

McGILL UNDERGRADUATE MATHEMATICS JOURNAL



the δ elta ϵ psilon

WINTER 2010

FOURTH ISSUE



McGill

THE δ ELTA- ϵ PSILON

MCGILL UNDERGRADUATE MATHEMATICS JOURNAL

CONTENTS

Distances between graphs and related problems <i>Hua Long Gervais</i>	3
May the Schwartz Be With You <i>Daniel Shapero</i>	11
Interview with Prof. David Wolfson <i>Ana Best</i>	18
Group Theory and Spectroscopy: A Chemist's View on Symmetry <i>Lai Chung Liu</i>	20
The Trepidation that is Math Anxiety <i>Maya Kaczorowski</i>	24
Écritures réduites de l'élément le plus long des groupes de Coxeter fini <i>Maxime Bergeron, Marco Robado et Maxime Scott</i>	28
Interview with Prof. Johanna Neslehova <i>Ana Best</i>	34
Spectral Theorem for Bounded Self-Adjoint Operators <i>Dana Mendelson and Alexandre Tomberg</i>	36
The Never-Ending Harmonic Series <i>Cyndie Cottrell</i>	42
Probabilistic Tracking of Pedestrians from Platform Force Data <i>Rishi Rajalingham</i>	47
Strange Cycling: Results and Examples of Delay Differential Equations as Dynamical Systems <i>Orianna DeMasi</i>	51
Interview with Prof. Russ Steele <i>Ana Best</i>	56
Rencontre numbers <i>Maya Kaczorowski</i>	59
Embeddings of Affine Lines in the Affine Plane and Polynomial Automorphisms <i>Hua Long Gervais</i>	62
White Noise <i>Daniel Shapero</i>	66

LETTER FROM THE EDITORS

For many of us, this year was both our first and our last on the δ elta- ϵ psilon editing team, and we had many ideas of what we considered to be a “great” journal. There have been countless discussions over the choice of articles, the size of the margins, appropriate cover art and which math jokes, if any, should be included in this year’s issue.

In the end, this issue satisfies our needs as editors, so we can only hope that you are also satisfied. (To next year’s editors, as you read this now for the first time the night before your deadline, keep going – it will be done soon.)

As in past years, the δ elta- ϵ psilon is once again in need of your help. As most of the editorial board graduates, this McGill math tradition must be kept alive. We need replacements, or there will be no more issues. Please send us an email or approach us in Burnside if you are interested – it’s a lot of fun, and an excellent way to learn the nuances of \LaTeX .

We encourage all of you undergraduates to keep contributing to the journal by sending in your articles. Summer research and independent studies are great ways to learn in depth about a particular topic that interests you.

Enjoy the articles and let us know any comments or suggestions you have.

Maya Kaczorowski
 Managing Editor
 The δ elta- ϵ psilon Editing Team
thedeltaepsilon@gmail.com



LETTER FROM SUMS

The Society of Undergraduate Mathematics Students (SUMS) would like to congratulate the δ elta- ϵ psilon on its fourth issue. We’re proud to sponsor such a successful undergraduate mathematics journal which showcases the accomplishments and research of undergraduate students at McGill University.

With yearly high-quality publications, the δ elta- ϵ psilon exemplifies that research is something accessible to undergraduates and supports McGill students in their effort to build on mathematical knowledge and foundations.

Sincerely,

Daphna Harel
 SUMS President
 (On behalf of SUMS council)
<http://sums.math.mcgill.ca/>



DISTANCES BETWEEN GRAPHS AND RELATED PROBLEMS

Hua Long Gervais

Given a fixed set of vertices, the problem considered here is to define a metric on the set of all k -regular graphs constructed from it. After basic notions such as Cayley graphs and adjacency matrices are introduced, various candidates for the sought metric are defined. We then derive some inequalities involving the various ideas of distance and look at whether they are equivalent or not.

1 BASIC NOTIONS AND FACTS

1.1 Types of graphs

Graphs are typically thought of as sets of vertices, some of which are linked by edges. An example of a graph in the real world is the internet, where the vertices are the various users, and the edges represent network connections between users. There are several types of graphs, which we classify according to whether they are simple graphs or multigraphs, and directed or undirected.

Definition. An *undirected simple graph* is an ordered pair $G = (V, E)$ where V is a set, and E is a subset of the subsets of size 2 of V .

The elements of V are called the vertices (or nodes) of G , and those of E are the edges of G . Rather than thinking about an ordered pair of sets, we typically visualize a graph by portraying the vertices as dots on a plane and the edges as lines that connect the dots. Note that our definition is made so that there is at most one edge between any two vertices and no vertex is connected to itself, this is what we mean by saying that the graph is simple. We say that for $u, v \in V$, u is adjacent to v if $\{u, v\} \in E$. Also, we call v a neighbour of u is v is adjacent to u . We assume that V is finite. Here is an example of an undirected simple graph:

Example.

$$V = \{1, 2, 3, 4, 5\}$$

$$E = \{\{1, 2\}, \{2, 3\}, \{2, 5\}, \{3, 5\}\}$$




Figure 1: This is an undirected simple graph.

Although we will not use this notion, we point out that if we define E to be a subset of the ordered subsets of size 2 of V , then we have a *directed graph*. In a directed graph, we think of the edge $\{u, v\}$ as having an orientation from u

to v , which is portrayed as an arrowhead on the edge when we visualize the graph.

Example.

$$V = \{1, 2, 3, 4, 5\}$$

$$E = \{\{1, 2\}, \{2, 3\}, \{2, 5\}, \{3, 5\}\}$$




Figure 2: This is a directed simple graph.

We can also allow multiple edges between vertices and that some vertices are adjacent to themselves. This type of graph is called a *multigraph* -as opposed to simple graphs. And again, we can decide whether edges are oriented or not for a multigraph.

Unless otherwise specified, we will use the word *graph* to designate undirected simple graphs.

Definition. The *degree* of a vertex of a graph G is the number of vertices of G to which it is adjacent. A graph G is *d-regular* if all of its vertices have degree d .

Definition. A graph G is *bipartite* if we can write its vertex set as the disjoint union of two sets R and L such that vertices in R are only adjacent to those in L and vice versa.

1.2 Walks and paths

A walk can be thought of as repeatedly moving from a vertex to an adjacent one in a graph by traversing the edges connecting them. A path is a walk where the same edge is never traversed twice. Formally, we have:

Definition. A *walk* in a graph $G = (V, E)$ is a sequence of vertices

$$x_1 x_2 \dots x_k$$

such that $\forall i \in \{1, 2, \dots, k - 1\}, \{x_i, x_{i+1}\} \in E$. The number $k - 1$ is the *length* of the walk. A *path* is a walk

with the additional condition that $\forall i, j \in 1, 2, \dots, k-1 \mid i \neq j, \{x_i, x_{i+1}\} \neq \{x_j, x_{j+1}\}$. A cycle of length $k-1$ is a path of length $k-1$ such that $x_1 = x_k$.

We will say that two vertices u and v are *connected by a path* if there exists a path $x_1 x_2 \dots x_k$ such that $x_1 = u$ and $x_k = v$. Note that the empty sequence is also a walk (or a path) and is interpreted as simply staying at the same vertex. Naturally, the empty walk has length 0.

Definition. Let u and v be two vertices of the graph $G = (V, E)$. The *distance* between u and v , denoted by $d(u, v)$ (or $d_G(u, v)$ if several graphs are involved), is the length of the shortest path connecting u and v . If there is no path connecting u and v , then $d(u, v) = \infty$.

It is straightforward to verify that this definition of distance is a metric on the vertices of a graph.

Definition. A graph G is *connected* if any two of its vertices are connected by some path.

Definition. Let $G = (V, E)$ be a graph, a *subgraph* is a graph $G' = (V', E')$ such that $V' \subset V$ and $E' \subset E$.

Definition. Let $G = (V, E)$ be a graph and $V' \subset V$, the *induced subgraph* of V' is the graph $G' = (V', E')$ where E' is the set of all edges of G that connect two vertices of V' .

Theorem 1. Any graph $G = (V, E)$ is the union of connected vertex disjoint subgraphs with no edge between them.

Proof. By induction on $n = |V|$. The case $n = 1$ is vacuously true. Suppose the claim is established for $n = k$. If G is connected, then G is the union of a single connected subgraph consisting of all of G . If not, there exist vertices u and v with no path connecting them. Let U be the set of all vertices that are connected to u by some path. Then V is the disjoint union of U and $V \setminus U$. Consider the induced subgraphs of U and $V \setminus U$, these have no vertices in common and no edge between them either or else there would be a path from u to an element of $V \setminus U$. Also, the induction hypothesis says that these subgraphs are the union of connected vertex disjoint subgraphs with no edge between them. Since G is the union of these induced subgraphs, we are done. \square

The “connected vertex disjoint subgraphs with no edge between them” of the last proposition are referred to as the components of the graph G .

The following definition is also useful.

Definition. Let $G = (V, E)$ be a graph and let $S \subset V$, $\partial S = \{v \in V \setminus S \mid \exists u \in S \text{ with } \{u, v\} \in E\}$.

1.3 Some graph invariants

Definition. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs, a *graph homomorphism* from G_1 to G_2 is a function $g : V_1 \rightarrow V_2$ such that $\{u, v\} \in E_1 \Rightarrow \{g(u), g(v)\} \in E_2$.

Definition. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs, an *isomorphism* of G_1 with G_2 is a bijection $f : V_1 \rightarrow V_2$ such that $\{u, v\} \in E_1 \Leftrightarrow \{f(u), f(v)\} \in E_2$. G_1 and G_2 are *isomorphic* if there exists an isomorphism between them.

Two graphs being isomorphic means that they are structurally the same and differ only by the labeling of their vertices. Note that if we remove the requirement that f be a bijection, then f can still be thought of as an embedding of G_1 into G_2 as an induced subgraph.

Definition. A graph invariant is a function $\Psi : \mathcal{G} \rightarrow \mathbb{R}$, where \mathcal{G} denotes the set of all graphs, such that $\Psi(G_1) = \Psi(G_2)$ whenever G_1 is isomorphic to G_2 .

Here are a few graph invariants:

- The *girth* of G is the length of its shortest cycle.
- The *average degree* of G is the average of the degrees of all its vertices.
- The *chromatic number* of G is the least number of colours needed to colour the vertices so that no adjacent vertices of G are of the same colour. (Note: Colouring a graph could be defined more precisely in terms of functions but it is not necessary for our purpose.)
- The *diameter* of G is the maximal distance between two of its vertices.

1.4 The adjacency matrix of a graph

Definition. Let $V = \{x_1, x_2, \dots, x_n\}$ and let $G = (V, E)$ be a graph. Define $A(G) \in M_{n \times n}(\mathbb{R})$ by:

$$a_{ij} = \begin{cases} 1 & \text{if } \{x_i, x_j\} \in E \\ 0 & \text{otherwise} \end{cases}$$

$A(G)$ is called the *adjacency matrix* of the graph G .

Since we are working with undirected graphs, $A(G)$ is symmetric and thus its eigenvalues are real. The multiset of eigenvalues of the adjacency matrix of a graph is called the *spectrum* of the graph. The spectrum gives us valuable information about the graph, as the following theorem shows.

Theorem 2. Let G be a graph with spectrum $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, then the following hold:

1. G is d -regular only if $d = \lambda_1$, the corresponding eigenvector is $\frac{1}{\sqrt{n}}(1, 1, \dots, 1)$.
2. If G is d -regular, then $|\lambda_i| \leq d$ for all i .
3. G is connected if and only if $\lambda_1 > \lambda_2$.
4. G is bipartite if and only if $\lambda_1 = -\lambda_n$.

Proof. See [1] p. 263-264. \square

Proposition 3. Let $V = \{x_1, x_2, \dots, x_n\}$ and let $G = (V, E)$ be a graph. Let $(b_{ij}) = A^n$. Then b_{ij} is the number of walks of length n that connect x_i to x_j .

Proof. We proceed by induction. The cases $n = 0$ and $n = 1$ are clear. Suppose that we know the claim to be true for $n = k$. Get all walks of length n from x_i to x_j by considering all walks of length $n - 1$ starting at x_i and connecting it to a neighbor of x_j . The numbers of walks of length $n - 1$ emanating from x_i reaching each of the vertices of G are given by the i th row of A^{n-1} (by the induction hypothesis). When multiplying this row by the j th column of A , the entries of the i th row that do not correspond to a neighbor of x_j will be multiplied by 0, whereas those that do correspond to a neighbor of x_j are multiplied by 1. Adding up, we get all walks of length n from x_i to x_j . \square

Definition. Suppose that $G = (V, E)$ is a d -regular graph. A $(k - 1)$ -random walk is a walk $x_1 x_2 \dots x_k$ where for each x_i , x_{i+1} is chosen uniformly at random among the d neighbours of x_i .

Thus, if we fix x_1 and initiate a $(k - 1)$ -random walk, each walk has a probability of $(\frac{1}{d})^{k-1}$ of occurring. However, several walks may end up at the same vertex. By the above proposition, the number of such walks is given by the entries of the adjacency matrix raised to the power of $k - 1$.

The following observations allow us to systematically calculate the distance between two vertices in a graph G using its adjacency matrix A :

1. The entries of the matrix $M(A, n) = A^0 + A^1 + A^2 + \dots + A^n$ are the number of walks of length less than or equal to n that connect the vertices of G .
2. If there exists a walk of length n between two vertices x_i and x_j of G , then there exists a path of length less than or equal to n between them. Thus, $d(x_i, x_j) \leq n$.
3. If there does not exist a walk of length n between x_i and x_j , then there cannot exist a path of the same length either and so $d(x_i, x_j) \neq n$.

We conclude that the distance between x_i and x_j in G is the least n such that the ij entry of $M(A, n)$ is not 0.

1.5 Cayley graphs

Cayley graphs are a nice way to ‘draw a picture of a finitely generated group. In such a graph, the nodes are the elements of the group and edges represent multiplication by a generator. The following definition is for the directed multigraph version of Cayley graphs.

Definition. Let G be a finitely generated group and let S be a set of generators for G (assume the identity of G does not belong to S). We define the Cayley graph of G on S to be the directed multigraph whose set of vertices is G , and elements g_1 and g_2 of G are adjacent if there exists some $s \in S$ such that $g_2 = sg_1$.

A directed graph gives a more complete picture of a group because we see precisely the effect of multiplication

by generators, however, we will mainly study Cayley graphs to produce a large class of interesting undirected simple graphs. We initiate this with the next definition.

Definition. Let G be a finitely generated group and S be a set of generators for G . Let $C(G, S)$ be the graph whose set of vertices is G and let g_1 and g_2 be connected if there exist some $s \in S$ such that either $g_1 = sg_2$ or $g_2 = sg_1$.

It is easy to verify that $C(G, S)$ satisfies our definition of an undirected simple graph. From now on, we will refer to $C(G, S)$ when speaking of a Cayley graph. An important feature of this type of graph is that an edge can either represent multiplication by a generator or by the inverse of a generator. Finally, the cardinality of G can be finite or infinite, so that we may have finite or infinite Cayley graphs. It turns out that geometric group theory is concerned with the study of infinite Cayley graphs.

Fact 4. The following are obvious implications of the fact that G is a group:

- In a directed Cayley graph, each vertex has one edge coming into it and going out of it for each $s \in S$.
- Both directed and undirected Cayley graphs are connected.
- Let $d = |S|$ and suppose that S contains k elements of order 2, then $C(G, S)$ is $(2d - k)$ -regular.

2 POSSIBILITIES FOR A DEFINITION OF DISTANCE BETWEEN GRAPHS

There are several candidates for the definition of distance between graphs, we present some of them in this section. We will assume that all graphs are on the same set of vertices $V = \{1, 2, \dots, n\}$ and that they are regular and connected.

Definition. Let G_1 and G_2 be two graphs with adjacency matrices A_1 and A_2 respectively. The *alpha distance* between G_1 and G_2 is defined to be:

$$\alpha(G_1, G_2) = \|A_1 - A_2\|$$

Where for $A \in M_{n \times n}(\mathbb{R})$, $\|A\| = \sup_{v \in \mathbb{R}^n} \frac{\|Av\|}{\|v\|}$. It is well known that this function is a norm on $M_{n \times n}(\mathbb{R})$. Moreover, since $A_1 - A_2$ is symmetrical, there exists an orthonormal basis that diagonalizes it and $\|A_1 - A_2\|$ is the maximal (in terms of magnitude) eigenvalue of $A_1 - A_2$.

Definition. Let G_1 and G_2 be graphs with the distance between vertices u and v denoted by $d_1(u, v)$ for G_1 and $d_2(u, v)$ for G_2 . Consider $B_i(v, r) = \{u \in V \mid d_i(u, v) \leq r\}$.

We define the *beta distance* between G_1 and G_2 to be:

$$\beta(G_1, G_2) = \sup_{v \in V, r \in \mathbb{N}} \left| |B_1(v, r)| - |B_2(v, r)| \right|$$

The set $B_i(v, r)$ is called a ball of radius r centered at v . Furthering the analogy with balls and spheres in \mathbb{R}^m , $\partial B_i(v, r-1) = \{u \in V \mid d_i(u, v) = r\}$ is obviously called the sphere or shell of radius r and center v . Intuitively speaking, the beta distance measures the difference in the growth of balls in the two graphs.

We make a few remarks on how to compute the beta distance. In the maximization process, we can restrict our search to $r \leq \max(\text{diam}(G_1), \text{diam}(G_2))$ because if we go past this upper bound, then $\|B_1(v, r) - B_2(v, r)\| = 0$ for all v since all vertices are within one diameter of any other vertex. In large graphs, it may be cumbersome to look for all paths connecting two vertices and then to decide which one is the shortest, especially if we must repeat this for all unordered pairs of vertices. However, there is a systematic way to avoid this using the adjacency matrices of our graphs. Recall the matrix $M(A_i, r)$ from section 1 and the observations that we made on it. Its entries are the number of walks of length less than or equal to r between the nodes of G_i , and a non-zero entry means that the distance between the corresponding nodes is less than or equal to r . Thus, the number of non-zero entries in the x th row is precisely $|B_i(v, r)|$. This allows us to compute $\|B_1(v, r) - B_2(v, r)\|$ for all $x \in V$ and $r \leq \max(\text{diam}(G_1), \text{diam}(G_2))$ and find the maximum without computing distances between all pairs of vertices.

Recall from elementary probability theory that two events A and B from some sample space are independent if $P(A)P(B) = P(A \cap B)$.

Definition. Let G_1 and G_2 be graphs and suppose that we pick two vertices u and v uniformly and independently from V . Consider the events $d_1(u, v) = a$ and $d_2(u, v) = b$ for $a, b \in \mathbb{N}$. We look at the parameter:

$$\gamma(G_1, G_2) = \left| \log \left(\sup_{(a,b) \in \mathbb{N}^2} \frac{P(d_1(u, v) = a \cap d_2(u, v) = b)}{P(d_1(u, v) = a)P(d_2(u, v) = b)} \right) \right|$$

We need to be careful with the last definition, if $\gamma(G_1, G_2)$ is close to 0, it means that the distance functions of the two graphs give nearly independent random variables so that the graphs are very different (contrary to the usual meaning of the distance being 0). In other words, knowing information about one graph does not allow to make any prediction about the distance between two vertices in the other graph.

Let us expand the formula for $\gamma(G_1, G_2)$. The random process carried out is the choice of an ordered pair (x, y) of elements of V . We have that $d(x, y)$ is a random variable. The probability that in the graph G_i , the distance between the chosen vertices is k is:

$$P(d_i(x, y) = k) = \frac{|\{(x, y) \in V \times V \mid d_i(x, y) = k\}|}{|V|^2}$$

In words, the number of ordered pairs of vertices that are k distance units away over the total number of ordered pairs. The above can also be rewritten as:

$$P(d_i(x, y) = k) = \frac{\sum_{x \in V} |\{y \in V \mid d_i(x, y) = k\}|}{|V|^2}$$

Incorporating this into the expression for $\gamma(G_1, G_2)$, we get:

$$\gamma(G_1, G_2) = \left| \log \left(\sup_{(a,b) \in \mathbb{N}^2} \frac{|V|^2 \sum_{x \in V} A}{(\sum_{x \in V} B)(\sum_{x \in V} C)} \right) \right|$$

$$\text{where } A = |\{y \in V \mid d_1(x, y) = a, d_2(x, y) = b\}|$$

$$B = |\{y \in V \mid d_1(x, y) = a\}|$$

$$C = |\{y \in V \mid d_2(x, y) = b\}|$$

3 SOME INEQUALITIES

3.1 The growth function of the Cayley graph of a free group on d generators

Suppose we take some vertex x of a d -regular graph $G = (V, E)$ and look at the various shells $S(x, r) = \{y \in V \mid d(x, y) = r\}$. The optimal growth of the size of those shells as r increases would occur in the case where for all r each node in $S(x, r)$ is adjacent to exactly one node in $S(x, r-1)$ and to $d-1$ nodes in $S(x, r+1)$. This is exactly how the (infinite) Cayley graph of a free group on d generators grows around its identity -or any other one of its elements actually. The goal of this subsection is to compute the number of vertices in the ball $B(e, r) := a_r$ of the infinite Cayley graph of a free group on d generators.

Example. Here is a picture of the Cayley graph of the free group on two generators. The nodes are the intersections of the straight lines, and each edge represents multiplication by a generator or its inverse. All branches actually extend to infinity.




Figure 3: The Cayley graph of $\langle a, b \rangle$.

It is interesting to solve this problem using a method that can be applied to any difference equation. The sequence a_r satisfies the recursive relation: $a_0 = 1$, $a_1 = d + 1$, $a_r = (d - 1)(a_{r-1} - a_{r-2}) + a_{r-1} = da_{r-1} - (d - 1)a_{r-2}$.

In matrix form:

$$\begin{pmatrix} a_r \\ a_{r-1} \end{pmatrix} = \begin{pmatrix} d & -(d-1) \\ 1 & 0 \end{pmatrix} \times \begin{pmatrix} a_{r-1} \\ a_{r-2} \end{pmatrix}$$

$$\begin{pmatrix} a_1 \\ a_0 \end{pmatrix} = \begin{pmatrix} d+1 \\ 1 \end{pmatrix}$$

$$\Rightarrow \begin{pmatrix} a_r \\ a_{r-1} \end{pmatrix} = \begin{pmatrix} d & -(d-1) \\ 1 & 0 \end{pmatrix}^{r-1} \begin{pmatrix} d+1 \\ 1 \end{pmatrix}$$

Let $A = \begin{pmatrix} d & -(d-1) \\ 1 & 0 \end{pmatrix}$. The eigenvalues of A are the roots of $\Delta(t) = (t - d)t - (-1)(d - 2) = t^2 - dt + (d - 1)$. These are: $\lambda_1 = \frac{d + \sqrt{d^2 - 4(d-1)}}{2} = d - 1$ and $\lambda_2 = \frac{d - \sqrt{d^2 - 4(d-1)}}{2} = 1$. The corresponding eigenvectors are $(d - 1, 1)$ and $(1, 1)$, respectively. Thus, we get that:

$$\begin{pmatrix} d-1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix}^{-1} \cdot \begin{pmatrix} d & 1-d \\ 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix}$$

Hence,

$$A^{r-1} = \left(\begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix} \right) \cdot \begin{pmatrix} (d-1)^{r-1} & 0 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix}^{-1}$$

Putting everything together and computing, we get:

$$\begin{pmatrix} a_r \\ a_{r-1} \end{pmatrix} = \begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} (d-1)^{r-1} & 0 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} d-1 & 1 \\ 1 & 1 \end{pmatrix}^{-1} \cdot \begin{pmatrix} d+1 \\ 1 \end{pmatrix}$$

$$a_r = \frac{1}{d-2} (d(d-1)^r - 2)$$

Since this corresponds to the maximal growth of balls centered at some vertex of a d -regular graph in general, we have an upper bound for the quantity $f(v, n)$ of definition 2:

$$|B(v, r)| \leq \frac{1}{d-2} (d(d-1)^r - 2)$$

where d is the degree of the regular graph G . Of course, this upper bound always holds but can only be attained when $r \leq \text{diam}(G)$.

3.2 A lower bound for $|B(v, r)|$

Definition. Define the expansion coefficient of G to be the number $h = \min_{|S| \leq \frac{|V|}{2}} \frac{|\partial S|}{|S|}$.

$$|B(v, r)| = |\partial B(v, r-1) \cup B(v, r-1)|$$

$$= |\partial B(v, r-1)| + |B(v, r-1)|$$

since the union is disjoint. So we have:

$$|B(v, r)| \geq |B(v, r-1)| + h|B(v, r-1)|$$

$$= |B(v, r-1)|(1+h)$$

Applying the last inequality repeatedly yields:

$$|B(v, r)| \geq |B(v, 0)|(1+h)^r = (1+h)^r$$

In [2], the authors relate the expansion coefficient of G to its spectrum with the following theorem.

Theorem 5. Let G be a d -regular graph with spectrum $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and expansion coefficient h . Then,

$$\frac{d - \lambda_2}{2} \leq h \leq \sqrt{2d(d - \lambda_2)}$$

Proof. See [2] p.454. □

For a d -regular graph, $\lambda_1 = d$ and $d - \lambda_2$ is called the spectral gap of G .

With this theorem, we can draw the following conclusion:

Proposition 6. Let $G = (V, G)$ be a d -regular graph with spectrum $d = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, and let $B(v, r) = \{u \in V \mid d(u, v) \leq r\}$, then for all $v \in V$, we have:

$$\left(1 + \frac{d - \lambda_2}{2}\right)^r \leq |B(v, r)| \leq \frac{1}{d-2} (d(d-1)^r - 2)$$

Proof. See the above discussion. □

4 INVESTIGATING THE IMPLICATIONS OF THE DEFINITIONS OF DISTANCE

In [3], the authors construct non isomorphic isospectral Cayley graphs of some simple groups. Similar questions that we answer are whether two isomorphic graphs necessarily have a beta distance of 0 and whether two graphs at beta distance 0 are necessarily isomorphic. In either case, the answer is easily shown to be “NO!!!” with a simple counterexample.

Consider the following two isomorphic graphs:

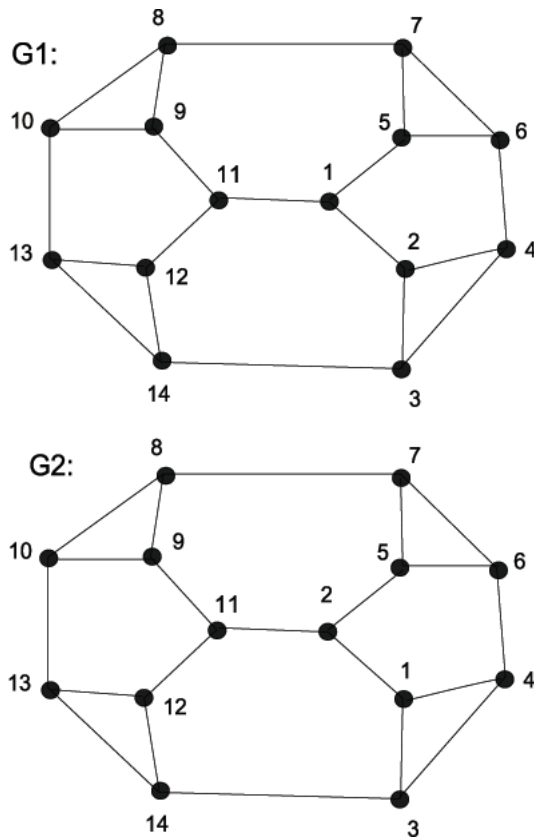


Figure 4: Two isomorphic graphs.

Computing the growth of balls, we get:

r	$ B_1(1,r) $	$ B_2(1,r) $	$ B_1(1,r) - B_2(1,r) $
0	1	1	0
1	4	4	0
2	10	8	2
3	14	12	2
4	14	14	0

$$\Rightarrow \beta(G_1, G_2) = 2 \neq 0$$




Figure 5: Two non-isomorphic graphs.

On the other hand, we can look at the two non-isomorphic graphs above: (G_1 contains triangles, but not G_2)

It is easy to check that all balls grow the same in those two graphs:

- Any ball of radius $r = 0$ contains only one vertex.
- Since the two graphs are 3-regular, all their balls of radius $r = 1$ contain four vertices.
- It is easy to verify that the diameter of both graphs is 2. Hence, any ball of radius $r = 2$ contains all the vertices of the graph in both G_1 and G_2 .

So that the beta distance of two graphs is 0 does not guarantee that they are isomorphic. It should come as no surprise that the beta distance does not classify isomorphic graphs. Indeed, the beta distance can be computed in polynomial time, but the problem of determining whether two graphs are isomorphic is conjectured to lie between P and NP (if $P \neq NP$).

Another question is whether we can construct graphs whose beta distance is 0 but whose alpha distance is not. The reverse question is quickly answered: since the alpha distance is defined using a matrix norm, it follows that $\alpha(G_1, G_2) = 0 \Rightarrow A(G_1) - A(G_2) = 0 \Rightarrow G_1 = G_2$ and so the beta distance of G_1 and G_2 is obviously 0.

We can use group actions to construct two Cayley graphs whose balls grow the same, but that have different adjacency matrices, which results in their alpha distance being non-zero.

Our procedure is based on the following proposition which precisely says that in any Cayley graph, balls of the same radius grow the same no matter what their center is.

Proposition 7. Let G be a finite group and let $C(G,S)$ be its Cayley graph for some set of generators S . Then for all $x, y \in G$ and all $r \in \mathbb{N}$, $f(x, r) = f(y, r)$.

Proof. It is enough to show that given any two $x_1, x_2 \in G$, we can find some bijection $\phi : G \rightarrow G$ such that $d(x_1, x) = d(x_2, \phi(x)) \forall x \in G$.

The function $* : G \times G \rightarrow G, g * v = vg$ is a transitive group action of G on itself. Define $\sigma_g : G \rightarrow G, \sigma_g(v) = vg$, for some $g \in G$. Let H be the a graph that is the same as $C(G, S)$ except for each vertex v being replaced by vg . Formally, H is the graph with set of vertices G , and set of edges $E = \{(x_1, x_2) \in G \times G | \{x_1g^{-1}, x_2g^{-1}\} \in E(C(G, S))\}$. We claim that σ_g is an isomorphism of $C(G, S)$ with H . Indeed, σ_g is clearly bijective and $(x_1, x_2) \in E(C(G, S)) \Leftrightarrow \exists s \in S | sx_1 = x_2 \text{ or } sx_2 = x_1 \Leftrightarrow \exists s \in S | sx_1g = x_2g \text{ or } sx_2g = x_1g \Leftrightarrow (\sigma_g(x_1), \sigma_g(x_2)) \in E(H)$.

Given any $x_1, x_2 \in G$, we can always choose some $g \in G$ such that $\sigma_g(x_1) = x_2$ because the action $*$ is transitive. Since σ_g is an isomorphism of graphs, it must satisfy $d(x, y) = d(\sigma_g(x), \sigma_g(y))$ for all $x, y \in G$. In particular, this means that $d(x_1, y) = d(x_2, \sigma_g(y))$ for all $y \in G$, as we needed. \square

Armed with proposition 7, let us outline a procedure to construct Cayley graphs with the same ball growth but different adjacency matrices. We know that balls grow the same in a Cayley graph, no matter which node they are centered at. So permuting the vertices of a Cayley graph with adjacency matrix A will yield another graph whose balls grow exactly the same as those of the original graph, but whose adjacency matrix has become MAM^{-1} where M is some permutation matrix. Note that as a byproduct we also get isomorphic graphs with different adjacency matrices.

5 THE COMMUTING ADJACENCY MATRICES PROBLEM

A solution to the following problem would help us find further links between definition 2 and definition 2.

Problem 8. Can you devise a systematic procedure to construct two graphs on the same set of vertices whose adjacency matrices commute?

We need to interpret what it means in terms of graphs if adjacency matrices commute. We assume that all graphs are on the same set of vertices $V = \{1, 2, \dots, n\}$. Let A_1 and A_2 be the adjacency matrices of G_1 and G_2 respectively. If we multiply the i th row of A_1 with the j th column of A_2 , we get the number of ways of going from node i to node j by first traversing an edge in G_1 , and then traversing an edge in G_2 . Conversely, the i th row of A_2 times the j th column of A_1 gives the number of ways of going from node i to node j by first going through an edge in G_2 and then through an edge in G_1 . Therefore, A_1 and A_2 commute if and only if for all nodes i and j , the number of ways of going from one to the other is the same whether we first traverse an edge of G_1 and then one of G_2 or do the opposite.

One approach would be to construct two Cayley graphs for the same group. The following lemma recasts the above in this context:

Lemma 9. Let G be a finite group. Let S and T be two sets of generators for G and denote $S \cup S^{-1}$ by S^* and $T \cup T^{-1}$ by T^* . Let A_S and A_T be the adjacency matrices of $C(G, S)$ and $C(G, T)$. Then A_S and A_T commute if and only if there exists a bijection $f : S^* \times T^* \rightarrow T^* \times S^*$ such that $\forall (s, t) \in S^* \times T^*$, $\prod f((s, t)) = \prod (s, t)$ where $\prod(x, y) = xy$.

Proof. Suppose we go from node g to node h of some Cayley graph $C(G, S)$ by traversing an edge between them. This means that there exists some $s \in S \cup S^{-1}$ such that $h = sg$ -we need to include S^{-1} because in the directed Cayley graph, the edge traversed could either go into g or out of it. If we return to our lemma, when going from g to h by first traversing an edge of $C(G, S)$ and then one of $C(G, T)$, we are saying that there exist $s \in S^*$ and $t \in T^*$ such that $h = ts$. Doing the same trip but this time taking an edge of $C(G, T)$ and then one of $C(G, S)$, we get that there exist $st \in S^*$ and $t \in T^*$ such that $h = sttg$. From our prior discussion, we

get that A_S and A_T commute if and only if for all $h, g \in G$, the number of distinct pairs $(s, t) \in S \times T$ such that $hg^{-1} = st$ is the same as the number of those $(t, s) \in T \times S$ such that $hg^{-1} = ts$. If we fix g and consider all products $h = stg$, we get that the aforementioned criterion is equivalent to the existence of the bijection described in the statement of the lemma. \square

Clearly, if G is an abelian group, then the function $f : (s, t) \rightarrow (t, s)$ satisfies the conditions of the lemma and we conclude that we can construct Cayley graphs with commuting adjacency matrices for all abelian groups. The question is then: Are there any other groups from which this can be done? We outline two approaches to this question.

The first idea was to start with two arbitrary sets S and T and to arbitrarily define a bijection as specified by lemma 9 with equality of products interpreted as equality in a word group on $S \cup T$. The equations arising from this equality condition would form a set of relations used in giving a presentation for a group G with $S \cup T$ as set of generators. Some more relations would be added so that G could be generated by either of S or T alone. It would remain to show that G is indeed a non-trivial finite group and to determine if it is isomorphic to some abelian group or not. Unfortunately, there is no systematic algorithm to perform these tasks and thus this procedure cannot be employed to systematically produce several groups with Cayley graphs whose adjacency matrices commute.

The other possibility is to take advantage of the fact that any finite group is a subgroup of S_n . Indeed, this allows us to design a routine that enumerates all finite groups by first enumerating all n and then finding all subgroups of S_n by computing the subgroup generated by each subset of S_n . One then checks to see if the adjacency matrices commute for the various sets of generators of a given group in our list.

6 SOME PROBLEMS ON GROUPS

The goal here is to present some problems that can be solved using the concepts of sections 1.

The general idea behind the first problem is to take some group G with presentation $\langle S | R \rangle$, given that G is finite. Suppose we want to multiply k elements of $S \cup S^{-1}$ and determine how many of those products will be equal to the identity. For example:

Problem 10. In the group D_3 , how many ways are there of writing the identity as a product of 1000 terms chosen from the set $\{s, r, r^2\}$?

This is solved by first constructing the Cayley graph of D_3 and using its adjacency matrix to find the number of walks of length 1000 from the identity to itself.

The second problem should be obvious if we bear the idea of distance between vertices of a graph in mind.

Problem 11. Consider S_4 and the set of generators $S = \{(1, 2), (1, 2, 3, 4)\}$. Define an *S-Jump* on g to be a multiplication of the element g of S_4 by some $s \in S \cup S^{-1}$. What is the least number of recursive S-Jumps required to go from the identity to the element $(1, 3, 2)(4, 1, 2)$?

Problem 12. Consider D_3 and the set of generators $S = \{(1, 2), (1, 2, 3, 4)\}$. Is it possible to partition D_3 into two disjoint subsets H_1 and H_2 such that $(S \cup S^{-1})H_1 \subset H_2$ and $(S \cup S^{-1})H_2 \subset H_1$?

This problem is equivalent to the question of whether the Cayley graph of D_3 on the given generators is bipartite or not and can be answered by computing its spectrum.

7 ACKNOWLEDGEMENTS

This article is a summary of the results obtained in a USRA project done during the summer 2008. I wish to thank Pro-

fessor Eyal Goren for his supervision.

REFERENCES

- [1] B. Bollobas. *Modern Graph Theory*. Springer-Verlag, New York, 1998.
- [2] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin (New Series) of the American Mathematical Society*, 43(4):43–561, October 2006.
- [3] A. Lubotsky, B. Samuels, and U. Vishne. Isospectral cayley graphs of some finite simple groups. *Duke Mathematical Journal*, 135:381–393, 2006.

JOKES

A math professor is talking to her little brother who just started his first year of graduate school in mathematics.

“What’s your favorite thing about mathematics?” the brother wants to know.

“Knot theory.”

“Yeah, me neither.” □

“Wasn’t yesterday you and your wife’s first wedding anniversary? What is it like having been married to a mathematician for a whole year?”

“She just filed for divorce...”

“I don’t believe it! Did you forget about your wedding day?”

“No. Actually, on my way back home from work, I stopped at a flower store and bought a bouquet of red roses for my wife. When I came home, I gave her the roses and said: ‘I love you.’ ”

“So, what happened?!”

“Well, she took the roses, slapped them around my face, kicked me in the groin, and threw me out of our apartment...”

“What a bitch!”

“No, no... it’s all my fault... I should have said: ‘I love you *and only you*.’ ” □

An American mathematician returns home from a conference in Moscow on real and complex analysis.

