbb{T})$ to a function in $l^2(\mathbb{Z})$, we have an inverse transform (the Fourier Series) which gives a decomposition of a function in terms of a set of complex exponentials, and we have a "spectral theorem" of the Laplacian since that set of complex exponentials is precisely the set of eigenfunctions of the Laplacian. For a more complete discussion of the Fourier Series, similar to the one in Section 4, the reader is encouraged to read the first five sections of Chapter 1 in Dym and McKean's text [11].

There is one very important tool that is normally seen in the context of the Fourier Series, and that is Poisson Summation. For a proof, consult Terras' [31] or Dym and McKean [11].

Theorem 5.1. (Poisson Summation Formula) If $f : \mathbb{R} \rightarrow \mathbb{C}$ is a Schwartz function, then

$$g(x) := \sum_{a \in \mathbb{Z}} f(x+a) = \sum_{a \in \mathbb{Z}} \hat{f}(a) e^{2\pi i a x}. \quad (22)$$

We will jump ahead at this point and begin to highlight properties of this construction that serve as a jumping off point for generalizing Fourier Analysis to broader class of spaces. One should note the following difference between the Fourier Transform on \mathbb{R} and the construction behind the Fourier Series: the Fourier Transform is giving an isometric isomorphism of $L^2(\mathbb{R})$ onto itself, whereas the Fourier series gives us an isometric isomorphism of $L^2(\mathbb{T})$ onto $l^2(\mathbb{Z})$. Our goal is to understand the algebraic structure

in play. We will also pay close attention to the role of the complex exponentials. The key to understanding various methods of generalizing this structure is understanding the various roles the the complex exponentials play. We will follow the presentation in Chapter 4 of Dym and McKean's classical text [11].

5.2. A Representation-Theoretic Viewpoint. We begin with a definition.

Definition 5.1. (Character) A *character* of a group G is a homomorphism of G into S^1 , $\chi : G \rightarrow S^1$:

- (a) For all $g \in G$, $|\chi(g)| = 1$.
- (b) For all $g_1, g_2 \in G$, $\chi(g_1 g_2) = \chi(g_1)\chi(g_2)$.

In the case of topological groups, we require the character be continuous.

The definition has some immediate consequences, the proofs of which are all elementary group-theoretic proofs.

Proposition 5.1. Let e be the identity element in G and g^{-1} denote the inverse of $g \in G$.

- (1) $\chi(e) = 1$, where e is the identity element in G .
- (2) $\chi(g^{-1}) = \chi(g)^{-1} = \overline{\chi(g)}$.

Definition 5.2. (Dual Group) Let \hat{G} be the collection of all characters χ of the group G . We call \hat{G} the dual group of G .

Proposition 5.2. \hat{G} is an abelian group under the multiplication defined by $\chi_1 \cdot \chi_2(g) = \chi_1(g)\chi_2(g)$. The identity of \hat{G} is the trivial character χ where $\chi(g) = 1$ for all $g \in G$.

And now we come to an important theorem. The set of complex exponentials in the Fourier series are precisely the characters of the group \mathbb{T} .

Theorem 5.2. $\chi_n(x) = e^{2\pi i n x}$ is a complete list of characters of S^1 .

And now, with a complete list of characters of \mathbb{T} , we can identify the dual group.

Theorem 5.3. $\hat{\mathbb{T}} \cong \mathbb{Z}$.

This tells us that \mathbb{Z} is the dual group and each $e^{2\pi i n x}$ is a character in the dual group (look back at the equation 20, the Fourier Series expansion of a periodic function). This is the piece we needed to see: the Fourier Series is summed over the elements of the dual group, with multiplication against elements in the dual group, the characters. The "Fourier Transform", equation 21, which gives us $\hat{f}(n)$, is a transform from $L^2(\mathbb{T})$ to $l^2(\mathbb{Z})$, where \mathbb{Z} is the dual group to \mathbb{T} .

Since we can see a shadow of Pontryagin Duality, we will point it out. The double dual of \mathbb{T} is isomorphic to itself.

Proposition 5.3. $\hat{\hat{\mathbb{T}}} \cong \mathbb{T}$.

5.3. Translation Invariant Subspaces of $L^2(\mathbb{T})$. Next, we will capture another property of these complex exponentials in $L^2(\mathbb{T})$, namely that they each generate a different translation invariant subspace.

Definition 5.3. (Translation Invariant Subspace) A closed subspace M of $L^2(\mathbb{T})$ is *translation invariant* if it is closed under translations, meaning if we define $f_y(x) = f(x + y)$ then f_y belongs to M for every $f \in M$ and every $y \in \mathbb{T}$.

Proposition 5.4. Let M_n be the closed, translation invariant subspace containing $e^{2\pi i n x}$. Then for $|n| \neq |m|$, $e^{2\pi i m x} \notin M_n$, and we have the following orthogonal decomposition of $L^2(\mathbb{T})$:

$$L^2(\mathbb{T}) = \bigoplus_{|n| < \infty} M_n \quad (23)$$

The reader may notice that this is an improved version of the Plancherel Identity given in Section 4, but applied to \mathbb{T} instead of \mathbb{R} .

5.4. Eigenfunctions of the Laplacian. Though we have computed this in Section 2, the computation rightfully belongs here. The eigenfunctions of the Laplacian are precisely the characters of \mathbb{T} . In light of the section 5.3 above, we can now describe the eigenfunctions more geometrically: M_n , the translation invariant subspace of $e^{2\pi i n x}$ and $e^{-2\pi i n x}$ is an eigenspace of the Laplacian.

5.5. Homomorphisms of the Convolution Algebra on $L^2(\mathbb{T})$. Next, we identify an algebraic structure on $L^1(\mathbb{T})$ which will give us another perspective on how the Fourier Transform is acting on $L^2(\mathbb{T})$.

Recall the definition of convolution for functions in \mathbb{R}^n . We can define the same operation on functions on \mathbb{T}^n .

Definition 5.4. (Convolution) Let f, g be in $L^1(\mathbb{T})$ and define the convolution operation as follows.

$$f * g(x) = \int_0^1 f(x-y)g(y)dy$$

Next, endow $L^2(\mathbb{T})$ with this multiplication (convolution), and define addition on the space as point-wise addition of functions. With this algebra on $L^1(\mathbb{T})$, we can define homomorphism of the algebra.

Definition 5.5. (Homomorphism of $L^1(\mathbb{T})$) A homomorphism of an algebra \mathcal{B} is a map $j : \mathcal{B} \rightarrow \mathbb{C}$ such that the map respects

- (a) complex multiplication: for all $\alpha \in \mathbb{C}$, $j(\alpha f) = \alpha j(f)$,
- (b) addition: for $f_1, f_2 \in \mathcal{B}$, $j(f_1 + f_2) = j(f_1) + j(f_2)$,
- (c) multiplication on \mathcal{B} : for $f_1, f_2 \in \mathcal{B}$, $j(f_1 * f_2) = j(f_1)j(f_2)$,
- (d) $|j(f)| \leq C\|f\|_{L^1}$, where C is a constant independent of f .

The first thing we should notice is that $j_n(f) := \hat{f}(n) = \int_0^1 f(x)e^{-2\pi i n x} dx$ is a homomorphism. Notice that we have a convolution theorem for the Fourier transform, whose proof can be readily adapted to give a convolution theorem for \hat{f} here, which means (c) is satisfied. In fact, this turns out that this is a complete list of homomorphisms!

Theorem 5.4. Define $j_n(f) := \hat{f}(n)$ for $f \in L^1(\mathbb{T})$. Then j_n for $n \in \mathbb{Z}$ is a complete list of homomorphisms of $L^1(\mathbb{T})$.

Now, we can extend this result to say something about $L^2(\mathbb{T})$. Notice that \mathbb{T} has a Haar measure, which is something we have glossed over until now. This will be discussed in more detail in Section 8. For now, we may think of this measure as the Lebesgue measure on $[0, 1]$. The point is, this is a finite measure, which tells us something about the relationship between $L^1(\mathbb{T})$ and $L^2(\mathbb{T})$, namely $L^2(\mathbb{T}) \subset L^1(\mathbb{T})$.

Theorem 5.5. Define $j_n(f) := \hat{f}(n)$ for $f \in L^2(\mathbb{T})$. Then j_n for $n \in \mathbb{Z}$ is a complete list of homomorphisms of $L^2(\mathbb{T})$.

5.6. Remark on the finite abelian case. We now digress to capture a bit of low-hanging fruit, often seen in an Algebra course. If we restrict our attention to finite abelian groups, using the observations above, we can prove the existence of a Fourier series, with convergence in L^2 , give a Plancherel Identity, and give a Poisson summation formula. We follow sections 4.5 and 4.6 in Dym and McKean's text [11].

We start by reminding the reader that finite abelian groups isomorphic to direct products of cyclic groups [10], and from this one could prove that G is isomorphic to its dual group \hat{G} .

We can define an inner product space on the functions of the group, $f : g \rightarrow \mathbb{C}$, and consider convergence in L^2 . Since the functions have finitely many values, any such choice of function is square-summable.

Definition 5.6. We define the inner product between two functions on a group G as follows.

$$(f_1, f_2) = \frac{1}{|G|} \sum_{g \in G} f_1(g) \overline{f_2(g)} \quad (24)$$

With this inner product, the norm of a function is given by $\|f\| = \sqrt{(f, f)}$. There is also a canonical isomorphism between the group and its double dual given by $g(\chi) := \chi(g)$.

