DYNAMICS OF PERTURBED ELLIPTICAL BILLIARD TABLES

PATRICK R. BISHOP, SUMMER CHENOWETH, EMMANUEL FLEURANTIN, EVELYN SANDER, AND JASON MIRELES JAMES

Abstract. Dynamical billiards consist of a particle on a two-dimensional table, bouncing elastically off a boundary curve. The state of the system is given by two numbers: one describing the location along the curve where the bounce occurs, and another describing the incoming angle of the bounce. Successive bounces define a two-dimensional area preserving map, and iterating this map gives a dynamical system first studied by Birkhoff. One of the simplest smooth table shapes is that of an ellipse, in which case the dynamics of the billiard map is completely integrable. The longstanding Birkhoff conjecture is that elliptical tables are the only smooth convex table for which complete integrability occurs. In this spirit, we present an implicit real analytic method for iterating billiard maps on perturbed elliptical tables. This method allows us to compute local stable and unstable manifolds of periodic orbits using the parameterization method. Globalizing these local manifolds numerically provides insight into the dynamics of the table.

Keywords: Dynamical Billiards, Birkhoff Conjecture, Stable/Unstable Manifolds, Parameterization Method
2020 AMS Subject Classifications: 37C83, 37M21, 37M05, 37D10

1. Introduction

The game of billiards, beloved since the 14th century, inspired the mathematical models proposed by Birkhoff in the 1920s [Bir27]. Since then, mathematical billiards has served as an especially rich proving ground for ideas in Hamiltonian systems theory. One of the motivations for Birkhoff's work was to solidify the foundations of thermodynamics by restricting attention to a single gas particle. Billiards (especially on polygonal tables) are moreover related to geodesic flow on surfaces, and have applications in classical mechanics and optics [BFB$^+$11, BDBT23, DHN03, BK96]. The reader will find more complete discussion of the history and application of these models, as well as many additional references, in any of the excellent sources [Gut86, Tab09, CM06, HW13, BFB$^+$11, Bun19, BFG$^+$22, Fra03, Gut12, LT07, MRRTS16, Mei92].

Mathematical billiards are introduced formally in Section 2.1 but, loosely speaking, the idea is to consider a massless particle moving without friction across a planar billiard table. The particle bounces off the table's boundary in an elastic collision and moves off in a new direction determined by Snell's law. If the boundary of the table is a simple closed curve, it can be parameterized with a single parameter. Since the angle of incidence is of course described by an angle, these two numbers describe the state of the system at collision. In between collisions the particle moves in a straight line with constant velocity, allowing one to reduce the system from a differential equation with impacts to a smooth map on a cylinder. The reduced system is known as a billiard map. The relation between the flow on the table and the billiard map is illustrated in Figure 1 for a pair of tables which, while simple, already suggest that billiard trajectories can be quite complicated.