The immigration officer at the airport glances at his landing card and says: “So, your trip to Russia was business related. What’s the nature of your business?”

“I am a professor of mathematics.”

“What kind of mathematics are you doing?”

The professor ponders for a split second, trying to come up with something that would sound specific enough without making the immigration officer suspicious, and replies: “I am an analyst.”

The immigration officer nods with approval: “I think it’s great that guys like you go to Russia to help those poor ex-commies to get their stock market on its feet...” □

“So how’s your boyfriend doing, the math student?”

“Don’t mention that crazy pervert to me anymore! We broke up.”

“How can you say such a nasty thing about him? He seemed to be such a nice boy.”

“Imagine! He was restless during the days and couldn’t sleep at night - always trying to solve his math problems. When he had finally done it, he wasn’t happy: he would call himself a complete idiot and throw all his notes into the garbage. One day, I couldn’t take it anymore, and I told him to drop math. You know what he told me?”

“No.”

“He said, he enjoyed it!!!” □

MAY THE SCHWARTZ BE WITH YOU

Daniel Shapero

Chances are that in studying differential equations or quantum mechanics you have come across a nasty function called δ whose defining property is that

$$\int_{-\infty}^{\infty} \delta(x)f(x) dx = f(0)$$

for any function f . What would such a function look like? For any $x \neq 0$, $\delta(x) = 0$. Then δ is almost everywhere 0 and hence we should have

$$\int_{-\infty}^{\infty} \delta(x) dx = 0,$$

yet by definition of δ this integral should be 1. And yet the delta “function” – you have doubtless also been told that it is not really a function – and its cousins, the distributions, are vital in many fields. How does one define these distributions rigorously and recover all the properties that we expect of them?

1 TEST FUNCTIONS AND DISTRIBUTIONS

Laurent Schwartz is credited as having made distribution theory rigorous by his work in the 1940s and 50s. His idea was to consider distributions as linear functionals rather than as true mappings on Euclidean space. In linear algebra classes we are taught about the dual space of a finite-dimensional vector space and various results about the zero sets of collections of functionals and so forth. One peculiar example of a linear functional that one typically sees is the map $T : C([a, b]) \rightarrow \mathbb{R}$ given by



Figure 1: Laurent Schwartz

$$T(\phi) = \int_a^b f(x)\phi(x) dx \tag{1}$$

where f is a fixed continuous function. This is a linear functional on an infinite-dimensional vector space, the likes of which are set aside in favor of studying finite-dimensional ones. Distributions are defined as elements of the dual space of a certain infinite-dimensional vector space \mathcal{D} , with the additional requirement that they be continuous in a certain sense. Much of the theory is inspired by the fact that most distributions should be representable by formula (1), as integration against some fixed function.

So, we will define a *test space* \mathcal{D} of functions $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with certain properties, and define distributions as linear functionals on \mathcal{D} . Since the test space only exists for the distributions to act on, we will take functions ϕ in \mathcal{D} to be as regular as possible. As such, the first requirement we will impose is that \mathcal{D} consist of functions which are infinitely differentiable. Next, in order to guarantee that as many integrals of the form (1) will converge, we will also require that any test function ϕ have compact support, so

that all integrals are over bounded regions. The existence of a C^∞ function of compact support is not quite obvious. One can check using the chain rule and induction that, for C a positive constant, the function

$$\Psi(x) = C \begin{cases} \exp\left(\frac{-1}{1-|x|^2}\right) & \|x\| < 1 \\ 0 & \|x\| \geq 1 \end{cases} \tag{2}$$

is smooth and has compact support. We choose the constant C so that $\int \Psi dx = 1$, for reasons that will become apparent later.

When dealing with infinite-dimensional vector spaces, one must worry about whether linear functionals are continuous. To even talk about the continuity of a function defined on \mathcal{D} , one must define what it means for a sequence in \mathcal{D} to converge. In our case, a sequence ϕ_n converges to a function ϕ in the sense of test functions if there is a fixed compact set K such that $\text{supp } \phi_n \subset K$ for all n , and such that for all multi-indices α , $\|\partial^\alpha(\phi_n - \phi)\|_\infty \rightarrow 0$. So, we require that ϕ_n and all its derivatives converge uniformly to those of ϕ . Since the convergence is uniform, the limit function ϕ will also be in \mathcal{D} .

Finally, a *distribution* is defined as a linear functional $T : \mathcal{D} \rightarrow \mathbb{R}$ such that, if ϕ_n converges to ϕ in the sense defined above, $T(\phi_n) \rightarrow T(\phi)$. To emphasize the dual pairing we will write $\langle T, \phi \rangle$ for $T(\phi)$. The space of distributions is denoted \mathcal{D}' .

Using the compactly supported smooth functions as test space is not the only choice possible. We will define another space \mathcal{S} which is also commonly used in place of \mathcal{D} , for the reason that \mathcal{S} will be much better behaved when we start looking at Fourier transforms. A function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ is called a *Schwartz function* if it is C^∞ and, for all multi-indices α and β , $\|x^\alpha \partial^\beta \phi\|_\infty < \infty$. These functions and their derivatives decay more rapidly at infinity than any polynomial. An example of a Schwartz function is the density of a Gaussian random variable,

$$\gamma(x) = \frac{1}{(2\pi)^{d/2}} e^{-\frac{\|x\|^2}{2}}. \tag{3}$$

The space of Schwartz functions is denoted by \mathcal{S} . A sequence of Schwartz functions ϕ_n converges to a function ϕ if, for all multi-indices α and β , $\|x^\alpha \partial^\beta (\phi_n - \phi)\|_\infty \rightarrow 0$. Again, since the convergence is uniform, the limit ϕ is then a Schwartz function.

We can then consider continuous linear functionals on \mathcal{S} . These are called *tempered distributions* and the space of tempered distributions is written \mathcal{S}' . Since every test function is also a Schwartz function, every tempered distribution is also an ordinary distribution, or in symbols $\mathcal{S}' \subset \mathcal{D}'$. The opposite inclusion however does not hold: there are distributions which are not tempered. For example, if we consider the functional

$$\langle T, \phi \rangle = \int e^{\|x\|^2} \phi(x) dx,$$

then T is a distribution that is not tempered. To see this, one could consider $\langle T, \gamma \rangle$ where γ is defined by (3); $\langle T, \gamma \rangle$ is not finite, even though $\langle T, \phi \rangle$ defines a continuous linear functional on \mathcal{D} .

We will later show that any distribution can be approximated by simpler ones, for which we require a notion of convergence of distributions. Accordingly, a sequence f_n of distributions converges to a distribution f if $\langle f_n, \phi \rangle \rightarrow \langle f, \phi \rangle$ for every $\phi \in \mathcal{D}$. If f_n converges to f , then f satisfies the required continuity properties of a tempered distribution as well; we will leave this fact unproven, but see [5]. This definition of convergence gives \mathcal{D}' a topology with which we can then talk about continuous operations on distributions.

We know about a wealth of distributions already: the Dirac delta for one, defined by

$$\langle \delta, \phi \rangle = \phi(0).$$

One can easily see that it has the required continuity properties. Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is locally integrable. Then we can define a distribution T_f by

$$\langle T_f, \phi \rangle = \int_{\mathbb{R}^d} f(x) \phi(x) dx, \quad (4)$$

which can be seen to have the required continuity properties by applying the dominated convergence theorem. One often identifies a locally integrable function f with the distribution T_f associated to it. Authors will often simply write that f is a distribution and use notation like $\langle f, \phi \rangle$ to mean (4). Since boundedness certainly amounts to local integrability, every test function ψ can be considered to be a distribution using this identification, so that $\mathcal{D} \subset \mathcal{D}'$. Similarly, any polynomial is a distribution, and so on and so forth. For tempered distributions, we need a slow growth condition on f as x grows large in order for (4) to define a tempered distribution, as we saw when we tried to take $f(x) = \exp(\|x\|^2)$. If f is bounded by some polynomial at infinity this will suffice, and one often speaks of functions of at most polynomial growth as being tempered distributions.

If a distribution T is given by integration against some function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ as in (4), the distribution T is said to be *regular* and f is called the *kernel*. Otherwise, T is called *singular*. The Dirac delta is an example of a singular distribution, by the argument given in the introduction. As we will soon see, every distribution can be approximated by regular distributions. This fact is comforting: while it seems that the space of tempered distributions is too big for one to understand, we can instead regard them as limits of functionals obtained by integrating against a smooth kernel. From now on if T is a regular distribution with kernel f we will instead refer to “the distribution f ”, although strictly speaking we are overloading terminology. Furthermore in light of the fact that every distribution is a limit of regular distributions we will denote distributions by f from now on as well.

2 SUPPORTS OF DISTRIBUTIONS

For a real function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we have a well-defined concept of the support of f as the set

$$\text{supp } f = \overline{\{x : f(x) \neq 0\}}.$$

Functions of compact support are particularly nice to deal with because they are often dense in various function spaces. We can also define this concept for distributions. Let f be a distribution. Suppose that U is a subset of \mathbb{R}^d such that, for all $\phi \in C_c^\infty(U)$,

$$\langle f, \phi \rangle = 0.$$

We say that f vanishes on U . Now suppose that $\{U_\alpha\}_{\alpha \in J}$ is an arbitrary collection of open sets such that f vanishes on each U_α . We claim that f vanishes on $\bigcup_{\alpha \in J} U_\alpha$. To see this, let $\phi \in C_c^\infty(\bigcup_{\alpha \in J} U_\alpha)$. Since ϕ has compact support, there exists a finite subcollection $U_{\alpha_1}, \dots, U_{\alpha_n}$ such that $\text{supp } \phi \subset \bigcup_{i=1}^n U_{\alpha_i}$. Now choose a partition of unity ψ_1, \dots, ψ_n subordinate to $U_{\alpha_1}, \dots, U_{\alpha_n}$, and write $\phi = \sum_{i=1}^n \psi_i \phi$. Then

$$\langle f, \phi \rangle = \sum_{i=1}^n \langle f, \psi_i \phi \rangle = 0,$$

since $\text{supp}(\psi_i \phi) \subset U_{\alpha_i}$ and f vanishes on U_{α_i} . With this fact proven we can then define

$$\text{supp } f = \left(\bigcup \{U : U \text{ open, } f \text{ vanishes on } U\} \right)^c.$$

We can then talk about distributions of compact support, the collection of which we will denote \mathcal{E}' .

3 OPERATIONS ON DISTRIBUTIONS

Operations on distributions are typically defined as follows: suppose that $A : C^\infty \rightarrow C^\infty$ is some linear operator such as

differentiation or the Fourier transform, and that A' is another operator such that $\int (A\phi)\psi dx = \int \phi(A'\psi) dx$. We should then define Af for f a distribution by the formula $\langle Af, \phi \rangle = \langle f, A'\phi \rangle$, since this formula already holds for the regular distributions. If this seems abstract, we will work out copious examples.

3.1 Derivatives

Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable and locally integrable. How does the derivative of f act as a distribution? Using the integration by parts formula,

$$\begin{aligned} \langle \partial_k f, \phi \rangle &= \int \partial_k f(x) \phi(x) dx \\ &= - \int f(x) \partial_k \phi(x) dx \\ &= \langle f, -\partial_k \phi \rangle, \end{aligned}$$

where the boundary terms are zero because either ϕ is compactly supported or its derivatives decay rapidly at infinity. By induction, for any multi-index α ,

$$\langle \partial^\alpha f, \phi \rangle = \langle f, (-1)^{|\alpha|} \partial^\alpha \phi \rangle. \tag{5}$$

Having shown that this formula holds for all regular distributions, we define the partial derivative $\partial^\alpha f$ of an arbitrary distribution f by formula (5). The operation of taking distributional derivatives is continuous with respect to the topology of \mathcal{D}' defined above: if $f_n \rightarrow f$,

$$\begin{aligned} \langle \partial^\alpha f_n, \phi \rangle &= \langle f_n, (-1)^{|\alpha|} \partial^\alpha \phi \rangle \\ \rightarrow \langle f, (-1)^{|\alpha|} \partial^\alpha \phi \rangle &= \langle \partial^\alpha f, \phi \rangle. \end{aligned}$$

All the other operations we define will also be continuous, and the proofs of continuity are the same as above.

3.2 Multiplication by Functions

If ψ is a smooth function and $\phi \in \mathcal{D}$, then $\psi\phi \in \mathcal{D}$, so that if f is locally integrable

$$\int (\psi(x)f(x))\phi(x) dx = \int f(x)(\psi(x)\phi(x)) dx.$$

If f is any distribution, we can multiply it by a smooth function ψ by the formula

$$\langle \psi f, \phi \rangle = \langle f, \psi\phi \rangle.$$

3.3 Affine Transformation

Fix some $y \in \mathbb{R}^d$, and let $(\tau_y f)(x) = f(x+y)$; τ_y is the operator of translation by y on functions. Making the substitution $w = x+y$,

$$\int f(x+y)\phi(x) dx = \int f(w)\phi(w-y) dw$$

so that we should define the translate by y of a distribution according to

$$\langle \tau_y f, \phi \rangle = \langle f, \tau_{-y} \phi \rangle.$$

What about reflection? If $Rf(x) = f(-x)$ is the operator of reflection on functions, we clearly have $\int Rf \cdot \phi dx = \int f \cdot R\phi dx$, so that we define the reflection of a distribution by

$$\langle Rf, \phi \rangle = \langle f, R\phi \rangle.$$

We can then note identities like $R\tau_{-y} = \tau_y R$.

Now let A be a non-singular linear map on \mathbb{R}^d , and for f locally integrable consider the function $f \circ A$.

$$\int f(Ax)\phi(x) dx = \int f(y)\phi(A^{-1}y) |\det A^{-1}| dy$$

so that in keeping with the above definitions

$$\langle f \circ A, \phi \rangle = \langle f, |\det A^{-1}| \phi \circ A^{-1} \rangle.$$

Together with translations we can transform distributions by applying linear maps to the independent variables.

3.4 Fourier Transform



Figure 2: How the Fourier transform works

A very important operation to carry over to distributions is the Fourier transform. We shall for the moment consider the tempered distributions, the reason for this being that every tempered distribution is the Fourier transform of another tempered distribution. Compactly supported smooth functions are never the Fourier transform of another compactly supported smooth function by a result called the Paley-Wiener theorem, found in [6], hence the preference. We will adopt the convention that the Fourier transform of a function $\phi \in \mathcal{S}$ is defined by

$$\widehat{\phi}(\xi) = \int e^{-2\pi i x \cdot \xi} \phi(x) dx.$$

The inversion formula for the Fourier transform, proven in [1], gives that

$$\phi(x) = \int e^{2\pi i x \cdot \xi} \widehat{\phi}(\xi) d\xi.$$

An application of Fubini's theorem gives that

$$\begin{aligned} \int \widehat{\phi}(\xi)\psi(\xi) d\xi &= \iint e^{-2\pi i x \cdot \xi} \phi(x)\psi(\xi) dx d\xi \\ &= \iint e^{-2\pi i x \cdot \xi} \psi(\xi)\phi(x) d\xi dx \\ &= \int \phi(x)\widehat{\psi}(x) dx. \end{aligned}$$

We then define for a tempered distribution f

$$\langle \widehat{f}, \phi \rangle = \langle f, \widehat{\phi} \rangle$$

where a similar formula holds for the inverse transform.

4 SOME FAMILIAR EXAMPLES

Let δ be the Dirac distribution in \mathbb{R} , and let $h(x) = \chi_{[0, \infty)}(x)$ be the Heaviside step function. We claim that $\delta = h'$ in the distributional sense. To see this, let $\phi \in \mathcal{D}$; by the fundamental theorem of calculus,

$$\begin{aligned} \langle h', \phi \rangle &= \langle h, -\phi' \rangle \\ &= - \int_0^\infty \phi'(x) dx \\ &= \phi(0) - \lim_{x \rightarrow \infty} \phi(x) \\ &= \phi(0) = \langle \delta, \phi \rangle \end{aligned}$$

in accordance with (5).

We can show easily that $\widehat{\delta} = 1$:

$$\langle \widehat{\delta}, \phi \rangle = \langle \delta, \widehat{\phi} \rangle = \widehat{\phi}(0) = \int \phi(x) dx = \langle 1, \phi \rangle.$$

Since we can take derivatives of distributions and multiply them by smooth functions, we can define linear differential operators on distributions. Let $L = \sum_{|\alpha| \leq k} a_\alpha(x) \partial^\alpha$ be a linear differential operator; we then have that

$$\begin{aligned} \langle Lf, \phi \rangle &= \left\langle \sum_{|\alpha| \leq k} a_\alpha \partial^\alpha f, \phi \right\rangle \\ &= \left\langle f, \sum_{|\alpha| \leq k} (-1)^{|\alpha|} \partial^\alpha (a_\alpha \phi) \right\rangle = \langle f, L^* \phi \rangle \end{aligned}$$

where $L^* = \sum \partial^\alpha a_\alpha$ is called the *adjoint* of L .

For L a linear differential operator as above, a *fundamental solution* for L is a distribution u such that $Lu = \delta$. We will see later when we study convolutions that fundamental solutions are indeed quite useful.

For $\phi \in \mathcal{S}$, one can prove using integration by parts that $\widehat{\partial^\alpha \phi} = (2\pi i \xi)^\alpha \widehat{\phi}$. Using this formula we then have

$$\langle \widehat{\partial^\alpha f}, \phi \rangle = \langle \partial^\alpha f, \widehat{\phi} \rangle = \langle f, (-1)^{|\alpha|} \partial^\alpha \widehat{\phi} \rangle.$$

Differentiating under the integral sign gives

$$\begin{aligned} \partial^\alpha \widehat{\phi} &= \partial^\alpha \int e^{-2\pi i x \cdot \xi} \phi(x) dx \\ &= (-1)^{|\alpha|} \int e^{-2\pi i x \cdot \xi} (2\pi i x)^\alpha \phi(x) dx \\ &= (-1)^{|\alpha|} \widehat{(2\pi i x)^\alpha \phi}, \end{aligned}$$

so pulling all the operators back to f gives that the formula $\widehat{\partial^\alpha f} = (2\pi i \xi)^\alpha \widehat{f}$ holds in a distributional sense as well.

Suppose now that L is a linear differential operator with constant coefficients and u is a tempered distribution. Taking the Fourier transform of Lu ,

$$\begin{aligned} \widehat{Lu} &= \sum_{|\alpha| \leq k} a_\alpha \widehat{\partial^\alpha u} \\ &= \left(\sum_{|\alpha| \leq k} a_\alpha (2\pi i \xi)^\alpha \right) \widehat{u} = P(2\pi i \xi) \widehat{u}, \end{aligned}$$

or in other words the differential operator L acts on tempered distributions by multiplication by a polynomial P , called the *symbol* of L , in Fourier space. More generally, a *Fourier multiplier* is an operator A on distributions given by function multiplication in the Fourier transform:

$$\widehat{Af} = m(\xi) \widehat{f},$$

so we allow more general functions than polynomials. One can immediately see that every constant-coefficient differential operator is a Fourier multiplier. But, this idea is quite powerful in that many other classical integral operators, the Hilbert transform being just one, are Fourier multipliers.

If f is a distribution, composing it with affine maps allows us to define invariance properties. For example, a distribution f is rotation-invariant if $f \circ Q = f$ for every orthogonal map Q . The distribution is even if $f \circ R = f$ for R the reflection map, and odd if $f \circ R = -f$. A distribution is positive-homogeneous of degree m if $f \circ tI = t^m f$ for every $t > 0$. As an illustration, the Dirac delta is rotation-invariant, even, and homogeneous of degree 0. If $y \in \mathbb{R}^d$ is a fixed vector such that $\tau_y f = f$, then f is periodic with period y . An example of a case where such analysis is fruitful is in trying to find fundamental solutions of differential operators. Consider the Laplace operator $\Delta = \sum_{k=1}^d \partial_k^2$. This operator commutes with rotations in the sense that, for any distribution f and any orthogonal matrix Q , $\Delta(f \circ Q) = (\Delta f) \circ Q$. This suggests that the fundamental solution of the Laplace operator should be rotation-invariant, and indeed it is as we shall see later.

5 CONVOLUTION AND APPROXIMATION

The last operation we would like to define for distributions is convolution. The convolution of two functions ψ and ϕ is the function defined by

$$(\psi * \phi)(x) = \int \psi(x-y) \phi(y) dy. \quad (6)$$

One can make a substitution in (6) to show that $\phi * \psi = \psi * \phi$. Suppose that ψ is differentiable and of compact support. The integral in (6) is taken over a compact set, so that we can safely differentiate under the integral sign to show that $\partial_k(\psi * \phi) = (\partial_k \psi) * \phi$. Convoluting a smooth function with any function then gives a smooth function as a result.

If ϕ is of compact support as well then $\psi * \phi$ is also of compact support: one can show in general that

$$\text{supp}(\psi * \phi) \subset \text{supp} \psi + \text{supp} \phi.$$

These two facts show that \mathcal{D} is closed under convolution.

Convolution also gives a very convenient way of approximating functions. Suppose that ψ is non-negative, smooth, has compact support containing 0 and $\int \psi dx = 1$. For definiteness we could use the function Ψ of (2). Define

$$\psi_t(x) = t^{-d} \psi\left(\frac{x}{t}\right).$$

Let f be bounded and uniformly continuous. We claim that $\psi_t * f \rightrightarrows f$ as $t \rightarrow 0$.

To see this, let $\varepsilon > 0$ be arbitrary and choose $\delta > 0$ so that, if $\|x - y\| < \delta$, $|f(x) - f(y)| < \varepsilon$. Since $\text{supp} \psi$ is compact, there exists $R > 0$ such that $\text{supp} \psi \subset B_R(0)$, the ball of radius R about 0. Now, suppose that $y \notin B_\delta(0)$. For $t < \delta/R$, $\|y/t\| \geq \|\delta/t\| > R$ and hence $\psi(y/t) = t^d \psi_t(y) = 0$. Then $\int_{B_\delta(0)^c} \psi_t(y) dy = 0$ for $t < \delta/R$, and since $\int \psi_t dy = 1$, given $x \in \mathbb{R}^d$,

$$\begin{aligned} |\psi_t * f - f| &= \left| \int \psi_t(y-x)f(y) dy - f(x) \right| \\ &= \left| \int \psi_t(y-x)(f(y) - f(x)) dy \right| \\ &\leq \int_{B_\delta(x)} \psi_t(x-y)|f(y) - f(x)| dy \\ &\quad + 2\|f\|_\infty \int_{B_\delta(x)^c} \psi_t(x-y) dy \\ &\leq \varepsilon \int_{B_\delta(0)} \psi_t dy + 2\|f\|_\infty \int_{B_\delta(0)^c} \psi_t dy \leq \varepsilon \end{aligned}$$

as required. Now, $\psi_t * f$ is a smooth function for each t , so that we have found a way to approximate a continuous function by a smooth one.

Suppose now that we apply this result on a test function ϕ . From (6),

$$\psi_t * \phi(x) = \langle \tau_x R \psi_t, \phi \rangle.$$

If we take $x = 0$, then we see that

$$\lim_{t \rightarrow 0} \langle R \psi_t, \phi \rangle = \phi(0) = \langle \delta, \phi \rangle,$$

hence $R \psi_t$ converges as a distribution to δ . In general, $\tau_x R \psi_t \rightarrow \delta_x$ as a distribution. Such functions ψ are called *approximations to the identity*. The classic example is the function Ψ from (2).

Our success so far suggests we proceed as follows: define the convolution of a test function ψ with a distribution f ; show that this distribution is in fact regular with smooth kernel; and finally, show that $\psi_t * f$ converges to f as a distribution where ψ_t is an approximate identity.

Again, (6) suggests that we define the convolution of a test function ψ with a distribution f pointwise by the formula

$$\psi * f(x) = \langle \tau_x R f, \psi \rangle = \langle f, \tau_x R \psi \rangle.$$

But, $\tau_y R \psi$ converges to $\tau_x R \psi$ in the sense of test functions as $y \rightarrow x$, so that by the continuity properties of distributions

$$\langle f, \tau_y R \psi \rangle \rightarrow \langle f, \tau_x R \psi \rangle$$

as $y \rightarrow x$. Hence $\psi * f(x)$ is a continuous function of x . But, we can do even better: $\psi * f$ is in fact differentiable and the formula $\partial_k(\psi * f) = (\partial_k \psi) * f$ holds as well. Let e_k be the k -th unit vector. Since the derivatives of ψ are all uniformly continuous,

$$\frac{\tau_{x+he_k} R \psi - \tau_x R \psi}{h} \rightarrow \tau_x R \partial_k \psi$$

as test functions. Since f is continuous,

$$\left\langle f, \frac{\tau_{x+he_k} R \psi - \tau_x R \psi}{h} \right\rangle \rightarrow \langle f, \tau_x R \partial_k \psi \rangle.$$

This implies that $\psi * f$ has a k -th partial derivative and

$$\partial_k(\psi * f) = (\partial_k \psi) * f.$$

Applying the same arguments as above, the partial derivatives of $\psi * f$ exist and are continuous, so $\psi * f$ is differentiable. Since in particular it is locally bounded, $\partial_k(\psi * f)$ is a distribution. Since ψ is infinitely differentiable we can apply induction to conclude that $\psi * f$ is also infinitely differentiable using the same arguments as above, and the formula $\partial^\alpha(\psi * f) = (\partial^\alpha \psi) * f$ holds for any multi-index α .

Another approach yields an equivalent formulation which we will also need. Supposing that f is locally integrable and using Fubini's theorem,

$$\begin{aligned} \langle \psi * f, \phi \rangle &= \int \psi * f(x) \phi(x) dx \\ &= \iint \psi(x-y) f(y) \phi(x) dy dx \\ &= \int f(y) \left[\int \psi(x-y) \phi(x) dx \right] dy \\ &= \langle f, (R\psi) * \phi \rangle. \end{aligned}$$

This suggests that we should also have

$$\langle \psi * f, \phi \rangle = \langle f, (R\psi) * \phi \rangle \tag{7}$$

for distributions using our other definition of convolution, which is true. While taking (7) as the definition of $\psi * f$ conceals the fact that $\psi * f$ is a regular distribution, it will aid us in defining the convolution of two distributions.

Before proceeding on this course, we can finally prove that an arbitrary distribution f is the limit of a sequence of regular distributions with C^∞ kernels. Let ψ_t be an approximate identity, and consider the sequence of distributions $\psi_t * f$. For ϕ a test function,

$$\langle \psi_t * f, \phi \rangle = \langle f, (R\psi_t) * \phi \rangle.$$

Note that if ψ_t is an approximate identity, so is $R\psi_t$. If we were using the function Ψ which we explicitly constructed earlier, then Ψ is even so that $R\Psi_t = \Psi_t$. Either way, $(R\psi_t) * \phi \rightarrow \phi$ as test functions. To see this, note that, for $t > 1$, $\text{supp } \psi_t \subset \text{supp } \psi_1$, and since

$$\begin{aligned} \text{supp}((R\psi_t) * \phi) &\subset \text{supp}(R\psi_t) + \text{supp } \phi \\ &\subset \text{supp}(R\psi_1) + \text{supp } \phi, \end{aligned}$$

the supports of $(R\psi_t) * \phi$ are all contained in a fixed compact set. For the uniform convergence of all derivatives,

$$\|\partial^\alpha((R\psi_t) * \phi - \phi)\|_\infty = \|(R\psi_t) * (\partial^\alpha \phi) - \partial^\alpha \phi\|_\infty \rightarrow 0$$

since $\partial^\alpha \phi$ is also a test function for any multi-index α . Since f is continuous,

$$\langle \psi_t * f, \phi \rangle = \langle f, (R\psi_t) * \phi \rangle \rightarrow \langle f, \phi \rangle$$

and hence $\psi_t * f \rightarrow f$ as a distribution. Since $\psi_t * f \in C^\infty$, we conclude that every distribution can be approximated by regular distributions with smooth kernels.

The same construction can be carried out when f is a tempered distribution, and ψ and ϕ are Schwartz functions. In this case $\psi * f$ is a regular tempered distribution with a C^∞ kernel. The approximate identities in the class of Schwartz functions are less restrictive, in the sense that they need not have compact support. We still require that $\int \psi_t dx = 1$ for all t but the condition on the support is now that for any $\varepsilon > 0$, $\psi_t \rightarrow 0$ on $B_\varepsilon(0)^c$. In particular, if γ is the Gaussian defined in (3), and

$$\gamma_\sigma(x) = \frac{1}{(2\pi)^{d/2}} \frac{1}{\sigma} \exp\left(-\frac{\|x\|^2}{2\sigma^2}\right),$$

then for any tempered distribution f , $\gamma_\sigma * f \rightarrow f$ in the distributional sense as $\sigma \rightarrow 0$, which can be given an interesting probabilistic interpretation.

6 MORE CONVOLUTION AND FUNDAMENTAL SOLUTIONS

Having defined the convolution of a test function with a distribution, we can now look at the convolution of two distributions, a more delicate issue. We will use (7): if f and u are distributions, a preliminary definition is that $f * u$ acts on test functions according to

$$\langle f * u, \phi \rangle = \langle u, (Rf) * \phi \rangle.$$

While we know that $(Rf) * \phi \in C^\infty$ by the reasoning given in the last section, we do not know that it is a test function because it may not have compact support. But, if we impose a restriction on f , namely that $f \in \mathcal{E}'$, then $f * \phi$ is compactly supported. To see this, note that $\text{supp}(\tau_x \phi) = \text{supp } \phi - x$. So, if $\|x\|$ is large enough, say greater than

$$\sup_{x \in \text{supp } f} \|x\| + \text{diam}(\text{supp } \phi),$$

then $\text{supp}(\tau_x R\phi) \cap \text{supp } Rf = \emptyset$ and thus

$$(Rf) * \phi(x) = \langle Rf, \tau_x R\phi \rangle = 0.$$

So, $(Rf) * \phi \in \mathcal{D}$ and so $f * u$ is defined by the above formula. While we have shown that restricting f to \mathcal{E}' is sufficient to guarantee that the convolution exists, is it necessary? We can relax the condition by considering what are called the *singular supports* of f and u , but unfortunately it is impossible to consistently define the convolution of two arbitrary distributions for precisely the consideration above.

All of this finds fruitful application in differential equations. Let $L = \sum_{|\alpha| \leq k} a_\alpha(x) \partial^\alpha$ be a linear partial differential operator with C^∞ coefficients. Recall that u is said to be a fundamental solution for L if $Lu = \delta$. Why should such a function be important? Suppose that we want to solve the inhomogeneous PDE $Lv = f$, where we will be vague about what f can be for now. Since L can be taken under convolutions, $L(u * f) = Lu * f = \delta * f = f$, so that $v = u * f$ is the solution to our problem. Here it becomes apparent why all the fuss was necessary about when distributions can be convolved: if f is a distribution rather than a smooth function we have to worry about supports, as $u * f$ may not even be defined.

Does every linear differential operator have a fundamental solution and if so how can they be found? Suppose now that L has constant coefficients, so that the a_α do not depend on x . We have seen that if u is a distribution, then

$$\widehat{Lu} = P(2\pi i \xi) \widehat{u},$$

where P is the symbol of L . To solve $Lu = f$ for $f \in \mathcal{S}$, if we take Fourier transforms we now have the problem $P\widehat{u} = \widehat{f}$, or

$$u(x) = \int_{\mathbb{R}^d} \frac{e^{2\pi i x \xi} \widehat{f}(\xi)}{P(2\pi i \xi)} d\xi.$$

The problem with this approach is that P may have zeroes. Anyone who has had his fair share of woes with integrals will no doubt fondly remember the day he learned about M. Cauchy's integral formula and the residue calculus, and indeed Augustin-Louis comes swooping down from the heavens to save us. One can show that \widehat{f} has an extension to a holomorphic function of several variables, deform the contour of integration to dodge the zeroes of P and thus define u , show that the integrand decays rapidly for $|\text{Im } \xi|$ large and differentiate under the integral sign to show that u is the required solution. This is the famed *Malgrange-Ehrenpreis Theorem*, which was one of the first results to convince people that distributions were good for much of anything at all. For a detailed proof see [1], [3].

7 CONCLUSION AND FURTHER READING

The topics we have covered above are only the first meter or two of a very deep rabbit hole. Other topics of interest in distribution theory:

Much as one can define continuous linear functionals, one can take tensor products of distributions and look at multilinear functionals. Then one can regard bilinear distributions as mappings from \mathcal{D} to \mathcal{D}' , where one has the famous *Schwartz kernel theorem* demonstrating that every map of \mathcal{D} to \mathcal{D}' is obtained from a distribution of several variables. This will have consequences in the follow-up article to this one, where one must look at the distribution kernel of a quadratic form that happens to be the variance of a generalized stochastic process.

We know how to apply affine-linear maps to the independent variables of a distribution. Can one also apply smooth maps? Using the change of variables formula it is indeed possible to compose a distribution f with a smooth function ψ which is a diffeomorphism on the support of f . Then one can look at distributions on manifolds.

We showed that if a distribution vanishes on a collection of open sets, then it vanishes on their union. This fact enabled us to define the support of a distribution as the complement of the largest open set on which it vanishes. Similarly, if a distribution f is given as integration against a smooth function for all test functions supported in U_α , then f also acts as integration against a smooth function on $\bigcup_\alpha U_\alpha$. We can then take the complement of the largest

open set on which f coincides with a smooth function to be the *singular support* of f . A powerful idea is to use the Fourier transform to examine the directions in Fourier space in which f is singular as well. This is one idea that underpins the gargantuan theory of *microlocal analysis*.

Hopefully this has been informative and will part the veil of mystery surrounding this oft-used yet rarely explained subject!

REFERENCES

- [1] G.B. Folland. *Introduction to Partial Differential Equations*. Princeton University Press, Princeton, NJ, 2nd edition, 1995.
- [2] I.M. Gel'fand and G. Shilov. *Generalized Functions*, volume I. Academic Press, New York, NY, 1964.
- [3] L. Hörmander. *The Analysis of Linear Partial Differential Operators*, volume I. Springer-Verlag, Berlin, 1990.
- [4] A.N. Kolmogorov and S.V. Fomin. *Introductory Real Analysis*. Dover Publications, New York, 1975.
- [5] M. Renardy and R.C. Rogers. *An Introduction to Partial Differential Equations*. Springer-Verlag, New York, 2nd edition, 2004.
- [6] F. Trèves. *Topological Vector Spaces, Distributions and Kernels*. Dover Publications, New York, 1995.

JOKES

Q: When did Bourbaki stop writing books?

A: When they realized that Serge Lang was a single person... \square

Q: What do you get if you add two apples and three apples?

A: A high school math problem! \square

Q: Why did the two vectors start an internet-based company?

A: Because they thought they had a good dot product. \square

Q: Did you hear that joke about the infinite line?

A: Don't worry, It doesn't have a point! \square

A mathematical limerick:

$$\left(\int_1^{\sqrt{3}} z^2 dz \right) \cdot \cos\left(\frac{3\pi}{9}\right) = \ln(\sqrt[3]{e})$$

Integral z -squared dz

from 1 to the square root of 3

times the cosine

of three pi over 9

equals log of the cube root of 'e'. \square

INTERVIEW WITH PROF. DAVID WOLFSON

Ana Best

DE: What is your name, and how long have you been at McGill?

My name is David Wolfson, and I've been at McGill since 1974.

DE: Tell me a bit about your background, both personal and academic.

I come from South Africa, and I did my undergraduate and Masters degrees at the University of Natal in South Africa, a B.Sc., an extra year called a B.Sc. Honours, and a one-year Masters. Then I went to the United States, to Purdue University, to do a Ph.D. in Statistics, and graduated in 1974, came to McGill, and have been here ever since.

DE: What are your favourite things about McGill and Montreal?

When I first came to Montreal, I thought the weather was really nice. It was neat, because I was a huge fan of *White Fang* and *The Call of the Wild* by Jack London, but over the years, the weather has sort of worn off as a novelty, so certainly not the weather. I guess the cosmopolitan nature of the city I find attractive. It's easy to get around, as it's not a huge city. I guess those are the main things that I really like.

As for McGill, its certainly nice being at a university that has a fine reputation, and the department is a very cordial place. I must say, for statisticians, it was not such a cordial place when I first arrived. But I think in the last 15 to 20 years, things have changed and it's become a much more friendly place for statisticians.

DE: How did you first become interested in statistics?


Well, that's a rather interesting story because when I did my undergraduate degree, the way it worked was, if you were doing a Bachelor of Science, you started out with four courses (the courses ran for a full year), so the first year you took four courses, then you took three courses in the second year, dropped one of those courses, and in your third year you retained two of the courses and they were called your major.

I was fanatical about chemistry when I started out. I had my own lab at home and did serious chemistry and so, like most people, I started with chemistry, physics, mathematics, and applied mathematics. And, within 6 months I hated chemistry, so chemistry was going to be gone in the second year. Physics - the physics department was quite weak, so physics was not an option for the second year, and I was stuck a bit because I needed three courses for the second year.

Now it turned out that mathematical statistics was a two-year major, because you were not allowed to take it until you'd had mathematics 1, because it had certain mathematical prerequisites. So, by default, statistics was the only two-year major that remained, and that's how I got into statistics. And I've been there ever since.

DE: What's your favorite probability distribution?

I guess the Weibull, because it's the parametric distribution that plays the biggest role in survival analysis.



Picture 1: Prof. David Wolfson

DE: What are you currently researching?

My main research areas are survival analysis and, to a lesser extent, Bayesian optimal design. I think that the main emphasis is really on survival analysis; survival analysis essentially deals with data that records events that occur in time, where there's a start date and an end date. In its simplest form, you're trying to say something about the distribution of the time between the start date and the end date of certain events.

What makes survival analysis different is that the data that you get are incomplete in two main ways: either you don't get to observe the end date because subjects drop out of the study or the study comes to an end. In such cases the data are called censored. Or in some cases, you have

subjects who don't live long enough, for instance, don't get to make it into your data set, and you tend to observe the longer survivors. The people with short survival tend not to be observed. And one of my main interests is in looking at such so-called left-truncated or length-biased data.

DE: What are your favorite and least favorite parts of research?

I think one of my favorite parts of research is suddenly having an insight into something that you've been looking at for a long time. There's a great deal of satisfaction in understanding something that you haven't been able to understand.