We can then highlight two properties these characters must have:

Theorem 5.6. (Orthogonality Principles)

(a) The characters form an orthogonal family:

$$(\chi_1, \chi_2) = \begin{cases} 1 & \chi_1 = \chi_2 \\ 0 & \chi_1 \neq \chi_2 \end{cases}$$

(b) The elements of the group can be identified with the corresponding element in the double dual $\hat{\hat{G}}$ using the isomorphism above. The elements of the double dual form an orthogonal family, when thought of as characters of the dual \hat{G} .

$$\sum_{\hat{G}} \chi(g_1) \overline{\chi(g_2)} = \begin{cases} 1 & g_1 = g_2 \\ 0 & g_1 \neq g_2 \end{cases}$$

Theorem 5.7. (Plancherel Theorem [11]) Let G be a finite abelian group, and f any function on G , $f : G \rightarrow \mathbb{C}$. Then f can be expanded into a Fourier series

$$f = \sum_{\hat{G}} \hat{f}(\chi) \chi \quad (25)$$

with coefficients

$$\hat{f}(\chi) = \frac{1}{|G|} (f, \chi) = \frac{1}{|G|} \sum_{g \in G} f(g) \overline{\chi(g)} \quad (26)$$

and an Plancherel identity:

$$\|f\|^2 = \sum_{g \in G} |f(g)|^2 = \sum_{\hat{G}} |\hat{f}(\chi)|^2 = \|\hat{f}\|^2 \quad (27)$$

Next, we show that we have a Poisson Summation formula. Before we generate this formula though, we need to identify a set of characteristic functions which are supported on a quotient of the group.

Theorem 5.8. (Poisson Summation Formula) Let f be a function on a finite abelian group G . Let H be a subgroup of G . Then we have the following.

$$\sum_H f(h) = \sum_{\hat{G}/\hat{H}} \hat{f}(\chi) \quad (28)$$

Now, we are beginning to see that this structure is transferable to other topological spaces. Finite abelian groups may not be that interesting as a topological space, but transferring this structure onto finite abelian groups has far-reaching ramifications. The interested reader should visit section 4.6 in Dym and McKean's text where using only Fourier analysis on finite abelian groups, the authors prove Gauss's Law of Quadratic Reciprocity.

6. KEY EXAMPLE: \mathbb{T}^2

Now we come across our first example of a translation surface. We may think of the 2-torus as a polygon in the plane, with opposite sides identified by translation, which fits nicely with one of the definitions of a translation surface [35] [33].

6.1. A Representation-Theoretic Viewpoint. Proceeding in a fashion similar to the previous section, we begin by identifying characters of \mathbb{T}^2 . The set of complex exponentials on \mathbb{T}^2 . The proof here is analogous to the proof referenced in Section 5 for the case of \mathbb{T} .

Theorem 6.1. $\chi_n(x) = e^{2\pi i(n \cdot x)}$ for $n \in \mathbb{Z}^2$, $x \in \mathbb{R}^2$, and \cdot being the usual dot product is a complete list of characters of \mathbb{T}^2 .

With a complete list of characters of \mathbb{T} , we can identify the dual group.

Theorem 6.2. $\hat{\mathbb{T}^2} \cong \mathbb{Z}^2$.

This tells us that \mathbb{Z}^2 is the dual group and each $e^{2\pi i(n \cdot x)}$ is a character in the dual group. This is the piece we needed to see, and we can use it to surmise what the correct Fourier Series would be in 2 dimensions:

$$f(x) = \sum_{n \in \mathbb{Z}^2} \hat{f}(n) e^{2\pi i(n \cdot x)}, \text{ where} \quad (29)$$

$$\hat{f}(n) = \int_{\mathbb{R}^2} f(x) e^{-2\pi i(n \cdot x)} dx \quad (30)$$

The Fourier Series is summed over the elements of the dual group, with multiplication against elements in the dual group, the characters. The "Fourier Transform", equation 30, which gives us $\hat{f}(n)$, is a transform from $L^2(\mathbb{T}^2)$ to $l^2(\mathbb{Z}^2)$, where \mathbb{Z}^2 is the dual group to \mathbb{T}^2 .

We will not prove convergence of this formula here. Instead, we will be able to see it as a consequence of the Peter-Weyl Theorem in section 8.

As before, we can see a shadow of Pontryagin Duality. The double dual of \mathbb{T}^2 is isomorphic to itself.

Proposition 6.1. $\hat{\hat{\mathbb{T}^2}} \cong \mathbb{T}^2$.

6.2. Translation Invariant Subspaces of $L^2(\mathbb{T}^2)$. Next, we will capture another property of these complex exponentials in $L^2(\mathbb{T}^2)$, namely that they each generate a different translation invariant subspace.

Proposition 6.2. Let M_n be the closed, translation invariant subspace containing $e^{2\pi i(n \cdot x)}$. Then for $|n| \neq |m|$, $e^{2\pi i(m \cdot x)} \notin M_n$, and we have the following orthogonal decomposition of $L^2(\mathbb{T}^2)$:

$$L^2(\mathbb{T}^2) = \bigoplus_{|n| < \infty} M_n \quad (31)$$

6.3. Eigenfunctions of the Laplacian. We can compute the eigenfunctions of the Laplacian using separation of variables, which lead us to conclude that the eigenfunctions of the Laplacian are precisely the characters of \mathbb{T}^2 . In light of the section 6.2 above, we can now describe the eigenfunctions more geometrically: for $n \in \mathbb{Z}^2$, M_n the translation invariant subspace of $e^{2\pi i(n \cdot x)}$ and $e^{-2\pi i(n \cdot x)}$ is an eigenspace of the Laplacian.

6.4. Homomorphisms of the Convolution Algebra on $L^2(\mathbb{T}^2)$. Next, we identify an algebraic structure on $L^1(\mathbb{T}^2)$ which will give us another perspective on how the Fourier Transform is acting on $L^2(\mathbb{T}^2)$.

Next, endow $L^2(\mathbb{T}^2)$ with multiplication (convolution as defined in section 6), and define addition on the space as pointwise addition of functions. With this algebra on $L^1(\mathbb{T}^2)$, we can define homomorphism of the algebra.

Definition 6.1. (Homomorphism of $L^1(\mathbb{T}^2)$) A homomorphism of an algebra \mathcal{B} is a map $j : \mathcal{B} \rightarrow \mathbb{C}$ such that the map respects

- (a) complex multiplication: for all $\alpha \in \mathbb{C}$, $j(\alpha f) = \alpha j(f)$.
- (b) addition: for $f_1, f_2 \in \mathcal{B}$, $j(f_1 + f_2) = j(f_1) + j(f_2)$.
- (c) multiplication on \mathcal{B} : for $f_1, f_2 \in \mathcal{B}$, $j(f_1 * f_2) = j(f_1)j(f_2)$.
- (d) $|j(f)| \leq C\|f\|_{L^1}$, where C is a constant independent of f .

The first thing we should notice is that $j_n(f) := \hat{f}(n) = \int_0^1 f(x)e^{-2\pi i(n \cdot x)} dx$ is a homomorphism. As before, we have a convolution theorem for the fourier transform, whose proof can be readily adapted to give a convolution theorem for \hat{f} here, which means (c) is satisfied.

Just as before, it turns out that this is a complete list of homomorphisms!

Theorem 6.3. Define $j_n(f) := \hat{f}(n)$ for $f \in L^1(\mathbb{T}^2)$. Then j_n for $n \in \mathbb{Z}^2$ is a complete list of homomorphisms of $L^1(\mathbb{T}^2)$.

And as before, we can extend this result to say something about $L^2(\mathbb{T}^2)$. The Haar measure on \mathbb{T}^2 is finite, so $L^2(\mathbb{T}^2) \subset L^1(\mathbb{T}^2)$, and the next theorem follows.

Theorem 6.4. Define $j_n(f) := \hat{f}(n)$ for $f \in L^2(\mathbb{T})$. Then j_n for $n \in \mathbb{Z}$ is a complete list of homomorphisms of $L^2(\mathbb{T})$.

7. ISSUES WITH THE NON-ABELIAN CASE

The work in the previous sections has provided us with a road-map for generalizing the Fourier Transform and inversion formula to any compact Abelian group, finite or otherwise.

However, we will run into problems if we are interested in leveraging this machinery on non-abelian groups. Recall, the dual group \hat{G} associated with a group G is an abelian group consisting of all of the characters of the group. This poses an issue if we are to study characters of non-abelian groups. The characters cannot capture any structure of the group that is not abelian. A key example is given by $SO(3)$. While it is a geometrically intuitive group being rigid motions of the sphere, the group is, in some sense, maximally nonabelian, as the following theorem shows. For a proof of this Theorem, consult Dym and McKean's text [11].

Theorem 7.1. The only character of $SO(3)$ is the trivial character $\chi \equiv 1$.

We see no information can be extracted from the study of representations of $SO(3)$ if we limit ourselves to characters. Since $SO(3)$ is the group of rigid motions of the sphere, this is severely limiting when it comes to studying the geometry underlying isometries of the sphere. In fact, we are faced with similar issues for the Euclidean Motion group $M(2)$ and the set of isometries of the upper half plane $PSL_2(\mathbb{R})$, both fundamental model spaces in geometry [21].