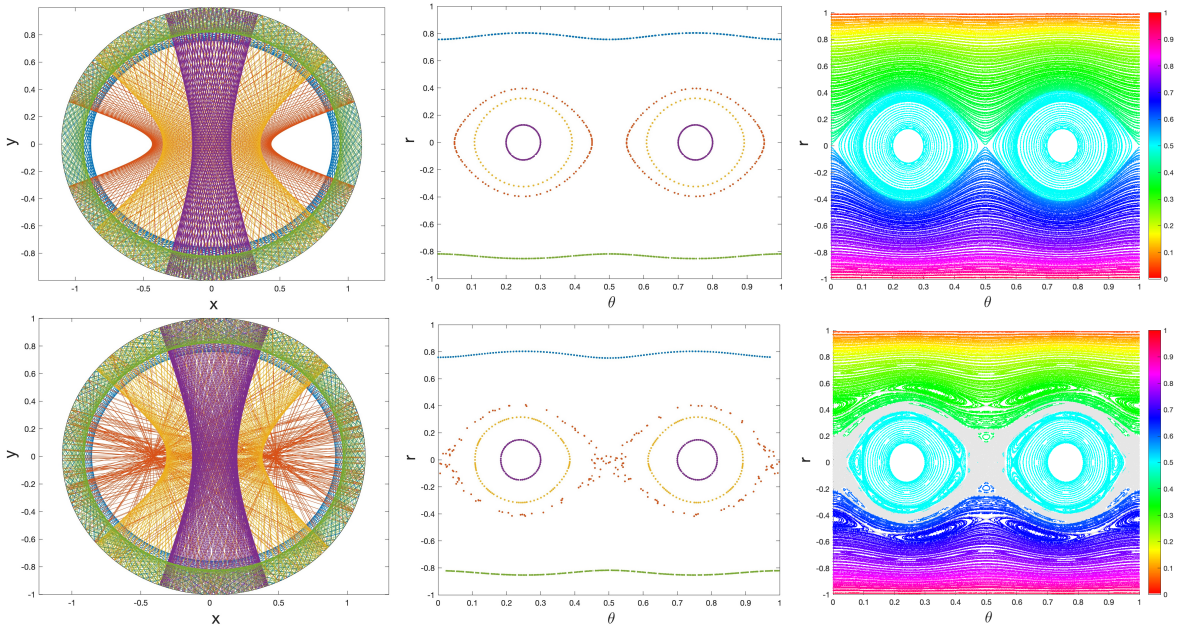

FIGURE 1. Top Left: Orbits on the physical table which is an unperturbed ellipse with eccentricity 0.4583. Each trajectory has a caustic, i.e. a curve which is tangent to every line. Top Middle: Phase space orbit corresponding to the trajectory on the top left. The value of θ indicates the location of bounce, and r measures the angle of the bounce. Top Right: The full phase space of orbits are quasiperiodic, colored by their frequency. Bottom Left, Middle: The trajectories and orbits for the same orbits, but for a small perturbation of the ellipse (Table A described in (2) and Table 1). Some caustics remain. However, the red trajectory no longer has a caustic; its corresponding orbit is chaotic in phase space. Bottom Right: For the full phase space, many quasiperiodic orbits persist and are colored by frequency, but some orbits are now chaotic, colored in gray.

Again, because particle trajectories follow straight lines between consecutive collisions, the dynamics are completely determined by the shape of the table. An important conjecture in this context is the Birkhoff conjecture, which claims that a billiard system is integrable if and only if the boundary of the table is an ellipse (this includes the case of the circle). From the beginning of the study of billiard problems in mathematics, this conjecture has occupied a central place, and work on the problem continues through the present day. Avila et al. [ADSK16] for example considered the case of infinitesimally perturbed ellipses, and Delshams et al. [DRR96, KS18a] proved local versions of the conjecture. In 2023, Baracco and Bernardi [BB24] proved that a totally integrable strictly-convex symplectic billiard table, whose boundary has strictly positive curvature, must be an ellipse. For more discussion of the theory and open problems, see [BFG$^+$22, Sch22].

Since the Birkhoff conjecture concerns integrable billiards, the question is closely related to the existence of chaotic motions. Chaos in billiards systems was first studied by

Bunimovich [Bun75, Bun79] in the 1970s, who showed that there exist tables with only focusing components (focusing trajectories like a lens) admitting chaotic dynamics. An overview of chaotic billiards is found in [CM06], and more recent work in this area includes [DHF19, Bun19]. If the Birkhoff conjecture holds, then we expect non-elliptical tables to admit chaotic motions.

In this paper, we numerically study a basic mechanism for generating chaotic dynamics; namely the homoclinic tangles first studied by Poincaré [Poi93a, Poi93b, Poi93c] and later formalized by Smale [Sma65]. These tangles are formed by transverse intersections between stable/unstable manifolds of hyperbolic fixed points and/or periodic orbits. To reliably locate these objects for billiard maps, we implement a parameterization method for numerically computing high order Taylor expansions of the local stable/unstable manifolds, and "grow" the local manifolds to find the desired intersections. Assuming the billiard table has a real analytic boundary, the series expansions are guaranteed to converge. Heuristic a-posteriori indicators provide practical guides to the useful domain for the truncated series. Figure 2, for example, illustrates stable and unstable manifolds computed for a variety of periodic orbits on the table of Figure 1. Figures 3 and 4 illustrate stable and unstable manifolds for periodic orbits of four other perturbations of elliptical tables. These manifolds intersect transversely, suggesting the existence of chaos in each case. We remark that the parametrization method is well suited for computer assisted proofs, and a follow-up paper will validate the findings presented in the present work.