And, I think least favorite thing is writing up results. After you've got something and you're convinced it's true, actually sitting down and writing the paper. But it's also one of my favorite things to do; it's a funny thing, but when you actually get down to do it and write a paper that conveys the ideas nicely, there's a good deal of satisfaction. But it's also a bit of a drudge contemplating that now you have to sit down and write a paper. Also one of the least favorite things is to get a devastating report from a referee. And we all do.

DE: When you were an undergraduate, what were your goals? Did you see yourself becoming a university professor?

No, I don't think so. When I started out, I anticipated doing

something in chemistry; research in chemistry. Then when I went into statistics, actually when I got my Masters degree, I looked around for a job in industry as a research statistician. And I think it was only when I started working on my Ph.D. that I came to think that, well, maybe I could be an academic.

DE: What advice do you have for undergraduates looking to go into statistics?

I think the main advice that I could give someone contemplating going into statistics is, first of all: get a strong math training. The undergraduate training in statistics is not vital. I think the vital thing is the strength in mathematics, with a basic training in statistics. A wide variety of statistics courses is not crucial.

Secondly, when you go to graduate school, I think that the successful statisticians nowadays should have an interest in data, or at least be motivated by real data problems. If you are going into statistics because you like mathematics and you like delta-epsilon arguments for the sake of delta-epsilon arguments, I'm not sure that statistics is for you, because there is comparatively little demand nowadays for the sort of 1950s-style navel-gazing statisticians.

So my main advice for students contemplating graduate studies in statistics is you should have an interest in problems that are motivated by real world data. It doesn't mean that you should just be a data analyst, but you should be interested in real world problems.

JOKES

Top Ten Reasons to Become a Statistician:

- Deviation is considered normal.
- We feel complete and sufficient.
- We are "mean" lovers.
- Statisticians do it discretely and continuously.
- We are right 95% of the time.
- We can legally comment on someone's posterior distribution.
- We may not be normal but we are transformable.
- We never have to say we are certain.
- We are honestly significantly different.
- No one wants our jobs.

Q: What do you get when you cross a sherpa and a mountain goat?

A: Nothing. you cant cross two scalars.

Chuck Norris can accept the null hypothesis. (He can also divide by zero.)

Q: What is a compact city?

A: It's a city that can be guarded by finitely many near-sighted policemen!

GROUP THEORY AND SPECTROSCOPY: A CHEMIST'S VIEW ON SYMMETRY

Lai Chung Liu

Symmetry considerations through an application of group theory is widely used in chemistry to understand the interaction of electromagnetic waves with matter. We give a brief introduction of mathematical groups and their properties. Then, we illustrate the use of groups as simplifying agents in the treatment of n-body quantum mechanical computations.

1 INTRODUCTION

To mathematicians, group theory is an immensely important and beautiful subject of study. Among other things, it provides a formal framework for the description of symmetry. It is this ability that has guaranteed it an important place in the works of physics and chemistry. As matter transforms and interacts with energy, it often rearranges itself into neat geometric structures. Although its constituent parts behave in a point-to-point manner, their symmetry leads to a degeneracy in allowed behaviours. Through this consideration, radical simplifications can be carried out in attempts at modelling complex systems, and once daunting problems are resolved with much more easily.

2 DEFINITIONS

2.1 Mathematical Groups

Simply, group theory is a branch of mathematics that studies algebraic structures called *groups*. By definition, a group \mathbf{G} is a set of elements A, B, C, \dots together with a binary operation between A and B denoted AB that satisfy the following properties:

1. Closure: if $A, B \in \mathbf{G}$, then $AB \in \mathbf{G}$.
2. Associativity: $A(BC) = (AB)C \forall A, B, C \in \mathbf{G}$.
3. Identity: $\exists I$ such that $IA = AI = A \forall A \in \mathbf{G}$.
4. Inverse: $\forall A \in \mathbf{G}, \exists A^{-1} \in \mathbf{G}$ such that $AA^{-1} = A^{-1}A = I$.

The number of elements in the group is called the *order* of the group. A subset of the group that is closed under the group operation and the inverse operation is called a *subgroup* of the group.

Two elements $A, B \in \mathbf{G}$ are said to be *conjugate* if they are related by a similarity transformation $C^{-1}AC = B$ where C is some element of the group. A *conjugacy class* is just a complete set of mutually conjugate elements. Note that each element of a group belongs to exactly one class and the identity operator I is always in its own class. It can be shown that the order of all class (and subgroups) must be integral factors of the order of the group.

2.2 Point Groups

Consider the set of all one-to-one, isometric transformations to the space of \mathbb{R}^3 that leave at least one point unchanged. Basically, there are five types of such *symmetry operations*:

1. Proper rotation of $2\pi/n$ about an axis, denoted C_n .
2. Reflection across a plane, denoted σ .
3. Inversion through a point, denoted i .
4. Improper rotation or rotation-reflection of $2\pi/n$ about an axis and across the perpendicular plane, denoted S_n .
5. Identity, denoted E .

From these symmetry operations, the symmetry of any object in \mathbb{R}^3 can be described as the set of all operations under which the object is invariant. Unsurprisingly, this set forms a group and is called the *point group* of the object. This definition allows us to neatly classify all objects based on the symmetry. Although there are infinitely many point groups in \mathbb{R}^3 , there are only a few which are common in chemistry and they are summarized in table 2.2.

Type	Point Groups
Non-axial	C_1, C_s, C_i
Cyclic	$C_n, C_{nh}, C_{nv}, C_{\infty v}$
Dihedral	$D_n, D_{nh}, D_{nd}, D_{\infty h}$
Improper rotation	S_n
Tetrahedral	T_d
Octohedral	O_h
Icosahedral	O_h

Table 1: List of common point groups.

Here, each point group is identified with its Schönflies symbol. Non-axial groups describe structures with very low symmetry, having only the identity operation (C_1), one reflection plane (C_s) or one inversion point (C_i). The subscript n indicates the order of the principal axis while h, v shows the presence of reflection plane, perpendicular and vertical to the principal axis respectively. When $n = \infty$, we have a linear structure. A dihedral group differs from its related cyclic group only by having perpendicular C_2 rotation axes.

The classification of a chemical species to its point group can be quite tedious. As a result, this task is usually done via a point group assignment scheme that can be found in any introductory textbook on spectroscopy.

2.3 Group Representations

A representation Γ of a group \mathbf{G} is a group action of \mathbf{G} on a vector space V by invertible linear maps. Most groups have infinitely many different representations. A convenient way of visualising a representation is to map each group element to a matrix. Since the group operation corresponds to composition of linear maps in the representation (by the definition of a group action), and since composition of linear maps corresponds to multiplication of matrices, the group operation becomes matrix multiplication. Thus, we can reduce a group-theoretic problem into a linear algebra one.

Now, two representations are said to be equivalent if they are similar to each other. This means that there is a single matrix such that conjugating any matrix of the first representation by this fixed matrix gives the corresponding matrix in the second one. Furthermore, a subspace W of V that is fixed under the group action is called a *subrepresentation*. As with matrices, a representation Γ can always be decomposed into a direct sum of subrepresentations Γ_i . If the direct sum is non-trivial, then the representation is said to be reducible, otherwise it is irreducible. It follows that the number of irreducible representations of a group is always equal to the number of conjugacy classes in the group.

2.4 Character of Group Representations

Let g be an element of group \mathbf{G} and Γ be a representation of \mathbf{G} . Define the *character* of g in the representation Γ as

$$\chi_{\Gamma}(g) = \text{Tr}(\Gamma(g)) \quad (1)$$

By this definition, the character of irreducible representations encode many important properties of the group in an even more compact form. In particular, it allows us to simply express the properties of these representations. Let h be the order of group \mathbf{G} , g an element of \mathbf{G} , $c(g)$ the order of the class containing g , Γ a reducible representation, and Γ_i the i th irreducible representation. Then, we have the following:

1. Direct sum: $\chi_{\Gamma_1 \oplus \Gamma_2}(g) = \chi_{\Gamma_1}(g) + \chi_{\Gamma_2}(g)$
2. Tensor product: $\chi_{\Gamma_1 \otimes \Gamma_2}(g) = \chi_{\Gamma_1}(g)\chi_{\Gamma_2}(g)$
3. Orthogonality theorem: $\sum_g \chi_i(g)\chi_j(g) = \delta_{ij}$
4. Decomposition theorem: $\chi_{\Gamma}(g) = \sum_i a_i \chi_{\Gamma_i}(g)$, where $a_i = \frac{1}{h} \sum_g c(g)\chi_{\Gamma_i}(g)\chi_{\Gamma}(g)$

To present information about a group, a *character table* is formed by listing the characters of the irreducible representation by rows while the group elements are listed by columns. An example is table 2.

C_{3v}	E	$2C_3$	$3\sigma_v$
A_1	1	1	-1
A_2	1	1	-1
E	2	-1	0

Table 2: Character table of the point group C_{3v} .

2.5 Mulliken Symbols

To identify irreducible representations of a group, they are named using the Mulliken symbols which are defined using the corresponding characters:

- A: symmetric under rotation, $\chi_{\Gamma}(C_n) = 1$
- B: symmetric under rotation, $\chi_{\Gamma}(C_n) = -1$
- E: doubly degenerate, $\chi_{\Gamma}(E) = 2$
- T: triply degenerate, $\chi_{\Gamma}(E) = 3$
- Γ_g : symmetric under inversion, $\chi_{\Gamma}(i) = 1$
- Γ_u : antisymmetric under inversion, $\chi_{\Gamma}(i) = -1$
- Γ'' : symmetric under reflection, $\chi_{\Gamma}(\sigma_h) = 1$
- Γ' : antisymmetric under reflection, $\chi_{\Gamma}(\sigma_h) = -1$

3 SYMMETRY AND SPECTROSCOPY

3.1 Notes on Molecular Spectroscopy

In general, spectroscopy is the study of the interaction between matter and electromagnetic waves. By probing the rotational, vibrational and electronic energy levels of molecules, it can reveal important information about their structure and dynamics. A relatively simple example is infrared spectroscopy which focuses on the infrared (IR) region of the electromagnetic spectrum (0.3 – 400 THz).

A typical IR experiment would involve the capture of the absorption spectrum of a sample by passing beams of IR light through it and measuring the frequency dependence of the transmitted intensity. Peaks in such spectra correspond to the natural frequencies at which molecules vibrate. Classically, a harmonic oscillator can be strongly driven by an periodic external force when they are in resonance. In quantum mechanics, the molecular Hamiltonian can be approximated by a harmonic oscillator in the neighborhood of the equilibrium molecular geometry under the Born-Oppenheimer and harmonic approximation. Then, the resonant frequencies are simply those associated the energy of transition from one normal mode of vibration to another.

3.2 Vibrations in Polyatomic Molecules

Consider a molecule of N atoms in \mathbb{R}^3 . Clearly, it can undergo much more complex vibrations than the simple oscillations of a diatomic molecule. However, all of this motion can be resolved into a superposition of a bound set of fundamental motions called normal modes of vibrations. By subtracting the rotational degrees of freedom from the total number of degree of freedoms, we get that the number of normal modes is $3N - 5$ for linear molecules, and $3N - 6$ otherwise. Physically, these vibrations can be recognizes as various forms of stretchings and bendings of the molecular bonds. A key observation on these modes can be made by assigning associated displacement vectors to each atom and considering their behaviour under the symmetry operators in the point group of the molecule. Then, it can be shown that:

Each set of normal modes of vibration forms a basis for an irreducible representation of the point group of the molecule.

In addition, each set of translation and rotation degrees of freedom forms a similar basis. With respect to spectroscopy, it is the symmetry of the modes that is of interest. To uncover the associated irreducible representations, one would need to construct the transformation matrices and compute their trace. However, for brevity, it suffices to proceed along the following steps:

1. Determine Γ_{stat} , the representation where each character is the number of atoms which are stationary under each symmetry operation.
2. Determine $\Gamma_{trans} \equiv \Gamma(x) \oplus \Gamma(y) \oplus \Gamma(z)$, the representation for the translational degrees of freedom.
3. Determine $\Gamma_{rot} \equiv \Gamma(R_x) \oplus \Gamma(R_y) \oplus \Gamma(R_z)$, the representation for the rotational degrees of freedom.
4. Determine $\Gamma_{tot} \equiv \Gamma_{stat} \otimes \Gamma_{trans}$, the representation for all degrees of freedom.
5. Determine $\Gamma_{vib} \equiv \Gamma_{tot} \ominus \Gamma_{trans} \ominus \Gamma_{rot}$, the representation for the vibrational degrees of freedom.
6. Apply the decomposition theorem to decompose Γ_{vib} into Γ_i , the irreducible representations.

Then, each Γ_i represents the symmetry of a particular normal mode v_i , denoted $\Gamma(v_i)$. To check the consistency of the calculations, it is useful to count the number of $\Gamma(v_i)$ from the decomposition and compare it with the total number of modes ($3N - 6, 3N - 5$). Unsurprisingly, this requirement is always satisfied when the multiplicity and degeneracy of component irreducible representations are considered.

3.3 Transition under Perturbation

Spectroscopy studies the effects of electromagnetic waves acting on the quantum state of electrons in matter. Given the time-varying nature of such waves, an application of time-dependent perturbation theory is necessary. First, consider the unperturbed Hamiltonian \hat{H}_0 and a chosen basis of orthonormal energy eigenkets $|n\rangle$ with their associated energy eigenvalues E_n :

$$\hat{H}_0|n\rangle = E_n|n\rangle \quad (2)$$

Introducing a time-dependent perturbation $\hat{V}(t)$, we get the Hamiltonian of the perturbed system.

$$\hat{H} = \hat{H}_0 + \hat{V}(t) \quad (3)$$

Let $|\psi(t)\rangle$ denote the quantum state of \hat{H} at time t . It obeys the time-dependent Schrödinger equation:

$$\hat{H}|\psi(t)\rangle = i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle \quad (4)$$

Since $|\psi(t)\rangle$ still belongs to the Hilbert space with basis $\{|n\rangle\}$, it can be written as

$$|\psi(t)\rangle = \sum_n c_n(t) e^{-iE_n t/\hbar} |n\rangle \quad (5)$$

where $c(t)$ are undetermined probability amplitudes. Combining Eq. (4) and (5), we obtain

$$\sum_n \left(i\hbar \frac{\partial}{\partial t} c_n(t) - c_n(t) \hat{V}(t) \right) e^{-iE_n t/\hbar} |n\rangle = 0 \quad (6)$$

By composing the expression with the bra $\langle k|$, it is reduced to a set of partial differential equations for the amplitudes.

$$\frac{\partial}{\partial t} c_n(t) = \frac{-i}{\hbar} \sum_k \langle n|\hat{V}(t)|k\rangle c_k(t) e^{-i(E_k - E_n)t/\hbar} \quad (7)$$

We define M_{kn} , the transition moment integral corresponding to the transition $k \rightarrow n$, as follows

$$M_{kn} = \langle n|\hat{V}(t)|k\rangle = \int \psi_n^* V(t) \psi_k d\tau \quad (8)$$

where ψ_j is the wavefunction of the eigenket $|j\rangle$.

3.4 Selection Rules

From section 3.3, it can be noted that $c_n(t) \equiv 0$ if $M_{kn} = 0$ and $c_n(0) = 0$. Physically, this implies that an initial state $|k\rangle$ cannot evolve into some final state $|n\rangle$ under the perturbation $\hat{V}(t)$, and the transition $k \rightarrow n$ is thus forbidden. By specifying conditions for which $M_{kn} \neq 0$, we can derive selection rules for allowed transitions.

For vibrational spectroscopy, the perturbing Hamiltonian is defined as

$$\hat{V}(t) = \hat{\mu} \cdot \hat{\mathbf{E}} \quad (9)$$

where $\hat{\mu} = \sum_i q_i \hat{\mathbf{r}}_i$ is the electric dipole moment operator and $\hat{\mathbf{E}}$ is the oscillating electric field operator of the electromagnetic wave. Then, we have

$$M_{kn} \propto \int \psi_n^* \mu \psi_k d\tau \quad (10)$$

Recall that

$$\int f(\mathbf{r})d\tau = \begin{cases} I & \text{if } f(\mathbf{r}) = f(-\mathbf{r}) \\ 0 & \text{if } f(\mathbf{r}) = -f(\mathbf{r}) \end{cases} \quad (11)$$

where $I \neq 0$. Thus, the problem of selection rules simply becomes a group-theoretic one:

$$\int \psi_n^* \mu \psi_k d\tau \neq 0 \text{ if } \Gamma(\psi_n^* \mu \psi_k) \subset A_g'' \quad (12)$$

where A_g'' is the totally symmetric irreducible representation. Since representations have the composition property, we can make this requirement more specific by considering the symmetry of the individual product elements.

From the definition, the electric dipole moment μ is a sum of linear terms in the Cartesian coordinates x, y, z . Then,

$$\Gamma(\mu) = \begin{pmatrix} \Gamma(x) \\ \Gamma(y) \\ \Gamma(z) \end{pmatrix} \quad (13)$$

Furthermore, from section 3.2, it is known that the vibrational wavefunctions are those of the harmonic oscillator. We denote by ψ_{nv_n} the wavefunction of the n^{th} normal mode of vibration v_n , at the v_n energy level. These functions are of the form $H_v(q)e^{-q^2}$, where H_v is the v^{th} Hermite polynomials and q is the normal coordinate. Then, it can be shown that

$$\Gamma(\psi_{nv_n}(q)) = \begin{cases} A_g'' & \text{if } v_n \text{ even} \\ \Gamma(v_n) & \text{if } v_n \text{ odd} \end{cases} \quad (14)$$

Finally, we have derived the selection rule for vibrational transitions by appealing to the symmetry of the problem. A similar derivation can be made for the selection rules associated with other types of transitions by swapping in another perturbing Hamiltonian (e.g. the polarization operator $\hat{\alpha}$ for Raman spectroscopy). Indeed, if the power of groups were not invoked, then we would have to evaluate the transition moment integral directly, a most unwelcome prospect to any chemist or mathematician.

4 PRACTICAL EXAMPLE

Using the derivation of the selection rule, we can perform a vibrational analysis of any molecule by checking the IR activity of all its possible transitions, thus allowing us to predict its IR spectrum. As an example, such an analysis is carried out here for the molecule dihydrogen monoxide, H_2O .

The molecule H_2O is bent, with two hydrogen attached at an angle to a central oxygen atom. Its structure thus belongs to the C_{2v} point group. Following the steps listed in section 3.2, we obtain table 3.

C_{2v}	E	C_2	$\sigma_v(xz)$	$\sigma'_v(yz)$	
A_1	1	1	1	1	z
A_2	1	1	-1	1	R_z
B_1	1	-1	1	1	x, R_y
B_2	1	-1	-1	1	y, R_x
Γ_{stat}	3	1	1	3	
Γ_{trans}	3	-1	1	1	
Γ_{rot}	3	-1	-1	3	
Γ_{tot}	9	-1	1	3	
Γ_{vib}	6	1	1	-1	

Table 3: Vibrational analysis of H_2O .

Applying the decomposition theorem,

$$\begin{aligned} \Gamma_{trans} &= \Gamma(x) + \Gamma(y) + \Gamma(z) \\ &= A_1 + B_1 + B_2 \\ \Gamma_{rot} &= \Gamma(R_x) + \Gamma(R_y) + \Gamma(R_z) \\ &= A_2 + B_1 + B_2 \\ \Gamma_{tot} &= \Gamma_{stat} \times \Gamma_{trans} \\ &= 3A_1 + A_2 + 2B_1 + 3B_2 \\ \Gamma_{vib} &= \Gamma_{tot} - \Gamma_{trans} - \Gamma_{rot} \\ &= 2A_1 + B_2 \end{aligned}$$

Let $\Gamma(v_1) = A_1, \Gamma(v_2) = A_1, \Gamma(v_3) = B_2$. We now consider the fundamental transition for each mode: $v = 0 \rightarrow v = 1$. Then,

$$\Gamma(\psi_{1,0}\mu\psi_{1,1}) = A_1 \times \begin{pmatrix} A_1 \\ B_1 \\ B_2 \end{pmatrix} \times A_1 = \begin{pmatrix} A_1 \\ B_1 \\ B_2 \end{pmatrix} \subset A_1 \quad (15)$$

$$\Gamma(\psi_{2,0}\mu\psi_{2,1}) = A_1 \times \begin{pmatrix} A_1 \\ B_1 \\ B_2 \end{pmatrix} \times A_1 = \begin{pmatrix} A_1 \\ B_1 \\ B_2 \end{pmatrix} \subset A_1 \quad (16)$$

$$\Gamma(\psi_{3,0}\mu\psi_{3,1}) = B_2 \times \begin{pmatrix} A_1 \\ B_1 \\ B_2 \end{pmatrix} \times A_1 = \begin{pmatrix} B_2 \\ A_2 \\ A_1 \end{pmatrix} \subset A_1 \quad (17)$$

Therefore, there are three normal modes of vibration in H_2O and they are all IR active, leading to three strong peaks in the IR spectrum of water.

REFERENCES

- [1] G.M. Barrow. *Introduction to Molecular Spectroscopy*. McGraw-Hill, New York, 1962.
- [2] D.C. Harris and M.D. Bertolucci. *Symmetry and Spectroscopy*. Oxford University Press, Oxford, 1978.
- [3] G. Herzberg. *Electronic Spectra of Polyatomic Molecules*. Van Nostrand Reinhold, New York, 1966.

THE TREPIDATION THAT IS MATH ANXIETY

Maya Kaczorowski

Math anxiety is a psychological phenomenon wherein individuals who are capable of solving a math problem because they possess the necessary skills are nonetheless unable to solve it due to frustration and lack of confidence in their own math abilities. Math anxiety is surprisingly wide-spread in our society, and often, for instance, affects students studying to become elementary school teachers, creating a vicious cycle which partially contributes to new students developing math anxiety. In this article, we briefly examine the current literature about the causes, effects, and those affected by math anxiety.

1 WHAT IS MATH ANXIETY?

Mathematics anxiety describes the feeling of helplessness and frustration experienced by individuals who, when confronted with math problems within their abilities, are nevertheless unable to solve them. Often, these individuals feel nervous prior to a math test, have trouble concentrating in class, or panic during an exam and blank on what they have learned. Math anxiety can become a concern as it limits university and career options, and furthermore can create uneasiness, emotional strain and embarrassment in social situations such as splitting a bill at a restaurant or calculating the amount of each ingredient required for a recipe. Sheila Tobias first described this phenomenon and coined the term “math anxiety” in a magazine article she wrote in 1976. She followed with a book two years later entitled *Overcoming Math Anxiety*.

2 IDENTIFYING OR DIAGNOSING MATH ANXIETY

Psychologists have developed several diagnostic tests based on simple, direct questions to determine who exhibits math anxiety. The most commonly used are Richardson and Suinn’s Mathematics Anxiety Rating Scale (MARS), Fennema and Sherman’s Mathematics Anxiety Scale (MAS) and Sandman’s Anxiety Toward Mathematics Scale (ATMS) [6, 7]. All of them ask questions about the feelings invoked when exposed to math, such as whether the individual feels uncomfortable and nervous or uneasy and confused during a math class or test [4]. All these tests have high re-test reliabilities. [For those interested, a short self-test can be found at <http://www.mathpower.com/anxtest.htm>.]

3 THOSE AFFECTED: RELATIONS TO GENDER AND AGE

Numerous studies have been conducted to determine who is more likely to be affected by math anxiety in terms of gender, age, math performance, perceptions about math, and math experience. It is impossible to know the prevalence of math anxiety as it is measured on a continuum, and there exists no set threshold which determines who does and does

not have math anxiety. Ashcraft and Kirk [2] chose their boundaries depending on the data set, so determining who qualified as having low-, medium- and high-math anxiety was relative. Out of 66 participants, they placed 18% at low-, 35% at medium- and 23% of participants at high-math anxiety, leaving 24% of participants described as non-math anxious. Another study testing the prevalence of math anxiety across educational backgrounds found that anywhere from one quarter to one half of all college students experience significant uneasiness in doing math, however the authors give no concrete estimates for those identified as having math anxiety [4].

The prevalence of math anxiety differs by gender. It has been shown in several studies that girls are more likely to experience math anxiety than boys, and girls also have much stronger negative affective reactions to math [4, 6, 9, 10, 12]. Boys have higher perceptions, expectations and intentions related to mathematics [10], however, math anxiety appears to be more strongly correlated with poor performance in mathematics in pre-university males than females [8].

The gender differences in math anxiety have some consistency across cultures, but the motivations may be different. In comparing math anxiety in the United States and Thailand, girls were more likely to exhibit math anxiety than boys in both cultures; however, only the girls’ performances in math were strongly related to their mother’s level of education. This suggests that a mother who performs poorly in math and who possibly has math anxiety is more likely to pass these traits on to her daughter [5].

Hembree’s study proposes the following reason for these gender differences:

“Across all grades, female students report higher mathematics anxiety levels than males. However, the higher levels do not seem to translate into more depressed performance or to greater mathematics avoidance on the part of female students. Indeed, male students in high school exhibit stronger negative behaviors in both these regards. This paradox may be explained along two lines: 1) Females may be more willing than males to admit their anxiety, in which case their higher levels are no more than a reflection of societal mores; 2) females may cope with anxiety better.” ([8], p.45)

Few studies have looked at the relation between the prevalence of math anxiety and different racial groups.

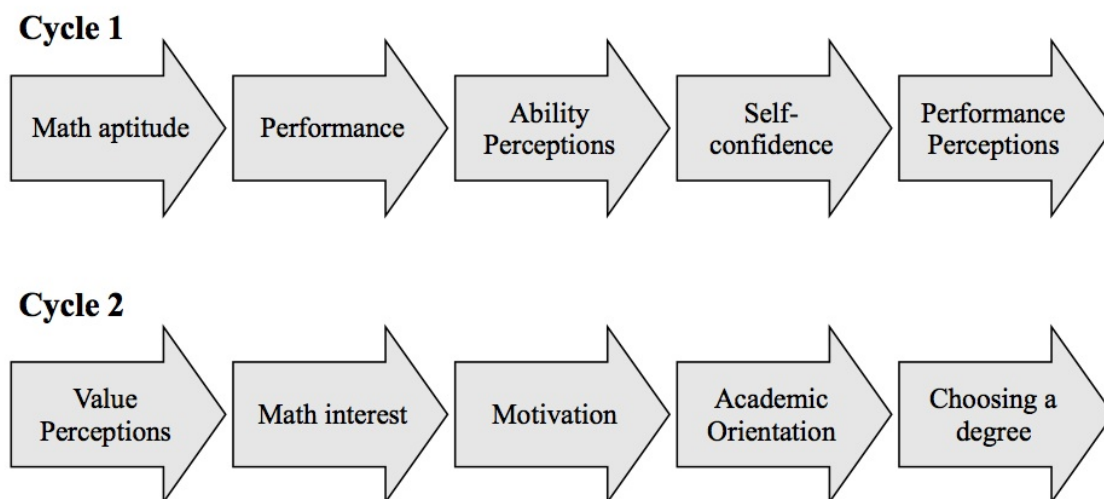


Figure 1: Cycles of the two proposed conceptual models of math anxiety

Hembree [8] found some correlations between the two. Although it seems that there is no difference between white and black student's likelihood of math anxiety, hispanic students on average had higher rates of math anxiety than any other racial group in the study.

There also exists a general progression in the development of math anxiety as individuals age. The prevalence of math anxiety among students in the 6th grade is very low, increasing to much higher levels in grade 9 students [12]. This is likely caused by the stress of beginning high school and more challenging math courses. Relatively speaking, freshman university students have low math anxiety compared to high school students [11]. However, there is still much variation within a group: in comparing female students entering university directly after high school with those who first took some time off, the older women were significantly more math anxious than those who had continuously been in school, possibly because of lack of practice in their time off [4].

4 A PROPOSED CONCEPTUAL MODEL OF MATH ANXIETY

Even though studies have been conducted over a wide range of age groups and therefore mathematical ken, which then makes the information we have on math anxiety difficult to synthesize, this can only help in providing the additional data and knowledge needed to track the development of math anxiety as individuals age. We now examine the theories of the progression of math anxiety as individuals age.

The progression of math anxiety due to math aptitude and value perceptions can be represented in two cycles. Several studies conducted according to this theoretical framework support this model due to the correlations be-

tween math anxiety and the cycle at each step. Both cycles are negatively correlated with math anxiety.

Cycle 1. It has been shown that higher achievement in mathematics is related to lower levels of math anxiety across several cultures including the North American culture [5] so that math aptitude and math anxiety are negatively correlated. Higher achievement and better academic performance leads students to form higher ability perceptions. In particular, the component of math anxiety which influences the individual's emotions and responsiveness to math, called the affective component, has been shown to be negatively correlated with math ability perceptions and performance expectancies [10, 12]. It has furthermore been found that negative self-perceptions about math due to math anxiety lead to lower performance [1]. This is likely because highly math anxious people end up with lower math competence and achievement as they intentionally select less exposure to math and learn less of the material to which they are exposed [1].

Cycle 2. The worry component of math anxiety, which makes the individual anxious during stressful situations, has been shown to be strongly negatively correlated to value perceptions and effort in mathematics [10, 12]. The perception that mathematics is a valuable skill leads to interest in math, which is also negatively correlated with math anxiety [11]. A highly math anxious student has low interest in math which gives little motivation to pursue math; therefore, math anxiety and motivation are also negatively correlated [1]. A lack of motivation and interest in mathematics and mathematics-related fields will influence what a student chooses to study in both high school and university. There exists a negative correlation between math anxiety and academic orientation [11]; in fact, studies show that

the more high school math and the higher levels of mathematics a student studies, the less likely they are to be math anxious [4, 11]. Those majoring in university in the physical sciences or mathematics have the least math anxiety, while students studying education have the most. The students with the highest math anxiety were those taking lower level math courses, such as remedial mathematics courses or mathematics for elementary school teachers [8]. All of the evidence supports the model's theorized negative correlation between math anxiety and academic orientation towards the physical sciences.

5 DEVELOPMENT OF MATH ANXIETY, AND POSSIBLE TREATMENT OR PREVENTION

Although it is uncertain what causes math anxiety, the examination of the relationships between math anxiety and test anxiety, working memory and physiological reactions may explain how math anxiety develops, and therefore pinpoint the best way to prevent it.

One study aimed to identify symptoms of math anxiety as physical behaviors; subjects were first tested for math anxiety on a variety of math anxiety test scales such as MARS, and then answered math questions [7]. Physiological reactions and avoidance were measured and very little relation was found between math anxiety and any physical symptoms of anxiety. The authors suggested that "it may be that individuals avoid math problem-solving situations, but once in them emit few avoidance behaviors." (p. 583)

Test anxiety can be separated into two components, worry, the cognitive component, and emotionality, the affective component [12]. These two components can also be applied to math anxiety. Math anxiety, like test anxiety, causes students to perform badly during a test because anxiety uses an individual's working memory span. A working memory, or an 'on-line' memory, is the limit on the amount of information an individual's brain can process at any given time. People who are very math anxious or test anxious will actually spend a lot of their working memory during a test worrying instead of thinking of the problem at hand; this is called a transitory disruption of the working memory [2]. Highly math anxious people will often sacrifice accuracy for speed, trying in a test situation to complete more problems but with more errors than when relaxed, especially as the difficulty of problems increase [2, 3]. So students with math anxiety will have a reduced working memory and therefore perform poorly.

A few methods have been tested to reduce levels of math anxiety in students and found to be fairly successful, such as systematic desensitization [13] and behavioral or cognitive-behavioral treatments [8]. Other treatments such as classroom interventions, relaxation training and group counseling have no effect on reducing math anxiety. However, like most medical or psychological problems, it is often more financially viable for society to prevent math anxiety rather

than treat it.

When examining methods to prevent math anxiety, an alarming trend becomes apparent. Out of a variety of university students tested for math anxiety, those who on average had the highest rates of math anxiety are the students preparing to become elementary school teachers [3]. This goes in tandem with the fact that students enrolled in lower level math courses, which includes math courses for education students, had more math anxiety than students enrolled in higher level courses [8, 11]. Math anxiety may be caused in part by highly demanding teachers who insist on correctness, yet provide little motivational support for students in their lessons [1]. This type of teaching style, more present in math classes, creates avoidance behaviors.

A parallel with parental influence can be seen: a mother with a lower educational background in mathematics and therefore more likely to have math anxiety will have children who are more likely to experience math anxiety [5]. If it is assumed that this is caused by learned behavior, in other words, that this is influenced by mothers teaching their children at home or helping them with their homework, what happens when a child's teacher experiences math anxiety? Is it passed on to the student? Studies have yet to be conducted analyzing this, but this may be the cause of the ineffectiveness of classroom interventions in treating math anxiety [8]. If the classroom is a major source of math anxiety, then without first dealing with the issues of math anxiety that affect educators, treatments for math anxiety in the classroom may be unsuccessful.

6 IMPACT ON INDIVIDUALS

It is clear that math anxiety can have a negative influence on an individual's life. Math anxiety relates to a student's performance, perceptions, motivation, and academic orientation. Those affected tend to have low math aptitude, are female, and may be Hispanic or have taken some time off prior to beginning post-secondary education. When examining the above cycles, we note that the only way to break a cycle is to successfully address one of its components; effectively, cut it out of the cycle. So which parts of the cycle should educators target to help these students? A few treatment methods for math anxiety have been shown to work, but it is much better to prevent rather than to treat. From previous findings of the relation between teaching methods and students' math anxiety, the best recommendation is that all elementary and grade school teachers be assessed for math anxiety and supported in trying new methods for teaching mathematics. As it is critical that all adults in this age of technology possess the basic skills of numeracy in order to effectively function as full-fledged members of society, math anxiety is a phenomenon which educators should take seriously and address as early as possible in a child's education.

REFERENCES

- [1] M.H. Ashcraft. Math anxiety: Personal, educational, and cognitive consequences. *Current Directions in Psychological Sciences*, 11(5):181–185, 2002.
- [2] M.H. Ashcraft and E.P. Kirk. The relationships among working memory, math anxiety, and performance. *Journal of Experimental Psychology: General*, 130(2):224–237, 2001.
- [3] M.H. Ashcraft and J.A. Krause. Working memory, math performance, and math anxiety. *Psychonomic Bulletin and Review*, 14(2):243–248, 2007.
- [4] N.E. Betz. Prevalence, distribution, and correlates of math anxiety in college students. *Journal of Counseling Psychology*, 25(5):441–448, 1978.
- [5] G. Engelhard. Math anxiety, mothers education, and the mathematics performance of adolescent boys and girls: Evidence from the united states and thailand. *The Journal of Psychology*, 124(3):289–298, 2001.
- [6] K.M. Harriss Dew, J.P. Galassi, and M.D. Galassi. Mathematics anxiety: Some basic issues. *Journal of Counseling Psychology*, 30(3):443–446, 1983.
- [7] K.M. Harriss Dew, J.P. Galassi, and M.D. Galassi. Math anxiety: Relation with situational test anxiety, performance, physiological arousal, and math avoidance behaviour. *Journal of Counseling Psychology*, 31(4):580–583, 1984.
- [8] R. Hembree. The nature, effects, and relief of mathematics anxiety. *Journal for Research in Mathematics Education*, 21(1):33–46, 1990.
- [9] M.M. Llabre and E. Suarez. Predicting math anxiety and course performance in college women and men. *Journal of Counseling Psychology*, 32(2):283–287, 1985.
- [10] J.L. Meece, A. Wigfield, and J.S. Eccles. Predictors of math anxiety and its influence on young adolescents course enrollment intentions and performance in mathematics. *Journal of Educational Psychology*, 82(1):60–70, 1990.
- [11] H. Resnick, J. Viehe, and S. Segal. Is math anxiety a local phenomenon? a study of prevalence and dimensionality. *Journal of Counseling Psychology*, 29(1):39–47, 1982.
- [12] A. Wigfield and J.L. Meece. Math anxiety in elementary and secondary school students. *Journal of Educational Psychology*, 80(2):210–216, 1988.
- [13] R.D. Zettle. Acceptance and commitment therapy (act) vs. systemic desensitization in treatment of mathematics anxiety. *The Psychological Record*, 53:197–215, 2003.

JOKES

Two mathematicians went out to lunch. Over lunch, one complained that most people don't understand even basic math. The other took a more optimistic view. A short time later, while the pessimist was in the bathroom, the other called the waitress over.

"I am going to call you over in a few minutes," he explained, "and I am going to ask you a question. I want you to answer $x^3/3$. OK?"

When the pessimist came back, he called the waitress over. "Look, I'll prove people understand math better than you think. OK, young lady, what is the integral of x^2 ?"

" $x^3/3$," she slowly repeated and walked away. Then she turned around and said, "plus a constant." □

On a blackboard at Concordia, " $1 + 1 = 3$, for large values of 1" □

A doctor, a lawyer and a mathematician were discussing the relative merits of having a wife or a mistress.

The lawyer says: "For sure a mistress is better. If you have a wife and want a divorce, it causes all sorts of legal problems."

The doctor says: "It's better to have a wife because the sense of security lowers your stress and is good for your health."

The mathematician says: "You're both wrong. It's best to have both so that when the wife thinks you're with the mistress and the mistress thinks you're with your wife – you can do some mathematics." □

A topologist is a person who doesn't know the difference between a coffee cup and a doughnut. □

Asked if he believes in one God, a mathematician answered:

"Yes, up to isomorphism." □

ÉCRITURES RÉDUITES DE L'ÉLÉMENT LE PLUS LONG DES GROUPES DE COXETER FINI

Maxime Bergeron, Marco Robado et Maxime Scott

We introduce group presentations by generators and relations to examine in further detail certain properties of the longest element of Coxeter groups. We are particularly interested in c -factorization and the c -sortability of such elements because of their connections to Generalized Associahedra, Lie Algebras and Cluster Algebras.

Nous introduisons les concepts de présentation par générateurs et relations d'un groupe afin d'étudier les propriétés particulières de l'élément le plus long de groupes de Coxeter. Nous sommes particulièrement intéressé par la c -factorisation de tels éléments et le fait qu'ils soient c -trié puisqu'il existe des liens entre ceux-ci et les associaèdres, les algèbres de Lie et les algèbres ammassées.