We should ask what precisely did the characters provide us with? A way to represent the underlying structure of the group. For non-abelian groups, we are interested in a similar sort of "non-abelian classifying space".

8. SOME REPRESENTATION THEORY: COMPACT GROUPS AND THE PETER-WEYL THEOREM

We will try to answer some of the questions posed at the end of the previous section. In order to generalize what we have done with characters, we will look for a "non-abelian classifying space". This search will lead us very naturally to the Representation Theory of Groups. Characters are a very specific case of representations, specifically suited for abelian groups. In this section, we follow the presentation given in [1].

8.1. Definitions and Schur's Lemma. As Bagchi, et al, [1] point out, we need another interpretation of \mathbb{S}^1 . It turns out, the way we should think of \mathbb{S}^1 is not actually as a subspace of \mathbb{C} , but as a subspace of the dual to \mathbb{C} , which turns out to be isomorphic to \mathbb{C} . Then \mathbb{S}^1 can be thought of as the set of unitary operators on \mathbb{C} . This is the right perspective to generalize, so we now define representations in the following way:

Definition 8.1. (Representation of a Group) Let G be a topological group, let H be a Hilbert space over \mathbb{C} , and let $\mathcal{B}(H)$ denote the Banach algebra of bounded linear operators on H . Let $\text{GL}(H)$ be the group of all invertible elements of \mathcal{B} . A *representation of a group G on H* is a group homomorphism

$$\pi : G \rightarrow \text{GL}(H) \tag{32}$$

such that for any $x \in H$, the map $g \rightarrow \pi(g)x$ from G to H is continuous. We will denote a representation of G as an ordered pair (π, H) .

Notice that this gives us a way of considering homomorphisms of non-abelian groups that can preserve the nonabelian structure. The remainder of this section will be focused on building the necessary machinery to prove a Fourier transform exists with a Plancherel identity, that there exists "invariant" subspaces of $L^2(G)$ and a decomposition similar to that of the decomposition we saw with $L^2(\mathbb{T})$. This theorem is called the Peter-Weyl Theorem. Along the way, we will extract a few other important properties underlying this machinery.

Definition 8.2 (Unitary). We say a representation (π, H) of a group G is unitary if $\pi(g)$ is a unitary operator for every $g \in G$.

We should notice that characters as we defined them are a simple example of unitary representations, where we allow the character to act on \mathbb{C} by complex multiplication, which amounts to a rotation for unitary operators.

We will now give a list the standard definitions required for basic representation theory.

Definition 8.3 (Uniformly bounded). We say a representation is uniformly bounded if $\sup_{g \in G} \|\pi(g)\| \leq K$ for some constant K . Here $\|\cdot\|$ denotes the usual operator norm.

Definition 8.4 (Invariant Subspace). Let (π, H) be a representation of a group G . We say a subspace M of H is an *invariant subspace* if for every $x \in M$ and $g \in G$, we have $\pi(g)x \in M$.

Definition 8.5 (Subrepresentation). Suppose M is an invariant subspace of a representation (π, H) of a group G . We say that $g \rightarrow \pi|_M(g)$ is a *subrepresentation*, we denote $(\pi|_M, M)$.

Definition 8.6 (Irreducible). Let (π, H) be a representation of a group G . We say (π, H) is an *irreducible representation* if the only invariant subspaces of H are $\{0\}$ and H .

Definition 8.7 (Completely reducible). We say a unitary representation (π, H) is *completely reducible* if there exists a family of closed, mutually orthogonal subspaces H_i such that each (π, H_i) is irreducible.

Now, we will introduce a special operator, the importance of which can be seen in the subsequent lemma: Schur's Lemma.

Definition 8.8 (Intertwining Operator). Let (π, H) be a unitary representation of a group G . Define \mathcal{I}_π as the set $\{T \in \mathcal{B}(H) \mid T\pi(g) = \pi(g)T \text{ for all } g \in G\}$. Any $T \in \mathcal{I}_\pi$ we call an *intertwining operator*.

Lemma 8.1 (Schur's Lemma). Let (π, H) be a unitary representation of a group G . (π, H) is irreducible if and only if $\mathcal{I}_\pi = \{\lambda I \mid \lambda \in \mathbb{C}\}$, where I denotes the identity operator on H .

Proof. Assume that $\mathcal{I}_\pi = \{\lambda I \mid \lambda \in \mathbb{C}\}$. If π is not irreducible, then we can find a non-trivial closed invariant subspace M , in other words $M \neq H$ and $M \neq \emptyset$. Let P be the projection operator $P : H \rightarrow M$, and we claim M is invariant for π if and only if $P\pi(g) = \pi(g)P$ for all $g \in G$. To see this, decompose

H into M and its perpendicular component (H is a Hilbert space). Then for any $x \in H$, $x = x_1 + x_2$ for $x_1 \in M$ and $x_2 \in M^\perp$. Then for all x ,

$$\begin{aligned} P\pi(g)(x) &= P\pi(g)(x_1 + x_2) \\ &= P(\pi(g)x_1 + \pi(g)x_2) \\ &= P\pi(g)x_1 \\ &= \pi(g)x_1 \\ &= \pi(g)P(x_1 + x_2) \\ &= \pi(g)P(x). \end{aligned}$$

Now, with the claim proven, we see that $P \in \mathcal{I}_\pi$, but P , being a projection, is not of the form $P = \alpha I$ for some $\alpha \in \mathbb{C}$, so this is a contradiction. We may conclude that there cannot be any non-trivial closed invariant subspaces, hence π is irreducible.

Conversely, assume π is irreducible. Let $T \in \mathcal{I}_\pi$, and notice that $T^* \in \mathcal{I}_\pi$, where T^* denotes the adjoint operator. Then, if we define $A = \frac{T+T^*}{2}$, and $B = \frac{T-T^*}{2i}$, we can write $T = A + iB$. Now, notice that $T \in \mathcal{I}_\pi$ if and only if $T^* \in \mathcal{I}_\pi$ if and only $A, B \in \mathcal{I}_\pi$. Since A and B are self-adjoint, we can reduce to the case where T is self-adjoint.

Now we will invoke a spectral theorem. We ask that the reader consult Rudin's text [26] for a proof of the theorem. The reader should notice that this is a weaker version of the theorem that Rudin proves. Before stating the theorem, however, we need a definition.

Definition 8.9 (Resolution of the Identity). Let \mathfrak{M} be a σ -algebra in a set Ω , and let H be a Hilbert Space. In this setting, a *resolution of the identity* (on \mathfrak{M}) is a mapping $E : \mathfrak{M} \rightarrow \mathcal{B}(H)$ with the following properties:

- (1) $E(\emptyset) = 0, E(\Omega) = 1$
- (2) Each $E(\omega)$ is a self-adjoint projection.
- (3) $E(\omega' \cap \omega'') = E(\omega')E(\omega'')$.
- (4) If $\omega' \cap \omega'' = \emptyset$, then $E(\omega' \cup \omega'') = E(\omega') + E(\omega'')$.
- (5) For every $x \in H$ and $y \in H$, the set function $E_{x,y}$ defined by $E_{x,y}(\omega) = \langle E(\omega)x, y \rangle$ is a complex measure on \mathfrak{M} .

Theorem 8.1 (Spectral Theorem for Self-Adjoint Operators). If $A \in \mathcal{B}(H)$ is self-adjoint, then there exists a unique resolution of the identity E on the Borel subsets of $\sigma(A)$ which satisfies

$$\int_{\sigma(A)} \lambda dE(\lambda) = A.$$

Furthermore, every projection $E(\omega)$ commutes with every $S \in \mathcal{B}(H)$ which commutes with A .

Since we are assuming T is self-adjoint, and its eigenvalues lie in the real line, we can apply Theorem 8.1 to our operator T and we have spectral decomposition of T as follows:

$$T = \int_{-\infty}^{\infty} \lambda dE(\lambda).$$

If there are two elements in the spectrum of T , call them λ_1 and λ_2 , then there exist disjoint sets in the Borel σ -algebra B_1 and B_2 such that $\lambda_1 \in B_1, \lambda_2 \in B_2$, and $B_1 \cup B_2 = \mathbb{R}$. Then, the spectral theorem tells us that the spectral projections $E(B_1)$ and $E(B_2)$ are in \mathcal{I}_π . But then, one of these projections must be trivial, otherwise $T \notin \mathcal{I}_\pi$. This means there cannot be two elements in the spectrum of T , so we see that $T = \lambda I$ for some λ , as desired. ■

Schur's Lemma is telling us that if you have an irreducible representation, the only operators that should commute with $\pi(g)$ for any $g \in G$ are the one's in the center of $\mathcal{GL}(H)$, when we think of $\mathcal{GL}(H)$

as a group. In some sense, we can interpret this as saying an irreducible representation should not cover up any “nonabelian” properties of the group; the structure of the group should be accurately displayed through the representation.

The reader interested in Lie Algebras should compare Schur’s Lemma as stated above to the version presented in Humphrey’s text on Lie Algebra [18], especially if the reader is familiar with the connection between Lie Groups and Lie Algebras.

With Schur’s Lemma, we can connect our representations more explicitly to the characters we used in the abelian setting.

Corollary 8.1. Let (π, H) be a irreducible unitary representation of an abelian group G . Then $\dim H = 1$.