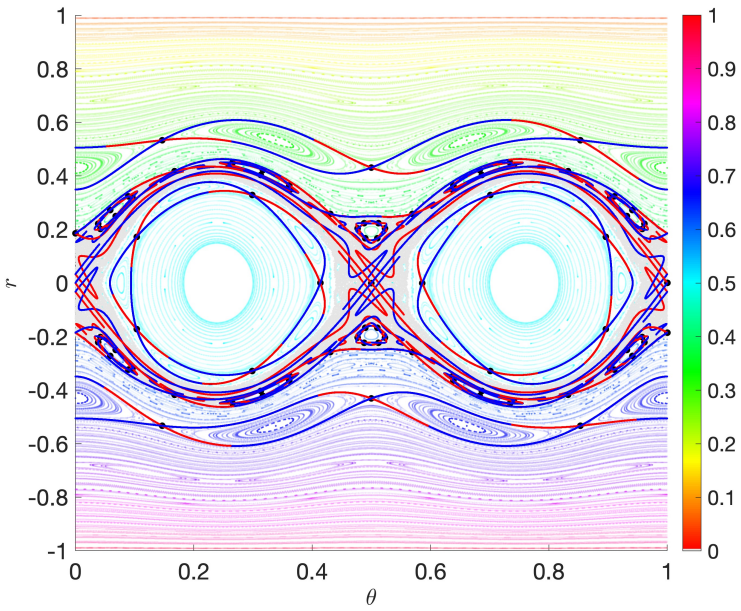

FIGURE 2. This figure shows the stable and unstable manifolds and their intersections for periodic orbits with a variety of periods. This data is for the same perturbed elliptical billiard table as shown in Fig. 1, (Table A in Tbl. 1).

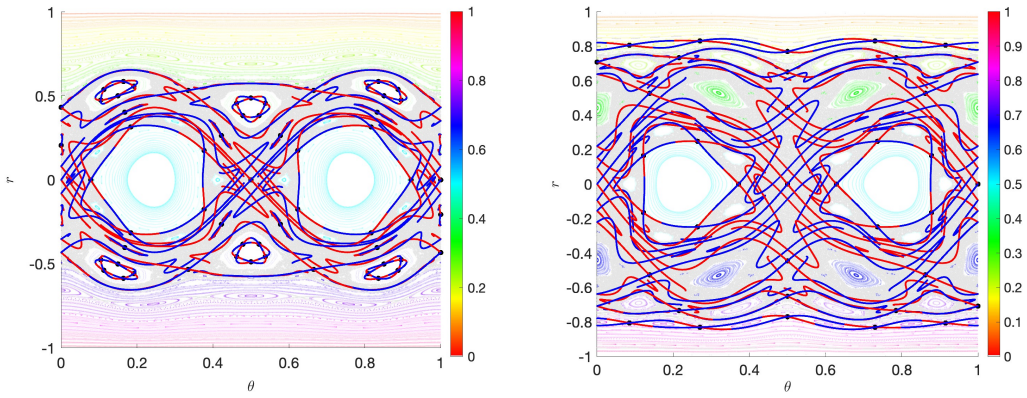

FIGURE 3. The figure shows stable and unstable manifolds for periodic orbits computed for Tables B and C (cf. Tbl. 1), larger perturbations of the same ellipse as in Figs. 1 and 2. The phase space in the background is colored by frequency for regular orbits and gray for chaotic orbits.