1 INTRODUCTION

La notion de groupe, fondamentale en mathématiques, est toujours des plus active. Toutefois, les groupes abstraits sont parfois difficiles à visualiser. C'est pourquoi, nous avons étudié les groupes de réflexions finis qui sont facilement représentés par la notion de symétrie. L'étude des groupes de symétrie trouve des applications dans de nombreux domaines, tels la théorie de Lie, la récente théorie des algèbres ammassées (Cluster Algebra), les associaèdres de même qu'en cristallographie, un domaine de la chimie.

Nous nous sommes intéressés plus particulièrement à la théorie des éléments c -triés (c -sortable en anglais) qui donne une nouvelle interprétation des associaèdres; cette notion a été introduite par Nathan Reading dans [5]. La construction des éléments c -triés a été réduite à l'étude d'une certaine écriture du mot le plus long dans un groupe de réflexion fini (voir [3]).

Exemple préliminaire

Pour illustrer les concepts abordés dans cet article, nous allons considérer l'exemple du groupe diédral \mathcal{D}_4 , le groupe des symétries qui préserve le carré. Rappelons que le groupe \mathcal{D}_4 est engendré par la symétrie s_1 d'axe une diagonale du carré et par la rotation r d'angle $\frac{\pi}{2}$. C'est en fait un groupe de réflexion, car il peut être engendré par $S = \{s_1, s_2\}$ où $s_2 = s_1 \circ r$; la réflexion obtenue en composant r avec s_1 . On peut noter les éléments de S comme suit:

$$e, s_1, s_2, s_1s_2, s_2s_1, s_1s_2s_1, s_2s_1s_2, s_1s_2s_1s_2$$

Notons que chacun de ces éléments peut s'écrire de plusieurs façons différentes à l'aide de relations telles que $e = r^4 = (s_1s_2)^4$ où e représente l'élément neutre du groupe diédral. Observons aussi que l'élément $w_0 := s_1s_2s_1s_2$ est le plus long mot sur \mathcal{D}_4 et qu'il s'écrit d'une seule autre manière: $s_2s_1s_2s_1$. Plus formellement, on peut voir les générateurs s_1 et s_2 comme des éléments d'un alphabet, et les éléments du groupe comme des mots sur cet alphabet.

¹nous noterons $\mathcal{M}(S) := \mathcal{M}(S \cup S^{-1})$ pour ne pas alourdir l'écriture

Notre objectif est de pouvoir, selon certaines hypothèses, choisir automatiquement l'une de ces expressions.

Nous commencerons d'abord par définir les groupes libres ainsi que la théorie de la présentation par générateurs et relations. Nous ferons ensuite un survol de la théorie des groupes de Coxeter finis. Enfin, nous aborderons le coeur de notre problématique en parlant de la c -factorisation des mots de Coxeter, plus particulièrement pour w_0 , le mot le plus long du groupe.

2 GROUPES LIBRES ET PRÉSENTATIONS DE GROUPES

Étant donné un groupe et un ensemble engendrant ce groupe (comme dans l'exemple de \mathcal{D}_4 avec l'ensemble $S = \{s_1, s_2\}$), l'intuition nous pousse à parler des éléments du groupe comme des mots sur l'alphabet des générateurs. Cette partie introduit les concepts permettant de parler des éléments de n'importe quel groupe en tant que mots sur un ensemble de générateurs, donnant ainsi un sens précis à l'intuition précédemment mentionnée. Pour plus de détails ainsi que les preuves, nous invitons les lecteurs à consulter n'importe quel bon livre d'algèbre abstraite.

2.1 Groupes libres

Soit S un ensemble de symboles appelé *ensemble générateur*. Posons l'ensemble S^{-1} , un ensemble en bijection avec S , tel que pour tout $s_i \in S$, on associe $s_i^{-1} \in S^{-1}$. L'ensemble des mots sur $S \cup S^{-1}$, noté $\mathcal{M}(S)$ ¹ forme alors un monoïde pour la concaténation des mots. Posons, $R \subset \mathcal{M}(S)$ appelé *relations d'adjacence*. L'ensemble R est choisi de sorte que nous puissions effectuer les réductions de forme $s_i s_i^{-1} = e_{F_S}$ où e_{F_S} représente l'élément neutre, associant effectivement un inverse à chaque élément de S . De cette manière, le quotient $\mathcal{M}(S)/R$ forme un groupe que l'on appelle le *groupe libre* sur l'alphabet S et on le note F_S .

Exemple 1. Soit $S = \{a, b, c\}$ un alphabet, alors $S^{-1} =$

$\{a^{-1}, b^{-1}, c^{-1}\}$. De plus,

$$abc, abb^{-1}ccba \text{ et } c^{-1}c^{-1}cc$$

sont des exemples d'éléments de $\mathcal{M}(S)$. Posons maintenant $F_S = \mathcal{M}(S)/R$ et notons


$$\pi: \mathcal{M}(S) \rightarrow \mathcal{M}(S)/R = F_S$$

la projection canonique. Alors abc est un représentant de la classe $\pi(abc)$, $accba$ est un représentant de la classe $\pi(abb^{-1}ccba)$ et le mot vide que nous notons e_{F_S} est un représentant de la classe $\pi(c^{-1}c^{-1}cc)$.

2.2 Présentation par générateurs et relations

Nous voulons maintenant pouvoir parler d'un groupe quelconque comme étant un ensemble de mots sur des générateurs. Pour ce faire, il nous faut rajouter un autre ensemble de relations par lequel nous quotienterons un groupe libre. Nous verrons que n'importe quel groupe est isomorphe au quotient d'un groupe libre. La *présentation* du groupe est alors la paire formée d'un ensemble générateur et de relations entre les générateurs. Nous tentons donc de trouver les relations nécessaires et suffisantes afin de pouvoir entièrement décrire un groupe donné.


Prenons G un groupe quelconque avec S un sous-ensemble générateur de G ; G est alors isomorphe à un quotient du groupe libre F_S . En particulier, il existe un ensemble minimal $R \subseteq F_S$ et un sous-groupe normal $N(R) \triangleleft F_S$ (l'intersection de tous les sous-groupes normaux dans F_S contenant R) tels que les éléments de $N(R)$ correspondent aux relations entre les éléments du groupe G . On considère alors la projection canonique $f: F_S \rightarrow G$ l'unique morphisme de groupe associant à chaque élément de S le générateur de G lui correspondant. Le noyau de cette projection correspondant précisément à $N(R)$, on peut voir que $F_S/N(R) \simeq G$ tel qu'illustré dans le diagramme suivant:



Exemple 2.

$$\mathcal{D}_m \simeq \langle r, s_1 \mid s_1^2, r^m, s_1 r s_1 r \rangle$$

Proof. Nous savons que $\mathcal{D}_m = \langle r, s_1 \rangle$ et donc en prenant $S = \{r, s\}$, on construit F_S l'ensemble des mots dans S et on obtiens le diagramme suivant où f représente la projection des mots de l'alphabet S dans le groupe \mathcal{D}_m :



Nous savons que dans \mathcal{D}_m , $s^2 = r^m = s r s r = e$ et donc pour $N(R) := N(s^2, r^m, s r s r)$, le sous-groupe normal de F_S engendré par ces éléments, nous avons $N(R) \subseteq \ker(f)$. D'autre part, en utilisant les relations de R , tout élément du groupe peut être écrit comme un produit $s_1^i r^j$ ou $1 \leq i \leq 2$ et $1 \leq j \leq m$. En effet, étant donné un élément, la relation $s_1 r = r^{-1} s_1$, nous permet de le réécrire tel qu'indiqué ci-dessus. On obtient donc la liste de $2m$ éléments :

$$(e, r, r^2, \dots, r^{m-1}, s_1, s_1 r, s_1 r^2, \dots, s_1 r^{m-1})$$

qui sont tous distincts puisque

$$r^j = r^k \iff j = k, \quad r \neq s_1$$

et

$$s_1 r^j = s_1 r^k \iff r^j = r^k \iff j = k$$

Il en suit que si $\ker(f)$ contient plus d'éléments, l'un de ceux énumérés ci-dessus qui n'y est pas déjà devrait y être ajouté, on aurait $|\mathcal{D}_m| < 2m$. Ceci étant une contradiction, on a que

$$\ker(f) = N(R)$$

et donc que

$$\mathcal{D}_m \simeq \langle r, s_1 \mid s_1^2, r^m, s_1 r s_1 r \rangle$$

□

À l'aide de méthodes similaires, on peut obtenir la présentation équivalente du groupe \mathcal{D}_m en tant que groupe de réflexions,

$$\mathcal{D}_m \simeq \langle s_1, s_2 \mid s_2^2, s_1^2, (s_1 s_2)^m \rangle$$

N.B. On appelle un mot m une *écriture* de $w \in G$ si $f(m) = w$ (ou f représente la projection canonique). Dans le reste du document, pour un élément $w \in W$ engendré par l'ensemble S , lorsque nous parlerons de m une écriture de w , il sera sous entendu que $m \in F_S$ et que $f(m) = w$ si cela ne porte pas à confusion.

3 GROUPES DE COXETER

Toutes les références, définitions et théorèmes ont été tirés du livre *Reflection groups and Coxeter groups* de James Humphreys [4].

Définition 1. Un *groupe de Coxeter* de rang n est un groupe engendré par un ensemble S de cardinalité n dont chaque paire d'éléments

$$s_i, s_j \in S$$

n'est sujette qu'aux relations de la forme $(s_i s_j)^{m(i,j)}$ où $m(i,i) = 1, m(i,j) = m(j,i) \geq 2$.

Autrement dit, tous les générateurs sont des involutions et si $(s_i s_j)^n = e$, alors $(s_j s_i)^n = e$. Ceci n'étant que la définition formelle, nous pouvons constater que la façon naturelle de penser à un groupe de Coxeter est de le considérer comme un groupe engendré par des réflexions agissant sur un espace vectoriel euclidien.

Exemple 3. Le groupe diédral

$$\begin{aligned} \mathcal{D}_{2m} &\simeq \langle r, s_1 \mid s_1^2, r^m, s_1 r s_1 r \rangle \\ &\simeq \langle s_1, s_2 \mid s_1^2 = s_2^2 = (s_1 s_2)^m \rangle \end{aligned}$$

Exemple 4. Le groupe symétrique \mathcal{S}_n a pour présentation,

$$\langle s_1, \dots, s_{n-1} \mid s_i^2, (s_i s_j)^2, (s_i s_{i+1})^3, |i - j| > 1 \rangle$$

Les groupes de Coxeter finis sont d'un intérêt particulier. En effet, on peut démontrer que les groupes de réflexion finis sont exactement les groupes de Coxeter finis.

3.1 Groupes de réflexions finis

Dans cette section, nous nous plaçons dans un espace vectoriel euclidien V de dimension n . Notre étude portera sur un groupe de réflexions W agissant sur l'espace V . Soit W un groupe fini engendré par un ensemble de réflexions. Chaque réflexion dans W définit un hyperplan de réflexion de dimension $n - 1$. Chacun de ces hyperplans définit à son tour un espace orthogonal de dimension 1, disons $\mathbb{R}\alpha$ avec $\|\alpha\| = 1$. Posons

$$\Phi = \bigcup_{s_\alpha \in W} \{\alpha, -\alpha\} \quad (s_\alpha \text{ une réflexion de } W)$$

Nous appelons Φ un *système de racines* pour W . Nous pouvons présumer sans perte de généralité que V est l'espace engendré par Φ . Il peut être démontré que l'ensemble Φ est caractérisé comme l'ensemble respectant les axiomes suivants:

1. $\forall \alpha \in \Phi, \Phi \cap \mathbb{R}\alpha = \{\alpha, -\alpha\}$
2. $\forall \alpha \in \Phi, s_\alpha \Phi = \Phi$

Définissons un *système de racines simples* $\Delta \subseteq \Phi$ comme étant une base de l'espace vectoriel V tel que tous vecteur dans Φ soit une combinaison linéaire des éléments de Δ à coefficients tous de même signe. De plus, nous imposons que l'ensemble $\{s_\alpha \mid \alpha \in \Delta\}$ engendre le groupe W . Il peut être montré que les systèmes simples existent.

Étant donné un vecteur $v = (x_1, \dots, x_k) \in V$ dans une certaine base de l'espace V , on dit de lui qu'il est positif si tous les $x_i \geq 0$. Puisque, dans la base Δ , chaque vecteur dans Φ s'écrit comme une combinaison linéaire à coefficients tous positifs ou tous négatifs, nous pouvons définir une partition de $\Phi = \Pi \cup -\Pi$ telle que tous les éléments de Π soient positifs. De plus, si $\alpha \in \Pi$, alors $-\alpha \notin \Pi$. Donc,

Π contient exactement la moitié des éléments de Φ , c'est-à-dire un par réflexion. Nous appelons Π un *système de racines positives* et avons donc une bijection entre l'ensemble des réflexions dans W et Π .

Évidemment, pour chaque choix de système simple de Φ , il existe un système positif différent. Il est étrange de constater que pour n'importe quelle paire de systèmes positifs $\Pi, \Pi' \subseteq \Phi$, il existe toujours un unique élément $w \in W$ tel que $w\Pi = \Pi'$. On dit donc que l'action de W sur les systèmes de racines positives est simplement transitive. Les systèmes simples héritent évidemment de cette propriété.

Exemple 5. La figure représente le groupe diédral \mathcal{D}_4 des symétries du carré. Les lignes pointillées représentent les hyperplans de réflexion (de simples droites dans \mathbb{R}^2). Dans cet exemple, l'ensemble des racines est

$$\Phi = \{\alpha_1, -\alpha_1, \alpha_2, -\alpha_2, \alpha_3, -\alpha_3, \alpha_4, -\alpha_4\}$$

les racines α_i et $-\alpha_i$ étant associés à l'hyperplan H_i et à la réflexion s_i . Un ensemble des racines simples serait

$$\Delta = \{\alpha_1, \alpha_2\}$$

Et le système de racines positives lui étant associé est

$$\Pi = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$$




Figure 1: Groupe diédral \mathcal{D}_4

Notation. Étant donné un groupe de Coxeter W et Δ le système simple lui étant associé, nous noterons $S := \{s_i \mid \alpha_i \in \Delta\}$ l'ensemble de ces générateurs de telle sorte que $W = \langle S \rangle$.

Notation. Nous nous permettrons de plus, lorsque cela ne portera pas à confusion, de noter les mots de la forme $s_{i_1} \cdots s_{i_k}$ comme $i_1 \cdots i_k$. Par exemple, le mot $s_1 s_2 s_3 s_4$ pourra aussi s'écrire 1234. Rappelons que les s_i sont des involutions ce qui nous évite d'avoir à écrire les s_i^{-1} sous cette notation.

3.2 L'élément le plus long

Considérant $w \in W$, on définit la longueur de w par rapport à un système simple Δ , $\ell(w)$, comme le plus petit nombre naturel r tel qu'il existe une écriture $w = s_1 \dots s_r$, $s_i \in \Delta$, et on attribue par convention $\ell(e) = 0$. Cette notion nous permet ensuite de définir une écriture de w comme étant réduite si et seulement si elle est de la forme $w = s_1 \dots s_{\ell(w)}$, $s_i \in \Delta$. Il découle directement de cette définition plusieurs propriétés de base telles que $\ell(w) = 1 \iff w = s_i$ ($s_i \in \Delta$), et $\ell(w) = \ell(w^{-1})$.

En cherchant une interprétation géométrique de cette propriété, on découvre que la longueur $\ell(w)$ d'un élément w est égale au nombre de vecteurs $\lambda \in \Pi$ tels que $w(\lambda) \in -\Pi$. Cette valeur correspond au nombre de racines positives dont l'image par l'action de w est une racine négative. Si l'on considère une paire de systèmes simples Δ et $-\Delta$ ainsi que Π et $-\Pi$, les systèmes positifs leur correspondant, nous savons qu'il existe un unique élément dans W noté w_0 tel que $w_0\Pi = -\Pi$. De plus, $\ell(w_0) = |\Pi|$, qui est un invariant pour le groupe W , et par la caractérisation géométrique qui vient d'être faite, c'est aussi la longueur maximale pour un élément de W . Nous voyons donc qu'il existe un unique élément de longueur maximale dans un groupe de Coxeter donné ayant évidemment plusieurs écritures différentes.

Le groupe \mathcal{D}_4 est un excellent exemple de la caractérisation géométrique de l'élément le plus long. Rappelons-nous de l'élément le plus long de ce groupe est $s_1s_2s_1s_2$ qui peut aussi s'écrire $s_2s_1s_2s_1$. Ce mot est en effet de longueur 4 ce qui correspond exactement au nombre de racines positives. Nous voyons plus précisément qu'en faisant agir cet élément sur l'ensemble des racines, toutes les racines positives sont envoyées sur des racines négatives.

Voici une caractérisation plus facile à manipuler de la longueur :

Définition 2. On appelle w_0 l'unique élément de W tel que $\forall s_i \in S, \ell(w_0) > \ell(w_0s_i)$.

Une propriété particulière de w_0 qui nous sera d'une grande utilité dans ce qui suit est la suivante.

Proposition 1. Soit W un groupe de Coxeter fini. L'élément le plus long $w_0 \in W$ a les propriétés suivantes:

1. $w_0 = w_0^{-1}$ (w_0 est une involution)
2. Tout mot réduit est préfixe d'une écriture réduite de w_0 .

3.3 Sous-groupes paraboliques

L'étude d'un groupe $W = \langle S \rangle = \langle s_\alpha \mid \alpha \in \Delta \rangle$ (ainsi que des Φ et Δ fixés qui lui sont associés) étant souvent facilitée par la connaissance de ses sous-groupes, on s'intéresse aux groupes engendrés par des sous ensembles de ses générateurs. On appelle donc *sous-groupe parabolique* et on note W_I le sous-groupe de W engendré par l'ensemble

de générateurs $I \subseteq S$. On note aussi Δ_I le système simple correspondant à W_I . Nous pouvons tout de suite voir les cas limites $W_\emptyset = e$ et $W_S = W$. Ces sous-groupes nous intéressent particulièrement puisque c'est à partir de ceux-ci que l'on obtiens tous les sous-groupes de W qui sont des groupes de Coxeter.

Bien entendu, pour un Δ fixé, la longueur d'un w dans W_I est la même que sa longueur dans W .

4 c-FACTORISATION

Nous introduisons ici la notion plus abstraite de c -factorisation. Étant donné un groupe de Coxeter W et son ensemble de générateurs S , afin de mieux caractériser les éléments $w \in W$, nous associons à chacun d'eux une écriture unique.

4.1 Éléments, mots et sous-mots de Coxeter

Définition 3. Soit W un groupe de Coxeter, on dit que c est un élément de Coxeter du groupe W si c est le produit de tous les générateurs (l'ensemble S) une et une seule fois.

Il peut bien entendu exister plusieurs éléments de Coxeter pour un même groupe W ainsi que plusieurs écritures pour chaque élément de Coxeter. Cette notion est d'ailleurs en correspondance avec les mots de Coxeter:

Définition 4. Un mot de Coxeter est une écriture particulière d'un élément de Coxeter. Deux mots de Coxeter différents m et m' peuvent correspondre au même élément de Coxeter par la projection f (définie dans la section 2.2), $f(m) = f(m') = c \in W$

Il existe plusieurs mots de Coxeter pour un groupe donné:

Exemple 6. Les mots de Coxeter $c_1 = 2134$, $c_2 = 2314$ et $c_3 = 2341$ pour le groupe \mathcal{S}_5 (le groupe symétrique avec la présentation usuelle) correspondent au même élément de Coxeter. Donc, c_1 est le même élément de Coxeter que c_2 et c_3 bien qu'ils aient des écritures différentes. D'autre part, les mots de Coxeter $c_1 = c_2 = c_3$ et $c_4 = 1234$ correspondent à des éléments de Coxeter différents.

Exemple 7. $c = s_1s_2$ et $c = s_2s_1$ sont deux mots de Coxeter différents du groupe \mathcal{D}_4 .

Exemple 8. $c_1 = 2134$ et $c_2 = 2314$ sont deux mots de Coxeter différents du groupe \mathcal{S}_5 .

Définition 5. Soit un alphabet S , et m un mot de S , on appelle m' un sous-mot de m si m' contient les mêmes lettres que m et qu'elles sont dans le même ordre que dans m . Cette notion est similaire à celle de sous-suite d'une suite.

Exemple 9. Soit $S = \{a, b, c\}$, $m = abc$, alors a, b, c, ac, ab et abc sont des sous-mots de m .

Cette notion est centrale à notre étude des mots de Coxeter. Étant donné c un mot de Coxeter du groupe W et I_j un sous ensemble des lettres de c , on note c_{I_j} le sous mot de c contenant les lettres de l'ensemble I_j . Autrement dit, on peut considérer c comme un suite de lettres, alors c_{I_j} est une sous-suite de c , si et seulement si c_{I_j} est un sous-mot de c .

4.2 C -factorisation

Soit c un mot de Coxeter du groupe W et $w \in W$. On dit alors que la c -factorisation de w est le sous mot réduit m de $c^\infty = ccccccc...$ le plus petit par rapport à l'ordre lexicographique tel que $f(m) = w$ dans W . Ici, l'ordre lexicographique correspond à l'ordre alphabétique sur l'alphabet ordonné S induit par c^∞ . Ce sous mot réduit m est unique pour une classe de mots de Coxeter équivalents (cette classe correspond à l'élément de Coxeter $f(c) \in W$ de $c \in F_S$).

De plus, on dit que m est une écriture c -factorisé de w s'il existe $c_{I_1}, c_{I_2}, \dots, c_{I_n}$ sous-mots de c tel que $m = c_{I_1}c_{I_2}\dots c_{I_n}$, est une écriture réduite minimale de w par rapport à l'ordre lexicographique de c^∞ . (Pour tout $m' = c'_{I_1}c'_{I_2}\dots c'_{I_n}$, $m \leq m'$ par rapport à l'ordre lexicographique dans c^∞ .)

Définition 6. Notons la relation "est un sous-mot" par \subseteq . On dit que $w = c_{I_1}c_{I_2}\dots c_{I_{p-1}}c_{I_p}$ est c -trié si $\emptyset \neq c_{I_p} \subseteq c_{I_{p-1}} \subseteq \dots \subseteq c_{I_1}$ ou $w = e$. On dit que w est toujours c -trié si w est c -trié pour tous les c , mots de Coxeter.

Exemple 10. Soit $c = 1234$, $w = 1234123121$ est c -trié où $c_{I_1} = 1234$, $c_{I_2} = 123$, $c_{I_3} = 12$ et $c_{I_4} = 1$. On voit bien que

$$S \supseteq c_{I_1} \supseteq c_{I_2} \supseteq c_{I_3} \supseteq c_{I_4}$$

4.3 w_0 toujours c -trié

La particularité de w_0 , que nous tentons d'éclaircir, est étroitement liée à sa c -factorisation. Le théorème de Nathan Reading qui nous intéresse ici dit que:

Théorème 1. Étant donné W un groupe de Coxeter, $w_0 \in W$ est toujours c -trié.

Bien qu'il existe une preuve de cet énoncé, elle fait appel à des notions fort peu élémentaires qui empêchent de comprendre totalement sa signification. Nous cherchons donc à établir une preuve plus élémentaire qui pourrait permettre de mieux comprendre cette propriété et ses liens avec les algèbres amassées.

5 RÉSULTATS

Grâce à des modélisation informatiques des logiciels Chevie [1] et Cambrian [2] du système d'algèbre computationnelle GAP [7], permettant de c -factoriser des mots

dans plusieurs groupes de Coxeters on obtenu des tables des éléments toujours c -triés de certains groupes et leurs diverses écritures réduites. On décèle ainsi plusieurs phénomènes récurrents dont le fait que les éléments toujours c -triés étaient tous des involutions. En approfondissant nos recherches, afin de mieux comprendre ce phénomène, nous avons mis au point plusieurs conjectures dont certaines restent encore à démontrer. Par contre, nous avons quand même réussi à catégoriser les mots toujours c -triés:

Théorème 2. Soit W un groupe de Coxeter fini, étant donné $I \subseteq S$ le sous ensemble de générateurs présents dans l'écriture de $w \in W$ toujours c -trié, w est l'élément le plus long de W_I .

Proof. Soit $w \in W$ un élément toujours c -trié et $I \subseteq S$ l'ensemble de générateurs présents dans l'écriture de w . Considérons dans le sous-groupe parabolique W_I un élément de Coxeter c pour lequel il existe une écriture où $s \in I$ est initial. Puisque w est c -trié par hypothèse, $w = sw'$ pour un certain $w' \in W_I$. Ainsi,

$$w^{-1} = w'^{-1}s$$

et

$$w^{-1}s = w'^{-1}s^2 = w'^{-1}$$

et

$$\forall s \in I, \ell(w^{-1}) > \ell(w^{-1}s)$$

Donc, w est l'élément le plus long du groupe W_I . \square

Corollaire 1. Les $w \in W$ toujours c -triés sont des involutions puisque w_0 est une involution.

Ce théorème, bien que simple, permet de généraliser le résultat de [6, p. 14] comme suit:

Théorème 3. Soit $I \subseteq S$ le sous ensemble de générateurs présents dans l'écriture de $w \in W$. Alors w est toujours c -trié si et seulement si w est l'élément le plus long de W_I .

6 CONCLUSION ET CONJECTURES

Notre étude des éléments toujours c -triés a permis de développer un nouvel outil qui, bien que fort simple, donne une nouvelle perspective sur cette propriété particulière de $w_0 \in W$. Le Théorème 2 qui dit que tout $w \in W$ toujours c -trié est nécessairement l'élément le plus long d'un $W_I \leq W$, nous permet d'envisager une nouvelle preuve que w_0 est toujours c -trié. Il s'agit maintenant de montrer que pour tout ensemble de générateurs S , il existe un $w \in \langle S \rangle$ toujours c -trié tel que tout $s \in S$ soit présent dans son écriture pour conclure que w_0 est toujours c -trié. Une telle approche pourrait offrir une preuve beaucoup plus simple et compréhensible du Théorème 1 et il s'agirait d'une voie prometteuse de recherche à suivre selon les auteurs. En effectuant plusieurs modélisations informatiques des groupes de Coxeter effectuées avec GAP-Chevie-Cambrian [1, 7, 2],

nous avons reformulé plusieurs conjectures énoncées auparavant. L'une d'elles particulièrement marquante serait que dans les groupes de Coxeter de rang impair, il existe un élément de Coxeter c et un nombre naturel m tel que $w_0 = c^m$. Il serait très intéressant de vérifier la véracité de cet énoncé et de ses applications. De plus, nous sommes toujours à la recherche d'un algorithme récursif permettant de construire w_0 et/ou sa c -factorisation étant donné c .

REMERCIEMENTS

Nous voudrions remercier le Laboratoire de Combinatoire et d'Informatique Mathématique (LaCIM)² de l'UQÀM et plus particulièrement Christophe Hohlweg³, notre directeur de recherche.

REFERENCES

[1] M. Geck, G. Hiss, F. Lübeck, G. Malle, and G. Pfeiffer. CHEVIE – A system for computing and processing generic character tables for finite groups of Lie type,

Weyl groups and Hecke algebras. *Appl. Algebra Engrg. Comm. Comput.*, 7:175–210, 1996.

- [2] C. Hohlweg. *Cambrian*. http://hohlweg.math.uqam.ca/?page_id=25.
- [3] Christophe Hohlweg, Carsten Lange, and Hugh Thomas. Permutahedra and generalized associahedra. *0709.4241*, September 2007.
- [4] James E Humphreys. *Reflection Groups and Coxeter Groups*. Number 29 in Cambridge studies in advanced mathematics. Cambridge University Press, Cambridge, Angleterre, 1990.
- [5] Nathan Reading. Cambrian lattices. *math/0402086*, February 2004. *Adv. Math.* 205 (2006), no. 2, 313-353.
- [6] Nathan Reading. Clusters, coxeter-sortable elements and noncrossing partitions. *math/0507186*, July 2005.
- [7] Martin Schönert et al. *GAP – Groups, Algorithms, and Programming – version 3 release 4 patchlevel 4*. Lehrstuhl D für Mathematik, Rheinisch Westfälische Technische Hochschule, Aachen, Germany, 1997.

JOKES

Que répond une mathématicienne venant d'accoucher à qui l'on demande "Avez-vous eu un garçon ou une fille ?"
Oui.

Qu'est-ce qu'un homme complexe dit à une femme réelle ?
Réponse: "viens danser !" (dans \mathbb{C})

"Mathematicians are like Frenchmen: whatever you say to them, they translate it into their own language, and forthwith it means something entirely different." – Goethe

Dans l'enfer topologique, la bière est contenue dans des bouteilles de Klein.

La limite quand n tend vers l'infini de $\frac{\sin x}{n}$ est 6.
Preuve : il suffit de simplifier par n au numérateur et au dénominateur.

Un mathématicien fou monte dans un bus et se met à menacer tout le monde : " Je vais vous intégrer! Je vais vous dériver!".
Tout le monde est effrayé et se sauve, sauf une jeune dame qui reste tranquille.
Le mathématicien fou arrive vers elle et dit : " Tu n'as pas peur ? Je vais t'intégrer! Je vais te dériver!".
La jeune dame répond : "Non, je n'ai pas peur, je suis e^x ."

Question : Combien faut-il de Bourbakistes pour changer une ampoule ?

R : Changer une ampoule est un cas particulier d'un problème plus général concernant l'entretien et la réparation d'un système électrique.

Pour déterminer un minorant et un majorant du nombre de personnes nécessaires, nous devons vérifier si les conditions du lemme 2.1 (disponibilité du personnel) et ceux du corollaire 2.3.55 (motivation du personnel) sont vérifiées.

Si et seulement si ces conditions sont réunies, on obtient le résultat en appliquant le théorème de la section 3.11.23.

Le majorant obtenu est, bien sûr, à prendre en compte dans un espace mesuré, muni de la topologie *-faible.

²<http://lacim.uqam.ca/>

³<http://hohlweg.math.uqam.ca/>

INTERVIEW WITH PROF. JOHANNA NESLEHOVA

Ana Best

DE: What is your name, and how long have you been at McGill?

My name is Johanna Nešlehová, and I have been here since July 2009.

DE: Tell me a bit about your background, both personal and academic.

I was born in Prague in 1977, and I went to school there. I started studying mathematics locally at Charles University and continued in Hamburg, Germany. I completed a Masters degree there and then I did my Ph.D. in statistics in Oldenburg, which is a little town west of Hamburg, about an hour and a half by car. After my Ph.D., I was offered a postdoctoral position in statistics and risk management at ETH Zurich. It was quite pleasant and I learned a lot. After two years, I became a lecturer there. I ended up staying three more years before coming to McGill.



Picture 1: Prof. Johanna Nešlehová

DE: What are your favourite things about McGill and Montreal?

I don't know Montreal all that well yet, but I find it really lively. The city seems to be thriving; there are lots of things happening, especially at the cultural level. I've been to a couple of great concerts already; there are many nice restaurants and I like the fact that it is multicultural too.

McGill is a great school. I have pleasant colleagues, and teaching here is enjoyable because the students are clever. The class that I had in the fall term was not too big, just seventeen students, but they were all quite interested so it was fun to teach them probability theory. The McGill community seems interesting as a whole, but I'm still discovering it.

DE: How did you first become interested in statistics?

Initially, I was exploring all sorts of things: analysis, combinatorial mathematics, graph theory, probability, and so on. But then when I came to Hamburg their statistics program caught my attention. When I started my studies there, I followed a class in nonparametric statistics that I really liked, but I had no background in statistics except for an introductory course. At the beginning, I didn't understand all that much, but there was something fascinating about the subject, so I worked hard and got to appreciate it more and more. I began to see how it connects many areas and things I had learned before. I also discovered that you could apply statistics in all sorts of ways. And I liked that combination, so in the end I became a statistician.

DE: What's your favorite probability distribution?

I think it would be the class of extreme-value distributions. It's quite a rich family; they arise as limits of properly scaled maxima. Just like the normal distribution, which appears in the central limit theorem as the limit of sums. Extreme-value distributions also exist in higher dimensions, where they have quite a nice geometric interpretation.

DE: What are you currently researching?

At the moment, I do research primarily in dependence modeling. I try to understand the relationship between variables, processes or other more complex phenomena. An approach that I am currently using involves copulas, which may be viewed as functional descriptions of dependence. It's not just one number summary but a whole function that can tell you a great deal about how variables are related. I am exploring these copula functions, and I try to estimate them statistically. What I like particularly are goodness-of-fit tests. In practice, specific parametric models are often used, and I am developing tests to check whether they fit the data well.

DE: What are your favorite and least favorite parts of research?

In my research, I often follow an idea intuitively, with a vague goal in mind. But then there comes a turning point:

things start to work out, you do a few calculations and suddenly you are beginning to see through your intuition. This is my favorite part; its quite fascinating to sense some structure behind and understand the message that you really want to tell. And I like to discuss with my colleagues throughout the process, too.

Once all the calculations are done and I know what the main message is, I like to step back and think of it as if it were a movie, with a storyline and characters. I enjoy writing articles as if they were little screenplays. What I like least are the final steps where everything is written but you go endlessly over it fishing for errors. And the review process of articles is sometimes really annoying.

DE: When you were an undergraduate, what were your goals? Did you see yourself becoming a university professor?

Actually, I was thinking about becoming an academic early on because I was familiar with that kind of life. My parents are not scientists, but my father was a painter and also a university professor. My mother is a researcher in art history at the Academy of Sciences in Prague, so I knew the academic environment quite well and liked it. I felt that if I became a university professor, I would have a lot of freedom. I could do really something creative while being helpful to the community through teaching and consulting. Universities are also full of young people; this makes for a lively environment.

Although I did not follow my parents footsteps, in my

heart I am a bit of an artist nonetheless. When I have a good stretch and lots of ideas, then I can do research day and night. On the other hand, there are times when I prefer to concentrate on other things, such as teaching, designing talks or developing course material. Above all, I find it important to have enough freedom to design my own schedule. An academic position is one of the rare jobs where you can really do this.

DE: What advice do you have for undergraduates looking to go into statistics?

I think it's a great choice because statistics is really broad. If you enjoy theoretical challenges, you have plenty of room for them. Statistics can be very abstract and complex, quite close to analysis, topology, algebra, geometry, etc. At the same time, statistics is rooted in applications, and many questions that we are trying to answer can help to solve interesting practical problems, for example in medicine.

Another important aspect is that if you try out statistical models on real data, you often discover that they do not fit reality perfectly. This helps to maintain a healthy balance and prevents you from drifting off into some abstract spaces and getting lost there.

From a practical perspective, statistics plays an increasing role in many fields and this generates quite a few interesting job opportunities outside academia. I have the feeling that sciences are getting more and more specialized and complex. As a result, there is a stronger need for statistics as it helps to validate things, and to see the bigger picture!

JOKES

A math professor is one who talks in someone else's sleep. □

When a statistician passes the airport security check, they discover a bomb in his bag. He explains.

"Statistics shows that the probability of a bomb being on an airplane is 1/1000. However, the chance that there are two bombs at one plane is 1/1000000. So, I am much safer..." □

"This is a one line proof... if we start sufficiently far to the left." □

Cantor did it diagonally.

Fermat tried to do it in the margin, but couldn't fit it in.

Galois did it the night before.

Möbius always does it on the same side.

Markov does it in chains.

Newton did it standing on the shoulders of giants.

Turing did it but couldn't decide if he'd finished. □

Q: Why did the chicken cross the Möbius strip?

A: To get to the other ... er, um ... □

"The problems for the exam will be similar to the discussed in the class. Of course, the numbers will be different. But not all of them. Pi will still be 3.14159..." □

SPECTRAL THEOREM FOR BOUNDED SELF-ADJOINT OPERATORS

Dana Mendelson and Alexandre Tomberg

We introduce the concepts of functional calculus and spectral measure for bounded linear operators on a Hilbert space. We state the spectral theorem for bounded self-adjoint operators in the cyclic case. We also compute the spectrum and the spectral measure in two concrete examples: a self-adjoint linear operator on a finite dimensional Hilbert space, and the discrete Laplacian operator on $\ell^2(\mathbb{Z})$.

1 INTRODUCTION

Diagonalization is one of the most important topics one learns in an elementary linear algebra course. Unfortunately, it only works on finite dimensional vector spaces, where linear operators can be represented by finite matrices.

Later, one encounters infinite dimensional vector spaces (spaces of sequences, for example), where linear operators can be thought of as "infinite matrices"¹. Extending the idea of diagonalization to these operators requires some new machinery. We present it below for the (relatively simple) case of bounded self-adjoint operators.

It is important to note that this generalization is not merely a heuristic desire: infinite dimensions are inescapable. Indeed, mathematical physics is necessarily done in an infinite dimensional setting. Moreover, quantum theory requires the careful study of functions of operators on these spaces – the functional calculus.

This may seem awfully abstract at first, but an example of a function of operators is known to anyone familiar with systems of linear ODEs. Given a system of ordinary linear differential equation of the form

$$x'(t) = Ax(t)$$

where A is a constant matrix, the solution is given by

$$x(t) = \exp(tA)x(0).$$

This is an instance of the matrix exponential, an operation that is well defined for finite dimensions.

Yet, quantum mechanics demands that we are able to define objects like this for any operator. In particular, the time evolution of a quantum mechanical state, ρ is expressed by conjugating the state by $\exp(itH)$ where H is the Hamiltonian of the system. This motivates the development of a functional calculus which allows us to define operator-valued equivalents of real functions.

But enough motivation, let us get on with the theory!

2 OPERATORS & SPECTRUM

2.1 Self-adjoint operators

Let \mathcal{H} be a Hilbert space and $A \in \mathcal{B}(\mathcal{H})$, the set of bounded linear operators on \mathcal{H} . In particular, in this ex-

¹Matrices with lots of dots!

position, we will focus on self-adjoint operators. In finite dimensions, an operator A is called self-adjoint if, as a matrix, $A = A^*$, where A^* denotes the conjugate transpose of A , i.e. $A^* = \overline{A}^T$.

Of course, in infinite dimensional space, this definition does not apply directly. We first need the notion of an adjoint operator in a Hilbert space. We begin by stating a result that we will use several times in this exposition.