Proof. Since G is abelian, $\pi(g)\pi(h) = \pi(gh) = \pi(hg) = \pi(h)\pi(g)$, so $\pi(g) \in \mathcal{I}_\pi$ for every g . Since π is irreducible, by Schur’s Lemma, we see $\pi = \lambda I$ for some $\lambda \in \mathbb{C}$. This means that π leaves any one-dimensional space invariant, so H can only be one-dimensional since π is irreducible. Thus $\dim H = 1$, as desired. ■

This means any irreducible unitary representation of an abelian group is a map into $GL(\mathbb{C}) \cong \mathbb{C}$. Unitary means we are mapping into S^1 , just as before. In other words, the group of characters is actually a group of *all* irreducible unitary representations of G ! This suggests it may be possible to generate Fourier series for compact groups if we can identify all irreducible representations of the group.

We now have need for the Haar measure. We will use the existence of the Haar measure without proof. For details additional details, consult the text by Bagchi, et al., [1], or alternatively Einseidler and Ward’s text [12]. A proof of the existence and uniqueness of the Haar measure on compact groups using fixed point theorems from functional analysis can be found in Zimmer’s text [34].

Definition 8.10 (Left Haar Measure). A left Haar measure on a locally compact group G is a positive regular Borel measure μ such that $\mu(gB) = \mu(B)$ for all $B \subset \mathcal{B}$ and $g \in G$.

Theorem 8.2 (Existence and Uniqueness of a Left Haar Measure). Let G be a locally compact group. There exists a left Haar measure μ on G . If ν is any other left Haar measure on G , then $\nu = C\mu$ for some constant C .

As Bagchi, et al. note, a compact group is unimodular, which means that a Left Haar measure is also a right Haar measure. Since our focus is currently on compact groups, we will say Haar measure instead of specifying left or right.

There are two important consequences of the existence of the Haar measure. For details of the proofs of these statements, refer to [1].

Theorem 8.3. Every irreducible unitary representation of a compact group is finite dimensional.

Theorem 8.4. Every unitary representation of a compact group is a direct sum of irreducible finite dimensional representations.

This is a powerful characterization: it tells us precisely what each unitary representation looks like. We now introduce some notation that will capture important notions from the work we have done thus far.

Definition 8.11 (\hat{G}). Let \hat{G} denote the set of equivalence classes of irreducible representations of G , where we say two representations (π_1, H_1) and (π_2, H_2) are equivalent if there exists a continuous linear isomorphism $T : H_1 \rightarrow H_2$ such that $\pi_2(g) = T\pi_1(g)T^{-1}$ for all $g \in G$. For unitary representations, such a T will necessarily be a unitary operator.

Definition 8.12 (Matrix Coefficients). Given a finite dimensional representation (π, H) of a compact group G and an orthonormal basis of H , we can define

$$\phi_{ij}(g) = \langle \pi(g)e_j, e_i \rangle, 1 \leq i, j \leq n \quad (33)$$

The ϕ_{ij} 's are continuous functions on G and are called the *matrix coefficients* of the representation π . Notice that since π is unitary, $\phi_{ij}(g)$ is a unitary matrix for each $g \in G$. When we write the ϕ_{ij} 's, it is understood that there is a choice of orthonormal basis made for H .

8.2. Schur's Orthogonality Relations. We now move onto capturing orthogonality relations between different representations. The reader is encouraged to review the orthogonality principles for finite abelian groups in Theorem 5.6.

Theorem 8.5 (Schur's Orthogonality Relations). Let π and ρ be two finite dimensional representations of G and $(\phi_{ij}(g)), (\psi_{kl}(g))$ the corresponding matrix coefficients of π, ρ , respectively, with respect to some fixed orthonormal bases in the respective Hilbert spaces. Then

- (1) $\langle \phi_{ij}, \psi_{kl} \rangle_{L^2(G)} = 0$ if π and ρ are not equivalent.
- (2) $\langle \phi_{ij}, \phi_{kl} \rangle_{L^2(G)} = \frac{1}{\dim H} \delta_{ik} \delta_{jl}$.

The interested reader is encouraged to consult Bagchi's text for a proof [1].

8.3. The Peter-Weyl Theorem. We will need one more definition before we can state the main theorem of this section.

Definition 8.13 (Hilbert-Schmidt norm). Let A be an operator. Then $\|A\|_{HS}$ denotes the *Hilbert-Schmidt norm*, where $\|A\|_{HS}^2 = \text{Tr}(AA^*)$, and Tr denotes the trace.

With this machinery, which largely generalizes much of what we have done in previous sections, we can state the Peter-Weyl theorem. The reader is encouraged to consult the text by Bagchi, et al., for accessible proof of the theorem.

Theorem 8.6. (Peter-Weyl Theorem)

- (1) Every irreducible unitary representation is equivalent to a subrepresentation of a right regular representation.
- (2) For each $\lambda \in \hat{G}$ there is a finite subspace E_λ such that E_λ is spanned by the matrix coefficients of a representation π in the equivalence class λ . E_λ is independent of the choice of π , and the following hold.
 - (a) Each E_λ is invariant under the right regular representation of G .
 - (b) If π is an irreducible unitary representation in the equivalence class λ , then the restriction of the right regular representation to E_λ is equivalent to the direct sum of d_π copies of π , where d_π is the dimension of the representation π , and consequently $\dim E_\lambda = d_\pi^2$. Additionally, d_π is independent of the choice of representation π , so we may define $d_\lambda = d_\pi$ for any choice of π .
 - (c) $L^2(G) = \bigoplus_{\lambda \in \hat{G}} E_\lambda$.
- (3) For $f \in L^2(G)$ and for each $\lambda \in \hat{G}$, choose $\pi_\lambda \in \lambda$. Define

$$\hat{f}(\lambda) = \int_G f(g) \pi_\lambda(g^{-1}) dg \quad (34)$$

Then $\hat{f} \in B(H_\lambda)$, and

$$f(g) = \sum_{\lambda \in \hat{G}} d_\lambda \text{Tr}(\hat{f}(\lambda) \pi_\lambda(g)) \quad (35)$$

where the series converges in $L^2(G)$, and

$$\|f\|_{L^2}^2 = \sum_{\lambda \in \hat{G}} d_\lambda \|\hat{f}(\lambda)\|_{HS}^2 \quad (36)$$

where $\|A\|_{HS}$ is the Hilbert-Schmidt norm of A .

- (4) The Fourier transform $\mathcal{F} : f \rightarrow \hat{f}$ is an isometry from $L^2(G)$ onto $L^2(\hat{G})$.

(5) $\mathcal{R}(G)$ is dense in $C(G)$, equipped with the sup norm.

Notice that the one assumption we carried throughout was that we had a topological group. Since this does not come with a smooth structure or metric, there is nothing we can say a priori about topological groups with regards to the Laplacian since no definition makes much sense without a metric on the space. However, we are making use of the Haar measure. So, we could also ask, "How does the Laplacian operate on a measurable function in $L^2(G)$?". Notice that for a measurable function, we very likely do not have differentiation (measurable functions are not necessarily even continuous!). However, we do have a way of defining a derivative on functions in $L^2(G)$, namely, by way of the Fourier Transform. We say then that the Laplacian acts on a function in L^2 by way of multiplication by a polynomial: apply the Fourier transform to the function, apply multiplication by the characteristic polynomial of the Laplacian (using properties 2 and 3 from Theorem 4.1), then apply the inverse Fourier Transform. So in some sense, there is a Laplacian available to us.

8.4. Remark on Locally Compact Groups. For the sake of completion, we digress momentarily to make a remark regarding locally compact groups. For locally compact abelian groups, we may consider \mathbb{R} as a model space, and generalize this framework as needed, the move to generalizing locally compact groups. Much easier said than done. This generalization was one of the crowning achievements of twentieth century mathematics, led in no small way by Harish-Chandra and others.

9. SYMMETRIC SPACES

We now move on to a more abstract setting: symmetric spaces. Symmetric spaces do not a priori come with a nice multiplication structure, so we cannot leverage the Peter-Weyl Theorem to get general results about these spaces. However, we do have other important pieces of geometric structure: a metric, a smooth structure on a manifold, and geodesic reversing isometries at each point. So while there may not be a group structure on the space itself, we can leverage the group structure of the isometries of the space.

Definition 9.1 (Symmetric Space). We say a Riemannian manifold is a symmetric space if at each point on the manifold, there is a geodesic reversing isometry.

9.1. Eigenfunctions of the Laplacian on a Compact Riemannian Manifold. Given a compact Lie Group G , we can apply the Peter-Weyl theorem and see that there is a decomposition of functions in terms of a group representation. The fact that we have a smooth structure on the group means that we can use the definition of the Laplacian given in Section 2, and show that the Inverse Fourier Transform gives a decomposition of functions in $L^2(G)$, where the basis consists of eigenfunctions of the Laplacian. This may lead one to ask whether or not such a decomposition exists for Riemannian manifolds. Our first order of business is to prove such a decomposition exists. In this section, we answer the question in the affirmative, following Buser's proof [2], which follows a proofs of Chavel [4] and Dodziuk [8]. The proof uses a method common in the study of Partial Differential Equations. We are able to define integral operators that give solutions to the heat equation with a special property: these operators form a semigroup. This, combined with a special connection the heat equation shares with Laplace's equation, enables us to prove the desired decomposition of L^2 of a compact smooth Riemannian manifold. Rosenberg uses a similar heat kernel technique to serve as a jumping off point for understanding the Laplacian on a Riemannian Manifold, and the study of Index Theory [25].