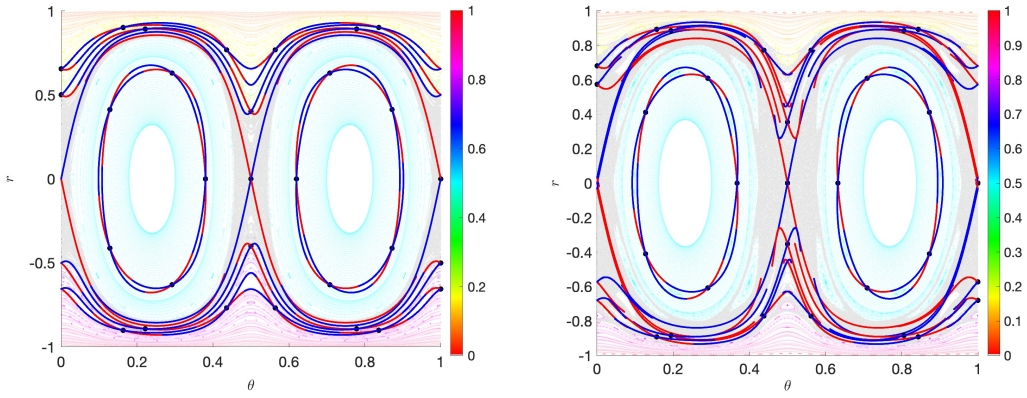

FIGURE 4. The figure shows stable and unstable manifolds for periodic orbits computed for Tables D and E (cf. Tbl. 1). These tables are perturbations of an ellipse with larger eccentricity. The phase space in the background is colored by frequency for regular orbits and gray for chaotic orbits.

The computational scheme is iterative, and a key step requires computing the power series of the billiard map composed with the current polynomial approximation. If the billiard map were explicitly polynomial, then such compositions could be worked out explicitly via Cauchy products, resulting in explicit recursion relations power series coefficients. See for example the worked problems in [GMJ17, FJ20], and also Appendix B of the current work. Non-polynomial nonlinearities are often dealt with by appending differential equations describing the nonlinear terms. Several examples of this procedure are given in the first reference just cited, and the procedure is reviewed in Appendix B. Moreover, formal series techniques for simple implicitly defined maps are also possible, as illustrated in [TMJ22].

However, when the implicit map is complicated enough, the techniques of the references just cited are unwieldy. Examples of such situations include Poincaré maps associated with differential equations, and also the billiard maps studied in the present work. In such situations, an alternative is to compute the desired composition (of a power series with the dynamical map) using methods of interpolation. Moreover, when interpolating analytic functions, the discrete Fourier transform (DFT) is known to provide especially accurate, stable, and efficient interpolation, especially when combined with the Fast Fourier Transform (FFT) algorithm. These techniques are also reviewed in Appendix B.

A complication with the DFT approach is that the map must must be evaluated in the complex plane, and this may involve non-trivial modifications of its numerical implementation. We refer to the recent works of Kumar [KAdlL25, KAdlL22, KAdlL21], where the parameterization method is implemented for Poincaré sections in problems from Celestial Mechanics, and where it is first necessary to extend the numerical integration, and then the Poincaré section, to complex variables via analytic continuation.

In the present work, we undertake a similar program for billiard maps. That is, to use the DFT, we first extend the billiard map into a complex domain (in both variables). Since the map is only implicitly defined, this involves numerical analytic continuation. Note, we do not claim the complex billiard map has any physical significance: it is simply a numerical convenience, as the manifolds we compute are real. This said, it is interesting to note that analytic continuation for billiards has appeared in the theory before, and we refer for example to the discussion following Remark 9 in the introduction of the paper [KS18b], where the authors explain that the main idea of their proof is to analyze the singularities in the complex plane associated with the action angle form of perturbed elliptic billiards. The coordinates used in the reference just cited are chosen for their theoretical properties, and do not appear to be easy to work with numerically. The coordinates used in the present work, on the other hand, are chosen to simplify the resulting numerical calculations. This work builds on results in the PhD Thesis of the first author [Bis24].