Let $T \in \mathcal{B}(\mathcal{H})$, for $y \in \mathcal{H}$, the map

$$x \mapsto \langle y | Tx \rangle$$

defines a bounded linear operator. Riesz's representation theorem for Hilbert spaces then tells us that $\exists! z \in \mathcal{H}$, such that

$$\phi(x) = \langle y | Tx \rangle = \langle z | x \rangle$$

We can now write $T^*(y) = z$ and *define* the adjoint T^* this way.

Definition. An operator $A \in \mathcal{B}(\mathcal{H})$ is said to be self-adjoint if

$$\langle Ax | y \rangle = \langle x | Ay \rangle$$

for all $x, y \in \mathcal{H}$, that is if $A = A^*$ with respect to our definition of the adjoint above.

Definition. λ is an eigenvalue of A if there exists $v \neq 0$, $v \in \mathcal{H}$ such that $Av = \lambda v$.

Equivalently, λ is an eigenvalue if and only if $(A - \lambda I)$ is not injective.

Several important properties of self-adjoint operators follow directly from our definition. First, the eigenvalues of a self-adjoint operator, A , are real. Indeed, let

$$Av = \lambda v$$

then,

$$\lambda \langle v | v \rangle = \langle Av | v \rangle = \langle v | Av \rangle = \overline{\lambda} \langle v | v \rangle$$

so $\lambda = \overline{\lambda}$. Moreover, if

$$Av = \lambda v, \quad Au = \mu u$$

then

$$\lambda \langle v | u \rangle = \langle Av | u \rangle = \langle v | Au \rangle = \overline{\mu} \langle v | u \rangle$$

Since $\lambda \neq \overline{\mu}$, we conclude that $\langle v | u \rangle = 0$. Which tells us that the eigenspaces of A corresponding to different eigenvalues are orthogonal.

These two simple facts are not only reassuring, but crucial for the study of quantum mechanical systems. In fact, for a quantum system, the Hamiltonian is a self-adjoint operator whose eigenvalues correspond to the energy levels of the bound states of the system. We can sleep well at night knowing that these energy levels are real values.

2.2 Spectrum

The spectrum of an infinite dimensional operator is an entirely different beast than just the sets of eigenvalues we are used to. To describe it, it is best to introduce some new terminology. We define this for a general operator T :

Definition. The resolvent set of T , $\rho(T)$ is the set of all complex numbers λ such that

$$R_\lambda(T) := (\lambda I - T)^{-1}$$

is a bijection with a bounded inverse. The spectrum of T , $\sigma(T)$ is then given by $\mathbb{C} \setminus \rho(T)$.

In general, the spectrum of a linear operator T is comprised of two disjoint components:

1. The set of eigenvalues is now called the *point spectrum*.
2. The remaining part is called the *continuous spectrum*.

Before we discuss some examples of continuous spectra, let us prove a simple result about $\sigma(T)$ that will be necessary later in the development of the functional calculus.

Lemma 1. The spectrum of a bounded linear operator is a closed and bounded subset of \mathbb{C} . In fact,

$$\sigma(T) \subseteq \{z \in \mathbb{C} : |z| \leq \|T\|\}$$

Proof. Recall that $\sigma(T) = \mathbb{C} \setminus \rho(T)$.

Closed Enough to show that $\rho(T)$ is open. Indeed, remark that by the convergence of the Neumann series, namely if $\|S\| < 1$ then $(I - S)$ is invertible and its inverse is given by

$$(I - S)^{-1} = \sum_{n=0}^{\infty} S^n.$$

Let $\lambda \in \rho(T)$. For any $\mu \in \mathbb{C}$,

$$\mu I - T = (\lambda I - T)^{-1} [(\mu - \lambda)(\lambda I - T)^{-1} - I]$$

exists if $|\mu - \lambda| \|(\lambda I - T)^{-1}\| < 1$.

Bounded Now, let $\lambda \in \mathbb{C}$ be such that $|\lambda| > \|T\|$. Then, $\exists \delta \in \mathbb{R}$,

$$|\lambda| > \delta > \|T\|$$

This means that $\forall x \in \mathcal{H}$,

$$\|Tx\| \leq \|T\| \|x\| < \|\delta x\| < \|\lambda x\|$$

And thus, $\forall x$,

$$0 < \|(\lambda I - T)^{-1}x\| < \|(\delta I - T)^{-1}x\| < \infty$$

so that $\lambda \in \rho(T)$. □

2.3 Examples of continuous spectra

The phenomenon of a purely continuous spectrum is uniquely found in infinite dimensional spaces, so for those who might never have ventured into these spaces before, this may seem a bit bizarre at first glance.

To offer a simple example, we consider the space $C([0, 1])$ of continuous functions defined on $[0, 1]$ and the operator A defined by

$$Ax(t) = tx(t).$$

Then $(A - \lambda I)x(t) = (t - \lambda)x(t)$ so

$$(A - \lambda I)^{-1}x(t) = \frac{1}{(t - \lambda)}x(t)$$

We cannot have that $tx(t) = \lambda x(t)$ so this operator has no eigenvalues. However, the spectrum is any value λ for which $t - \lambda$ vanished. Thus, the whole interval $[0, 1]$ is in the spectrum of A . Hence A has purely continuous spectrum.

Consider now a much more realistic example that will arise later in our treatment. We let $\mathcal{H} = l^2(\mathbb{Z})$, the Hilbert space of doubly infinite, square summable sequences and we let $A = \Delta$, the discrete Laplacian. If $x = (x_n)$,

$$(Ax)(n) = x_{n+1} + x_{n-1} - 2x_n.$$

Then A is self-adjoint and has no eigenvalues. We will later see that its continuous spectrum is the entire interval $[0, 4]$.

3 FUNCTIONAL CALCULUS

3.1 Operator-valued functions

In finite dimensional case, there is a natural way to write down the formula of a linear operator with solely the knowledge of its eigenvalues (i.e. spectrum) and eigenvectors. In fact, if $M \in M_n(\mathbb{C})$ with eigenvalues $\{\lambda_k\}$ and associated eigenvectors $\{v_k\}$, then

$$M = \sum_{k=1}^K \lambda_k P_k$$

where P_k is the orthogonal projection on v_k .

We can view this linear combination as an operator-valued function defined on the spectrum of M :

$$\sigma(M) \rightarrow M_n(\mathbb{C}), \quad \lambda_k \mapsto \lambda_k P_k$$

We can use this idea to define functions of operators. Indeed, if $f : \sigma(M) \rightarrow \mathbb{C}$ we can set

$$f(M) := \sum_{k=1}^K f(\lambda_k) P_k$$

Note that as the spectrum consists of finitely many points, this construction allows us to define $f(M)$ for *any* complex-valued f defined on the spectrum. For instance, in the case of the matrix exponential, mentioned in the introduction, we obtain

$$\exp(M) = \sum_{n=0}^{\infty} \frac{M^n}{n!} = \sum_{k=1}^K e^{\lambda_k} P_k.$$

We can think of this definition as a mapping associating an operator-valued equivalent to functions on $\sigma(M)$:

$$f(z) \mapsto f(M)$$

However, as we pointed out above, in infinite dimensional case, the spectrum need not be pure point. Hence, we need to extend this idea to a larger class of functions.

For this section, our goal is to extend the mapping ϕ above to *all* continuous functions defined on the spectrum of a bounded self-adjoint operator A . Before we begin, let us introduce the notion of an algebra-morphism.

Definition. An algebra-morphism is a map

$$\phi : X \rightarrow Y$$

preserving scalar multiplication, addition and multiplication in the spaces X and Y . In other words, $\forall x \in X, y \in Y$ and all scalars α ,

1. $\phi(\alpha x) = \alpha \phi(x)$
2. $\phi(x + y) = \phi(x) + \phi(y)$
3. $\phi(xy) = \phi(x)\phi(y)$

Note that these properties simply reflect our notions of pointwise addition and multiplication of functions. Indeed, we want the operator-valued equivalents defined by ϕ to obey these notions, and so, requiring ϕ to be an algebra-morphism is a natural constraint.

3.2 Continuous functional calculus

For this section, we let A be a bounded, self-adjoint operator. Let P be a polynomial, with

$$P(x) = \sum_{k=0}^n \alpha_k x^k$$

then we define

$$P(A) := \sum_{k=0}^n \alpha_k A^k$$

We thus have a map $\phi : \mathbb{C}[x] \rightarrow \mathcal{B}(\mathcal{H})$ with

$$P \mapsto P(A)$$

This ϕ is an algebra-morphism and satisfies $\phi(\bar{P}) = \phi(P)^*$. Moreover, if λ is an eigenvalue of A then $P(\lambda)$ is an eigenvalue of $P(A)$. This fact can be reformulated as the following

Lemma 2 (Spectral Mapping Theorem).

$$\sigma(P(A)) = \{P(\lambda) : \lambda \in \sigma(A)\}$$

Proof. Let $\lambda \in \sigma(A)$ and consider $Q(x) = P(x) - P(\lambda)$ then λ is a root of $Q(x)$ and so there is a polynomial $R \in \mathbb{C}[x]$ such that $Q(x) = (x - \lambda)R(x)$. Thus

$$P(A) - P(\lambda) = (A - \lambda I)R(A) = R(A)(A - \lambda I)$$

Since $(A - \lambda I)$ is not invertible for $\lambda \in \sigma(A)$, $P(A) - P(\lambda)$ is not invertible, so $P(\lambda) \in \sigma(P(A))$. Conversely, let $\mu \in \sigma(P(A))$ then, by factoring we obtain that

$$P(x) - \mu = \alpha(x - \lambda_1) \cdots (x - \lambda_n)$$

and

$$P(A) - \mu = \alpha(A - \lambda_1) \cdots (A - \lambda_n)$$

Since $\mu \in \sigma(P(A))$, $P(A) - \mu$ is not invertible and so there is some $i \in \{1, \dots, n\}$ such that $A - \lambda_i$ is not invertible. This $\lambda_i \in \sigma(A)$ but $P(\lambda_i) = \mu$ by the first part of this discussion. \square

Before developing a continuous functional calculus, we require one more simple technical claim.

Lemma 3.

$$\|P(A)\| = \sup\{|P(\lambda)| : \lambda \in \sigma(A)\}$$

Proof.

$$\begin{aligned} \|P(A)\|^2 &= \|P(A)^* P(A)\| = \|\bar{P}P(A)\| = \\ &= r(\bar{P}P(A)) \end{aligned}$$

Now, by the above lemma we have that

$$\begin{aligned} r(\bar{P}P(A)) &= \sup\{|\mu| : \mu \in \sigma(\bar{P}P(A))\} \\ &= \sup\{|\bar{P}P(\lambda)| : \lambda \in \sigma(A)\} \\ &= \sup\{|P^2(\lambda)| : \lambda \in \sigma(A)\} \end{aligned} \quad \square$$

This brings us to the main result of this section:

Theorem 4 (Continuous Functional Calculus). Let $C(\sigma(A))$ be the continuous functions defined on the spectrum of A . There exists a unique map $\phi : C(\sigma(A)) \rightarrow \mathcal{B}(\mathcal{H})$,

$$\phi(f) = f(A)$$

such that

1. ϕ is an algebra-morphism.
2. $f(A)^* = \overline{f(A)}$.

3. If $f(x) = x$ then $f(A) = A$.
4. $\|f(A)\| = \|f\|_\infty$.
5. $\sigma(f(A)) = \{f(\lambda) : \lambda \in \sigma(A)\}$ and if λ is an eigenvalue of A then $f(\lambda)$ is an eigenvalue of $f(A)$.

Proof. Due to 1 and 3, φ must coincide with our previously defined map on the polynomials. We only need to extend it uniquely to $C(\sigma(A))$, the space of continuous functions

$$\sigma(A) \rightarrow \mathbb{C}$$

Recall the Stone-Weierstrass theorem:

If X is compact, the set of polynomials over X is dense in $C(X)$.

By lemma 1, $\sigma(A)$ is closed and bounded in \mathbb{C} , and hence, compact by Heine-Borel theorem. Thus, the map φ is densely defined on $C(\sigma(A))$ and can be extended by continuity. The uniqueness of such an extension is guaranteed by the isometry from lemma 3.

Finally, taking limits in lemmas 2 and 3 proves the properties 5 and 4 respectively. \square

So far, this has been a natural extension, but what follows is a miracle.

4 SPECTRAL MEASURES & BOREL FUNCTIONAL CALCULUS

We let A be as before and, given any $\psi \in \mathcal{H}$, we define

$$L : C(\sigma(A)) \longrightarrow \mathbb{C}$$

$$f \longmapsto \langle \psi | f(A) \psi \rangle$$

L is a continuous linear functional with

$$|\langle \psi | f(A) \psi \rangle|^2 \leq \|f(A)\| \|\psi\|^2 \leq \|f\|_\infty \|\psi\|^2$$

Moreover, L is positive. Indeed, if $f \geq 0$, by theorem 4(5),

$$\sigma(f(A)) = f(\sigma(A)) \subseteq (0, \infty)$$

Also, by continuity of f , $\sigma(f(A))$ is compact. Recall the Riesz-Markov theorem for a locally compact Hausdorff space X :

For any positive linear functional Φ on $C_c(X)$, there exists a unique Borel measure μ on X such that $\forall f \in C_c(X)$,

$$\Phi(f) = \int_X f d\mu$$

As we have seen before, $\sigma(A)$ is compact, so every continuous function on it is compactly supported, i.e. $C_c(\sigma(A)) = C(\sigma(A))$.

Thus, by the Riesz-Markov theorem, there is a positive Borel measure μ_ψ such that for all $f \in C(\sigma(A))$,

$$L(f) = \langle \psi | f(A) \psi \rangle = \int f d\mu_\psi$$

Something interesting has just happened here. The right hand side makes sense even if f is not continuous, just measurable. So we can extend our definition of L to an arbitrary measurable g by setting

$$\langle \varphi | g(A) \psi \rangle := \int g d\mu_\psi$$

We can push our luck a bit further. Using the polarization identity, we define for all $\varphi, \psi \in \mathcal{H}$

$$\begin{aligned} \langle \varphi | g(A) \psi \rangle := & \frac{1}{4} [\langle \varphi + \psi | g(A) \varphi + \psi \rangle - \\ & - \langle \varphi - \psi | g(A) \varphi - \psi \rangle + \\ & + i \langle \varphi + i\psi | g(A) \varphi + i\psi \rangle - \\ & - i \langle \varphi - i\psi | g(A) \varphi - i\psi \rangle] \end{aligned}$$

Given g and ψ , $\varphi \mapsto \langle \varphi | g(A) \psi \rangle$ is linear and continuous. We can (yet again!) apply Riesz's representation theorem. It states that for such a bounded linear functional, there is a unique $h_{g,\psi} \in \mathcal{H}$ such that

$$\langle \varphi | g(A) \psi \rangle = \langle \varphi | h_{g,\psi} \rangle$$

for all $\varphi \in \mathcal{H}$.

Remark. Be careful to note that the left-hand side is merely notation for our extended definition of the inner product, while the right-hand side is the true inner product of our Hilbert space.

We should also note that the above is sheer magic. It is a good day when we can unite both forms of the Riesz representation theorem in one proof.

We now let $g(A) : \mathcal{H} \longrightarrow \mathcal{H}$ with $\psi \mapsto h_{g,\psi}$.

Theorem 5. All the properties of $g(A)$ that hold for a continuous g also hold for a measurable g .

The proof is left as exercise for the reader.

Special Case. If $g(x) = \chi_A(x)$, then $g(A)$ is an orthogonal projection.

Proof. By the algebra property, $g(A)^2 = g^2(A)$, and as $\chi_A^2 = \chi_A$, we get that $g(A)^2 = g(A)$. Hence, $g(A)$ is a projection.

$$g(A)^* = \overline{g(A)} = g(A)$$

Thus, $g(A)$ is an orthogonal projection. \square

Definition. $\psi \in \mathcal{H}$ is called a cyclic vector of A if

$$\{P(A)\psi : P \in \mathbb{C}[x]\}$$

is dense in \mathcal{H} .

Theorem 6 (Spectral theorem for self-adjoint operators). Let A be a self-adjoint operators on a Hilbert space \mathcal{H} and suppose that ψ is a cyclic vector of A . Then there is a measurable function $f: \sigma(A) \rightarrow \mathbb{R}$ and a unitary map

$$U: \mathcal{H} \rightarrow L^2(\sigma(A), \mu_\psi)$$

so that

$$(UAU^{-1}f)(\lambda) = \lambda f(\lambda).$$

Remark. In finite dimensions, a unitary matrix is a change of basis matrix from one orthonormal basis to another. Thus, the unitary operator U can be thought of as a change of basis operator. In fact, diagonalization of finite matrices produces a similar result:

$$D = PAP^{-1}$$

where P is the matrix changing standard basis to the eigenbasis of A .

Now, if A does not admit of a cyclic vector on the entire space, we are not entirely out of luck. The decomposition theorem for Hilbert spaces (see [2]) comes to our aid. It assures us that we may decompose our Hilbert space into orthogonal subspaces \mathcal{H}_n and our operator into corresponding components A_n , such that A_n has a cyclic vector on \mathcal{H}_n . The proof of this theorem as well as the proof of the general case of the spectral theorem will be omitted here.

5 EXAMPLES

5.1 Finite dimensional case

Let $\mathcal{H} = \mathbb{C}^n$, and, as in section 3, set

$$A = \sum_{j=1}^k \lambda_j P_j \quad f(A) = \sum_{j=1}^k f(\lambda_j) P_j.$$

This can always be done, as any self-adjoint operator on a finite dimensional vector space is necessarily diagonalizable. Fix $\psi \in \mathbb{C}^n$, then

$$\begin{aligned} \int f d\mu_\psi &= \langle \psi | f(A) \psi \rangle \\ &= \sum_{j=1}^k f(\lambda_j) \langle \psi | P_j \psi \rangle \\ &= \sum_{j=1}^k f(\lambda_j) \|P_j \psi\|^2. \end{aligned}$$

$$\implies \forall f, \int f d\mu_\psi = \sum_{j=1}^k f(\lambda_j) \|P_j \psi\|^2$$

Thus, using the second equation in theorem 6, we can conclude that

$$\mu_\psi = \sum_{j=1}^k \|P_j \psi\|^2 \delta(\lambda - \lambda_j).$$

²Many thanks to Prof. Jakobson.

That is, the spectral measure is a counting measure, where each λ_j is weighted according to the norm of the corresponding $P_j \psi$ vector.

→ When does A admit a cyclic vector?

By definition above, ψ is cyclic if

$$\begin{aligned} \{P(A)\psi : P \in \mathbb{C}[x]\} = \mathbb{C}^n &\iff \\ \iff \left\{ \sum_{j=1}^k \overbrace{P(\lambda_j)}^{\dim k} P_j \psi : P \in \mathbb{C}[x] \right\} = \mathbb{C}^n &\iff \dim n \end{aligned}$$

Hence, ψ is cyclic if $k = n$, i.e. there are n distinct (simple) eigenvalues.

ψ is cyclic \iff spectrum is simple.

5.2 Discrete Laplacian

Let Δ be the discrete Laplacian on $\ell^2(\mathbb{Z})$. For $\psi \in \ell^2(\mathbb{Z})$, Δ acts as

$$(\Delta\psi)(n) = \psi_{n+1} + \psi_{n-1} - 2\psi_n$$

Given $\psi \in \ell^2(\mathbb{Z})$ we can define $\hat{\psi} \in L^2([0, 2\pi))$ by

$$\hat{\psi}(\xi) = \sum_{n \in \mathbb{Z}} e^{in\xi} \psi_n$$

with

$$\psi_n = \frac{1}{2\pi} \int_0^{2\pi} e^{-in\xi} \hat{\psi} d\xi$$

then the map

$$\begin{aligned} U: \ell^2(\mathbb{Z}) &\rightarrow L^2([0, 2\pi), \frac{d\xi}{2\pi}) \\ \psi &\xrightarrow{U} \hat{\psi} \end{aligned}$$

is unitary. We wish to investigate the behaviour of Δ under this map, that is what does $U\Delta U^{-1}$ give us. Take $f(\xi) \in L^2([0, 2\pi))$ then $(\Delta U^{-1}f)(n) =$

$$\begin{aligned} &= \frac{1}{2\pi} \int_0^{2\pi} \left[e^{-i(n+1)\xi} + e^{-i(n-1)\xi} + e^{-in\xi} \right] f(\xi) d\xi \\ &= \frac{1}{2\pi} \int_0^{2\pi} e^{-in\xi} (e^{-i\xi} + e^{i\xi} - 2) f(\xi) d\xi. \end{aligned}$$

Hence $(U\Delta U^{-1})(\xi) = (2 \cos \xi - 2)f(\xi)$.

Remark (about analysis, the universe and life in general²). Note that the unitary map U defined above is the discrete Fourier transform on $\ell^2(\mathbb{Z})$. Thus, this provides us with an example of the Fourier transform mapping a complicated looking operator on a sequence space, to a multiplication operator on the familiar space $L^2([0, 2\pi), \frac{d\xi}{2\pi})$. Fairy tales really do come true!

Spectrum Now that we have the form of our operator in Fourier space, we are equipped to determine what its spectrum is. We look at the resolvent set

$$\begin{aligned}\rho(\Delta) &= \{\lambda : (\lambda - \Delta) \text{ is invertible}\} \\ &= \{\lambda : \lambda - 2 \cos \xi + 2 \neq 0 \forall \xi \in [0, 2\pi)\} \\ &= \mathbb{C} \setminus [-4, 0]\end{aligned}$$

Hence $\sigma(\Delta) = [-4, 0]$.

Spectral measure Lastly, for $f \in C([-4, 0])$, in Fourier space $f(\Delta)$ is just multiplication by $f(2 \cos \xi - 2)$. Fix $\psi \in l^2(\mathbb{Z})$ and let $\hat{\psi}$ be the corresponding function in $L^2([0, 2\pi))$. Then

$$\begin{aligned}\int f(\Delta) d\mu_\psi &= \langle \psi | f(\Delta) \psi \rangle_{\ell^2} \\ &= \langle U^{-1} U \psi | f(\Delta) U^{-1} U \psi \rangle_{\ell^2} \\ &= \left\langle U \psi \left| \left(U f(\Delta) U^{-1} \right) U \psi \right\rangle_{L^2} \\ &= \langle \hat{\psi} | f(2 \cos \xi - 2) \hat{\psi} \rangle_{L^2} \\ &= \frac{1}{2\pi} \int_0^{2\pi} \bar{\hat{\psi}} f(2 \cos \xi - 2) \hat{\psi} d\xi \\ &= \frac{1}{2\pi} \int_0^{2\pi} f(2 \cos \xi - 2) |\hat{\psi}(\xi)|^2 d\xi.\end{aligned}$$

Now splitting the region of integration into two equal parts and applying the change of variables

$$\lambda = 2 \cos \xi - 2$$

yields

$$\begin{aligned}\int f d\mu_\psi &= \int_{-4}^0 f(\lambda) \left[\left| \hat{\psi} \left(\arccos \left(\frac{\lambda+2}{2} \right) \right) \right|^2 + \right. \\ &\quad \left. + \left| \hat{\psi} \left(-\arccos \left(\frac{\lambda+2}{2} + 2\pi \right) \right) \right|^2 \right] \frac{d\lambda}{\sqrt{-\lambda^2 - 4\lambda}}\end{aligned}$$

Thus,

$$d\mu_\psi = [\dots] \frac{d\lambda}{\sqrt{-\lambda^2 - 4\lambda}}$$

6 CONCLUSION

We conclude this article with an interesting result. We recall Lebesgue's decomposition theorem that states that any measure μ on \mathbb{R} has a unique decomposition into

$$\mu = \mu_{pp} + \mu_{ac} + \mu_{sing},$$

the pure-point, absolutely continuous and singularly continuous parts. Moreover, these three measures are mutually singular.

Given a self-adjoint operator, $A \in \mathcal{B}(\mathcal{H})$, we define

$$\mathcal{H}_{pp} = \{\psi | \mu_\psi \text{ is pure point}\}$$

and similarly for \mathcal{H}_{ac} and \mathcal{H}_{sing} . We then have that

$$\mathcal{H} = \mathcal{H}_{pp} \oplus \mathcal{H}_{ac} \oplus \mathcal{H}_{sing}$$

where each subspace is invariant under A . We can now define $\sigma_{pp}(A)$ to be the spectrum of A restricted to \mathcal{H}_{pp} and we further have that

$$\sigma(A) = \sigma_{pp}(A) \cup \sigma_{ac}(A) \cup \sigma_{sing}(A)$$

where this union might not be disjoint. In quantum mechanics, in particular, self-adjoint operators represent physical observables of a given system and their spectra correspond to the outcomes of measurements. Roughly speaking, the absolutely continuous spectrum corresponds to free states while the pure point corresponds to bound states.

However, the observables are not necessarily bounded (take the momentum and position operators, say). Hence, in order to fully appreciate and apply spectral theory in a quantum mechanical setting, one must turn to spectral theory for unbounded operators.

REFERENCES

- [1] Laurent Bruneau. Spectral theorem for bounded self-adjoint operators, 2009. Lecture notes from a mini-course given at Cergy Pontoise University.
- [2] Vojkan Jaksic. Topics in spectral theory. In *Open Quantum Systems I. The Hamiltonian Approach.*, volume 1880 of *Lecture Notes in Mathematics*, pages 235–312. Springer, 2006.
- [3] Andrej Nikolaevich Kolmogorov and S V Fomin. *Introductory real analysis*. Prentice-Hall, Englewood Cliffs, NJ, 1970.
- [4] Michael Reed and Barry Simon. *Methods of modern mathematical physics*. Academic Press, San Diego, CA, 1980.
- [5] Walter Rudin. *Real and complex analysis: mathematics series; 3rd ed.*, volume I. (Functional Analysis) of *Mathematics Series*. McGraw-Hill, New York, NY, 1987.

THE NEVER-ENDING HARMONIC SERIES

Cyndie Cottrell

Mathematicians were surprised when they first discovered the divergence of the harmonic series. We discuss why, and present elementary proofs of this fact.

1 WHAT IS IT?

By definition, $\sum_{n=1}^{\infty} \frac{1}{n}$ is the harmonic series. It is the sum of the reciprocals of every natural number. The harmonic sequence converges, as $\lim_{n \rightarrow \infty} \frac{1}{n} = 0$, and it is well-established that the alternating harmonic series converges $\sum_{n=1}^{\infty} (-1)^n \frac{1}{n}$. On the other hand, at the beginning of any calculus course, students are taught that the harmonic series diverges. One may toy with this idea and casually accept that $\frac{1}{n}$ does not get “small enough fast enough” to converge, i.e., $\sum_{n=1}^{\infty} \frac{1}{n}$ loses the race to convergence. Yet, after having accepted that a series such as $\sum_{n=1}^{\infty} \frac{1}{n^2}$ and the geometric series ($\sum_{n=1}^{\infty} ar^{n-1}$ where a is some constant and $|r| < 1$) converge, this fact becomes surprising. Another interesting aspect of this series is that when one revolves the graph of $y = \frac{1}{x}$ around the positive half of the x-axis, one would expect that the link between this function and the harmonic series would cause this volume to be infinite. Instead, one obtains a shape called Gabriel’s horn with the following volume [2]:

$$\begin{aligned} V &= \int_1^{\infty} \pi y^2 dx \\ &= \pi \int_1^{\infty} \frac{1}{x^2} dx \\ &= \pi \left[\frac{-1}{x} \right]_1^{\infty} \\ &= \pi \left[\left(\lim_{x \rightarrow \infty} \frac{-1}{x} \right) - \frac{-1}{1} \right] \\ &= \pi \end{aligned}$$




Figure 1: Gabriel’s horn

The harmonic series is indeed an astonishing series. In order to better understand it, and more specifically its divergence, we review ten elementary proofs of this well-established truth. But first, a bit of history.

2 THEN AND NOW

It is well-accepted that the first scholar to prove the divergence of the harmonic series was Nicholas Oresme. His proof was presented in the mid-1300s, and he used one of the simplest methods of grouping terms [3]. This proof was lost for many years, but in 1647, Pietro Mengoli published another proof of divergence, and roughly 30 years later the Bernoulli brothers each discovered a proof [7]. Since then, a number of other proofs have been developed.

This series has a very simple structure, and is easily understood. Despite this simplicity, it plays an important role in the study of mathematics. When considering whether or not a certain series is convergent or divergent, one of the most common series used with the comparison test to prove divergence is the harmonic series.

In 1735, Euler defined γ , also known as Euler’s constant, as intricately linked with the harmonic series. It was first defined as the limit of the difference between the n^{th} partial sum of the harmonic series and $\ln(n)$. This involvement with γ is one of the most noteworthy uses of the harmonic series. The significance of Euler’s constant, γ , should be enough to convince anyone that the divergence of the harmonic series is an important truth. If $\sum_{n=1}^{\infty} \frac{1}{n}$ converged, γ , would not exist [3], as $\gamma = 0$ in this case.

Finally, the divergence of this series is a reminder that rigorous proofs can lead us to interesting counter-intuitive results, and such results have been surprising mathematicians for centuries.

Proof 1. Appropriately, we begin with Oresme’s proof as shown by Havil [3]:

First,

$$\begin{aligned} n \in \mathbb{N} &\Rightarrow n + 1 > n \\ &\Rightarrow \frac{n+1}{n} > 1 \\ &\Rightarrow \frac{1}{n} > \frac{1}{n+1} \end{aligned}$$

Now,

$$\begin{aligned} \sum_{n=1}^{\infty} \frac{1}{n} &= 1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \frac{1}{6} + \frac{1}{7} + \frac{1}{8}\right) + \dots \\ &> 1 + \frac{1}{2} + \left(\frac{1}{4} + \frac{1}{4}\right) + \left(\frac{1}{8} + \frac{1}{8} + \frac{1}{8} + \frac{1}{8}\right) + \dots \\ &= 1 + \frac{1}{2} + \left(\frac{2}{4}\right) + \left(\frac{4}{8}\right) + \dots \\ &= 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \dots \end{aligned}$$

$$\begin{aligned} \frac{1}{k+1} + \frac{1}{k+2} + \dots + \frac{1}{k^2} &\geq \overbrace{\left(\frac{1}{k^2} + \dots + \frac{1}{k^2}\right)}^{k^2 - k \text{ times}} \\ &= \frac{k^2}{k^2} - \frac{k}{k^2} \\ &= 1 - \frac{1}{k} \end{aligned}$$

$$\text{So, } \frac{1}{k} + \frac{1}{k+1} + \dots + \frac{1}{k^2} \geq \frac{1}{k} + 1 - \frac{1}{k} = 1.$$

is divergent, so the larger series $\sum_{n=1}^{\infty} \frac{1}{n}$ is also divergent.

Proof 2. Now we present Mengoli’s proof from the 17th century, which involves grouping three terms of the series at a time [4]:

First of all, for $n \in \mathbb{N} \setminus \{1\}$,

$$\begin{aligned} \frac{1}{n-1} + \frac{1}{n+1} &= \frac{n+1+n-1}{(n-1)(n+1)} \\ \frac{2n}{n^2-1} &> \frac{2n}{n^2} = \frac{2}{n} \\ \Rightarrow \frac{1}{n-1} + \frac{1}{n+1} + \frac{1}{n} &> \frac{2}{n} + \frac{1}{n} = \frac{3}{n} \end{aligned}$$

Now, suppose that the harmonic series converges to S , then

$$\begin{aligned} S &= 1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5} + \frac{1}{6} + \frac{1}{7} + \frac{1}{8} + \dots \\ &= 1 + \left(\frac{1}{2} + \frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \frac{1}{6} + \frac{1}{7}\right) \\ &\quad + \left(\frac{1}{8} + \frac{1}{9} + \frac{1}{10}\right) + \dots \\ &> 1 + \left(\frac{3}{3}\right) + \left(\frac{3}{6}\right) + \left(\frac{3}{9}\right) + \dots \\ &= 1 + \left(\frac{3}{3 \cdot 1} + \frac{3}{3 \cdot 2} + \frac{3}{3 \cdot 3} + \dots\right) \\ &= 1 + \left(1 + \frac{1}{2} + \frac{1}{3} + \dots\right) \\ &= 1 + S \\ \Rightarrow S &> 1 + S \end{aligned}$$

We have a contradiction. Our assumption of the convergence of the harmonic series is false.

Proof 3.

This proof was developed by Jacob Bernoulli and groups terms in an interesting manner [4]. Choose some k such that $1 < k \in \mathbb{N}$.

We group terms as follows:

$$\begin{aligned} \sum_{n=1}^{\infty} \frac{1}{n} &= 1 + \left(\frac{1}{2} + \frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \dots + \frac{1}{25}\right) + \dots \\ &= 1 + \left(\frac{1}{2} + \dots + \frac{1}{2^2}\right) + \left(\frac{1}{5} + \dots + \frac{1}{5^2}\right) + \dots \\ &\geq 1 + 1 + 1 + \dots \end{aligned}$$

So, the series diverges.

Proof 4. This proof was developed by Jacob Bernoulli’s brother, Johann [4]. We first note that

$$\sum_{n=1}^{\infty} \frac{1}{n(n+1)} = \sum_{n=1}^{\infty} \left(\frac{1}{n} - \frac{1}{n+1}\right)$$

is a telescoping series and has k^{th} partial sum

$$T_k = 1 - \frac{1}{k+1} \Rightarrow \lim_{k \rightarrow \infty} T_k = \lim_{k \rightarrow \infty} 1 - \frac{1}{k+1} = 1$$

So this series converges to 1. Also,

$$\begin{aligned} \sum_{n=k}^m \frac{1}{n(n+1)} &= \frac{1}{k} - \frac{1}{m+1} \\ \Rightarrow \sum_{n=k}^{\infty} \frac{1}{n(n+1)} &= \lim_{m \rightarrow \infty} \left\{ \frac{1}{k} - \frac{1}{m+1} \right\} = \frac{1}{k} \end{aligned}$$

Assume that the harmonic series converges to S :

$$\begin{aligned}
 S &= 1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \frac{1}{5} + \frac{1}{6} + \dots \\
 &= 1 + \frac{1 \cdot 1}{2 \cdot 1} + \frac{1 \cdot 2}{3 \cdot 2} + \frac{1 \cdot 3}{4 \cdot 3} + \frac{1 \cdot 4}{5 \cdot 4} + \frac{1 \cdot 5}{6 \cdot 5} + \dots \\
 &= 1 + \frac{1}{2} + \frac{2}{6} + \frac{3}{12} + \frac{4}{20} + \frac{5}{30} + \dots \\
 &= 1 + \left(\frac{1}{2} + \frac{1}{6} + \frac{1}{12} + \dots \right) \\
 &\quad + \left(\frac{1}{6} + \frac{1}{12} + \frac{1}{20} + \dots \right) \\
 &\quad + \left(\frac{1}{12} + \frac{1}{20} + \frac{1}{30} + \dots \right) \\
 &= 1 + \sum_{n=1}^{\infty} \frac{1}{n(n+1)} + \sum_{n=2}^{\infty} \frac{1}{n(n+1)} \\
 &\quad + \sum_{n=3}^{\infty} \frac{1}{n(n+1)} + \dots \\
 &= 1 + 1 + \frac{1}{2} + \frac{1}{3} + \dots \\
 &= 1 + S
 \end{aligned}$$

We have a contradiction, so our assumption fails.

Proof 5. The next proof uses the integral test, and can be found in any basic calculus textbook [5]: First, we let $f(x) = \frac{1}{x}$ and estimate $\int_1^{\infty} f(x)dx$:

$$\begin{aligned}
 \int_1^{\infty} f(x)dx &= \int_1^{\infty} \frac{1}{x} dx = [\ln(x)]_1^{\infty} \\
 &= \lim_{x \rightarrow \infty} \ln(x) - \ln(1) \\
 &= \lim_{x \rightarrow \infty} \ln(x)
 \end{aligned}$$

which is infinite.




Figure 2: Integral test

Also, it is clear from Figure 2 that $\sum_{n=1}^{\infty} \frac{1}{n} > \int_1^{\infty} f(x)dx$ so the series diverges.

Proof 6. This short proof separates the odd terms from the even terms to contradict the convergence of the harmonic series [4]. We first suppose that $\sum_{n=1}^{\infty} \frac{1}{n}$ converges to S .

$$\begin{aligned}
 S &= \sum_{n=1}^{\infty} \frac{1}{n} \\
 &= \sum_{n=1}^{\infty} \frac{1}{2n} + \sum_{n=1}^{\infty} \frac{1}{2n-1} \\
 &= \frac{1}{2} \sum_{n=1}^{\infty} \frac{1}{n} + \sum_{n=1}^{\infty} \frac{1}{2n-1} \\
 &= \frac{1}{2}S + \sum_{n=1}^{\infty} \frac{1}{2n-1} \\
 \Rightarrow \frac{1}{2}S &= \sum_{n=1}^{\infty} \frac{1}{2n-1}
 \end{aligned}$$

Yet, $\frac{1}{2n} < \frac{1}{2n-1}$ is true for all n , so

$$\begin{aligned}
 \sum_{n=1}^{\infty} \frac{1}{2n} &< \sum_{n=1}^{\infty} \frac{1}{2n-1} \\
 \Rightarrow \frac{1}{2}S &< \frac{1}{2}S
 \end{aligned}$$

Contradiction.

Proof 7. This proof is by contradiction as well. [6]. We first show the following:

$$\begin{aligned}
 \sum_{n=1}^k \frac{1}{n} + \sum_{n=1}^{2k} \frac{1}{n} &= 2 \sum_{n=1}^k \frac{1}{n} + \sum_{n=k+1}^{2k} \frac{1}{n} \\
 &= 2 \sum_{n=1}^k \frac{1}{n} + \underbrace{\frac{1}{k+1} + \dots + \frac{1}{2k}}_{k \text{ elements}} \\
 &> 2 \sum_{n=1}^k \frac{1}{n} + \frac{k}{2k} \\
 &= 2 \sum_{n=1}^k \frac{1}{n} + \frac{1}{2}
 \end{aligned}$$

Now, we suppose that the harmonic series converges to S . So,

$$\begin{aligned}
 \lim_{k \rightarrow \infty} \sum_{n=1}^k \frac{1}{n} + \lim_{k \rightarrow \infty} \sum_{n=1}^{2k} \frac{1}{n} &= \lim_{k \rightarrow \infty} \left(\sum_{n=1}^k \frac{1}{n} + \sum_{n=1}^{2k} \frac{1}{n} \right) \\
 &\geq \lim_{k \rightarrow \infty} \left(2 \sum_{n=1}^k \frac{1}{n} + \frac{1}{2} \right) \text{ by the above inequality} \\
 &= 2S + \frac{1}{2}
 \end{aligned}$$

Yet,

$$\begin{aligned} \lim_{k \rightarrow \infty} \sum_{n=1}^k \frac{1}{n} + \lim_{k \rightarrow \infty} \sum_{n=1}^{2k} \frac{1}{n} &= 2S. \\ \Rightarrow 2S &\geq 2S + \frac{1}{2} \end{aligned}$$

But this is impossible, thus concluding the proof.