Recall from Section 2 that we have defined the Laplacian on a Riemannian Manifold (M, g) . Since the heat equation will play a central role in this subsection, we will need a definition of a "fundamental solution to the heat equation", and the subsequent existence and uniqueness theorem.

Definition 9.2. (Fundamental solution to the heat equation on M) Let M be any Riemannian manifold without boundary. A continuous function $p = p(x, y, t) : M \times M \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$ is called a fundamental solution to the heat equation on M if it belongs to $C^{2,1}(M \times M \times \mathbb{R}_{>0})$ and satisfies:

- (a) $\frac{\partial p}{\partial t} = -\Delta_x p$, where Δ_x denotes the Laplacian with respect to the first argument,
 (b) $p(x, y, t) = p(y, x, t)$, and
 (c) $\lim_{t \rightarrow 0^+} \int_M p(x, y, t) f(y) dM(y) = f(x)$, where the convergence is locally uniform for test functions f , and by test function we mean a function with compact support on M .

The third property should be reminiscent of a Dirac sequence in distribution theory, where the limit of the sequence is the Dirac- δ integrated against f . In fact, this *is* a Dirac sequence of positive type [31]. In the theory of distributions, the Dirac- δ is also called an approximate identity.

Next, we state the existence and uniqueness theorem for the fundamental solution of the heat equation on a compact connected Riemannian manifold without boundary. We will use this without proof. The interested reader should refer to Chavel's text [4].

Theorem 9.1 (Existence and Uniqueness of the Heat Kernel on a Manifold). Let M be any n -dimensional manifold without boundary. Then M has a unique fundamental solution of the heat equation, denoted p_M . The function p_M is in $C^\infty(M \times M \times \mathbb{R}_{>0})$. For $0 < t < 1$, p_M has the following bounds, where the constant c_M depends on M .

$$0 \leq p_M(x, y, t) \leq c_M t^{\frac{1}{2}}$$

Now we can state the main theorem of this subsection.

Theorem 9.2 (Spectral Theorem of the Laplacian on a Compact Riemannian Manifold). Let M be a compact connected Riemannian manifold without boundary. The eigenvalue problem

$$\Delta f = \lambda f \tag{37}$$

has a complete orthonormal system of C^∞ -functions ϕ_0, ϕ_1, \dots in $L^2(M)$ with corresponding eigenvalues $\lambda_0, \lambda_1, \dots$. These have the following properties.

- (a) $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$, where $\lambda_n \rightarrow \infty$ as $n \rightarrow \infty$.
 (b) $p_M(x, y, t) = \sum_{n=0}^{\infty} e^{-\lambda_n t} \phi_n(x) \phi_n(y)$

In order to prove this theorem, we will black box two theorems from spectral geometry, the Hilbert-Schmidt Theorem and Mercer's Theorem, which we will state below.

Theorem 9.3 (Hilbert-Schmidt Theorem). Let M be a compact connected Riemannian manifold and let \mathcal{K} be the integral operator defined by

$$\mathcal{K}[f](x) = \int_M K(x, y) f(y) dM(y), f \in L^2(M) \tag{38}$$

where $K : M \times M \rightarrow \mathbb{R}$ is a continuous function which is symmetric: $K(x, y) = K(y, x)$. Then the eigenvalue problem $\mathcal{K}[\phi] = \eta \phi$ has a complete orthonormal system of eigenfunctions ϕ_0, ϕ_1, \dots in $L^2(M)$ with corresponding eigenvalues η_0, η_1, \dots , where $\eta_n \rightarrow 0$ as $n \rightarrow \infty$. The kernel K has the following expansion in the L^2 -sense:

$$K(x, y) = \sum_{n=0}^{\infty} \eta_n \phi_n(x) \phi_n(y) \tag{39}$$

Theorem 9.4 (Mercer's Theorem). Let M , K , and \mathcal{K} be as stated in the Hilbert-Schmidt Theorem above. Assume that almost all eigenvalues η_0, η_1, \dots are non-negative. Then K has the expansion

$$K(x, y) = \sum_{n=0}^{\infty} \eta_n \phi_n(x) \phi_n(y) \tag{40}$$

where the convergence of the series is uniform on $M \times M$.

With all of our conveniently black-boxed theorems, we can follow through the proof in a way that highlights the relationship between the Laplacian and the heat kernel. The proof of the Spectral Theorem above will follow from six Lemma's. We will start with a definition of the heat equation on a manifold and what we mean by a solution.

Definition 9.3 (Solution to the Heat Equation with Initial Condition). Let M be a compact, connected Riemannian manifold without boundary and $f : M \rightarrow \mathbb{R}$ be a continuous function. A continuous function $u = u(x, t) : M \times \mathbb{R}_{\geq 0}$ is called a *solution to the heat equation with initial condition $f(x)$* if $u \in C^{2,1}(M \times \mathbb{R}_{>0})$, and if u satisfies the heat equation:

(a) $\frac{\partial u}{\partial t} = -\Delta u$, where Δ is applied to the first variable $x \in M$.

and satisfies the initial conditions

(b) $u(x, 0) = f(x)$ for all $x \in M$.

Now, we will use the existence of a fundamental solution to show that for any initial condition f , where $f : M \rightarrow \mathbb{R}$ is a continuous function, there exist a solution to the heat equation.

Lemma 9.1 (Existence Theorem). Let $p = p(x, y, t)$ be a fundamental solution to the heat equation on a compact, connected Riemannian manifold without boundary M . Let $f : M \rightarrow \mathbb{R}$ be a continuous function. Then

$$u(x, t) := \int_M p(x, y, t) f(y) dM(y), \text{ for } t > 0 \quad (41)$$

has a continuous extension to $t = 0$, and this extended function is a solution to the heat equation for initial condition $u(x, 0) = f(x)$.

Proof. First notice that the integral converges. Since p and f are continuous, their product is continuous, and the product achieves a maximum on M since M is compact. Then the integral is bounded by this maximum and the volume of M . Additionally, since M is compact and partial derivatives of $p(x, y, t)f(y)$ with respect to t and x are continuous, we may pass the derivatives as needed.

$$\begin{aligned} \frac{\partial u}{\partial t} &= \frac{\partial}{\partial t} \int_M p(x, y, t) f(y) dy \\ &= \int_M \frac{\partial p}{\partial t}(x, y, t) f(y) dy \\ &= \int_M -\Delta_x p(x, y, t) f(y) dy \\ &= -\Delta_x \int_M p(x, y, t) f(y) dy \\ &= -\Delta u(x, t) \end{aligned}$$

The continuous extension to $t = 0$ follows from property (c) of the fundamental solution $p(x, y, t)$. (This is a property common among Dirac-sequences.) ■

Next, we show that this solution u has a few valuable properties, and prove uniqueness of the solution.

Lemma 9.2 (Uniqueness Theorem). The solution to the heat equation in the preceding Lemma is unique, and has the following properties.

(a) $\frac{d}{dt} \int_M u dM = 0$ for $t > 0$.

(b) $\frac{d}{dt} \int_M u^2 dM \leq 0$ for $t > 0$.

Proof. Let u be any solution. The reader should be aware that this integration is happening locally over finitely many charts on M . Refer to [22] for details. For $t > 0$, we have the following.

$$\begin{aligned}
\frac{d}{dt} \int_M u dM &= \int_M \frac{\partial u}{\partial t} dM \\
&= - \int_M \Delta u dM \\
&= 0, \text{ by integration by parts applied to } \Delta u \cdot 1.
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
\frac{d}{dt} \int_M u^2 dM &= \int_M u \frac{\partial u}{\partial t} + \frac{\partial u}{\partial t} u dM \\
&= -2 \int_M u \Delta u dM \\
&= -2 \int_M \langle \text{grad } u, \text{grad } u \rangle dM, \text{ by integration by parts} \\
&= -2 \int_M \|\text{grad } u\|^2 dM \\
&\leq 0.
\end{aligned}$$

With these properties, uniqueness follows. ■

In fact, while this proves uniqueness of the solution, it is not difficult to show the fundamental solution is also unique. Simply let p and q be fundamental solutions, and notice that by the uniqueness we just proved, we have

$$\int_M p(x, y, t) f(y) dM(y) = \int_M q(x, y, t) f(y) dM(y)$$

for all $t > 0$. The result follows by linearity of the integral and positivity of the fundamental solution. Consult Buser's text for details [2].

Next, we will show the fundamental solution has a very nice property, making it a kernel.

Lemma 9.3. $\int_M p(x, y, t) dM(y) = 1$

Proof. Notice that $f(y) = 1$ is a continuous function on the manifold M . By our existence theorem, there is a solution $u(x, t) = \int_M p(x, y, t) dM(y)$ with initial condition $f(y) = 1$. Next, notice that $u(x, t) = 1$ is a solution to the heat equation with initial condition $f(x) = 1$. By uniqueness, we see that $1 = u(x, t) = \int_M p(x, y, t) dM(y)$ is the unique solution, as desired. ■

The sequence of lemmas above works nicely in the case of a compact space. For a non-compact space, such as \mathbb{R} , the heat kernel has the same property, but one would need to work a bit more to prove it. See [31].

We will now re-interpret our machinery: we can think of the integral in the existence theorem as an operator which takes a continuous function on the manifold as an input. We will extend this notion as follows.