The remainder of the paper is structured as follows: Section 2 introduces the setup, definitions, notation, and prior results. Section 3 gives the details of our numerical methods for iterating the billiard map on any perturbed elliptical table with smooth boundaries that is convex. In Section 4, we discuss how we compute the stable and unstable manifolds for periodic orbits. In Section 5, we detail our results, presenting computed stable and unstable manifolds for periodic orbits on five different billiard tables shown in Figures 9–15. Section 6 contains our conclusions and an outline of future work.

2. Definitions, background, and prior results

This section introduces the definitions and notation along with a review of prior results.

2.1. **Dynamical billiards – from physical table to phase space.**

We focus on connected, strictly-convex, planar domains with smooth (in fact real analytic) boundaries, as done in [HW13]. Such a billiard system is given by the following data.

Definition 2.1. *A billiard system consists of a point particle confined to move in a connected, strictly convex, and closed planar domain $\mathcal{D} \subset \mathbb{R}^2$ with boundary curve Γ. We assume that Γ is a smooth non-self-intersecting curve. In particular, Γ is sufficiently smooth that it can*

be parameterized by a real analytic map $B : [0,1] \to \mathbb{R}^2$, where $B(0) = B(1)$, and $B'(\theta)$ is a nonvanishing tangent vector. We assume that B has counter-clockwise orientation on Γ.

The point particle moves within domain $\mathcal{D}$ according to the following criteria.

1. *The trajectory curve $G(t) \subset \mathcal{D}$ for all $t \geq 0$. We require $G(t)$ to be a continuous union of line segments, such that for any time interval for which G is in the interior of $\mathcal{D}$, $\mathbf{v}(t) = G'(t)$ is constant, and $|\mathbf{v}(t)| = 1$.*
2. *The trajectory of the particle begins at $t = 0$ at a point $G(0) \in \Gamma$ on the boundary of the domain.*
3. *Whenever the trajectory $G(t)$ reaches the boundary, it is reflected elastically back into the interior of $\mathcal{D}$ according to the standard Snell's law of reflection.*

Based on the above definition, the particle trajectory is fully determined by the location of the boundary collisions, along with the direction of vector $\mathbf{v}$ after the collision. Let $\{t_k\}_{k=0}^{\infty}$ be a discrete increasing sequence of times at which the particle hits the boundary Γ. Define θ_k by $G(t_k) = B(\theta_k)$. Let γ_k be the angle that $\mathbf{v}(t)$ makes with the tangent vector $B'(\theta_k)$ immediately after the bounce, and let

$$r = \cos\gamma, \text{ where } -1 < r < 1.$$

Then the vector (θ_k, r_k) defines the k-th point of collision[1]. The restriction on r is due to the fact that $|r| \geq 1$ corresponds to the particle leaving the convex billiard table.

For a given boundary parametrization $B(\cdot)$, there is a uniquely determined diffeomorphism f on $S^1 \times (-1, 1)$, which is as smooth as the boundary, such that

$$f\begin{pmatrix}\theta_k \\ r_k\end{pmatrix} = \begin{pmatrix}\theta_{k+1} \\ r_{k+1}\end{pmatrix}.$$

Before moving to the general case, we review the simplest smooth tables.

Lemma 2.1 (Circular Table)**.** *The billiard map on a circular table with boundary parameterized by*

$$B(\theta) = \begin{pmatrix}\cos(2\pi\theta) \\ \sin(2\pi\theta)\end{pmatrix},\ 0 \leq \theta < 1$$

has the following explicit solution. Starting at the point (θ_0, r_0), it is given by

$$f\begin{pmatrix}\theta \\ r\end{pmatrix} = \begin{pmatrix}\theta + C \\ r_0\end{pmatrix},$$

where $C = \arccos(r_0)/\pi$ depends only on the initial condition.

This implies that for the circle, orbits consist of rigid rotation on horizontal lines in the phase space θ versus r. The proof of this result is due to Birkhoff [Bir27].

Lemma 2.2 (Ellipse)**.** *For any ellipse, the function f is fully integrable. All orbits are quasiperiodic or periodic orbits, and all solutions lie along smooth curves.*

We will give more details on how to compute f in the next section.