Proof 8. We first prove that $e^x > 1 + x \forall x > 0$:
We know that e^x and $\ln(x)$ increase as x increases and $\ln\left(\frac{e^x}{1+x}\right)$ is defined and positive for all positive $x \in \mathbb{R}$, so

$$\begin{aligned} \ln\left(\frac{e^x}{1+x}\right) &> 0 \\ \ln(e^x) - \ln(1+x) &> 0 \\ \ln(e^x) &> \ln(1+x) \\ e^x &> 1+x \end{aligned}$$

Next, we shortly discuss infinite products in order to better understand when they converge [1]. Partial products are defined as:

$$P_k = \prod_{n=1}^k a_n = a_1 \cdot a_2 \cdots a_k$$

and an infinite product exists if the sequence of partial products $(P_k)_{k=1}^\infty$ converges. If the terms are non-positive, the product must converge to zero. On the other hand, when $a_n > 0 \forall n \in \mathbb{N}$, we see that

$$P_k = \exp\left(\ln \prod_{n=1}^k a_n\right) = \exp\left(\sum_{n=1}^k \ln(a_n)\right).$$

Thus, it is clear that the infinite product converges if and only if $\sum_{k=1}^\infty \ln(a_k)$ converges to a finite number or diverges to $-\infty$ where $\exp\left(\sum_{n=1}^k \ln(a_n)\right) = 0$.

In 1976, Honsberger published a proof using infinite products and the above inequality that was similar to the following [4]:

$$\begin{aligned} \text{Consider } \sum_{n=1}^\infty e^{\frac{1}{n}} &= \exp\left(\sum_{n=1}^\infty \frac{1}{n}\right) \\ &= \prod_{n=1}^\infty \exp\left(\frac{1}{n}\right) \\ &> \prod_{n=1}^\infty \left(1 + \frac{1}{n}\right) \\ &= \prod_{n=1}^\infty \left(\frac{n+1}{n}\right) \end{aligned}$$

Yet, $\frac{n+1}{n}$ is positive $\forall n \in \mathbb{N}$ and does not converge to 0, so we can claim the following:

$$\begin{aligned} \prod_{n=1}^k \frac{n+1}{n} &= \left(\frac{2}{1}\right) \left(\frac{3}{2}\right) \left(\frac{4}{3}\right) \left(\frac{5}{4}\right) \cdots \left(\frac{k+1}{k}\right) \\ &= \frac{(k+1)!}{k!} = \frac{k!(k+1)}{k!} = k+1 \end{aligned}$$

Thus,

$$\begin{aligned} \lim_{k \rightarrow \infty} \exp \sum_{n=1}^k \frac{1}{n} &> \lim_{k \rightarrow \infty} (k+1) \\ \Rightarrow \exp\left(\sum_{n=1}^\infty \frac{1}{n}\right) &\text{ is unbounded} \\ \Rightarrow \sum_{n=1}^\infty \frac{1}{n} &\text{ is unbounded.} \end{aligned}$$

Proof 9. Without claiming credit for it, Havil presents the following proof involving a geometric series [6]:
We note that $\sum_{n=1}^\infty ar^{n-1}$, where a is a constant and $|r| < 1$, is a geometric series which converges to $\frac{a}{1-r}$. If $e^x < 1$, then $(1 - e^x)^{-1}$ is the sum of a geometric series with $a = 1$ and $r = e^x$, and so

$$(1 - e^x)^{-1} = \sum_{n=1}^\infty (e^x)^{n-1}$$

We then use this equality in the following integral where $x \leq 0$,

$$\begin{aligned} \int_{-\infty}^0 \left(\frac{e^x}{1-e^x}\right) dx &= \int_{-\infty}^0 (e^x(1-e^x)^{-1}) dx \\ &= \int_{-\infty}^0 \left(e^x \sum_{n=1}^\infty (e^x)^{n-1}\right) dx \\ &= \int_{-\infty}^0 e^x(1+e^x+e^{2x}+\dots) dx \\ &= \int_{-\infty}^0 (e^x+e^{2x}+e^{3x}+\dots) dx \\ &= \left[e^x + \frac{1}{2}e^{2x} + \frac{1}{3}e^{3x} + \dots\right]_{-\infty}^0 \\ &= \left(1 + \frac{1}{2} + \dots\right) \\ &\quad - \lim_{x \rightarrow -\infty} \left(e^x + \frac{1}{2}e^{2x} + \dots\right) \\ &= \left(1 + \frac{1}{2} + \dots\right) - 0 \\ &= \sum_{n=1}^\infty \frac{1}{n} \end{aligned}$$

Yet, this integral can also be evaluated using substitution:

$$(*) \int_{-\infty}^0 \left(\frac{e^x}{1-e^x}\right) dx = - \int_{-\infty}^0 (-e^x(1-e^x)^{-1}) dx$$

We let $u = 1 - e^x$, so

$$\begin{aligned}
(*) &= -\int_{x=-\infty}^{x=0} \left(\frac{1}{u}\right) du \\
&= \left[-\ln(u)\right]_{x=-\infty}^0 \\
&= \left[-\ln(1-e^x)\right]_{-\infty}^0 \\
&= -\ln(1-e^0) - \lim_{x \rightarrow -\infty} (-\ln(1-e^x))
\end{aligned}$$

Now,

$$\begin{aligned}
\lim_{x \rightarrow -\infty} (-\ln(1-e^x)) &= -\ln\left(1 - \lim_{x \rightarrow -\infty} e^x\right) \\
&= -\ln(1-0) = 0
\end{aligned}$$

and $-\ln(1-e^0) = -\ln(1-1)$; the \ln function is undefined for 0. This implies that the value of this integral is infinite, so the sum of the harmonic series is also infinite.

Proof 10.

For this final proof, we use the Fibonacci numbers [6]. These well-known numbers are recursively defined as $f_0 = 1$, $f_1 = 1$ and $\forall n \in \mathbb{N}$, $f_{n+1} = f_n + f_{n-1}$.

Though it can be shown using calculus that

$$\lim_{n \rightarrow \infty} \frac{f_{n+1}}{f_n} = \frac{1 + \sqrt{5}}{2},$$

the golden ratio, we will not show it here.

Now,

$$\begin{aligned}
f_{n+1} &= f_n + f_{n-1} \\
\Rightarrow f_{n-1} &= f_{n+1} - f_n \\
\text{So, } \lim_{n \rightarrow \infty} \frac{f_{n-1}}{f_{n+1}} &= \lim_{n \rightarrow \infty} \frac{f_{n+1} - f_n}{f_{n+1}} \\
&= \lim_{n \rightarrow \infty} 1 - \frac{f_n}{f_{n+1}} \\
&= 1 - \lim_{n \rightarrow \infty} \frac{f_n}{f_{n+1}} \\
&= 1 - \frac{2}{1 + \sqrt{5}} \\
&= \frac{\sqrt{5} - 1}{1 + \sqrt{5}} \neq 0
\end{aligned}$$

Thus, the series $\sum_{n=1}^{\infty} \frac{f_{n-1}}{f_{n+1}}$ diverges as the limit of its terms is not 0. Yet,

$$\begin{aligned}
\sum_{n=1}^{\infty} \frac{1}{n} &= 1 + \frac{1}{2} + \frac{1}{3} + \left(\frac{1}{4} + \frac{1}{5}\right) \\
&\quad + \left(\frac{1}{6} + \frac{1}{7} + \frac{1}{8}\right) + \left(\frac{1}{9} + \dots + \frac{1}{13}\right) + \dots \\
&\geq 1 + \frac{1}{2} + \frac{1}{3} + \left(\frac{1}{5} + \frac{1}{5}\right) \\
&\quad + \left(\frac{1}{8} + \frac{1}{8} + \frac{1}{8}\right) + \underbrace{\left(\frac{1}{13} + \dots + \frac{1}{13}\right)}_{5 \text{ times}} + \dots
\end{aligned}$$

$$\Rightarrow \sum_{n=1}^{\infty} \frac{1}{n} = 1 + \frac{1}{2} + \frac{1}{3} + \frac{2}{5} + \frac{3}{8} + \frac{5}{13} + \dots$$

We note that the first few terms of the Fibonacci sequence are the following: 1, 1, 2, 3, 5, 8, 13... We have thus grouped the terms of the harmonic series such that:

$$\begin{aligned}
\sum_{n=1}^{\infty} \frac{1}{n} &\geq 1 + \frac{f_0}{f_2} + \frac{f_1}{f_3} + \frac{f_2}{f_4} + \frac{f_3}{f_5} + \frac{f_4}{f_6} + \dots \\
&= 1 + \sum_{n=1}^{\infty} \frac{f_{n-1}}{f_{n+1}} \text{ which diverges.}
\end{aligned}$$

Hence, the harmonic series does not converge.

3 IT NEVER ENDS

Another intriguing series that diverges is the sum of the reciprocals of each prime. This series is even smaller than the harmonic series, but does not converge. A short discussion can be found in section 3.2 of [3] by the interested reader.

And so it is, the harmonic series diverges to positive infinity. No matter how surprising this is, it is clear that this fact must be accepted. The simple structure of this series has allowed mathematicians to come up with these pretty little proofs, and many more. Perhaps there are infinitely many of these; perhaps the “series of harmonic-series-divergence proofs” diverges as well.

REFERENCES

- [1] S.W. Drury. Class notes for math 255. Unpublished class notes, McGill University, 2006.
- [2] Julian F. Floron. Gabriel’s wedding cake. *The College Mathematics Journal*, 30(1):35–38, 1999.
- [3] Julian Havil. *Gamma: Exploring Euler’s Constant*. Princeton University Press, Princeton, NJ, 2003.
- [4] I.I. Hirschman. *Infinite Series*. Holt, Rinehart and Winston, New York, 1962.
- [5] S.J. Kifowit. More proofs of the divergence of the harmonic series. Unpublished, available at faculty.prairiestate.edu/skifowit/htdocs/harm2.pdf.
- [6] S.J. Kifowit and T.A. Stamps. The harmonic series diverges again and again. *The AMATYC Review*, 27:31–43, 2006.
- [7] Eric W. Weisstein. Harmonic series, math world, 14 Sep 2009. <http://mathworld.wolfram.com/HarmonicSeries.htm>.

PROBABILISTIC TRACKING OF PEDESTRIANS FROM PLATFORM FORCE DATA

Rishi Rajalingham

This article presents a probabilistic approach to track and estimate lower body pose of users interacting with a virtual ground surface. Based on the low resolution platform force observations arising over time from pedestrians walking on a tiled floor, we use a Bayes filter, specifically a particle filter, to approximate the distribution of parameters corresponding to a human pose. The work presented is ongoing research as part of Natural Interactive Walking (NIW), which investigates multimodal interaction with a virtual ground surface. Applications for such a method include foot sensing upon contact, as well as controlling virtual skeletons (avatars) via natural foot interactions.

1 INTRODUCTION

Pose estimation is a classic problem in computer vision where one is interested in automatically inferring posture of individuals from image data; traditional solutions involve training a classifier with image features to identify distinct poses. Significantly, the loss of information in the planar projection of a 3-dimensional scene onto an image sequence and the artifacts caused by lighting and inter-subject variability already render it a difficult, if not ill-posed problem. This paper presents a coarse pose estimation algorithm based not on visual data, but rather the foot-floor contact forces of pedestrians interacting naturally with a floor surface.

Many of the afore-mentioned challenges remain relevant, but are compounded by the limitations of the sensing infrastructure, the difficult mapping of footprints to posture, the complex dynamics of human motion. As a result, we aim only to coarsely estimate lower body pose, corresponding to feet positions and orientations, and track it over the course of arbitrary walking motion.

We present a probabilistic approach to track lower body pose of users interacting with a virtual floor surface. This work was done as part of the Natural Interactive Walking project, which investigates virtual and augmented reality paradigms, specifically multi-modal interaction with a virtual floor. The remainder of this paper is sectioned as follows: we first present the infrastructure and define the problem at hand, then lay out the theoretical background and motivation for Bayes filters, in particular particle filters, and conclude with the specific system design. Preliminary results are available online, as the following is part of ongoing work.

2 PROBLEM STATEMENT

The NIW infrastructure consists of a floor surface, a 6x6 array of rigid tiles, each 1 square foot of area. An actuator is bolted to the under belly of each tile for vibro-tactile feedback. Force sensing resistors (FSRs) are positioned under elastic suspension mounted at the four corners of each tile (see Figure 1).




Figure 1: Schematic of NIW infrastructure.

As a result of the sensing configuration of the floor system, incoming force data is at a very low spatial resolution. Indeed, the entire floor surface corresponds to a 12×12 data matrix. However, previous NIW work has shown that estimating feature points corresponding to tiles' centers of pressures can drastically increase the sensing resolution. In particular, we are able to estimate lower body pose reasonably well in the static case, by clustering center of pressure feature points, labeling them as feet and coarsely reconstructing the skeleton. In the dynamic case, it is important to consider the temporal continuity of observations. This motivates the use of Bayes filters in considering the dynamic system that is arbitrary walking.

Another motivation to use Bayesian filtering in tracking pedestrian pose is that they have been shown to be effective despite occluded or noisy observations. In the context of pose estimation and tracking from platform forces, we note that the discontinuities of force observations, due to the slice-projection of a dynamic pose onto a platform, are akin to occlusions in images, while the low resolution of sensor data calls for a robust estimation method.

3 THEORY

3.1 Bayes Filters

Let x_t, z_t be the time indexed states and corresponding observations, respectively, of a general dynamic system as illustrated in Fig 2 . Bayes filters probabilistically estimate at any time t the state x_t by approximating the probability

distribution $Bel(x_t)$. $Bel(x_t)$ is called the *belief* of state x and is defined by the posterior probability of x_t conditioned on the history of available observations:

$$Bel(x_t) = p(x_t | z_t, z_{t-1}, z_{t-2}, \dots) = p(x_t | z_{0:t})$$




Figure 2: Dynamic system: states x_t evolving as a Markov process, with corresponding observations z_t .

Assume that the dynamic system is Markov, i.e. the states comprise a Markov process (the current state depends at most on one past state). Then, we have $p(x_t | x_{0:t-1}) = p(x_t | x_{t-1})$. Similarly, the observations comprise a Hidden Markov Model $p(z_t | x_{0:t-1}) = p(z_t | x_{t-1})$. To compute the posterior density, simply use Bayes' Theorem on $p(x_t \cap z_{0:t-1} \cap z_{0:t})$:

$$p(x_t | z_{0:t}) p(z_{0:t} | z_{0:t-1}) = p(z_{0:t} | x_t) p(x_t | z_{0:t-1})$$

Marginalizing over x_t, x_{t-1} respectively,

$$LHS = p(x_t | z_{0:t}) \int_0^0 p(z_{0:t} | x_t) p(x_t | z_{0:t-1}) dx_t,$$

$$RHS = p(z_{0:t} | x_t) \int_0^0 p(x_t | x_{t-1}) p(x_{t-1} | z_{0:t-1}) dx_{t-1}$$

Since $p(z_{0:t} | z_{0:t-1})$, a prior over observations, is constant relative to x_t ,

$$p(x_t | z_{0:t}) = \frac{p(z_{0:t} | x_t) \int_0^0 p(x_t | x_{t-1}) p(x_{t-1} | z_{0:t-1}) dx_{t-1}}{\int_0^0 p(z_{0:t} | x_t) p(x_t | z_{0:t-1}) dx_t}$$

$$Bel(x_t) \propto p(z_{0:t} | x_t) \int_0^0 p(x_t | x_{t-1}) Bel(x_{t-1}) dx_{t-1}$$

This can be thought of as a predict-and-update process:

1. Using the *motion probability* described by $p(x_t | x_{t-1})$, predict posterior density states x as $Bel^*(x_t) = \int_0^0 p(x_t | x_{t-1}) Bel(x_{t-1}) dx_{t-1}$
2. Using the *likelihood* prescribed by $p(z_t | x_t)$, update posterior based on new observation: $Bel(x_t) = \alpha p(z_t | x_t) Bel^*(x_t)$

Thus, given *observation* and *dynamic models* to describe the likelihood and motion probability of a dynamic system respectively, one can estimate its state with the mean of the posterior density, $x_t = E[Bel(x_t)]$.

Common implementations of Bayes filters are the Kalman filter, extended Kalman Filter (EKF), and the particle filter. Since the first two make strong assumptions about the dynamic system (for instance, linear observation and dynamic models and Gaussian posteriors), the particle filter is favoured in applications where these may not apply.

3.2 The Particle Filter

The particle filter, also known as bootstrap filter, condensation algorithm, survival of the fittest, Sampling Importance Resampling Filter, etc..., uses a discrete set of weighted samples (particles) $S_t = (X_t^i, w_t^i)_i$ to approximate $Bel(x_t)$.

As the number of particles becomes large, S_t becomes a better approximation of $(x_t, p(x_t | z_{0:t}))$. The Sampling-Importance-Resampling algorithm for the particle filter is shown below, and illustrated in Figure 3:

1. Initialize particles (ex: as a normal distribution, with mean corresponding to initial estimate, and appropriate variance).
2. For each observation z_t ,
 - (a) Sample and predict particles based on dynamic model.
 - (b) Set weight of particles based on importance, computed as normalized likelihood: $w_t^i = \frac{p(z_t | x_t^i)}{\sum_{j=0}^n p(z_t | x_t^j)}$.
 - (c) Resample particles based on weights (Create new set of particles S_t^* by sampling S_t with selection probability proportional to weight).




Figure 3: Particle filter algorithm: (1) Set weight based on likelihood (2) Resample (3) Project forward with motion model.

Note that the resampling step distinguishes SIR from SIS (Sampling Importance Sampling), and is meant to eliminate degenerate particles. Moreover, particle filters allow for priors to further eliminate unlikely states.

4 DESIGN

In the context of pedestrian tracking and pose estimation based on platform force observations, we define the Bayesian filter variables as:

- \mathbf{z}_t : the time indexed observations as normalized platform force measurement vector.
- \mathbf{x}_t : the time indexed states as a vector describing kinematic poses. The state is implemented as a 25-dimensional vector concatenating 3-dimensional position, velocity and acceleration vectors of three skeletal nodes (right foot, left foot, center of mass) and a discrete-valued foot-floor contact variable (no/left/right/left&right foot contact). $\mathbf{x}_t = (lf, \dot{lf}, \ddot{lf}, rf, \dot{rf}, \ddot{rf}, com, \dot{com}, \ddot{com}, \beta)$. See Figure 4.




Figure 4: Schematic of state representation, parameters corresponding to a pose.

We then define a particle filter by its observation and dynamic models.

4.1 Observation Model

As previously noted, the observation model defines the likelihood function $L(\mathbf{x}_t) = p(\mathbf{z}_t | \mathbf{x}_t)$ describing the probability of observation of captured force data \mathbf{z}_t given a pose \mathbf{x}_t . We compute the likelihood function as follows:

Assume $\mathbf{z}_t = \mathbf{H}(\mathbf{x}_t) + \mathbf{V}_t$, where H is the observation model which maps the true state space into the observed space and \mathbf{V}_t is the observation noise. We may assume zero-mean Gaussian white noise with covariance R_t , $\mathbf{V}_t \sim N(0, R_t)$.

Our observation model first generates the expected observations $\mathbf{H}(\mathbf{x}_t)$ for state \mathbf{x}_t using simulated mechanic poses. A foot is simulated as centers of pressure located in the contact area. Figure 5 shows the steps to map a state (a foot pose) into the observation space (force sensor values), via the estimation of centers of pressure. Each center

of pressure is then linearly distributed as a point mass across a rigid tile towards the four force sensors:

$$f_i = \frac{1}{d_i} \sum_{j=0}^4 \frac{1}{d_j}$$




Figure 5: Schematic of observation model, computing expected observations from state.

Using this observation model, we compute $p(\mathbf{z}_t | \mathbf{x}_t) = N(\mathbf{z}_t; \mathbf{H}(\mathbf{x}_t), \mathbf{V}_t)$.

The likelihood function is modified to consider human postural constraints, namely to consider only those postures which are humanly possible within the context of walking. Let γ^A be a set of kinematic postural constraints for action A (arbitrary walking). The postural prior $p(\mathbf{x}_t | \gamma^A)$ is evaluated as a product of single variable Gaussians considering stance size and stance angle constraints. Let the stance size ζ_s be the distance between the two feet, and the stance angle ϕ_s the relative angle difference between them. Then $p(\mathbf{x}_t | \gamma^A) = N(\zeta_s; \mu_{\zeta_s}, \sigma_{\zeta_s}) \times N(\phi_s; \mu_{\phi_s}, \sigma_{\phi_s})$.

Then, assuming independence of these constraints relative to the observation model, we have $L(\mathbf{x}_t) = p(\mathbf{x}_t | \gamma^A) p(\mathbf{z}_t | \mathbf{x}_t)$.

4.2 Dynamic Model

The dynamic model describes system dynamics via the motion probability $p(\mathbf{x}_t | \mathbf{x}_{t-1})$, the probability of state \mathbf{x}_t given the previous state \mathbf{x}_{t-1} . The particle filter motion model propagates particles \mathbf{x}_t^i forward in time in between observations.

Recall that $\mathbf{x}_t = (lf_t, \dot{lf}_t, \ddot{lf}_t, rf_t, \dot{rf}_t, \ddot{rf}_t, com_t, \dot{com}_t, \ddot{com}_t, \beta_t)$, concatenating position, velocity and acceleration

vectors of three skeletal nodes (right foot, left foot, center of mass) and a discrete-valued foot-floor contact variable. Note also that arbitrary walking is a highly complex motion. We approximate its state dynamics with first order motion and discrete state transitions.

4.2.1 First Order Motion Model

Let ξ_t be a free skeletal node, i.e. right foot, left foot, or center of mass not in contact with the floor. Consider the following first order motion, corresponding to its linear drift. It is illustrated in Figure 6

$$\begin{aligned}\xi_t &= \xi_{t-1} + \dot{\xi}_{t-1} dt \\ \dot{\xi}_t &= \dot{\xi}_{t-1} + \eta_t, \eta_t \sim N(0, \sigma_v)\end{aligned}$$




Figure 6: Schematic of first-order motion model— linear drift of free skeletal nodes.

4.2.2 State transition

Recall that the discrete-valued contact variable $\beta \in \{NONE, LF, RF, LRF\}$ relates to the various foot-floor contact configurations. The dynamic model includes transitions from the various foot-floor contact states via the following stochastic process (Figure 7), which approximates discrete stepping motion. δ is an empirically determined probability of no change in β .

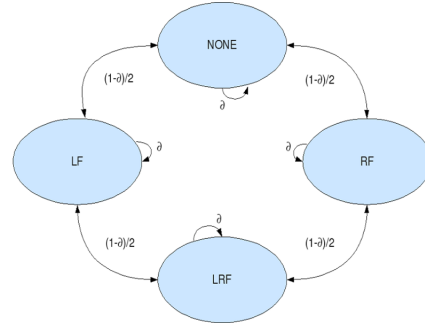


Figure 7: Stochastic state transition diagram for stepping motion approximation.

5 RESULTS AND CONCLUSION

To the best of our knowledge, tracking and estimating pose based on platform forces is a novel idea. It is in fact part of ongoing work at the Shared Reality Lab at McGill University, with preliminary results suggesting that the approach taken in this paper is promising. For details and updates, please visit see <http://www.cim.mcgill.ca/~rishi/>.

6 ACKNOWLEDGEMENTS

The above article was a brief account of a summer project completed by ECSE undergraduate student Rishi Rajalingham as part of the Natural Interactive Walking (NIW) project. This work was conducted under the supervision of Yon Visell and Dr. Jeremy Cooperstock from McGill University's Center for Intelligent Machines (CIM), with funding from NSERC USRA. For more information about NIW, visit <http://www.niwproject.eu/>.

REFERENCES

- [1] I. Rius, J. Varona, X. Roca, and J. Gonzalez. Posture constraints for bayesian human motion tracking. Lecture Notes in Computer Science, 2006.
- [2] Y. Visell, A. Law, J. Ip, R. Rajalingham, S. Smith, and J. Cooperstock. Interaction capture in immersive virtual environments via an intelligent floor surface. IEEE VR 2010 (Under review), 2010.
- [3] K. Yin and D.K. Pai. Footsec: an interactive animation system. In Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, 2003.

STRANGE CYCLING: RESULTS AND EXAMPLES OF DELAY DIFFERENTIAL EQUATIONS AS DYNAMICAL SYSTEMS

Orianna DeMasi

Delay differential equations (DDEs) are similar to ordinary differential equations (ODEs) but are a larger class of differential equations that have are more versatile for modeling physical situations and have richer behavior as dynamical systems. To explore the use of DDE's in modeling, two classic examples from biology are presented where including delays created realistic behavior in the model. Then, as an example of complex dynamics, a test DDE with state-dependent delays is introduced and recent results on its bifurcation properties are presented.

1 INTRODUCTION

Delay differential equations (DDEs) allow the derivative of a system to be affected by past states. Such dependency gives DDEs richer dynamics and more flexibility as models than ODEs. As a result, there is much interest in DDEs for applications in experimental sciences. For example DDE's are used in ecology as population models where the delay is maturation time from juvenile to reproductive adult, in engineering for remote control of machines where the delay is time between sensor and computer response, and in physiology for modeling neurons where delay is time between transferring chemical and response. In spite of their utility, many intriguing and bizarre dynamics that arise from DDEs are not well understood as the mathematical theory governing DDEs is delicate and still in need of much development.

DDE's are most generally described as equations of the type

$$y'(t) = f(t, y(t), y_t) \quad y_t = y(t - \tau), \tau > 0 \quad (1)$$

where y_t is a function mapping the interval $[-\tau, 0]$ into \mathbb{R}^d and $f: \Omega \in \mathbb{R} \times C \rightarrow \mathbb{R}^d$; the derivative depends on the current state and past states. τ is known as the delay and can take progressively more complicated forms, which cause various problems for solving or even numerically approximating the equations. To solve or approximate, (1) must be restated as an initial value problem with initial conditions. While in first order ODE's, only a simple initial value y_0 is given, due to continuous dependence on past states, DDE's knowledge is required of the function on an interval in the past so DDE's require a continuous initial function $\phi(t)$ to be given as history on an entire interval $[t_0 - \tau, t_0]$. As an initial value problem with n delays (1) becomes

$$\begin{cases} y'(t) = f(t, y(t), y(t - \tau_1), \dots, y(t - \tau_n)) & t_0 \leq t \\ y(t) = \phi(t), & t \leq t_0 \end{cases} \quad (2)$$

Recall that a scalar first order ODE can be rewritten as a linear system of equations by declaring a new variable $z(t) = y'(t)$. It then becomes obvious that the ODE is two dimensional. Similarly, an n th order ODE is of $n+1$ dimensions and requires n initial conditions. Thus ODE's of finite order define systems of finite dimensions. In contrast, as DDE's require a function over an entire interval, they are of infinite dimension. The dimensionality is one of the most

fascinating differences between ODE's and DDE's and it prevents much ODE theory from applying to DDE's; it also opens lots of space for interesting behavior to occur when DDE's are considered as dynamical systems.

2 DELAYS IMPROVE LIFE TWO FAMOUS EXAMPLES

There are two classic examples where delays significantly improved models. Both are taken from biological applications, the first from population dynamics and the second from physiology.

2.1 Ecology- Logistic Equation

The logistic equation is a common tool for modeling population levels [2]. This equation takes into account average population growth and death in a coefficient r and the growth or death dependent on population density and stress on resources. It has been know for quite some time that its discrete form

$$u_{n+1} = -ru_n(1 - u_n) \quad (3)$$

can give rise to chaos. On the other hand, the continuous model $u(t) = -ru(t)(1 - u(t))$ is much less versatile; it only generates monotonic solutions. The delayed form

$$u(t) = -ru(t)(1 - u(t - \tau)) \quad (4)$$

improves the continuous model by delaying the density term $(1 - u(t - \tau))$ which considers growth when population is small enough and mortality if population exceeds environmental resources. The delay allows the model to not monotonically approach a steady state, which is unrealistic in observed populations. (2.1) has solutions which are monotonic for $r \in (0, 1/e)$, oscillate as they approach a steady state for $r \in [1/e, \pi/2)$ and periodic orbits for $r > \pi/2$.




Figure 1: Delayed logistic model in phase plane. Shows one solutions monotonically approaching steady state, the other oscillating to steady state. Picture from [2]

2.2 Physiology – Mackey-Glass Equation

Another example of DDE success is one of the most famous differential delay equations. The Mackey-Glass equation

$$\frac{du}{dt} = \beta \frac{u(t-\tau)}{1+u(t-\tau)^n} - \gamma, \quad (5)$$

with $\beta, \gamma, \tau, n \in \mathbb{R}^+$ was developed in the 1980's by Michael Mackey and Leon Glass, from McGill University, to model levels of circulating white blood cell in patients. Considering the physiological basis of leukocyte production, Mackey and Glass used (5) with the delay τ as cell maturation time before circulation, to study why some individuals suffer from Cyclic Neutropenia, a disease which causes blood cell levels to fluctuate rapidly and violently. It was found that for certain parameter ranges, the equation cycled similar to normal cell cycling levels. However, when parameters were perturbed sufficiently, the mellow dynamics broke into chaotic fluctuations (see figure 2).




Figure 2: Mackey-Glass equation in phase space. Shows chaotic behavior equation can produce. Picture from [3]

The cycling found received much accolade as it lent insight into physiology through mathematics and it showed

chaotic behavior from a single equation rather than a system of equations. Note that for chaotic behavior in ODE's at least three equations (dimensions) are needed.

3 STATE DEPENDENT DDE AS A DYNAMICAL SYSTEM

Now, we present some results from a current research project. This projects investigates the rich dynamics of a DDE with 2 delays which is of the form

$$\begin{aligned} \varepsilon \dot{u}(t) = & -\gamma u(t) - \kappa_1 u(t - a_1 - c_1 u(t)) \\ & - \kappa_2 u(t - a_2 - c_2 u(t)) \end{aligned} \quad (6)$$

Note that the delays $\tau = \tau(t, u(t))$ are state-dependent. Such delays are difficult because it is not clear where the delays will fall and where the breaking points will decrease order of continuity and how of if it's possible to bound the delays. The above form is chosen as the linear dependence indicates that it's the simplest form of a state-dependent delay that can be considered.

3.1 Parameter Conditions

Some parameter conditions must be established to ensure the equation does not become advanced but remains delayed. Without loss of generality it can be assumed

$$-\frac{a_2}{c_2} \leq -\frac{a_1}{c_1} < 0$$

In order for (6) to remain delayed, both delays must remain less than the given time; by the ordering of parameters this holds when

$$-\frac{a_1}{c_1} \leq u(t)$$

The above condition gives a lower bound on the DDE which can be used to find an upper bound on the solution using

$$0 \leq \varepsilon \dot{u}(t) \leq -\gamma u(t) + \frac{a_1}{c_1} (\kappa_1 + \kappa_2)$$

With the solution's upper bound and forcing \dot{u} to vanish or become positive at the lower bound $\dot{u}(-\frac{a_1}{c_1}) \geq 0$, yields a second condition on the parameters of the DDE. These two conditions are summarized by

$$\begin{aligned} (1) \quad & 0 > -\frac{a_1}{c_1} > -\frac{a_2}{c_2} \\ (2) \quad & \gamma \geq \kappa_2 + \kappa_1 \end{aligned} \quad (7)$$

and ensure that $u(t) > -\frac{a_1}{c_1}$ for all t and that (6) remains delayed.

3.2 Linearization

Consider that $u(t) = 0$ satisfies (6), is clearly a steady state solution, and in general it is the only steady state solution except for rare cases with specific parameters. Similar to ODE theory, we can linearize the above equation around

$u(t)$ for u near 0 to gain intuition about stability and dynamics of the system. Expanding in a Taylor Series, each delayed term becomes of the form

$$u(t - a_i - c_i u(t)) \Big|_{u \approx 0} = u(t - a_i) + u'(t - a_i)(-c_i u(t)) + \frac{u''(t - a_i)}{2!}(-c_i u(t))^2 + \dots$$

Ignoring terms of higher orders near $u = 0$ leaves constant delay approximations.

$$u(t - a_i - c_i u(t)) \approx u(t - a_i), i = 1, 2$$

Substituting in (6) we are left with a linear problem.

$$\varepsilon \dot{u}(t) = -\gamma u(t) - \kappa_1 u(t - a_1) - \kappa_2 u(t - a_2) \quad (8)$$

To consider the characteristic values λ of equation (8), substitute $u = e^{\lambda t}$ which yields the characteristic values to be zeros of

$$f(\lambda) = \varepsilon \lambda + \gamma + \kappa_1 e^{-\lambda a_1} + \kappa_2 e^{-\lambda a_2}$$

As f is a transcendental equation, there are infinitely many (complex) roots. Note that if $\lambda, \varepsilon, \gamma, \kappa_1, \kappa_2 > 0$ then there are no positive real roots.

To look at the behavior of such roots, substitute $\lambda = x + iy$ and separate parts with Euler's equation to see the real part of the equation must satisfy

$$Re(f(\lambda)) = \varepsilon x + \gamma + \kappa_1 e^{-a_1 x} \cos(a_1 y) + \kappa_2 e^{-a_2 x} \cos(a_2 y) = 0$$

And the imaginary part must satisfy

$$Im(f(\lambda)) = \varepsilon y - \kappa_1 e^{-a_1 x} \sin(a_1 y) - \kappa_2 e^{-a_2 x} \sin(a_2 y) = 0$$

Rearranging terms, squaring, and summing; these two equations describe a curve

$$(\varepsilon x + \gamma)^2 + (\varepsilon y)^2 = \kappa_1^2 e^{2a_1 x} + \kappa_2^2 e^{-2a_2 x} + 2\kappa_1 \kappa_2 e^{-(a_1 + a_2)x} \cos((a_1 - a_2)y) \quad (9)$$

on which all characteristic values must lie. Examples of such curves are in figure 3. To establish the location of characteristic values on this curve we can also consider that at such points (roots of f) we would need

$$g(\lambda) = (Re(f))^2 + (Im(f))^2 = 0$$




Figure 3: For $\gamma = 3, \varepsilon = 1, \kappa_1 = 4, a_1 = 1.5, \kappa_2 = 1, a_2 = .5$ the curve described by (9) is indicated. The actual location of characteristic values is indicated by a contour plot of the roots of function $g(\lambda)$. Hopf bifurcations occur as characteristic values pass into right half plane.

3.3 Bifurcations and Dynamics

For different parameter values (9) deforms. When some parameters are increased, the curve moves right as it does, characteristic values pass into the right half plane (see figure 4). Considering Euler's equation we note

$$u(t) = e^{\lambda t} = e^{(x+iy)t} = e^{xt} (\cos(yt) + i \sin(yt))$$

When all characteristic values have negative real parts, then perturbations will decay and the solution will evolve back to the steady state. If any characteristic values exist with positive real parts, then dynamics will not decay back to the steady state but will instead persist (or grow) into a cycle with frequency given by the imaginary coefficient. When the characteristic values cross into the right half plan, there is a sudden change in dynamics as the steady state loses stability and instead of evolving to the steady state, solutions evolve to periodic orbits. This type of change in dynamics is referred to as a Hopf bifurcation. Figure 5 expresses the important features of a hopf bifurcation causing the system to go from a steady state into an orbit of growing amplitude as the bifurcation parameter changes.

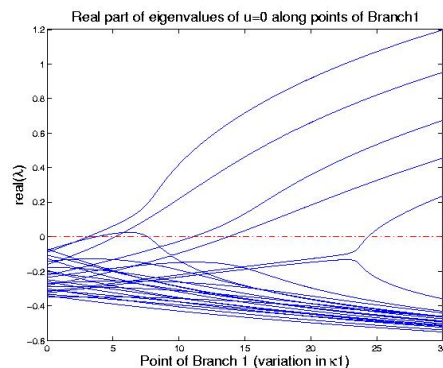


Figure 4: Real parts of characteristic values passing into right half plane as κ_1 is increased. Values crossing the imaginary axis cause a destabilization of the steady state.




Figure 5: Simplified view of what happens when a single Hopf bifurcation occurs. The steady state turns into a periodic orbit with increasing amplitude.

Consider κ_1 as the bifurcation parameter in equation (6). As this parameter is increased, the system is destabilized. Successive Hopf bifurcations occur indicating the system becomes more unstable or unstable in more dimensions. What is interesting is how the system becomes unstable for different values of κ_2 . Refer to figures 6,7 which show the amplitude of periodic orbits. Note that each intersection with the horizontal axis indicates a point of zero amplitude which is where the Hopf bifurcation occurs. As κ_1 is increased, the amplitude of the period increases. When $\kappa_2 = .1$ the leading branch of orbits (out of the first bifurcation) increases monotonically (figure 6). This behavior is to be expected. What is really interesting is that we do not see the same behavior with $\kappa_2 = 3$ (figure 7). Instead, at the larger value of κ_2 the leading branch increases smoothly for a bit, then winds back over itself before continuing.

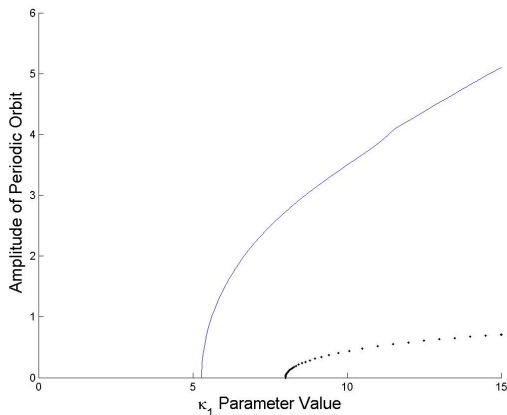


Figure 6: $\kappa_2 = .1$ Plot indicates the number of and amplitude of periodic orbits out of Hopf bifurcations. Note smooth increase in amplitude indicated by first branch. This is expected behavior of system.




