Preface

The theory of dynamical systems has its origins in mechanics. A basic motivating problem, explored by the dynamical pioneers Henri Poincaré and George Birkhoff around the turn of the 20th century, was to predict the motion of heavenly bodies. Today, dynamics is one of the most lively areas in mathematics, and the "dynamical approach" is used to solve problems in a range of mathematical areas, including number theory, geometry, and analysis. To give one example, dynamical techniques have led to significant recent progress on the Littlewood conjecture on simultaneous Diophantine approximations, which dates back to the 1930s. A recent breakthrough in three-dimensional topology, the proof of the surface subgroup conjecture, relies in part on dynamical properties of geodesic flows. Some of these techniques are touched on in these lectures, in particular Bryce Weaver's exposition of Margulis's method. An application of complex dynamics to astrophysics appears in Roland Roeder's lectures: a single light source can have at most $5n - 5$ images when lensed by n point masses. And the list goes on.

This monograph is intended to introduce the reader to the field of dynamical systems by emphasizing elementary examples, exercises, and bare-hands constructions. These notes were written for the Undergraduate Summer School of the thematic program "Boundaries and Dynamics" held in 2015 at the University of Notre Dame, in partnership with the NSF and the French GDR Platon 3341 CNRS.

Roughly speaking, a dynamical system is a space that can be transformed by a fixed set of rules (classically these rules are deterministic, but in the last chapter, random dynamics is explored). By applying these rules repeatedly, under a process called iteration, the space evolves over a discrete set of time intervals. In a slight variation of this definition, the system evolves over a continuous time interval such as the real numbers, still subject to the rules given by, for example, an ordinary differential equation. In both settings, the object of the game is to understand the future states of the system. Starting at a particular point in the space and following its future iterations gives an *orbit* or *trajectory* of the system. Many questions arise, depending on the system. Are there bounded trajectories? *Periodic* orbits, which return to their starting point after a finite period of time? Orbits that fill the space densely? When do two systems have the same orbits in some sense, and what are invariants of a system that can detect this type of equivalence? Can such invariants be computed using periodic trajectories?

Billiard flows are a basic type of continuous time dynamical system and arise naturally as models in physics. Here the table itself gives a framework for the space, and a point in the space, imagined as a ball with a specified velocity, travels in the space over time by reflecting off the sides of the table. In the first chapter, Diana Davis starts with a square billiard table and takes us to more exotic tables, the *rational* polygons (i.e., tables where all corner angles are equal to a fractional part of π). Tables give rise to surfaces: unfolding the billiard table, the trajectory of the ball can be reimagined as a "straight line" curve on a *translation surface* created from a polygon where parallel edges are identified by translations. For the square, the associated surface is a *flat torus* and resembles a donut. This particular surface is rich in symmetries, like reflections across verticals and horizontals. A class of transformations of this surface called *shear* maps can be used to understand the behavior of the straight line curves on the torus and hence to describe the trajectories on a square billiard table (Chapter 1, Theorem 5.5). Periodic trajectories are characterized by an initial velocity with rational slope (Chapter 1, Exercise 2.1). Such trajectories can be grouped into families of *cylinders* of parallel periodic trajectories with equal lengths. Analogous techniques can be used to analyze rational billiards, with highly symmetric tables, known as Veech tables, admitting a particularly complete understanding.

This deceptively simple-sounding class of dynamics continues to captivate researchers, with many open problems remaining.

For example, let $N(r)$ be the number of cylinders in a rational billiard table corresponding to periodic trajectories of length at most r. A famous conjecture states that the limit of $N(r)/r^2$ exists and is nonzero. The best result to date was given by Alex Eskin, Maryam Mirzakhani (Fields Medal 2014!), and Amir Mohammadi. The beauty of their work comes from the method: they deduce asymptotic properties of the counting function $N(r)$ from the description of trajectories of a dynamical system defined on a very big space, the *moduli space*, in which the initial billiard is just a point! For billiards on irrational polygons, few tools are available and not much is known. For example, the existence of periodic trajectories is an open problem even for triangles.

In the second chapter, Bryce Weaver restricts his attention to a discrete dynamical system on the flat torus defined by a 2×2 matrix A with integer coefficients and determinant 1. The eigenvalues of such a matrix are multiplicative inverses of each other; to avoid trivial dynamics, we assume one of these eigenvalues is real and bigger than 1.

This example, introduced in the 1960s by Vladimir Arnold (and playfully termed the "cat map"), is the quintessential model for the class of *uniformly hyperbolic* dynamical systems that are highly sensitive to initial conditions. Different invariants express the unpredictable behavior of such systems. One of them, *topological entropy*, measures the exponential rate at which points separate. For the process generated by the matrix A, this invariant equals the logarithm of the biggest eigenvalue $\lambda > 1$ of the matrix (Chapter 2, Proposition 3.9). The entropy is connected to the asymptotic behavior of the counting function $P_n^{\mathcal{O}}$ defined by the number of A-periodic points of period less than n. Roughly speaking, this invariant corresponds to the growth rate of $P_n^{\mathcal{O}}$. More precisely, the limit of $P_n^{\mathcal{O}} \times n/\lambda^n$ exists and equals $\lambda/(\lambda - 1)$ (Chapter 2, Theorem 4.2). The chapter provides two proofs of this fundamental relationship. The first one relies on the algebraic nature of the transformation A

and is elementary. The first proof even gives an explicit formula for the number of A-periodic points of period n (Chapter 2, Theorem 4.5). The second proof uses an argument based on the presence of expanding and contracting directions attached to A (Chapter 2, Section 4.2) and on the fact that this transformation "mixes" the sets (Chapter 2, Proposition 4.15). This approach, which combines geometry and ergodic theory, was developed by Gregory Margulis. It is longer and less precise than the first one (Chapter 2, Theorem 4.5 is replaced by a weaker version, Theorem 4.11) but it applies to a vastly more general class of systems: those for which fine algebraic information (such as the eigenvalues of a matrix) is not available. Indeed, Margulis's method can be used to deduce geometrical information about negatively curved compact manifolds, in particular the growth rate of the number of closed geodesics as a function of their length.