[1]To avoid cumbersome notation, when we talk about a vector, unless otherwise specified, we implicitly assume that it is a column vector.

2.2. **Birkhoff conjecture.** Billiards on ellipses (including the special case of a circle) have simple nonchaotic behavior. Birkhoff conjectured that these are the only shapes for which this is true.

Theorem 2.1 (Birkhoff Conjecture)**.** *For a convex smooth boundary, if the billiard system is completely integrable, it must be an ellipse.*

Birkhoff first discussed the two-dimensional billiard system with convex boundaries [Bir27]. In 1995, Delshams and Ramírez-Ros [DRR96] proved that any non-trivial, symmetric perturbation of the ellipse is not integrable. More recently, Avila et al. [ADSK16], proved the Birkhoff conjecture for tables that are perturbed ellipses with small eccentricity. This was extended by Kaloshin and Sorrentino [KS18b]. Proofs of local versions of the conjecture are given in [DRR96, KS18a]. Bialy and Mironov [BM22] proved the Birkhoff conjecture for centrally-symmetric $\mathcal{C}^2$-smooth convex planar billiards. Baracco and Bernardi [BB24] proved that a totally integrable strictly-convex symplectic billiard table, whose boundary has strictly positive curvature, must be an ellipse. For more discussion on recent results and open problems in Birkhoff billiards, refer to [BFG$^+$22, KS22, LT07].

2.3. **Related numerical work.** There are many existing methods for numerically computing billiard maps. Levi and Tabachnikov [LT07] develop a method for convex tables that utilizes the minimization property to find the point of next contact of the particle with the boundary. An issue with this method is that the minimization problem has two solutions. For trajectories that are close to tangent, it is hard to distinguish the spurious from the correct solution.

Lansel and Porter [LP04] and [Tur16] simulate classical billiard systems that are not necessarily convex. Both methods are computationally slow, as they directly track the full trajectory of the ball. Solanpää et al. [SLR16] consider more general nonconvex billiard systems for two-dimensional tables. They simulate models for various billiards and diffusion models including those with multiple charged particles and those subject to magnetic fields. In [dCHSL22], da Costa et al. describe billiard dynamics on oval-like tables using a polar equation, with a circle approximation to estimate the next point of contact with the boundary.

There are a number of other works with methods for computing billiards maps. See [AV21, BFB$^+$11, CF21, Dat17, Kni98, MRRTS16, PRK22, PTZ23, RGK21]. None of these previous works includes computation of invariant manifolds for billiard maps, and they are not ideally suited for that purpose, meaning that we had to create a new numerical method rather than being able to use one of the existing methods.

3. Numerical computation of the billiard map

We now discuss our numerical procedure for evaluating the real billiard map.

3.1. **Table shapes.** Note that an ellipse with eccentricity $\sqrt{1 - b_1^2/a_1^2}$ has parameterization

$$E(\theta) = \begin{pmatrix} a_1 \cos(2\pi\theta) \\ b_1 \sin(2\pi\theta) \end{pmatrix}. \tag{1}$$

We focus on perturbed ellipses of the form $B(\theta) = (x(\theta), y(\theta))$ with Fourier series given by

$$
\begin{aligned}
x(\theta) &= a_{1,1}\cos(2\pi\theta) + \sum_{k=2}^{n} a_{k,1}\cos(2k\pi\theta) + b_{k,1}\sin(2k\pi\theta) \\
y(\theta) &= b_{1,2}\sin(2\pi\theta) + \sum_{k=2}^{n} a_{k,2}\sin(2k\pi\theta) + b_{k,2}\sin(2k\pi\theta).
\end{aligned} \tag{2}
$$

Table	A	B	C	D	E
$(a_{1,1}, a_{2,1}, a_{3,1})$	(1.1,0.03,0)	(1.1,0.05, 0.00015)	(1.1, 0.08, 0.0002)	(2, 0.04,0)	(2, 0.05,0)
$(b_{1,2}, b_{2,2}, b_{3,2})$	(1,0.03,0)	(1,0.035, 0.0001)	(1, 0.095, 0.0001)	(1, 0.035,0)	(1, 0.065,0)
Eccentricity	0.4583	0.4583	0.4583	1.7321	1.7321

TABLE 1. The coefficients used for our choice of perturbed elliptical billiard tables, referred to as Billiard Tables A–E. Any unspecified coefficients are zero. The bottom row states the eccentricity of the associated unperturbed ellipse.