Definition 9.4 (\mathcal{P}_t). For every f in $L^2(M)$ and for every $t > 0$, we define

$$\mathcal{P}_t[f](x) := \int_M p(x, y, t) f(y) dM(y), \text{ where } x \in M \tag{42}$$

Now, we can show that this is actually a very nice operator, with smoothing properties. This should not come as a surprise, since kernels produce functions with properties similar to the kernel.

Lemma 9.4. Each \mathcal{P}_t is a compact positive self-adjoint operator. For any f in $L^2(M)$, $\mathcal{P}_t[f]$ is in $C^\infty(M)$.

Proof. First notice that $\mathcal{P}_t[f]$ is well-defined since the integral converges. To see this, apply Hölder's Inequality. Then, as before, since M is compact, we may pass partial derivatives of x and t through the integral. Since $p(x, y, t)$ is a smooth function, we see that $\int_M p(x, y, t) f(y) dM(y)$ is a smooth function in x and t .

Now, recall that a operator $T : E \rightarrow F$ is compact if it is a bounded operator such that $\overline{T(E)}$ is compact in F . The operator is clearly linear by linearity of the integral. Boundedness of this operator is immediate from Hölder's inequality. Compactness of the operator follows from the fact that M is compact and has finite volume, and the result follows. See [34] for details.

Next, from the second property of the fundamental solution to the heat equation, we know p is a symmetric kernel. By Fubini,

$$\begin{aligned} \langle f, \mathcal{P}_t[g] \rangle &= \int_M \int_M f(x) p(x, y, t) g(y) dM(y) dM(x) \\ &= \int_M \int_M f(x) p(y, x, t) g(y) dM(x) dM(y) \\ &= \langle \mathcal{P}_t[g], f \rangle. \end{aligned}$$

Positivity of the operator follows from the next lemma. ■

Now, we can identify one of the central properties we need to leverage to prove the Spectral Theorem for the Laplacian: these integral operators have the semigroup property.

Lemma 9.5. For $s > 0$ and $t > 0$ we have the semigroup property:

$$\mathcal{P}_s \circ \mathcal{P}_t = \mathcal{P}_{s+t} \tag{43}$$

In particular, this means $\mathcal{P}_s \circ \mathcal{P}_t = \mathcal{P}_t \circ \mathcal{P}_s$.

Proof. Since $C^0(M)$ is dense in $L^2(M)$, and \mathcal{P}_t is a continuous operator for any t , we can prove the statement for any $f \in C^0(M)$ and bootstrap to attain the desired result. Fix $t > 0$ and let $f \in C^0(M)$. First, we show that $\mathcal{P}_s \circ \mathcal{P}_t$ is a fundamental solution to the heat equation on M . By applying Fubini's Theorem, we can see that

$$\mathcal{P}_s \circ \mathcal{P}_t[f] = \int_M \left(\int_M p(z, x, s) p(x, y, t) dM(x) \right) f(y) dM(y),$$

so in particular, we can express $\mathcal{P}_s \circ \mathcal{P}_t$ as

$$\mathcal{P}_s \circ \mathcal{P}_t = \int_M p(z, x, t' - t) p(x, y, t) dM(x)$$

where we make the change of variables $s + t = t'$. The reader can check to verify that properties (a) and (b) of Definition 9.2 are satisfied. To see that the third property is satisfied, consider the following:

$$\begin{aligned}
\lim_{t' \rightarrow 0^+} \mathcal{P}_s \circ \mathcal{P}_t[f] &= \int_M \left(\int_M p(z, x, t' - t) p(x, y, t) dM(x) \right) f(y) dM(y) \\
&= \lim_{t' \rightarrow 0^+} \int_M p(z, x, t' - t) \int_M p(x, y, t) f(y) dM(y) dM(x) \\
&= \lim_{t' \rightarrow 0^+} \int_M p(z, x, t') \lim_{t \rightarrow 0^+} \int_M p(x, y, t) f(y) dM(y) dM(x) \\
&= \lim_{t' \rightarrow 0^+} \int_M p(z, x, t') f(x) dM(x) \\
&= f(z).
\end{aligned}$$

Since the integral is over a compact space, and the functions continuous, we may pass the limit. Notice that as t' goes to 0, so must t . This gives us the desired result.

Then, the fact that this kernel is equivalent to \mathcal{P}_{s+t} follows from the uniqueness of the kernel, Theorem 9.1, and uniqueness of the solutions, Lemma 9.2. \blacksquare

Next, we prove one more statement that we will need. Notice that the statement holds for f continuous by the existence lemma. We will need to do a bit of bootstrapping to achieve the following.

Lemma 9.6. For any $f \in L^2(M)$ we have

$$\lim_{t \rightarrow 0^+} \mathcal{P}_t[f] = f \quad (44)$$

where convergence is with respect to the L^2 norm.

Proof. Since continuous functions are dense in $L^2(M)$ (M is compact), we will prove the statement for continuous functions and the result will follow from a bootstrapping argument. Let f be continuous, and we know $\mathcal{P}_t[f]$ is a solution to the heat equation with initial condition f . Then, by Lemma 9.2,

$$\begin{aligned}
\frac{d}{dt} \|\mathcal{P}_t[f]\|_{L^2} &= \frac{d}{dt} \left(\int_M \mathcal{P}_t[f]^2 dM \right)^{\frac{1}{2}} \\
&= \frac{1}{2} \left(\int_M \mathcal{P}_t[f]^2 dM \right)^{-\frac{1}{2}} \frac{d}{dt} \int_M \mathcal{P}_t[f]^2 dM \\
&\geq 0.
\end{aligned}$$

Then, since property (c) of the fundamental solution tells us that $\lim_{t \rightarrow 0^+} \int_M p(x, y, t) f(y) dM(y) = f(x)$ where the convergence is uniform on M , compact, we have that $\|\mathcal{P}_t[f]\|_{L^2} \leq \|f\|_{L^2}$. The statement follows from the reverse triangle inequality. For the general case, let f be in $L^2(M)$. There exists a sequence $f_n \rightarrow f$ in L^2 , where f_n for $n \in \mathbb{N}$ is continuous. Fix any $\epsilon > 0$. Pick n such that $\|f_n - f\|_{L^2} < \frac{\epsilon}{3}$. Pick t such that $\|\mathcal{P}_t[f_n] - f_n\|_{L^2} < \frac{\epsilon}{3}$.

$$\begin{aligned}
\|\mathcal{P}_t[f] - f\|_{L^2} &= \|\mathcal{P}_t[f] - \mathcal{P}_t[f_n] + \mathcal{P}_t[f_n] - f_n + f_n - f\|_{L^2} \\
&\leq \|\mathcal{P}_t[f] - \mathcal{P}_t[f_n]\|_{L^2} + \|\mathcal{P}_t[f_n] - f_n\|_{L^2} + \|f_n - f\|_{L^2} \\
&< \|f_n - f\|_{L^2} + \|\mathcal{P}_t[f_n] - f_n\|_{L^2} + \|f_n - f\|_{L^2}, \text{ by Hölder's Inequality} \\
&< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \\
&= \epsilon.
\end{aligned}$$

\blacksquare

Now we have all of the necessary machinery to prove our spectral theorem.

Proof of Theorem 9.2. We start by applying the Hilbert-Schmidt theorem to each of the operators \mathcal{P}_t . For $t = 1$, let $\phi_0, \phi_1, \phi_2, \dots$ be the eigenfunctions of \mathcal{P}_1 that form a complete orthonormal system in $L^2(M)$ with the corresponding eigenvalues $\eta_0, \eta_1, \eta_2, \dots \geq 0$, where $\eta_j \rightarrow \infty$ as $j \rightarrow \infty$.

Next, notice that $\phi_0, \phi_1, \phi_2, \dots$ are eigenfunctions of \mathcal{P}_t for any $t > 0$. To see this, we use the semigroup property. First, let $t = \frac{1}{k}$ and notice $\mathcal{P}_1 = (\mathcal{P}_{\frac{1}{k}})^k$. If ϕ is an eigenfunction of $\mathcal{P}_{\frac{1}{k}}$ with eigenvalue η , then ϕ is also an eigenfunction of \mathcal{P}_1 , but with eigenvalue η^k . Then, since we have a complete set of eigenfunctions for $\mathcal{P}_{\frac{1}{k}}$, \mathcal{P}_1 and \mathcal{P}_1 have the same orthogonal system of eigenspaces. This means that $\phi_0, \phi_1, \phi_2, \dots$ are eigenfunctions for both \mathcal{P}_1 and $\mathcal{P}_{\frac{1}{k}}$. The corresponding eigenvalues for $\mathcal{P}_{\frac{1}{k}}$ are $\eta_0^{\frac{1}{k}}, \eta_1^{\frac{1}{k}}, \eta_2^{\frac{1}{k}}, \dots$.

Next, we apply the semigroup property once more to get that ϕ_j are eigenfunctions for \mathcal{P}_t with eigenvalues η_j^t for any positive $t \in \mathbb{Q}$. Then, we can use the continuity of $p(x, y, t)$ to get that the statement holds for all $t > 0$.