In the third chapter, Roland Roeder considers a dynamical system on the complex plane $\mathbb{C}$, governed by a quadratic map $p_c(z) = z^2 + c$, where c is a complex number, a *parameter* that can be changed to vary the dynamics of the system. The study of this family of maps $\{p_c : c \in \mathbb{C}\}$ was initiated by two founders of holomorphic dynamics, Pierre Fatou and Gaston Julia. This area came to explosive life in the 1980s with the introduction of so-called quasiconformal methods on the theoretical side and, on the experimental side, with the blossoming of the personal computer as a mathematical tool.

Despite the simplicity of their defining formula, the maps p_c exhibit complicated dynamics. Points which are far from the origin O in $\mathbb{C}$ escape to infinity. Among the other points, there is at most one periodic orbit around which spiral the orbits of nearby points. The set of parameters c for which p_c has a single *attracting fixed point* is contained inside a cardioid (Chapter 3, Lemma 2.9). Increasing the size (i.e., period) of the periodic orbit attaches open blobs to this cardioid in the complex parameter plane. Collecting all of these blobs together, the set M_0 of parameters c for which the map p_c admits such *attracting periodic points* has a rich topological and combinatorial structure. This set M_0 is contained in the famous Mandelbrot set M, defined as the set of parameters for which the orbit of O remains bounded. (This set was named for Benoît Mandelbrot, who brought public attention to this class of dynamical systems and its vivid computer images in the 1980s). A central open question in complex dynamics asks when the closure of M_0 coincides with M. Although it is unsolved, it was proved in 1997 for a thin slice of the parameter plane: the restriction of M_0 and M to the real line (i.e., for the parameters c being real numbers). The set M is still mysterious, but many interesting properties are known. In particular, it is self-similar: M contains arbitrarily small copies of itself. Adrien Douady and John Hubbard have proved that M is connected, and it is conjectured that M is locally connected. There is a deep relationship between M and the dynamics of an individual map p_c through the shape of the *filled Julia set* K_c of points having a bounded orbit under p_c. Namely, K_c is connected if and only if c belongs to

M (Chapter 3, Section 3). This principle is exploited to transfer information from the dynamics of an individual member of M back to geometric information about M itself. In particular, it is used to show that the boundary of M has amazing complexity: unlike the boundary of a disk, which is a smooth one-dimensional curve, the boundary of M has (Hausdorff) dimension 2. To show the limits of our understanding, it is not known whether the boundary of M might even have positive area!

The last chapter deals with a type of nondeterministic dynamical system: a *random walk*, for which the iteration at each step is governed by a probability law. In a sense, this is a classical dynamical system in which the transformation rules are allowed to include a roll of the dice. The theory of random walks, which mixes geometry and probability, was first developed in the 1920s. It has incredibly broad applicability and today, less than a century later, it is nearly ubiquitous in science and engineering. Pablo Lessa concentrates on simple walks on a combinatorial object, an *infinite graph*, obtained by starting at a vertex and choosing a random neighbor at each step. The central question concerns the *recurrence property*: does the walk visit "almost surely" every vertex infinitely many times? These walks are one of the most classical examples of how the geometry of the underlying space influences the behavior of stochastic processes on that space. The first result in this direction was obtained by George Pólya for grids: the simple walk on the two-dimensional grid $\mathbb{Z}^2$ is recurrent but on the *three-dimensional grid* $\mathbb{Z}^3$ the walk is not recurrent—it is *transient* (Chapter 4, Section 2.7). The study of simple random walks on $\mathbb{Z}^d$ is a first step in understanding a more complicated object: a continuous time stochastic process on $\mathbb{R}^d$ (or on Riemannian manifolds) called *Brownian motion*. The "wire mesh" in $\mathbb{R}^d$ with vertices in the grid $\mathbb{Z}^d$ is an example of a *Cayley graph*, which encodes the structure of a finitely generated group G and its generators (in this case $\mathbb{Z}^d$ with the standard generating set). For this special class of graphs, Nicholas Varopoulos proved that recurrence of the random walk can be characterized entirely by certain algebraic/geometric properties of the group G. In particular, if we define the counting function $f_G(n)$ of G to be the number of words of length at most n (relative to a set of generators), then the random walk on the Cayley graph of G is recurrent if and only if $f_G(n)$ is bounded above by a polynomial function of degree at most 2 (Chapter 4, Section 4). The field of geometric group theory grew in the 1980s to study the relationship between the algebraic properties of groups and the geometric properties of their Cayley graphs. A theorem proved by Mikhail Gromov illustrates the deep relationship between these two objects: the counting function $f_G(n)$ has polynomial growth if and only if G admits a finite index nilpotent subgroup.

Françoise Dal'Bo September 2015
François Ledrappier
Amie Wilkinson