Notice that for a table specified by (2), if $a_{1,1}$ and $b_{1,2}$ are the only nonzero coefficients, then the table is an ellipse. In particular, for any table, the *ellipse associated with this table* is the ellipse with the same $a_{1,1}$ and $b_{1,2}$ values (all other coefficients set to zero). The particular tables we have used for our numerics are referred to as Billiard Tables A–E. The coefficients $a_{k,i}$ and $b_{k,i}$ are given in Tbl. 1. For Tables A-C, the associated ellipse has $a_{1,1} = 1.1$ and $b_{1,2} = 1$. For Tables D and E the associated ellipse as $a_{1,1} = 2$ and $b_{1,2} = 1$. Tables A-E and the associated ellipses are depicted in Figure 5.

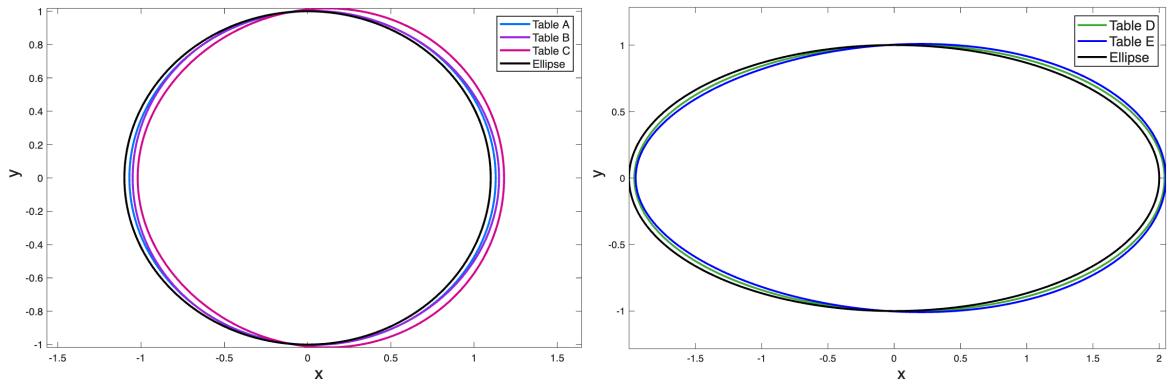

FIGURE 5. Shape of billiard tables used in the numerics. (Left) Billiard Tables A-C are perturbations of an ellipse with eccentricity 0.4583. (Right) Billiard Tables D and E are perturbations of an ellipse with eccentricity of 1.7321.

Since ellipses are convex and since $B(\theta)$ varies continuously with respect to the coefficients, the table remains convex when the higher-order Fourier coefficients are sufficiently small – specifically, if for $i = 1, 2$ and for all $k > 1$, $|a_{k,i}|$ and $|b_{k,i}|$ are significantly smaller than $|a_{1,1}|$

and $|b_{1,2}|$, the table remains convex. To illustrate how convexity is lost as these pertubations grow, consider the parameterized table with coefficients given by:

$$(a_{1,1}, a_{2,1}) = (1.1, 0.03\epsilon)$$
$$(b_{1,2}, b_{2,2}) = (1, 0.025\epsilon). \tag{3}$$

For $\epsilon = 0$, this table is an ellipse. Figure 6 (left panel) shows how the table shape changes as ϵ increases, while the right panel shows the minimum value of the signed curvature measured along the boundary of the table. At $\epsilon \approx 9$, the minimum signed curvature passes through zero, and the table loses convexity. The tables given by (3) have the same associated ellipse as Billiard Tables A, B, and C; in particular Billiard Tables A-C have coefficients $a_{2,1}$ and $b_{2,2}$ with $\epsilon < 4$, which is clearly before convexity is lost.