By Lemma 9.4, we see that ϕ_j is smooth for all j . By Lemma 9.6, $\lim_{t \rightarrow 0^+} \mathcal{P}_t[\phi_j] = \phi_j$, so $\eta_j^t \rightarrow 1$ as $t \rightarrow 0^+$. Additionally, we see that $\eta_j > 0$ for all $j \in \mathbb{N}$. Then, by the compactness of \mathcal{P}_1 , we can conclude all of the eigenspaces are finite dimensional, and we can arrange the eigenvalues in decreasing order.

Now, we claim $1 = \eta_0 > \eta_1 \geq \eta_2 \geq \dots \geq 0$. To see this, first notice that Lemma 9.3 can be interpreted as telling us that 1 is an eigenvalue for the constant function. Next, let ϕ_j be a non-constant eigenfunction of \mathcal{P}_1 . Then, using this slight improvement to Lemma 9.2,

$$\begin{aligned} \frac{d}{dt} \|\mathcal{P}_t[\phi_j]\|^2 &= -2 \int_M \langle \text{grad } \mathcal{P}_t[\phi_j], \text{grad } \mathcal{P}_t[\phi_j] \rangle dM \\ &= -2\eta_j^{2t} \int_M \langle \text{grad } \phi_j, \text{grad } \phi_j \rangle dM \\ &< 0 \end{aligned}$$

since $\eta_j > 0$. This is a slight variation on Lemma 9.6, and gives us the following strict inequality: $\|\mathcal{P}_t[\phi_j]\| < \|\phi_j\|$, from whence we can conclude $\eta_j < 1$.

Now, we can finish the proof. Let

$$\lambda_j = -\log \eta_j \text{ for } j = 0, 1, 2, \dots$$

Then since $\mathcal{P}_t[\phi_j]$ is a solution to the heat equation, we have $\frac{\partial}{\partial t} \mathcal{P}_t[\phi_j] = -\Delta \mathcal{P}_t[\phi_j]$. Then,

$$\begin{aligned} 0 &= \Delta \mathcal{P}_t[\phi_j] + \frac{\partial}{\partial t} \mathcal{P}_t[\phi_j] \\ &= \Delta \eta_j^t \phi_j + \frac{\partial}{\partial t} \eta_j^t \phi_j \\ &= \eta_j^t \Delta \phi_j + (\log \eta_j) \eta_j^t \phi_j \\ &= \eta_j^t (\Delta \phi_j + (\log \eta_j) \phi_j) \\ &= e^{-t\lambda_j} (\Delta \phi_j - \lambda_j \phi_j), \end{aligned}$$

and we see that the ϕ_j are also eigenfunctions of the Laplacian with eigenvalues λ_j ! In fact, the λ_j have the desired properties. To see the second statement in the theorem, apply Mercer's Theorem to the compact positive operator \mathcal{P}_t . ■

9.2. **The flat plane, \mathbb{R}^2 .** Now, we return to the discussion on symmetric spaces. The groundwork has been set: any compact quotient of one of the model symmetric spaces results in a space where the eigenfunctions of the Laplacian decompose L^2 . Thus, we can define fourier series on these spaces. We will define transforms on the covering space here, where the quotients have the induced series.

The case of the flat plane is handled by section 4 above, where we let $n = 2$. The resulting Fourier Transform has all of the desired properties.

9.3. **The sphere, S^2 .** For the sphere, we can leverage its isometry group, $SO(3)$. Since $SO(3)$ is a compact group, we can use the Peter-Weyl theorem to generate spherical harmonics [31]. Alternatively, we can construct a transform by computing the eigenfunctions of the sphere on the Laplacian. The result is the Hankel Transform. See [31] for details.

9.4. **The Hyperbolic Plane, \mathbb{H} .** For the Hyperbolic plane, we cannot use Peter-Weyl since $SL_2(\mathbb{R})$ is not a compact group. However, we could begin the study of semisimple Lie Groups, and look for ways to leverage the representation theory of $SL_2(\mathbb{R})$ to generate transforms. For now, though, we will suggest a more straightforward path. Using separation of variables, we can compute the eigenfunctions of the Laplacian on \mathbb{H} , and use these to generate a transform. See [31] for details. The resulting transform is the Kontorovich-Lebedev Transform.

10. CONNECTIONS TO GEODESICS

10.1. **Eigenvalues and Isospectrality.** In this section, we follow Buser's text [2] and connect the spectrum of the Laplacian to the length spectrum. Additionally, we state Sunada's Theorem, which gives us a technique to answer the following question: can you hear the shape of a Riemann Surface?

The reader should be aware though. While translation surfaces are all topologically equivalent to a Riemann Surface (refer to Section 11), the surfaces in this section are assumed to have a hyperbolic metric. Some of the following theorems do not require constant curvature, and the author will attempt to point this out when it becomes relevant. That being the case, there may be an opportunity to adapt proofs of these theorems to translation surfaces. One may hope that our definition of a Laplacian on a translation surface will lead to similar consequences.

10.2. **Huber's Theorem.** First, we need a few definitions.

Definition 10.1 (Oriented closed geodesics). We say two parametrized closed geodesics $\gamma, \gamma' : S^1 \rightarrow M$ are *equivalent* if there is a homeomorphism $f : S^1 \rightarrow S^1$ of the form $f(t) = t + c$ for some real constant c such that $\gamma' = \gamma \circ f$. We then say that a *oriented closed geodesic* is an equivalence class of closed parametrized geodesics.

Definition 10.2 (Primitive oriented closed geodesics). Let γ and δ be closed geodesics and let $m \in \mathbb{Z} \setminus \{0\}$. We say that γ is the *m-fold iterate* of δ if $\gamma(t) = \delta(mt)$ for $t \in S^1$. A closed geodesic (oriented or not) is said to be *primitive* if it is not the *m-fold iterate* of another closed geodesic for some $m \geq 2$. We also call an *orientated primitive closed geodesic* a *prime geodesic*.

Definition 10.3 (Length spectrum). The sequence arranged in ascending order of all lengths of oriented closed geodesics on a surface is called the *length spectrum*.

Definition 10.4 (Primitive length spectrum). The sequence arranged in ascending order of all lengths of primitive oriented closed geodesics is called the *primitive length spectrum*. Notice that this is a subsequence of the length spectrum.

Now we can state Huber's Theorem.

Theorem 10.1 (Huber's Theorem). Two compact Riemann Surfaces of genus $g \geq 2$ have the same spectrum of the Laplacian if and only if they have the same length spectrum.

10.3. Sunada's Theorem.

Definition 10.5 (Almost Conjugate or Gassmann equivalent). Let G be a finite group. Two subgroups H_1, H_2 of G are called *almost conjugate* or *Gassman equivalent* if for all $g \in G$

$$|[g] \cap H_1| = |[g] \cap H_2|. \quad (45)$$

where $[g]$ denotes the conjugacy class of an element, $\{\sigma g \sigma^{-1} | \sigma \in G\}$.

Theorem 10.2 (Sunada's Theorem). Let M be a complete Riemannian manifold and let G be a finite group acting on M by isometries with at most finitely many fixed points. If H_1 and H_2 are almost conjugate subgroups of G acting freely on M , then the quotients of $H_1 \backslash M$ and $H_2 \backslash M$ are isospectral.

Sunada's original proof followed what he calls a rather routine technique from number theory, adapted to a geometric setting [29].

10.4. Can you hear the shape of a Translation Surface? This suggests a fascinating question, reminiscent of the elegant question posed by Kac [20]: can you hear the shape of a translation surface? Kac's original question was "Can you hear the shape of a drum?", a reference to polygonal domains on \mathbb{R}^2 . Interestingly, Kac does mention that Bochner had asked him essentially the same question ten years prior. In fact, before Kac asked the question, Milnor found two 16-dimensional isospectral flat tori that are not isometric [24] (because of course he did). Kac notes this in his original paper. Later, Ikeda found examples of isospectral, but not isometric, lens spaces, quotients of S^3 by a \mathbb{Z}/p -action [19]. Vignéras found an example of isospectral, but not isometric, Riemann surfaces with constant negative curvature [32].

Gordon, Webb, and Wolpert answered Kac's question about polygonal domains in \mathbb{R}^2 negatively 26 years later [14]. They used Sunada's construction to create non-isomorphic surfaces that are isospectral. Buser, Conway, Doyle, and Semmler identified several polygonal domains with the same property only a couple years later [3].

It is, however, possible that the translation structure required of the polygonal domains of a translation surface provide the necessary structure to eliminate these examples. Moon, D., et al, have recently answered a variation of the question in the affirmative for Billiard tables, where they have put an equivalence on the space by affine transformations [9]. Billiard tables unfold to become a small class of translation surfaces. They have constructed a new spectrum (not the spectrum of the Laplacian), they call the *bounce spectrum*. The techniques they use stem from symbolic dynamics. They label the edges of a polygon by letters, and generate sequences of letters as they keep track of what edge the geodesic hits. We note immediately, as did they, that the bounce spectrum is unchanged by dilations, rotations, and translations. This is interesting because the spectrum of the Laplacian can hear "area", noted by Kac [20], but the bounce spectrum cannot. We should pose the question though: is it necessary to hear the area? What information are we trying to glean from the spectrum?

10.5. Eigenfunctions and Quantum Ergodicity. In physics, there is a phenomenon that occurs in quantum chaos that physicists have named "scarring". The non-rigorous version of this is that nodal sets (zero sets) of the eigenfunctions of the Laplacian tend to localize around unstable orbits.