Figure 7: $\kappa_2 = 3$ The branch of periodic solutions out of the first Hopf point changes stability in six different bifurcation points. These points are indicated by arrows. Particular interest is the apparent lack of stable object in region between bifurcation # 1 and # 2.

This is not standard behavior and it indicates much more is going on within the system than mere Hopf bifurcations. Consider the region of κ_1 values between bifurcations numbers #4 and #3. There are two stable portions of branch above this region indicating that there are two stable, co-existing periodic orbits. Co-existing stable orbits mean that long-term behavior is non-unique and where the system ends up depends on the initial conditions. Also, as an unstable orbit lives between the stable orbits (indicated with the dotted branch) co-existing stable orbits imply that different perturbations from this unstable orbit will converge to different stable orbits. Perhaps the most interesting object that has been found thus far in (6) lives in the region between bifurcation # 1 and # 2. Note that in this region there is apparently no stable object. It is easy to show that the system is bounded so solutions cannot be unbounded but must converge to something. Looking at Floquet multipliers to see how stability is lost, it becomes clear that in this region the periodic cycle bifurcates itself at # 1. Instead of orbiting around on a single path and always reaching each point an infinite number of times, the system evolves on an orbit around the orbit; it evolves on a surface around the orbit and carves out an invariant torus. The cartoon in figure 8 indicates how one can think of the solution evolving on a torus surrounding the periodic solution which existed before bifurcation #1.




Figure 8: Solution evolves in time over the surface of a torus. The torus surrounds the periodic orbit which lost stability; solutions perturbed off the unstable orbit converge to living on the surface of the torus.




Figure 9: Simulations of the torus found for (6) with $\kappa_2 = 3$.

4 SUMMARY

DDE's are capable of modeling strange and wonderful behavior. As seen in the two examples of chaos above, DDE's can describe quite complicated systems with what appear as relatively simple equations. Further, DDE's require dif-

ferent mathematical theory than ODE's and must be approached with delicacy so as to not lose any interesting underlying behavior of the system in numerical simulations.

5 ACKNOWLEDGEMENTS

Thank you to Tony Humphries for his guidance and tutelage on DDE's, Felica Magpanty for her editing, and McGill University for financial support.

REFERENCES

- [1] A. Bellen and Z. Marino. *Numerical Methods for Delay Differential Equations*. Clarendon Press, 2003.
- [2] R. Knapp. Delay logistic equation, wolfram demonstration project, 2009. <http://demonstrations.wolfram.com/DelayLogisticEquation/>.
- [3] M. Mackey and L. Glass. Mackey-glass equation, scholarpedia, 2009. http://www.scholarpedia.org/article/Mackey-Glass_equation?url.

JOKES

Q: Why couldn't the Möbius strip enroll at the school?

A: The school required an orientation. □

Q: What is the world's longest song?

A: "Aleph-nought Bottles of Beer on the Wall." □

You know how dumb the average guy is? Well, by definition, half of them are even dumber than that. □

The difference between an economist and a statistician: people believe what economists say about the future, but not what statisticians say about the past. □

An algebraist had to move to a new place. His wife didn't trust him very much, so when they stood down on the street with all their things, she asked him to watch their ten trunks, while she got a taxi. Some minutes later she returned. The husband said:

"I thought you said there were ten trunks, but I've only counted nine."

The wife said, "No, there are TEN!"

"...but I have counted them: 0, 1, 2, ..." □

"I bet you can't divide fourteen sugar cubes into three cups of coffee so that each cup has an odd number of sugar cubes in it."

"That's easy: one, one, and twelve."

"But twelve isn't odd!"

"Twelve is an odd number of sugar cubes to put into a cup of coffee..." □

"What is the area of a circle?"

" πr^2 ?"

"No, pies are not square. Pies are round. Cornbread is square." □ (Editor's note: what?)

INTERVIEW WITH PROF. RUSS STEELE

Ana Best



Picture 1: Prof. Russ Steele

DE: What is your name, and how long have you been at McGill?

I'm Russ Steele, and this is my eighth year at McGill; I started in fall 2002.

DE: Tell me a bit about your background, both personal and academic.

I grew up in Pittsburgh, Pennsylvania, and did my undergraduate degree at Carnegie Mellon University, which is an engineering school in Pittsburgh, and then I did my Masters degree there as well, both in statistics. Then I went to the University of Washington in Seattle and did my Ph.D. there; I was there for four years. And then, after I finished my Ph.D., I came to McGill in 2002. I think I defended my thesis in July, and was here a week later.

DE: What are your favourite things about McGill and Montreal?

McGill – It definitely is the students first. There were three things that I wanted when I was looking at jobs. The top three criteria I had were: it had to be in a big city, it had to be on the east coast, and it had to have very good undergraduate students.

So, I would say that for McGill, the undergraduates are just really really good considering the size of the university.

At Carnegie Mellon, there were about 4 or 5 thousand undergraduates at the time I went there, so it's a tiny private school. And then Washington is enormous, and I found that the quality of the undergraduates was not the same as I was used to at Carnegie Mellon.

I have to say that, coming to McGill, even though McGill is still a fairly large school, the level is very high. Not just the really good students being really good, which you would expect from any school that has a lot of students, it's also the average student that I see in a majors course is just really good. So I would say that that's probably my favorite thing.

I do a lot of collaborative work in medical research, and the medical school here is just off the charts ridiculously good, so I've been very fortunate to collaborate with some of the best people in the world in their disciplines, so that makes collaboration really exciting and you kind of get to work on really hard problems, which is really good.

The statistics group is really cohesive; it was hard for me to come to a math and statistics department because that was very different for me. Carnegie Mellon has a separate statistics department, the University of Washington has a separate statistics department. I was kind of worried at times about what it would be like to come here. But I find that our group functions very well, and everybody gets along very well.

We all have somewhat related but very different interests, which is basically what you want to have. Not everybody is doing exactly the same thing, but we all know roughly what everybody else is doing. I don't interact as much with the mathematics faculty or applied math faculty on curriculum-related issues, but I can say that I like the fact that every one of the faculty members in our statistics group is very much oriented towards teaching statistics. I think that's one of the reasons that our programs are so good.

In terms of Montreal, I love the fact that it's a big city, I love the fact that it's on the east coast. The first winter was rough, but that was mostly because I didn't know how to dress. I honestly still prefer the winters in Seattle because it was always nice to be able to play soccer outside in January in your shorts, but it is nice to see the sun in November through March, which we did not get so much in Seattle. So I find that the weather here is not so bad.

Everybody plays up the whole "European side" of Montreal, which I don't buy into as much, but I will say that the bilingual stuff that goes on, with the anglophone and francophone cultural stuff plus the high level of diversity and immigration definitely makes this a much more interesting place to live than many places in North America.

DE: How did you first become interested in statistics?

I had one nine-week period during my 11th-grade math

course where we did a little bit of very basic statistics; mean, median, standard deviation, the normal distribution, and a bit of combinatorics/probability stuff, and I really liked it. As we were doing that section, I was trying to think about what I would do. I was kind of interested in but honestly not that good at math, but I was pretty good and I liked math a lot more than everybody else did. So I kind of figured that that would be a reasonable thing to try to look for a job in.

It was just lucky that I had these nine weeks of statistics. I read a lot more about it, and so when I graduated high school, I was kind of interested in becoming an actuary. I'd read about it, it seemed like they made a lot of money, and they always are ranked high in terms of job satisfaction. So when I went to Carnegie Mellon, I intended to study statistics.

The first or second class that I took at Carnegie Mellon, I did an introductory statistics course with Joel Greenhouse, who is editor of *Statistics in Medicine* and does a lot of biostatistical stuff. Joel came down the stairs and started writing on the board about this problem he was working on with his medical collaborators, and I basically thought that this guy had the best job I had ever seen. It was obvious that he had just gotten to work, he was very relaxed, he was very excited about what he was doing, and I was hooked. From then on I knew that was what I wanted to do.

DE: What's your favorite probability distribution?

I guess I would have to say, just off the top of my head, the Beta. The Beta is very flexible, it's very useful in a lot of contexts. As a Bayesian statistician, I love the Beta because it's a conjugate prior for the Binomial distribution and it turns out that a lot of things have a Beta distribution, and it turns up in a lot of places where you wouldn't expect. But it's a nice distribution, it lives in $[0, 1]$, but at the same time, with the shape and scale parameters you can get a lot of different looking variations.

DE: What are you currently researching?

I do a lot of different things, but most of my methodological research time right now is being devoted to missing data problems. The big thing that I'm really interested in is what people should do whenever there are missing values in the data set that they're collecting.

This tends to be a very difficult problem. Many times in statistics, we have things like the central limit theorem, and we talk about having i.i.d observations, and about maximum likelihood estimators, and assume regularity about the problem. But missing data basically messes everything up. It wreaks havoc both when it's missing for reasons you don't care about, and when it's missing for reasons that you do care about. It can be difficult to tell the difference between the two and determine how sensitive your model conditions are. It's really a good intersection of being a very hard

problem but being at the same time something that every researcher has to deal with.

In terms of where in missing data I'm most interested in working, it's kind of at the intersection of hard statistical methodology but trying to boil it into something that people would actually use. You can come up with all sorts of complicated things for missing data, but the problem is that 99% of people who actually analyze data are unable to do them. So you have to come up with solutions that give correct answers that people will actually use.

Probably the hottest topic I've got going right now in my collaborative work is in scleroderma research. It's a rheumatic disease, in the same class as rheumatoid arthritis or lupus. It's autoimmune, with all sorts of stuff going on in the body. What I love about the disease, from a research perspective, is that no-one knows what is going on. People don't know very much about what causes it, and it's very hard to characterize.

One of the things as a collaborative statistician is that there's nothing better than a situation where people know almost nothing. So they're collecting lots of data, but we're able to do a lot of interesting things because people don't have a firm preconceived notion of how to analyze the data, because they don't even know the different biological mechanisms that are involved. So we get to do a lot of new stuff in those particular problems because we don't have to analyze our data in the same way that other people have analyzed their data in order to get it published. We can do whatever we want and throw super-advanced methods at the problems.

DE: What are your favorite and least favorite parts of research?

The least favorite part, let's be honest, the hardest part of research is the review process. I'm willing to admit in print that I'm a pretty sensitive guy. So when you spend an enormous amount of time on a paper and you get a rough review back, and in statistics in particular the reviews can be a bit rough, it can be very difficult sometimes.

The review process in statistics is a bit longer than in many other disciplines, so what happens is that you work on something really hard, and you send it off. And three to six months later you get back a blistering review of someone who really doesn't like what you've done, and asks you to redo it in a different way. But of course the time lag is such that you've already forgotten what you've done, so you have to go back, remember what it is that you've done, get back into it, make the changes that they've asked for, and send it back and see what happens. So I find that this is the toughest part of research.

The best part by far is when you "see the gap." So there's a moment when you're doing research, and you're trying to solve a particular problem, and you've read as much of the stuff as you can read about the problem. And then you have his moment where you realize that the idea

that you've had has never been had before.

Even publishing the paper is not as exciting to me as that moment. Although, I don't know if that's because after the review process you're so tired at the end that it's just a relief to get it published, but I think that's not it. I think that once you have the idea, and you realize that it's a good idea, it's like a drug. That's what keeps you going back, I think, is to just keep having these little ideas and keep seeing these ways to improve what people are doing. And that, to me, is probably the best part.

DE: When you were an undergraduate, what were your goals? Did you see yourself becoming a university professor?

Unfortunately yes. I wish that I could give better guidance to students but, like I said, I had that moment at the beginning of Joel Greenhouse's class, and that was it for me. He basically had the job that I wanted, and I was very lucky to have people like Joel Greenhouse as extremely good mentors at Carnegie Mellon. They really kind of helped to move me along the path toward being a professor.

I remember there were a couple of times that I wasn't quite sure. Everybody reaches a point, around their second to last year of undergraduate, where it's just your most brutal year and you think "there's no way I could put in for another five years of this after I finish my degree." But really it was only a momentary lapse of faith, and by the time I was through my most brutal part of my junior year at Carnegie Mellon, I was really sure that being a professor was what I wanted to do.

Also, this was confirmed to me in my undergraduate research with Steve Fienberg at Carnegie Mellon. I really realized that I loved to do research, and I loved to work on problems, to read papers, and learn about new things outside of my classes, so I think that that really helped as well.

DE: What advice do you have for undergraduates looking to go into statistics?

I think the best advice I could give them is to be patient. One of the things that can be frustrating about statistics is that it is a very dichotomous world of mathematical and clean problems as well as data analysis and real problems that are very messy. It's, in my opinion, impossible for someone to

be the best at both.

I think particularly for math students it's difficult to be thrown into a real statistical problem and realize that there's not a right answer. There are wrong answers, but there's not one right way to approach the problem. I think that they're comfortable with the fact that there is an answer and no matter how long they bang their head against the wall, there is an answer out there, because the professor would not have asked for the proof if it were not true. Whereas with statistics, there's not that certainty that there's an answer at the end, and I think that sometimes math students can get a bit frustrated by this.

And simultaneously, some people are really gifted at applied statistics because they're good at this grey area, and wonderfully good at constructing models. But at the same time they need to be patient because it's important that you have at least a certain level of comfort with the mathematical machinery because in the end you need to know how this stuff that you're using works.

That's what really separates a statistician from someone in another discipline, is knowing how this machinery really works, because in the event that something does not work, you need to be able to identify what went wrong, which is really difficult to get without a background in at least some mathematics. In fact, I wish that I had more courses in mathematics when I was an undergrad.

Another part of the patience is I think that it's hard to have a really good GPA as a statistics student. Again, because of this issue that it's really hard to be really good at both sides of it, so you have to accept that maybe you're not going to get A's in all of your courses. This says nothing about you as a person except that maybe there are some things that you are better at than others.

One of the nice things about statistics is that the competition, in my mind, is much lower than, say, for medical school. And so no-one sees having a lower GPA as a weakness when you're out in the graduate statistics market. As long as you have a 3.5, someone will seriously consider you at even the best graduate schools. So that's something that students need to think about, because I think sometimes they fear taking courses outside of their comfort zone, but that's the worst thing you can do because then you're missing things that you need.

JOKES

Why don't statisticians like to model new clothes? Lack of fit.

Did you hear about the statistician who was thrown in jail? He now has zero degrees of freedom.

The last few available graves in a cemetery are called residual plots.

Relations between pure and applied mathematicians are based on trust and understanding. Namely, pure mathematicians do not trust applied mathematicians, and applied mathematicians do not understand pure mathematicians.

Rencontre NUMBERS

Maya Kaczorowski

How many ways can we permute N points such that no points remain fixed in the permutation? Furthermore, how many permutations of N points exist such that n points are fixed? While seemingly simple, the solution is not obvious, so we discuss several ways of calculating this. Further, we examine the interesting behaviour with larger N .

1 MOTIVATION

This problem is first attributed to Montmort, who wondered how many matching cards could be found in two mixed-up decks if we matched each card to its corresponding one in the other deck [4]. Other examples include the possibilities of correctly placing pages of a manuscript back in order after it was blown apart by the wind; and if we took a random substitution cipher, if any letters would be matched to themselves in the code.

2 TERMINOLOGY AND SIMPLEST CASES

A *rencontre* number is the number of permutations such that out of N points, n remain fixed. We will denote the *rencontre* number for N and n by $D_{(N,n)}$. We call a *dérangement* a permutation for a particular N and n in which no points remain fixed (sometimes this is called a *subfactorial* [2]). In our notation, there are then $D_{(N,0)}$ *dérangements* for a particular N .

Let's start with the simplest cases. We label the points alphabetically, and bold all fixed points so that it is easier to refer to a particular permutation. When $N = 0$, there is only one way to fix $n = 0$ points in a permutation (recall that $n = 0$ points fixed means that no points are fixed). When $N = 1$, there is no way to fix $n = 0$ points, but there is one way to fix $n = 1$ points, the permutation **a**.

When $N = 2$, there are only two possible permutations of these points, the permutations **ab** and **ba**. There is one way to fix $n = 0$ points and one way to fix $n = 2$ points. Note that there is no way to fix $n = 1$ points (or else we must fix all of them).

When $N = 3$, for a,b,c, we have the permutations

abc **acb** **bca**
 cba **cab**
 bac

Here, we find that there is still only one way to fix $N = n = 3$ points and no way to fix $N = N - 1 = n = 2$ points. This will clearly always be the case, as the only permutation fixing all points is the original permutation, so there is only one way; and it is impossible to fix $N - 1$ points without implying that the N th point must also be fixed.

We find that there are 3 ways to fix $n = 1$ points, which

can be explained probabilistically as:

$$\binom{3}{1} \cdot D_{(2,0)}$$

This is the probability of “choosing” one of the three points to fix, and rearranging the remaining points in any way so long as they are not fixed, i.e. they can be chosen to be any *dérangements*.

That leaves the remaining permutations as *dérangements*, so there are two *dérangements* for $N = 3$, i.e. $D_{(3,0)} = 2$.

The last relatively easy case is for $N = 4$, where we have the permutations:

abcd	abdc	acdb	bcda
	adcb	adbc	bdac
	acbd	cbda	badc
	dbca	dbac	cdab
	cbad	bdca	cdba
	bacd	dacb	cadb
		bcad	dabc
		cabd	dcba
			dcab

And so we find:

$$\begin{aligned} D_{(4,4)} &= \binom{4}{4} = 1 \\ D_{(4,3)} &= 0 \\ D_{(4,2)} &= \binom{4}{2} \cdot D_{(2,0)} = 6 \cdot 1 = 6 \\ D_{(4,1)} &= \binom{4}{1} \cdot D_{(3,0)} = 4 \cdot 2 = 8 \\ D_{(4,0)} &= 9 \end{aligned}$$

It is clear that there exists a recursive relation between the permutations with no fixed points for N and those with one fixed point for $N + 1$. You can choose the fixed point $\binom{N+1}{1} = N + 1$ ways, then the remaining N points must have no fixed points. This relation occurs for any $D_{(N,n)}$ and $D_{(N-n,0)}$. Generally,

$$D_{(N,n)} = \binom{N}{n} D_{(N-n,0)}$$

And so the first few *rencontre* numbers are

$N \setminus n$	0	1	2	3	4
0	1				
1	0	1			
2	1	0	1		
3	2	3	0	1	
4	9	8	6	0	1

3 CALCULATING THE NUMBER OF *dérangements* $D_{(N,0)}$

It is still impractical to find the number of *dérangements* $D_{(N,0)}$ by calculating all possible permutations of a set of points. However, recall that the total number of unique permutations of N points is $N!$, where the first point is chosen out of N possibilities, the second out of $N - 1$ possibilities, etc. Then, we can define the relation:

$$\begin{aligned} D_{(N,0)} &= N! - \sum_{i=1}^N D_{(N,i)} \\ &= N! - \sum_{i=1}^N \binom{N}{i} D_{(N-i,0)} \end{aligned}$$

There are other methods of calculating these *dérangements*. For instance, we could use alternating sums. Suppose we wanted to count the number of permutations of N elements having at least one point fixed. We first count the permutations fixing just one point: we choose the point that gets fixed, $\binom{N}{1}$ ways, and then rearrange the other points $(N - 1)!$ ways, so there are $\binom{N}{1}(N - 1)!$ such permutations. But we have over-counted, as the permutation fixing '1' and '2' is counted under both those fixing '1' and those fixing '2'. And so we end up with the following alternating sum equation [1,3]:

$$\begin{aligned} N! - D_{(N,0)} &= \text{\#fixing 1} - \text{\#fixing 2} + \\ &\quad \text{\#fixing 3} - \dots + (-1)^{N-1} \text{\#fixing N} \\ &= \binom{N}{1}(N - 1)! - \binom{N}{2}(N - 2)! + \\ &\quad \dots + (-1)^{N-1} \binom{N}{N}(N - N)! \\ &= \sum_{i=1}^N (-1)^{i-1} \binom{N}{i} (N - i)! \\ D_{(N,0)} &= N! + \sum_{i=1}^N (-1)^i \binom{N}{i} (N - i)! \\ D_{(N,0)} &= \sum_{i=0}^N (-1)^i \binom{N}{i} (N - i)! \end{aligned}$$

This is our first deterministic equation. Note that this

can also be written as

$$\begin{aligned} D_{(N,0)} &= \sum_{i=0}^N (-1)^i \binom{N}{i} (N - i)! \\ &= \sum_{i=0}^N (N - i)! (-1)^i \binom{N}{i} \\ &= \sum_{i=0}^N i! (-1)^{N-i} \binom{N}{i} \end{aligned}$$

Indeed, this equation holds; we verify it for $N = 4$ and $n = 0$, and we find:

$$\begin{aligned} D_{(4,0)} &= 4! - \binom{4}{1}(4 - 1)! + \binom{4}{2}(4 - 2)! \\ &\quad - \binom{4}{3}(4 - 3)! + \binom{4}{4}(4 - 4)! \\ &= 4! - 4! + 6 \cdot 2 - 4 \cdot 1 + 1 \\ &= 12 - 4 + 1 \\ &= 9 \text{ as we previously obtained} \end{aligned}$$

We could also calculate the number of *dérangements* by looking at the number of elements fixed in a permutation when it is written in cyclic form [1]. Recall that all permutations can be written as a unique product of cycles. So we only need to count the number of products of disjoint cycles of length of at least 2.

For example, if we wanted once again to calculate the number of permutations with no fixed points for 4 elements, then we would consider only the possibility of two 2-cycles, or of one 4-cycle, else some elements are fixed. If we have two 2-cycles, then there are 4 choices for the first element and 3 for the second. But these could be written (ab) or (ba), so we must divide by 2. Similarly for the second cycle, the third element has two choices and the choice of the last element is forced, and these can be permuted 2 ways. However, we could also permute the two cycles, as (ab)(cd) and (cd)(ab), and so must once again divide by 2. So for two 2-cycles, only 3 leave no elements fixed:

$$\frac{(4 \cdot 3)}{2} \cdot \frac{(2 \cdot 1)}{2} \cdot \frac{1}{2} = 3$$

For one 4-cycle, there are four ways to choose the first, three to choose the second, etc. But these elements can be written 4 ways to represent the same cycle, i.e. (abcd), (bcda), (cdab), and (dabc); so we divide by 4. Then there are 6 ways to leave no elements fixed in a 4-cycle:

$$\frac{(4 \cdot 3 \cdot 2 \cdot 1)}{4} = 6$$

Then we once again find that for 4 elements, there are $3 + 6 = 9$ permutations with no fixed points. However, this is not a very practical method to go about counting *dérangements*, as it first requires the calculation of all possible decompositions into cycles of length at least 2, which is much more complicated for large N .

$N \setminus n$	0	1	2	3	4	5	6	7	8	9	10
0	1										
1	0	1									
2	1	0	1								
3	2	3	0	1							
4	9	8	6	0	1						
5	44	45	20	10	0	1					
6	265	264	135	40	15	0	1				
7	1854	1855	924	315	70	21	0	1			
8	14833	14832	7420	2464	630	112	28	0	1		
9	133496	133497	66744	22260	5544	1134	168	36	0	1	
10	1334961	1334960	667485	222480	55650	11088	1890	240	45	0	1

Table 1: A larger table of *rencontre* numbers, taken from [4], p.65

4 APPROXIMATIONS FOR LARGE N AND APPLICATIONS

From our deterministic formula,

$$\begin{aligned}
 D_{(N,0)} &= \sum_{i=0}^N (N-i)! (-1)^i \binom{N}{i} \\
 &= \sum_{i=0}^N (-1)^i \frac{N!}{i!} \\
 &= N! \sum_{i=0}^N \frac{(-1)^i}{i!}
 \end{aligned}$$

But recall the power series for e is

$$e^x = \sum_{i=0}^{\infty} \frac{x^i}{i!}$$

And so we take $x = -1$ to find [3]:

$$\begin{aligned}
 D_{(N,0)} &\approx \frac{N!}{e} \text{ for large } N \\
 D_{(N,0)} &= \left\lfloor \frac{N!}{e} + \frac{1}{2} \right\rfloor \text{ the nearest integer}
 \end{aligned}$$

And so, if we have a large N , approximately $\frac{N!/e}{N!} = \frac{1}{e} \approx \frac{1}{3}$ of all permutations will have no fixed elements.

This question, about the number of fixed elements in a permutation, is particularly applicable in cryptography. It is important that a key used in an encryption be ‘random enough’, that is to say, the key is random enough that nothing additional about the message could be determined without knowledge of the key. For example, consider that we have a simple substitution cipher amongst letters of the alphabet, so that every letter of the alphabet is replaced by an-

other letter of the alphabet. Without considering the methods cryptanalysts would usually try to break the cipher, ideally the cipher would have no fixed letters, or as few as possible, to give away no information about the message without decoding it. Cryptographers could then use *rencontre* numbers to determine the probability that a random permutation of the letters created a cipher with 0 letters fixed, 1 letter fixed, etc.

Although other results exist concerning the recurrence relationships between *dérangements* that were not mentioned in this article, they do not aid in the calculations of *rencontre* numbers. The deterministic equations we have provided above and the behaviour described for large N are enough to calculate *rencontre* numbers and *dérangements* in most situations.

REFERENCES

[1] Paul Garrett. Overheads on fixed point of permutations, crypto and number theory course. University of Minnesota, online at http://www.math.umn.edu/~garrett/crypto/Overheads/06_perms_otp.pdf, January 2007.

[2] Bernd C. Kellner. Some remarks on kurepa’s left factorial. Preprint. arXiv:math.NT/0410477, 2009.

[3] Ken Ribet. Discrete mathematics class notes, mathematics 55. Class Notes at <http://math.berkeley.edu/~ribet/Math55/>, Fall 1997.

[4] John Riordan. *Introduction to Combinatorial Analysis*. Wiley, New York, 1958.

EMBEDDINGS OF AFFINE LINES IN THE AFFINE PLANE AND POLYNOMIAL AUTOMORPHISMS

Hua Long Gervais

In affine algebraic geometry, one defines an affine line to be an algebraic variety whose coordinate ring is a polynomial ring in one variable. If one considers affine plane curves, then it turns out that the question of whether a curve is a line or not is related to the characterization of all polynomial automorphisms of the ring $k[x, y]$, which is completely solved by the Jung-van der Kulk theorem. We will give a very terse outline of this connection.

1 BACKGROUND NOTIONS FROM ALGEBRAIC GEOMETRY

In affine algebraic geometry, an affine line is defined as an algebraic variety with coordinate ring isomorphic to a polynomial ring in one indeterminate. Naturally, one is interested in finding concrete examples of such an object, and we address this issue by trying to understand what subsets of the affine plane are affine lines. It turns out this question can be answered by describing the automorphisms of a polynomial ring in two variables, a task that is accomplished by the Jung-van der Kulk theorem. We will start by briefly presenting the background material.

Although we will only consider the affine plane and the projective plane, it is of no harm to give the general definition here.

Definition. Let k be a field and define $\mathbb{A}^n(k) = \{(a_1, \dots, a_n) \mid a_i \in k \forall i = 1, \dots, n\}$. We call $\mathbb{A}^n(k)$ affine n -space over k .

In particular, $\mathbb{A}^2(k)$ is called the affine plane. Of course, $\mathbb{A}^n(k)$ can be identified with the n -dimensional vector space k^n . This allows us to formulate the next definition.

Definition. Let k be a field, the projective n -space (over k) $\mathbb{P}^n(k)$ is the set of all one-dimensional subspaces of $\mathbb{A}^{n+1}(k)$.

It is often more convenient to identify projective space with the set of all $(n+1)$ -tuples (x_1, \dots, x_{n+1}) where not all x_i are zero and we identify (x_1, \dots, x_{n+1}) with (y_1, \dots, y_{n+1}) if there exists $\lambda \in k^*$ such that $(x_1, \dots, x_{n+1}) = \lambda(y_1, \dots, y_{n+1})$. Since projective space is then seen as a set of equivalence classes, we denote its “points” by $[x_1 : \dots : x_{n+1}]$. We can view $\mathbb{A}^n(k)$ as a subset of $\mathbb{P}^n(k)$ by considering the set U_i of all points such that $x_i \neq 0$. We have $[x_1 : \dots : x_i : \dots : x_{n+1}] = [\frac{x_1}{x_i} : \dots : 1 : \dots : \frac{x_{n+1}}{x_i}]$ and $U_i = \{[x_1 : \dots : 1 : \dots : x_{n+1}] \mid x_j \in k \forall j \neq i\}$, from which the bijection $U_i \rightarrow \mathbb{A}^n$ is obvious. In fact, we see that projective n -space is covered by n affine pieces U_i , $i = 1, \dots, n$.

When $n = 2$, we have the projective plane, and from now on, we will only deal with the affine plane and the projective plane, which we will simply denote by \mathbb{A}^2 and \mathbb{P}^2 where the field k will be understood from the context.

Definition. Let k be a field and let $F(X, Y) \in k[X, Y]$. A point $P = (x, y) \in \mathbb{A}^2$ is a zero of F if $F(x, y) = 0$, which

we write more succinctly as $F(P) = 0$. If $I \subset k[X, Y]$, we let $V(I) = \{P \in \mathbb{A}^2 \mid F(P) = 0 \forall F \in I\}$ and call it the zero set of the subset I . We write $V(F)$ for the zero set of the ideal generated by a single polynomial F . Conversely, given $S \subset \mathbb{A}^2$, we let $I(S) = \{F \in k[X, Y] \mid F(P) = 0 \forall P \in S\}$. A subset S of \mathbb{A}^2 is called an affine algebraic set if $S = V(I)$ for some subset I of $k[X, Y]$.

We have analogous definitions for the projective plane, but care must be taken in defining a zero of a polynomial. Let $P = [x : y : z] \in \mathbb{P}^2$ and let $F \in k[X, Y, Z]$, then we would want to define P to be a zero of F if $F(x, y, z) = 0$, but we have to make sure that F also vanishes on all the representatives of P , that is, that $F(\lambda x, \lambda y, \lambda z) = 0$ for all $\lambda \in k^*$. This is what we adopt as our definition of a zero of F .

Let us focus on the affine plane for a moment. It is straightforward to verify the following properties of zero sets:

1. If I is the ideal of $k[X, Y]$ generated by S , then $V(I) = V(S)$. So a zero set is always the zero set of an ideal $I \subset k[X, Y]$.
2. If $\{I_\alpha\}$ is any collection of ideals, then $V(\bigcup_\alpha I_\alpha) = \bigcap_\alpha V(I_\alpha)$.
3. If $I \subset J$, then $V(I) \supset V(J)$.
4. $V(I) \cup V(J) = V(\{FG \mid F \in I, G \in J\})$.
5. $V(0) = \mathbb{A}^2$, $V(1) = \emptyset$, and $V(X - a, Y - b) = \{(a, b)\}$.

We note in passing that the above show that we can define a topology on the affine plane by declaring the closed sets to be those subsets that are the zero set of some ideal of $k[X, Y]$. This topology is called the Zariski topology, but do note that we have considered it in a very restrictive case.

We also have the following easy properties of ideals of sets of points:

1. If $X \subset Y$, the $I(X) \supset I(Y)$.
2. $I(\emptyset) = k[X, Y]$ and $I(\mathbb{A}^2) = (0)$ if k is an infinite field.
3. $I(V(S)) \supset S$ for any set S of polynomials, and $V(I(X)) \supset X$ for any set X of points.
4. If V is an algebraic set, then $V = V(I(V))$. If I is the ideal of an algebraic set, then $I = I(V(I))$.

We need an extra definition to be able to say more.

Definition. Let R be a commutative ring, and let $I \subset R$ be an ideal. The radical of I is the set $Rad(I) = \{a \in R \mid a^n \in I \text{ for some integer } n\}$. Then $Rad(I)$ is an ideal. If $I = Rad(I)$, then I is called a radical ideal.

Here are more properties:

1. $I(X)$ is a radical ideal for all sets of points X .
2. For any ideal I of $k[X, Y]$, $V(I) = V(Rad(I))$ and $Rad(I) \subset I(V(I))$

An algebraic set is called irreducible if it cannot be written as the union of two smaller algebraic sets. Using the fact that $k[X, Y]$ is a Noetherian ring, it is possible to show that any algebraic set can be written uniquely up to the order of the components as a union of irreducible algebraic sets. Moreover, an algebraic set is irreducible if and only if its ideal is prime. We shall call an irreducible algebraic set an affine variety.

We put an end to this affine algebraic sets melodrama with a very important theorem in commutative algebra:

Theorem 1 (The Hilbert Nullstellensatz in two variables). Let k be an algebraically closed field and let $I \subset k[X, Y]$ be an ideal, then $I(V(I)) = Rad(I)$.

This allows us to establish a one-to-one correspondence between radical ideals and algebraic sets:

$$\begin{aligned} V &\longmapsto I(V) \\ I &\longmapsto V(I) \end{aligned}$$

The basic properties also hold for projective algebraic sets and ideals of subsets of the projective plane. The Nullstellensatz also holds provided we are careful about excluding the ideal (X, Y, Z) because of the exclusion of the point $(0, 0, 0)$ from the projective plane. We will call irreducible projective algebraic sets projective varieties. We shall not delve further into those details.

As final remark, we mention that the notion of an algebraic set can be generalized to multiprojective spaces $\mathbb{P}^{n_1} \times \dots \times \mathbb{P}^{n_r} \times \mathbb{A}^m$. One then defines the Zariski topology on this space by declaring the closed sets to be algebraic sets. A quasiprojective variety is an open subset of an irreducible algebraic set. This is overkill for our purposes however. There are even further generalizations of this in modern algebraic geometry, and these are definitely major overkill.

The interested reader can look at [1] for an excellent systematic treatment of this material.

2 AFFINE LINES IN THE PLANE

Definition. Let $V \subset \mathbb{A}^2$ be an affine variety, then the coordinate ring of V is the quotient ring $\Gamma(V) = k[X, Y]/I(V)$.

Definition. An affine plane curve is subset of \mathbb{A}^2 of the form $V(F)$ for some $F \in k[X, Y]$.

In the theory of plane curves, F is typically allowed to have multiple components; in practice, this just means that F need not be an irreducible polynomial. See figure 1 for examples.




Figure 1: a) $F(X, Y) = (X + Y)(X - Y)$ has multiple components b) $F(X, Y) = Y^2 - X^3 - X^2$ is irreducible

Also, F and λF , $\lambda \in k^*$ are considered as the same curve.

Definition. An affine line is an algebraic variety X whose coordinate ring $\Gamma(X)$ is a polynomial ring in one variable, that is: $\Gamma(X) \simeq k[T]$, where T is some indeterminate.

Although the above definition refers to the general notion of an algebraic variety and its coordinate ring, restricting it to the affine varieties we have considered so far is enough for our purposes.

We wanted to understand which subsets of the affine plane are lines. So let $C = V(F)$ be an irreducible affine plane curve, then our problem is to determine when there exists a $t \in k[X, Y]/(F(X, Y))$ such that $\Gamma(C) = k[X, Y]/(F(X, Y)) = k[t]$ (we will use capital letters to denote elements of $k[X, Y]$ and lower case letters to denote their residues modulo F). This questions turns out to be related to considerations about automorphisms of $k[X, Y]$.

Let

$$\begin{aligned} \alpha : k[X, Y] &\longrightarrow k[X, Y] \\ X &\longmapsto F_1(X, Y) \\ Y &\longmapsto F_2(X, Y) \end{aligned}$$

be an automorphism. We say that (F_1, F_2) is an automorphic pair. Now, suppose that the F of the previous paragraph is part of an automorphic pair (F, G) arising from an automorphism α . Then, $k[X, Y] = k[F, G]$, that is, $X = \varphi(F, G)$ and $Y = \phi(F, G)$ for some polynomials φ, ϕ in $k[X, Y]$. Therefore, $\Gamma(C) = k[X, Y]/(F) = k[F, G]/(F) = k[g]$. So it is enough for F to be part of some automorphic pair in order for C to be a line. Naturally, one then asks about the converse: if C is a line, then is F part of an automorphic pair?

It was proven in 1971 that the answer depends only on the characteristic of k . If $char(k) = 0$, then that $C = V(F)$ is a line does guarantee that we can find a G such that (F, G) is an automorphic pair. If $char(k) > 0$, then the answer is no in general.

We know there is a bijective correspondence between automorphisms of $k[X, Y]$ and affine lines in characteristic zero, so in that case our problem is reduced to finding out how to construct all the automorphisms of $k[X, Y]$. The answer to this problem is given by the Jung-van der Kulk theorem.

3 THE JUNG-VAN DER KULK THEOREM

Let $GA_2(k) = \text{Aut}_k(k[X, Y])$ and consider the following particular elements of $GA_2(k)$:

$$\begin{aligned} \alpha_{a,b,c,d,e,f} : k[X, Y] &\longrightarrow k[X, Y] \\ X &\longmapsto aX + bY + e \\ Y &\longmapsto cX + dY + f \\ a, b, c, d, e, f \in k \quad \det \begin{pmatrix} a & b \\ c & d \end{pmatrix} &\neq 0 \\ \beta_F : k[X, Y] &\longrightarrow k[X, Y] \\ X &\longrightarrow X \\ Y &\longrightarrow Y + F(X) \\ F &\in k[X] \end{aligned}$$

Let $T_2(k)$ be the subgroup of $\text{Aut}_k(k[X, Y])$ generated by all automorphisms of the above form, $T_2(k)$ is called the subgroup of tame automorphisms.

With these notions, we can now state the Jung-van der Kulk theorem:

Theorem 2. Let k be any field, then $GA_2(k) = T_2(k)$

This says that all automorphisms of $k[X, Y]$ can be obtained by composing tame automorphisms repeatedly. Notice that in two variables, the theorem holds for k having any characteristic.

This theorem was first proved by Jung in 1942 for fields of characteristic zero and then by van der Kulk in 1953 for arbitrary characteristic. Several different proofs have been given afterwards, and a full one appears in chapter 1 of [2]. We will only briefly describe the progression of one possible approach to prove this theorem.

The important technique consists in blowing up points in the affine plane. Roughly speaking, we take a point in \mathbb{A}^2 , remove it from the plane, and replace it by a projective line whose points represent all the possible tangent directions at the removed point. The result is an algebraic variety (in the general sense) that can be covered by two pieces isomorphic (as algebraic varieties) to the affine plane.

More precisely, the blow-up of \mathbb{A}^2 is

$$B = \{(x, y, [x, y]) \mid (x, y) \neq (0, 0)\} \cup \{(0, 0, [x, y]) \mid [x, y] \in \mathbb{P}^1\} \subset \mathbb{A}^2 \times \mathbb{P}^1$$

We have the morphism (think of it as a sort of homomorphism for algebraic varieties):

$$\begin{aligned} \pi : B &\longrightarrow \mathbb{A}^2 \\ (x, y, [x, y]) &\longmapsto (x, y) \end{aligned}$$

Note that $\pi|_{B \setminus (0, 0, \mathbb{P}^2)} : B \setminus (0, 0, \mathbb{P}^2) \longrightarrow \mathbb{A}^2 \setminus \{(0, 0)\}$ is a bijection and $\pi^{-1}((0, 0)) = (0, 0, \mathbb{P}^2)$. These express the fact that the blow-up of \mathbb{A}^2 consists of an affine plane with the origin replaced by a projective line with one point for each direction.

To describe the use of blow-ups in proving the Jung-van der Kulk theorem, we need to define what the singularities of a curve are. Let F be an affine plane curve, then F can be written uniquely as $F = F_m + F_{m+1} + \dots + F_n$ where F_i is a homogeneous polynomial of degree i and the F_i are arranged in order of increasing degree. Then $(0, 0)$ is a singularity of F if and only if $m > 1$, or equivalently if $\frac{\partial F}{\partial X}(0, 0) = \frac{\partial F}{\partial Y}(0, 0) = 0$. The degree m is called the multiplicity of the singularity. F has a singularity at an arbitrary point P if, given a translation T taking $(0, 0)$ to P , $(0, 0)$ is a singularity of $F \circ T$, and its multiplicity is defined as the number m appearing in the decomposition of $F \circ T$ as a sum of homogeneous polynomials as above. It is a well-known fact that a homogeneous polynomial in two variables factors into linear homogeneous polynomials. The factors of F_m are the tangent lines of F , and the multiplicity of a tangent line is defined as its exponent in the factorization of F_m . At a singular point, F will either have several distinct tangent lines of multiplicity 1, or tangent lines of multiplicity greater than 1. For example, in the following figure, curve a) has a double tangent at the origin and curve b) has two single tangents at the origin.




Figure 2: a) $F(X, Y) = Y^2 - X^3$ b) $F(X, Y) = Y^2 - X^3 - X^2$

Now, given a curve C in the affine plane passing through the origin, one considers the set $\pi^{-1}(C \setminus \{(0, 0)\}) \subset B$ and takes its closure in B in the Zariski topology sense, that is, the smallest open subvariety of B containing $\pi^{-1}(C \setminus \{(0, 0)\})$. This curve in B is called the blow-up of C at the origin. If C had a singular point at the origin, then its blow-up has better singular points in the sense that the singularity at the origin has split into several singular points of lower multiplicities. Intuitively, this is because the fiber above the origin consists of those points on the projective line above $(0, 0)$ which correspond to the tangent directions to C at the

origin. Now, each of those singularities lies in one of the affine plane pieces that cover B . So the natural thing to do is to blow up each of those affine pieces at the singular points of the blown up curve to get yet another blown up curve and so on. Repeating this ultimately yields a non-singular curve, called the non-singular model of the original curve C .

Back to the Jung-van der Kulk theorem, we let $F = l.o.t. + F_n$ and $G = l.o.t. + G_m$ be any automorphic pair, where F_n and G_m are the homogeneous polynomials of highest degree of F and G respectively, and “l.o.t.” means “lower order terms.” Then F and G can meet only once in the affine plane. Let us give a quick argument for this.

Given any affine varieties V and W , we have a one-to-one correspondence between homomorphisms of their coordinate rings and polynomial maps from W to V :

$$\alpha : \Gamma(V) \longrightarrow \Gamma(W) \longleftarrow \alpha^* : W \longrightarrow V$$

where $\alpha^*(P) = (\alpha(x)(P), \alpha(y)(P))$. Moreover, α^* is an isomorphism (i.e. is bijective with inverse a polynomial map) if and only if α is an isomorphism.

In our case, we have $V = W = \mathbb{A}^2$ and $\Gamma(V) = \Gamma(W) = k[X, Y]$. To the automorphism $X \mapsto F(X, Y), Y \mapsto G(X, Y)$, we therefore associate the isomorphism of the affine plane with itself $P \mapsto (F(P), G(P))$. Suppose now that F and G meet at two different points P and Q in \mathbb{A}^2 , both P and Q are mapped to the origin by our associated isomorphism of affine varieties, which contradicts the fact that it must be injective. So F and G only meet at one point in the affine plane and with a suitable change of coordinates, we might as well assume it is the origin.

The rest of the proof goes through by induction on $m + n$. We embed \mathbb{A}^2 in \mathbb{P}^2 and consider the projective closures of F and G . Then by Bezout’s theorem, F and G must meet mn times in \mathbb{P}^2 , counted with multiplicities. Hence, they must meet $mn - 1$ times in the hyperplane at infinity $Z = 0$. After dealing with some cases, one concludes that there is actually only one point of intersection with multiplicity $mn - 1$ and therefore, after choosing suitable coordinates, one may assume that $F_n = Y^n$ and $G_m = Y^m$. We then blow up F and G repeatedly until we obtain nonsingular models for both of these curves. From this point on, technical algebraic arguments show that $m|n$ or $n|m$. So we can

reduce the degrees of F if $m|n$ or G if $n|m$, by composing our starting automorphism with a suitable tame automorphism. Since the result is assumed to hold for lower values of $m + n$ and trivially holds for $m + n = 2$, this shows that the composed automorphism is tame, and hence so is the original one.

If we return to the original problem of understanding when a curve is a line, then we have an answer to the characteristic zero case. A curve $C = V(F)$ is a line if it can be obtained by repeatedly applying transformations of the α and β types described above to X or Y .

Naturally, one might ask what happens with automorphisms of polynomial rings in three variables. It turns out not much is known about this. In 2002, Shestakov and Umibaev proved that the subgroup of tame automorphisms of $k[X, Y, Z]$ is properly contained in the whole group of automorphisms for k a field of characteristic 0. More specifically, they showed that the automorphism constructed by Nagata is not tame. It appears even experts consider their proof as not being widely understood. Another problem pertaining to polynomial automorphisms is the famous Jacobian conjecture. Given a polynomial automorphism, it is a necessary condition that its Jacobian matrix is invertible. The Jacobian conjecture states that this condition is in fact sufficient, and despite its apparent simplicity, it remains largely intractable to this day.

4 ACKNOWLEDGEMENTS

I learned about algebraic geometry and polynomial automorphism in the course of a USRA project. I wish to thank Professor Peter Russell for his diligent supervision.

REFERENCES

- [1] W. Fulton. *Algebraic Curves: An Introduction to Algebraic Geometry*. Addison Wesley Publishing Company, Reading, MA, 1998.
- [2] A. Van den Essen, H. Huy Vui, H. Kraft, P. Russell, and D. Wright. *Polynomial Automorphisms and Related Topics*. Publishing House for Science and Technology, Hanoi, Vietnam, 2007.

JOKES

“The world is everywhere dense with idiots.” □

These days, even the most pure and abstract mathematics is in danger to be applied. □

Some mathematicians become so tense these days that they do not go to sleep during seminars. □

A lecturer:

“Now we’ll prove the theorem. In fact I’ll prove it all by myself.” □

WHITE NOISE

Daniel Shapero

In the 1930s and 40s physicists began to use differential equations with stochastic forcing terms to describe systems driven by random fluctuations or with randomly changing parameters. In doing so, their differential equations typically included a term of the form $W_t dt$, where W_t is *white noise*, a random signal having equal power distributed across all frequencies. While these methods were not completely rigorous, in a mere 6-page paper, Kiyoshi Itô, with the critical insight that a differential of the form $W_t dt$ is equal to a differential dB_t where B_t is the Wiener process, gave the theory a systematic foundation. Furthermore, his eponymous formula allowed for the adaptation of the usual solution methods of deterministic differential equations to stochastic ones.

But, in what sense is white noise the derivative of the Wiener process? Since the sample paths of Brownian motion are almost surely nowhere-differentiable, it is immediately clear that some mischief is afoot. That mischief is distribution theory, which we will use extensively, along with some general facts about the Wiener process, to give a proof of Itô's insight.

1 BROWNIAN MOTION AND THE WIENER PROCESS



Figure 1: Robert Brown

The seeds of modern stochastic analysis were sown in the 19th century by botanist Robert Brown, who studied the random and erratic motions of a grain of pollen immersed in water. The realization that these motions were due to neighbouring water molecules' random bombardments was due to none

other than Albert Einstein in his *anno mirabilis* 1905, the same year that he wrote his famous paper on special relativity. Einstein suggested that the average of these collisions via the central limit theorem produces a normal distribution, so that the process is essentially Gaussian.

It was not until the 1930s, after Kolmogorov axiomatized probability theory, that the subject was treated in complete rigour by Norbert Wiener. The Wiener process, also referred to interchangeably as Brownian motion, is a continuous-time stochastic process B_t such that:

- (1) For $s, t \in [0, \infty)$, $B_t - B_s$ is normally distributed with variance $|t - s|$, or $B_t - B_s \sim N(0, |t - s|)$, and $B_0 = 0$ almost surely.
- (2) For $0 < t_1 < \dots < t_k$, the random variables $B_{t_{j+1}} - B_{t_j}$ are independent.
- (3) The map $t \mapsto B_t$ is continuous almost surely.

These properties characterize Brownian motion uniquely, up to almost-sure equality. For a proof, see [3].

We have been somewhat vague in stating (3): almost surely with respect to what probability measure on what sample space? The choice of sample space is somewhat arbitrary, and there is more than one on which such a stochastic process can be defined. The standard is for Ω to be the

space of all continuous functions from $[0, \infty)$ to \mathbb{R} , and the probability measure P , called the *Wiener measure*, is obtained from using Kolmogorov's extension theorem.

Another important property is that the sample paths of Brownian motion are almost surely nowhere-differentiable, a fact which underlies the goal of this entire article. The proof is due to Kakutani, Dvoretzki and Erdős, and can be found in [1]; while it is rather difficult, we cannot in good conscience leave it out and thus will sketch it. If f is a continuous function such that $f'(s)$ exists for some $s \in [0, \infty)$ and $|f'(s)| < \beta$, then there is some δ -neighbourhood of s on which f is Lipschitz continuous, namely

$$|f(t) - f(s)| < 2\beta|t - s|$$

for $t \in (s - \delta, s + \delta)$. Replacing δ with $3/n$ for n natural, we can then look at the collection of all continuous functions which are 2β -Lipschitz on intervals $(\frac{k-1}{n}, \frac{k+2}{n})$ where $k \in \mathbb{N}$; call this collection of functions A_n . Any function which is differentiable at some point must then be contained in every A_n for n greater than some fixed N , so if we want to show that the differentiable functions have probability zero all we must show is that $P(\liminf A_n) = 0$, where P is the probability law of Brownian motion. Using the fact that $B_{t+h} - B_t$ has the same distribution for any h , this reduces, after a slew of inequalities, to the statement that

$$\lim_{n \rightarrow \infty} n \left[\frac{1}{\sqrt{2\pi n}} \int_{-c\beta}^{c\beta} e^{-\frac{x^2}{2n}} \right]^3 = 0,$$

where c is some constant. This implies that the probability that sample paths of Brownian motion have a derivative somewhere which is less in absolute value than some arbitrary β is zero, and taking a countable sequence β_n tending to infinity gives the result.

As a consequence, the sample paths of Brownian motion are almost surely of unbounded variation on any interval. If a sample path of B had bounded variation on some interval, it could then be written as the difference of two monotone functions, and a monotone function is almost-everywhere differentiable, in contradiction to the non-differentiability of paths. This fact would lead one to believe that one cannot

define Riemann-Stieltjes integration with respect to Brownian motion, and yet the Itô integral does precisely that!




Figure 2: two approximate sample paths of the Wiener process on $[0, 1]$

2 GAUSSIAN PROCESSES AND WHITE NOISE

Suppose that X_t is a stochastic process taking values in \mathbb{R} . The process X_t is called *Gaussian* if, for every finite collection t_1, \dots, t_k of real numbers, $(X_{t_1}, \dots, X_{t_k})$ is a jointly Gaussian random variable. The process is called *stationary* in the wide sense if $E(X_t) = m$ is constant in time, and the covariance of the process at two sample points, defined to be

$$\text{Cov}(X_s, X_t) = E[(X_s - m)(X_t - m)],$$

is a function only of $t - s$. We then write

$$\text{Cov}(X_s, X_t) = \text{Cov}_X(t - s).$$

The subscript X will be dropped when it is clear which process we are talking about. While $E(B_t) = 0$ for all t , Brownian motion is not wide-sense stationary because

$$\text{Cov}(B_s, B_t) = E(B_s B_t) = \min\{s, t\} = \frac{s + t - |s - t|}{2}$$

is not a function of $t - s$. From now on we will refer to wide-sense stationary processes simply as stationary, but there are other closely related concepts such as strict stationarity which have different definitions.

What would we imagine a white noise process Z to be like if we had to pin down a definition? If $s \neq t$, we would expect that Z_s and Z_t are independent. It should be equally likely to take positive and negative values, suggesting that $E(Z_t) = 0$. Furthermore, its fluctuations at any given instant of time are very large, so taking them to be infinite is physically reasonable; in other words, $\text{Var}(Z_t) = \infty$. In analogy with white light, if we analyze Z in the frequency domain, we would expect it to have equal energy distributed at every frequency. We should then have that $\widehat{\text{Cov}}_Z = 1$, the Fourier transform of the covariance function of white noise is identically equal to 1. Since the Fourier transform of the Dirac delta is 1, this implies that the covariance function of the process is given by $\text{Cov}_Z(t - s) = \delta(t - s)$. The *spectral density* of a stationary process is defined as the Fourier transform of its covariance function, and it illuminates many properties of the process. The spectral density is widely used in signal processing applications. Our condition above that $\widehat{\text{Cov}}_Z = 1$ states that white noise has flat spectral density.

With these ideas in mind, we can give a bad definition of a white noise process: a stationary, Gaussian stochastic process Z_t with $E(Z_t) = 0$ and $E(Z_s Z_t) = \delta(t - s)$, where δ is the Dirac delta function, is called *white noise*. How could a definition possibly deserve the unabashed description of being bad? For one, there is no such object which satisfies it. The Dirac delta is not a function, only a generalized function. Furthermore, if Z_t is to be Gaussian, then $E(Z_t^2) = \text{Var}(Z_t) = \infty$. Gaussian random variables have finite variance by definition. In fact, there is no such stochastic process Z_t defined in the usual sense, as a family of random variables taking real values.

How do we reconcile the many signals which we would describe as white noise with the nonexistence of white noise according to the definition above? We defined a white noise signal to have flat spectral density across all frequencies. The spectral density of a real signal which we would describe as white noise will be identically equal to 1 across some broad frequency range, but it will then decay to 0. For example, if one examines the velocity of a particle undergoing Brownian motion, at a small enough time scale the collisions with the surrounding fluid are elastic and there is a maximum frequency with which they can occur. But, we do not enforce this decay rate of the Fourier transform in our mathematical model because our observations are so coarsely grained in time and space, far above the scale at which one can discern a frequency cutoff, that for all intents and purposes the Fourier transform is constant. To give another example, suppose you were to hear an audio signal with flat spectral density across the entire range of

your hearing: so far as your ear knows, the signal has flat spectral density across the entire frequency spectrum.

While this definition of white noise is vacuous, in a way we are actually not far off. Our stated goal was to show that the derivative of Brownian motion is white noise. While we cannot take derivatives of Brownian motion, we can take difference quotients and look at what happens in the limit. Following the treatment in [4], consider the stochastic process

$$\Delta_\tau B_t = \frac{B_{t+\tau} - B_t}{\tau},$$

where $\tau > 0$, and evaluate

$$C_\tau(s, t) = \text{Cov}(\Delta_\tau B_s, \Delta_\tau B_t).$$

Since $(B_t, B_{t+\tau})$ is multi-normal then so is $\Delta_\tau B_t$, and $E(\Delta_\tau B_t) = 0$. To evaluate C_τ , suppose that $s < t$, $\tau > 0$.

$$C_\tau(s, t) = \frac{E[B_{s+\tau}(B_{t+\tau} - B_t)] - E[B_s(B_{t+\tau} - B_t)]}{\tau^2}$$

But, $B_{t+\tau} - B_t$ is independent of $B_s = B_s - B_0$, so the second expectation is zero. For the first expectation we use that $\min\{s + \tau, t + \tau\} = s + \tau$, so that

$$\begin{aligned} C_\tau(s, t) &= \tau^{-2}[s + \tau - E(B_{s+\tau}B_t)] \\ &= \tau^{-2}[s + \tau - \min\{s + \tau, t\}]. \end{aligned}$$

We can then pull s out of the minimum to write

$$\begin{aligned} s + \tau - \min\{s + \tau, t\} &= \tau - \min\{\tau, t - s\} \\ &= \tau - \tau \min\left\{1, \frac{t-s}{\tau}\right\} \end{aligned}$$

which gives

$$C_\tau(s, t) = \frac{1 - \min\left\{1, \frac{t-s}{\tau}\right\}}{\tau} = \frac{\max\{1 - \frac{t-s}{\tau}, 0\}}{\tau}.$$

But, we notice that the covariance function of the process $\Delta_\tau B$ can be regarded as a function of a single variable, evaluated at $|t - s|$, which we will do from now on. In particular, $\Delta_\tau B$ is stationary. A more convenient form of the covariance function is

$$\text{Cov}_{\Delta_\tau B}(t) = C_\tau(t) = \frac{1 - \frac{|t|}{\tau}}{\tau} \chi_{[-\tau, \tau]}(t).$$

Now, for each τ , $C_\tau \geq 0$, $\int C_\tau dt = 1$ for all τ , and on any set not containing a neighbourhood of 0, C_τ converges uniformly to 0. These properties show that the sequence $\{C_\tau\}$ is an approximation to the identity, or

$$\lim_{\tau \rightarrow 0} \int C_\tau(t) \phi(t) dt = \phi(0)$$

for any function ϕ in the Schwartz space \mathcal{S} . To see this, note that $\text{supp } C_\tau \subset [-\tau, \tau]$. If ϕ is a Schwartz function on

\mathbb{R} , it is uniformly continuous, so that, given $\varepsilon > 0$, there exists δ such that, if $|s - t| < \delta$, $|\phi(s) - \phi(t)| < \varepsilon$. Choosing $\tau < \delta$,

$$\begin{aligned} |C_\tau * \phi(t) - \phi(t)| &= \left| \int_{-\infty}^{\infty} C_\tau(s) \phi(s - t) ds - \phi(t) \right| \\ &= \left| \int_{-\tau}^{\tau} C_\tau(s) (\phi(s - t) - \phi(t)) ds \right| \\ &\leq \sup_{s \in [-\tau, \tau]} |\phi(s - t) - \phi(t)| \int_{-\tau}^{\tau} C_\tau(s) ds < \varepsilon. \end{aligned}$$

Another way of putting it is that C_τ converges to the Dirac delta distribution as $\tau \rightarrow 0$ in the sense of distributions.

To give another justification we can compute the Fourier transform of C_τ and show that it converges to the constant function 1. Since C_τ is an even function supported on $[-\tau, \tau]$, the integral defining \widehat{C}_τ can be taken from $-\tau$ to τ . We can also ignore the imaginary part, which is an odd function and hence integrates to 0.

$$\begin{aligned} \widehat{C}_\tau(\xi) &= 2 \int_0^\tau \left(\frac{1}{\tau} - \frac{t}{\tau^2} \right) \cos(2\pi t \xi) dt \\ &= \frac{\sin(2\pi \tau \xi)}{\pi \tau \xi} - \frac{2}{\tau^2} \int_0^\tau t \cos(2\pi t \xi) dt \\ &= \frac{1}{\pi h^2 \xi} \int_0^\tau \sin(2\pi t \xi) dt = \frac{1 - \cos(2\pi \tau \xi)}{2(\pi \tau \xi)^2} \\ &= \frac{1 - \cos^2(\pi \tau \xi) + \sin^2(\pi \tau \xi)}{2(\pi \tau \xi)^2} = \left(\frac{\sin(\pi \tau \xi)}{\pi \tau \xi} \right)^2 \end{aligned}$$

Noting that $\widehat{C}_\tau \rightarrow 1$ as $\tau \rightarrow 0$ for any ξ , we then have that $\widehat{C}_\tau \rightarrow \widehat{\delta}$ as well. So, we have shown that the limit of power spectral density of $\Delta_\tau B$ at each point of the frequency spectrum approaches 1 as $\tau \rightarrow 0$, another defining quality of white noise.

The above two computations show that the limit of finite-difference approximations $\Delta_\tau B$ to the derivative of B is what we would expect the corresponding value to be for a white noise process, if it existed. In a very palpable sense then white noise is the derivative of the Wiener process, but can we make precise the way in which this statement is true? Indeed we can and shall, but doing so will of necessity take into account what the delta function truly is: a distribution.

3 DISTRIBUTION-VALUED STOCHASTIC PROCESSES

We will make ample use of distribution theory in what follows, as it will enable us to make sense of the idea that Brownian motion can have a derivative at all. For the relevant discussion and definitions, see the companion article to this one, which is devoted entirely to the subject.

The first viewpoint of stochastic processes is to regard them as an indexed family $\{X_t\}_{t \in T}$ of real-valued random variables defined on some probability space (Ω, \mathcal{F}, P) .

However, if we let \mathbb{R}^T be the collection of all functions from T to \mathbb{R} , we can instead consider a stochastic process as a measurable function from Ω to \mathbb{R}^T , which then induces a measure on \mathbb{R}^T . While \mathbb{R}^T can be a huge space, especially when $T = [0, \infty)$, typically the induced measure is concentrated on continuous or at least càdlàg functions.

The ideas we will present instead require that we consider mappings that take values in the space \mathcal{S}' of tempered distributions. In order for this to not seem completely insane note that we can easily regard the sample paths of Brownian motion as being tempered distributions. To prove this, we need to show that the sample paths are bounded by polynomials almost surely with respect to their probability law.

A heuristic proof relies on the construction of Brownian motion as an appropriate scaling limit of a symmetric random walk, as proved in [4]. A random walk is a sum of Bernoulli random variables, and the sum of n Bernoulli random variables which take the values ± 1 can be at most $n(n+1)/2 = \mathcal{O}(n^2)$.

For a more rigorous argument, let B_t be a sample path of the process. For any fixed t , B_t is Gaussian with variance t and mean 0, so it has the density $\frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{x^2}{2t}\right)$. Hence,

$$E[|B_t|] = \frac{1}{\sqrt{2\pi t}} \int_{-\infty}^{\infty} |x| e^{-\frac{x^2}{2t}} dx = \sqrt{\frac{2}{\pi}} \sqrt{t}.$$

Since $\int_0^\infty \sqrt{t}(1+t^2)^{-1} dt < \infty$, we have that

$$\int_0^\infty \frac{E[|B_t|]}{1+t^2} dt < \infty.$$

The integrand is positive and jointly measurable in t and ω , so by Fubini's theorem, we can exchange the expectation and the t -integral:

$$E \left[\int_0^\infty \frac{|B_t|}{1+t^2} dt \right] < \infty.$$

Then

$$P \left[\int_0^\infty \frac{|B_t|}{1+t^2} dt = \infty \right] = 0,$$

so that even the most blunt estimate that almost all sample paths of Brownian motion are $\mathcal{O}(t^2)$ as $t \rightarrow \infty$ will suffice for our purposes. (The Law of the Iterated Logarithm gives the sharp estimate that $|B_t| = \mathcal{O}(\sqrt{2t \log \log t})$ as $t \rightarrow \infty$.) The sample paths of Brownian motion then define tempered distributions according to the formula

$$\langle B, \phi \rangle = \int_0^\infty B_t \phi(t) dt$$

where ϕ decreases rapidly as $t \rightarrow \infty$.

4 GENERALIZED GAUSSIAN PROCESSES

Suppose we consider stochastic processes X taking values in the space of tempered distributions \mathcal{S}' . We can then

consider the random variable $\langle X, \phi \rangle$ for ϕ some Schwartz function. The process X is called *generalized Gaussian* if $\langle X, \phi \rangle$ is a Gaussian random variable in the usual sense for every $\phi \in \mathcal{S}$. We can then compute quantities like the mean $E(\langle X, \phi \rangle)$, the variance $\text{Var}(\langle X, \phi \rangle)$, the characteristic function $E(e^{i\xi \langle X, \phi \rangle})$ and so forth.

How are we to define stationarity in this new framework? One condition is easy: $E(\langle X, \phi \rangle) = m$, a constant not depending on ϕ . Recall that τ_u is the operator of translation by the vector u on \mathcal{S} defined by

$$\tau_u \phi(t) = \phi(t+u).$$

If for every $u \in \mathbb{R}$ and every $\phi, \psi \in \mathcal{S}$,

$$E(\langle X, \tau_u \phi \rangle \langle X, \tau_u \psi \rangle) = E(\langle X, \phi \rangle \langle X, \psi \rangle),$$

then X is called a *generalized wide-sense stationary process*. From now on when we use the word stationary it will be understood that we mean wide-sense stationary.

Suppose we look at the Wiener process as a random variable taking values in the space \mathcal{S}' of tempered distributions. Given a Schwartz function ϕ , what is the distribution of $\langle B, \phi \rangle$? Enter the workhorse of calculus, the Riemann integral. Let

$$R_n = \frac{1}{n} \sum_{k=1}^{n^2} B_{\frac{k}{n}} \phi\left(\frac{k}{n}\right), \tag{1}$$

be a Riemann sum for the integral of $B_t \phi(t)$, so that

$$\lim_{n \rightarrow \infty} R_n = \int_0^\infty B_t \phi(t) dt. \tag{2}$$

To give another justification of this limit, let

$$f_n(t) = \sum_{k=1}^{n^2} B_{\frac{k}{n}} \phi\left(\frac{k}{n}\right) \chi_{\left[\frac{k}{n}, \frac{k+1}{n}\right)}(t).$$

Each f_n is a simple measurable function, and f_n converges pointwise to $B \cdot \phi$. Furthermore, $\int f_n d\mu = R_n$. Since ϕ decays rapidly at infinity, we can bound $B_t \phi(t)$ by something integrable, say t^{-2} , for large t . We can then invoke the dominated convergence theorem:

$$\lim_{n \rightarrow \infty} \int_0^\infty f_n(t) dt = \lim_{n \rightarrow \infty} R_n = \int_0^\infty B_t \phi(t) dt.$$

Now, since each $B_{\frac{k}{n}}$ has mean zero, $E[R_n] = 0$ also, and taking the limit

$$E \left[\int_0^\infty B_t \phi(t) dt \right] = E[\langle B, \phi \rangle] = 0$$

as well.

Keeping in mind that the random vector $(B_{\frac{k}{n}})_{1 \leq k \leq n^2}$ is Gaussian with mean 0 and covariances $\min(\frac{i}{n}, \frac{j}{n})$, we can

then compute the variance of R_n as

$$\begin{aligned}\text{Var}(R_n) &= \frac{1}{n^2} \sum_{i=1}^{n^2} \sum_{j=1}^{n^2} \phi\left(\frac{i}{n}\right) \phi\left(\frac{j}{n}\right) E\left(B_{\frac{i}{n}} B_{\frac{j}{n}}\right) \\ &= \frac{1}{n^2} \sum_{i=1}^{n^2} \sum_{j=1}^{n^2} \phi\left(\frac{i}{n}\right) \phi\left(\frac{j}{n}\right) \min\left(\frac{i}{n}, \frac{j}{n}\right).\end{aligned}$$

We can then regard this as the Riemann approximating sum of the double integral

$$\int_0^\infty \int_0^\infty \min(s, t) \phi(s) \phi(t) \, ds \, dt,$$

where the same argument as before justifies the convergence of the sum. If we draw a picture of the function $\min(s, t) \phi(s) \phi(t)$ on the s, t -plane and use its symmetry in reflecting over the line $s = t$, we can see that the variance can be written more conveniently as

$$\text{Var}(\langle B, \phi \rangle) = 2 \int_0^\infty \int_0^t s \phi(s) \phi(t) \, ds \, dt \quad (3)$$

and it is in this form that we will compute the variance of the derivative of the Wiener process. Having chosen some $\phi \in \mathcal{L}$, it follows from (1) that the random variable $\langle B, \phi \rangle$ is Gaussian with mean 0 and variance given by (3), since a pointwise limit of Gaussian random variables whose means and variances converge is also Gaussian.

We should note that we have swept something slightly delicate under the rug here. The argument outlined above uses that

$$\lim_{n \rightarrow \infty} E[R_n] = E\left[\int_0^\infty B_t \phi(t) \, dt\right],$$

but we must be careful: the expectation is with respect to the Wiener measure on $C([0, \infty))$, and Wiener measure was constructed in a rather mysterious fashion, one not particularly amenable to explicit computations. The expectation of R_n can be easily computed from the finite-dimensional distributions of the Wiener process, since R_n depends on only finitely many times. But, does the expectation of R_n approximate the expectation of $\int B_t \phi(t) \, dt$ with respect to the full Wiener measure? Indeed it does, but a full justification requires examining how the Wiener measure is defined by the Riesz representation theorem. In keeping with the elementary spirit of the above calculations we will omit it, but the reader is referred to [6].

With the additional machinery of generalized stochastic processes under our belt we can now give the right definition of a white noise process Z : a generalized stochastic process which is both stationary and Gaussian, having zero mean and variance

$$E(\langle Z, \phi \rangle^2) = \int_0^\infty \int_0^\infty \delta(s-t) \phi(s) \phi(t) \, ds \, dt.$$

5 WHITE NOISE IS THE DERIVATIVE OF BROWNIAN MOTION

Now we come to the meaning of all this sound and fury: computing the variance of the generalized derivative of Brownian motion. We will follow closely the treatment in [6]. Recall that if f is a distribution on the line and ϕ is a Schwartz function, then the distributional derivative of f is defined by the formula

$$\langle f', \phi \rangle = -\langle f, \phi' \rangle.$$

Accordingly, if we consider the Wiener process, its distributional derivative acts on Schwartz functions by the formula

$$\langle B', \phi \rangle = -\int_0^\infty B_t \phi'(t) \, dt.$$

We can compute the variance of $\langle B', \phi \rangle$ from (3): it is simply the variance of $\langle B, -\phi' \rangle$, or

$$\text{Var}(\langle B', \phi \rangle) = 2 \int_0^\infty \int_0^t s \phi'(s) \phi'(t) \, ds \, dt. \quad (4)$$

What we want to do now is beat this integral into a convenient form so we can recognize that $\text{Var}(\langle B', \phi \rangle)$ is given by a familiar distribution acting on ϕ . Use integration by parts to evaluate the s -integral in (4):

$$\begin{aligned}\int_0^t s \phi'(s) \, ds &= s \phi(s) \Big|_0^t - \int_0^t \phi(s) \, ds \\ &= t \phi(t) - \int_0^t \phi(s) \, ds.\end{aligned}$$

The t -integrals of each term can be evaluated separately. For the first summand we use integration by parts again.

$$\begin{aligned}\int_0^\infty t \phi(t) \phi'(t) \, dt &= \frac{1}{2} \int_0^\infty t \frac{d}{dt} \phi(t)^2 \, dt \\ &= \frac{1}{2} t \phi(t)^2 \Big|_0^\infty - \frac{1}{2} \int_0^\infty \phi(t)^2 \, dt \\ &= -\frac{1}{2} \int_0^\infty \phi(t)^2 \, dt.\end{aligned}$$

The second integral is again evaluated via integration by parts.

$$\begin{aligned}-\int_0^\infty \left(\int_0^t \phi(s) \, ds \right) \phi'(t) \, dt &= -\phi(t) \int_0^t \phi(s) \, ds \Big|_0^\infty \\ &\quad + \int_0^\infty \phi(t)^2 \, dt \\ &= \int_0^\infty \phi(t)^2 \, dt\end{aligned}$$

Adding up the last two integrals yields

$$\text{Var}(\langle B', \phi \rangle) = \int_0^\infty \phi(t)^2 \, dt$$

but, by a clever trick, this can be rewritten as

$$\int_0^\infty \int_0^\infty \delta(s-t) \phi(s) \phi(t) \, ds \, dt.$$

So, the distribution kernel of $\text{Var}(\langle B', \phi \rangle)$ is $\delta(s-t)$, and so B' , interpreted as a tempered distribution, is white noise.

6 STOCHASTIC DIFFERENTIAL EQUATIONS

Finally, we can hint at the first step necessary to defining the Itô integral. The intuitive notion of a stochastic differential equation as conceived by physicists was a differential equation of the form

$$\frac{dX}{dt} = b_t + \sigma_t Z_t$$

where Z is white noise, b is the *drift* coefficient and σ the *diffusion* coefficient. We allow for b and σ to be random as well. We can then reformulate this as an integral equation:

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s Z_s ds,$$

but since $Z_t dt = dB_t$,

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dB_s.$$

One must then give a meaning to the integral with respect to Brownian motion, determine what class of functions can be integrated, and set about solving these integral equations. A form of the chain rule holds for this Itô calculus, called the Itô formula, which in many cases does allow us to find explicit solutions: if X_t is a solution to the previous stochastic differential equation and $Y_t = f(t, X_t)$ for f twice continuously differentiable, then

$$Y_t = Y_0 + \int_0^t \left[\frac{\partial f}{\partial t}(s, X_s) + \frac{1}{2} \frac{\partial^2 f}{\partial x^2}(s, X_s) \right] ds + \int_0^t \frac{\partial f}{\partial x}(s, X_s) dX_s.$$

When b , σ and f are fairly nice functions we can solve most of these with comparable difficulty to solving ordinary differential equations. Real life not being exactly solvable, the sample paths of these stochastic processes can be simulated numerically.

JOKES

A stats major was completely hung over the day of his final exam. It was a True or False test, so he decided to flip a coin for the answers. The stats professor watched the student the entire two hours as he was flipping the coin...writing the answer...flipping the coin...writing the answer. At the end of the two hours, everyone else had left the final except for the one student. The professor walks up to his desk and interrupts the student, saying:

“Listen, I have seen that you did not study for this statistics test, you didn’t even open the exam. If you are just flipping a coin for your answer, what is taking you so long?”

The student replies bitterly (as he is still flipping the coin):

“Shhh! I am checking my answers!” □

A Neanderthal child rode to school with a boy from Hamilton. When his mother found out she said, “What did I tell you? If you commute with a Hamiltonian you’ll never evolve!” □

Q: How can you tell that you are in the hands of the Mathematical Mafia?

A: They make you an offer that you can’t understand. □

7 CONCLUSION

Stochastic analysis is a burgeoning theory with many facets. The study of stochastic differential equations has ballooned in the past few decades in light of its success as a model for the values of financial assets. Additionally, stochastic differential equations are seeing growing use in all fields of science as many realize that they can incorporate fluctuations in physical parameters or random forcing, for example in turbulence theory. The white noise process is particularly satisfying to study because it was the genesis of the entire subject of stochastic analysis, in a heuristic form reliant mostly on physical intuition. But, we can now look back, reinterpret and justify rigorously its uses in the early days of the field.

REFERENCES

- [1] R.B. Ash and C. Doléans-Dade. *Probability and Measure Theory*. Academic Press, 2nd edition, 1999.
- [2] L.B. Koralov and Y.G. Sinai. *Theory of Probability and Random Processes*. Springer, 2nd edition, 2007.
- [3] N.V. Krylov. *Introduction to the Theory of Random Processes*. American Mathematical Society, 2002.
- [4] H-H. Kuo. *Introduction to Stochastic Integration*. Springer, 2006.
- [5] B. Øksendal. *Introduction to Stochastic Differential Equations*. Springer, 5th edition, 2002.
- [6] S. Sternberg. *Theory of Functions of a Real Variable*. 2005.

CREDITS

*in alphabetical order**The δ -epsilon Editing Team*

- David Ayele
- Ana Best
- Deepy Chahal
- Orianna DeMasi
- Maya Kaczorowski
- Dana Mendelson
- Daniel Shapero
- Alexandre Tomberg
- Katherine Woodard

Cover Art and Design

- Paul Koosis
- Dana Mendelson

Reviewing Board

- Tayeb Aissiou
- Tigran Atoyan
- Sara Froehlich
- Mark Hagen
- Tony Li
- Jeremy Macdonald
- Felicia Magpantay
- Marc Masdeu
- Amaan Mehrabian
- Jason Polák
- Mireille Schnitzer
- Michael Wong

We would like to give credit to the following websites for the math jokes in this year's issue:

<http://www.math.ualberta.ca/~runde/jokes.html>
<http://www.mathjokes.org/>
<http://www.cms.math.ca/Recreation/mjokes-1.html>
http://varatek.com/scott/math_jokes.html
<http://www.math.utah.edu/~cherk/mathjokes.html>
http://trucsmaths.free.fr/blagues_math.htm
<http://www.btinternet.com/~se16/hgb/statjoke.htm>

ACKNOWLEDGEMENTS

We would first like to thank Professors David Wolfson, Johanna Nešlehová and Russell Steele for taking the time to speak with us, and for the great advice they have given to the undergraduates. Their presence at McGill and their involvement with the δ -epsilon greatly contributes to the undergraduate math community at McGill. A special thanks goes to Professor Paul Koosis for letting us use his drawings on our cover.

We must thank our Reviewing Board, the math graduate students at McGill who so graciously gave their time and reviewed the articles received this year. This allowed the δ -epsilon to be peer-reviewed and greatly improved the caliber of our articles.

We are grateful to the SUMS, the Science Undergraduate Society, the Arts Undergraduate Society and the Department of Mathematics and Statistics for providing the funding which enabled us to publish this year's issue.

Of course, we must thank the authors for their interesting articles, and you, the reader, for the thing you are defined to do!

x & y axes and lying, with
that $f(z)$ is analytic in D
 $\in D$, take a square Q (with
 axes) lying, with its interior
 at z_0 . Let ϵ be the largest