This is an open area of research in mathematics, where one of the primary open questions concerns scarring on the Bunimovich stadium. The reader is encouraged to read Terence Tao's blog post as an introduction to the topic. One fascinating connection is to quantum ergodicity, and Rudnick and Sarnak's conjecture on Quantum Unique Ergodicity. We define quantum ergodicity as the property that some spaces have - the square of orthonormal eigenfunctions of the Laplacian tend to the volume form on the space, almost always. Quantum unique ergodicity removes the almost always.

For the interested reader, there is a route into this area of mathematics via Translation surfaces. It begins with the equilateral triangle that unfolds into a hexagon [33]. We have explicit eigenfunctions on this triangle thanks to Gabriel Lamé, a French mathematician. We can then study questions regarding nodal sets.

11. TERMINUS A QUO: LAPLACIAN ON A TRANSLATION SURFACE

We end where we began. How does one go about defining the Laplacian? We have found the tip of many rather large icebergs: a connection between the Fourier Transform and the Laplacian, the role of representation theory.

In the introductory remarks, we hinted at two different definitions of a translation surface: a polygon in the plane with parallel sides identified and a Riemannian manifold with singularities whose transition functions consist of translations. It is not hard to see these two things are the same. There are in fact three equivalent definitions for a *translation surface* commonly in use by researchers [33], but we will focus only on these two.

Definition 11.1 (Translation Surface as a Polygon). A translation surface is an equivalence class of polygons in the plane with edge identifications: Each translation surface is a finite union of polygons in \mathbb{C} , together with a choice of pairing of parallel sides of equal length that are on “opposite sides.” (So for example two horizontal edges of the same length can be identified only if one is on the top of a polygon, and one is on the bottom. Each edge must be paired with exactly one other edge. These conditions are exactly what is required so that the result of identifying pairs of edges via translations is a closed surface.) Two such collections of polygons are considered to define the same translation surface if one can be cut into pieces along straight lines and these pieces can be translated and re-glued to form the other collection of polygons. When a polygon is cut in two, the two new boundary components must be paired, and two polygons can be glued together along a pair of edges only if these edges are paired.

Definition 11.2 (Translation Surface as a Riemannian Manifold with Singularities). A translation surface is a closed topological surface X , together with a finite set of points Σ and an atlas of charts to $\text{Con } X \setminus \Sigma$ whose transition maps are translations, such that at each point p_0 of Σ there is some $k > 0$ and a homeomorphism of a neighborhood of p_0 to a neighbourhood of the origin in the $2k + 2$ half plane construction that is an isometry away from p_0 . The singularity at p_0 is said to have cone angle $2\pi(k + 1)$, since it can be obtained by gluing $2k + 2$ half planes, each with an angle of π at the origin. The term “cone point” is another synonym for “singularity.”

For a proof of the equivalence, the reader should see Wright’s survey [33].

It makes sense then to consider a Laplacian on each object, and ensure that the definition agrees.

11.1. Translation Surface as a Polygon. Here, the somewhat obvious choice is to use the usual Laplacian in \mathbb{C} , where we invoke a set of boundary conditions based on the identifications via translations.

11.2. Translation Surface as a Riemannian Manifold with Cone Points. To define a Laplacian on a Riemannian manifold with singularities would take a substantial amount of work. Luckily, Jeff Cheeger did this work in the late 1970s and early 1980s [5] [6]. His work generalized existing structures on smooth manifolds, effectively enabling analysis on cone manifolds. His methodology required that he define an L^2 – *cohomology* of the space, and with this, generalize Hodge Theory. The end result was the ability to define a coordinate-free version of the Laplacian on cone manifolds. This work will be a phenomenal starting point for defining a Laplacian on a translation surface.

12. CONCLUSION AND NEXT STEPS

The possible next steps are numerous, but all involve starting with a definition. Once we thoroughly investigate Cheeger’s construction, we can land on a definition. And once we understand the definition, we can ask the following concrete questions.

- (1) Does our definition yield useful results on square-tiled surfaces (covers of \mathbb{T}^2)? How does this relate to the work of Hillairet [17]?

- (2) How does this relate to the work of Hillairet and Judge [16]? Can we say something about the Golden L?

Our work has also led us to the following more abstract, big-picture questions.

- (1) Can you hear the shape of a translation surface?
- (2) Is there a useful formulation of Selberg's Trace formula?
- (3) Is there a Prime Number Theorem for translation surfaces?
- (4) Can we relate the nodal sets of the eigenfunctions of the Laplacian to geodesics? (Quantum Scarring)
- (5) Does quantum ergodicity hold for translation surfaces? Quantum unique ergodicity?

Clearly, Professor Athreya's question has led us to many more questions, and in the coming years, we hope to see progress in this area.

REFERENCES

- [1] S.C. Bagchi, S. Madan, A. Sitaram, and U.B. Tewari, *A First Course on Representation Theory and Linear Lie Groups*, Universities Press, Hyderabad, 2000.
- [2] P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Birkhäuser, Boston, 1992.
- [3] P. Buser, J. Conway, P. Doyle, and K.D. Semmler, *Some Planar Isospectral Domains*, Math. Res. Not. **9** (1994), 391-400.
- [4] I. Chavel, *Eigenvalues in Riemannian Geometry*, Academic Press, Orlando-San Diego-New York, 1984.
- [5] J. Cheeger, *On the Spectral Geometry of Spaces with Cone-like singularities*, Proc. Nat. Academy of Sciences (1979).
- [6] ———, *Spectral Geometry of Singular Riemannian Spaces*, J. Differential Geometry (1983).
- [7] F. Chung, *Spectral Graph Theory*, 2nd ed., American Mathematical Society, Providence, 1997.
- [8] J. Dodziuk, *Eigenvalues of the Laplacian and the Heat Equation*, American Mathematical Monthly **88** (1981).
- [9] M. Duchin, V. Erlandsson, C. Leininger, and C. Sadanand, *You can hear the shape of a billiard table: Symbolic dynamics and rigidity for flat surfaces* (2018).
- [10] D. Dummit and R. Foote, *Abstract Algebra*, 3rd ed., John Wiley & Sons, USA, 2004.
- [11] H. Dym and H.P. McKean, *Fourier Series and Integrals*, Academic Press, Inc., New York, 1972.
- [12] M. Einsiedler and T. Ward, *Ergodic Theory with a view towards Number Theory*, Springer-Verlag, London, 2011.
- [13] G. Folland, *Real Analysis: Modern Techniques and Their Applications*, 2nd ed., John Wiley & Sons, Inc., New York, 1999.
- [14] C. Gordon, D. Webb, and S. Wolpert, *One cannot hear the shape of a drum*, Bulletin of the American Mathematical Society **27** (1992), no. 1, 134-138.
- [15] J. Hadamard, *An Essay on the Psychology of Invention in the Mathematical Field*, Princeton University Press, Princeton, 1945.
- [16] L. Hillairet and C. Judge, *Generic Spectral Simplicity of Polygons*, Proceedings of the American Mathematical Society **v. 137** (January 2009).
- [17] L. Hillairet, *Spectral Decomposition of square-tiled Surfaces*, Mathematische Zeitschrift **vol. 260, no. 2** (2008).
- [18] J. Humphreys, *Introduction to Lie Algebras and Representation Theory*, Springer-Verlag, New York, 1972.
- [19] A. Ikeda, *On spherical space forms which are isospectral but not isometric*, J. Math. Soc. **35** (1983), 437-444.
- [20] M. Kac, *Can one hear the shape of a drum?*, Amer. Math. Monthly **73** (1966), 1-23.
- [21] J. Lee, *Introduction to Riemannian Geometry*, 2nd ed., Springer, New York, 2019.
- [22] ———, *Introduction to Smooth Manifolds*, 2nd ed., Springer, New York, 2013.
- [23] ———, *Introduction to Topological Manifolds*, 2nd ed., Springer, New York, 2011.
- [24] J. Milnor, *Eigenvalues of the Laplace operator on certain manifolds*, Proc. Nat. Acad. Sci. **51** (1964), 542.
- [25] S. Rosenberg, *The Laplacian on a Riemannian Manifold*, Cambridge University Press, Cambridge, 1992.
- [26] W. Rudin, *Functional Analysis*, McGraw-Hill, United States of America, 1973.
- [27] W. Schlag, *A Course in Complex Analysis and Riemann Surfaces*, American Mathematical Society, Providence, 2014.
- [28] J. Stewart, *Calculus*, 8th ed., Cengage Learning, Boston, 2015.
- [29] T. Sunada, *Riemannian coverings and isospectral manifolds*, Ann. of Math. (2) **121** (1985), 169-186.
- [30] W. Strauss, *Partial Differential Equations: An Introduction*, 2nd ed., John Wiley & Sons, Inc., Hoboken, 2008.
- [31] A. Terras, *Harmonic Analysis on Symmetric Spaces and Applications I*, 2nd ed., Springer, New York, 2013.
- [32] M. F. Vignéras, *Variétés riemanniennes isospectrales et non isométriques*, Ann. of Math. **112** (1980), 21-32.
- [33] A. Wright, *Translation Surfaces and their Orbit Closures: An Introduction for a Broad Audience*, EMS Surv. Math. Sci. (2015).
- [34] R. Zimmer, *Essential Results of Functional Analysis*, The University of Chicago Press, Chicago, 1990.
- [35] A. Zorich, *Flat Surfaces*, Frontiers in Number Theory, Physics, and Geometry **Vol. 1** (2006).