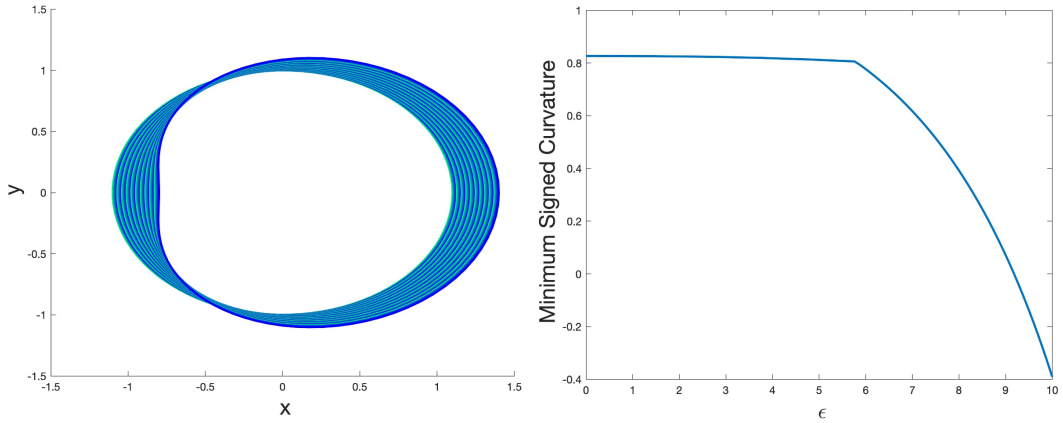

FIGURE 6. (Left) The shape of a table with the nonzero coefficients given by (3) for values of ϵ between 0 and 10. (Right) The minimum signed curvature of the table. When $\epsilon \approx 9$, the signed curvature is zero at a point on the boundary, implying that the table has lost convexity.

3.2. **Iteration for elliptical tables.** In order to compute the billiard map for general tables, we use the fact that we can directly compute the billiard map for an ellipse. We compute the location of each bounce on the table's corresponding ellipse, then use this as the initial guess for Newton's method to find the location of the bounce on the perturbed table. In this section we explain how to compute the ellipse case.

Let $B(\theta) = (x(\theta), y(\theta)) = (a_1 \cos(2\pi\theta), b_1 \sin(2\pi\theta))$ parameterize an ellipse, as introduced in (1). Given an input (θ, r), we seek $(\hat{\theta}, \hat{r})$ such that $f(\theta, r) = (\hat{\theta}, \hat{r})$. Based on r, we can use the arccosine function to find the direction vector $\mathbf{v}$. Since $B(\theta)$ and $B(\hat{\theta})$ must lie on the line parallel to the vector $\mathbf{v} = (v_1, v_2)$, there exists s_* such that

$$B(\hat{\theta}) = B(\theta) + s_*\mathbf{v}. \tag{4}$$

Let

$$d_1 = v_1/a_1 \text{ and } d_2 = v_2/b_1.$$

Combining these equations with the equation for an ellipse and solving for s_* gives

$$s_* = -2\left(\frac{d_1\cos(2\pi\theta) + d_2\sin(2\pi\theta)}{d_1^2 + d_2^2}\right). \tag{5}$$

Plugging this value back into (4) to get $B(\hat{\theta}) = (\hat{x}, \hat{y})$. We find $\hat{\theta}$ from $B(\hat{\theta})$ since

$$\frac{b_1}{a_1}\tan(2\pi\hat{\theta}) = \frac{\hat{y}}{\hat{x}}.$$

This last statement is where the simplification of working with the ellipse is really clear, since for a perturbed ellipse, we cannot solve directly and must instead use a two-dimensional root finding method. Finally, from this point finding $\hat{r}$ is exactly the same as for the perturbed case, and it will be explained in the next section.

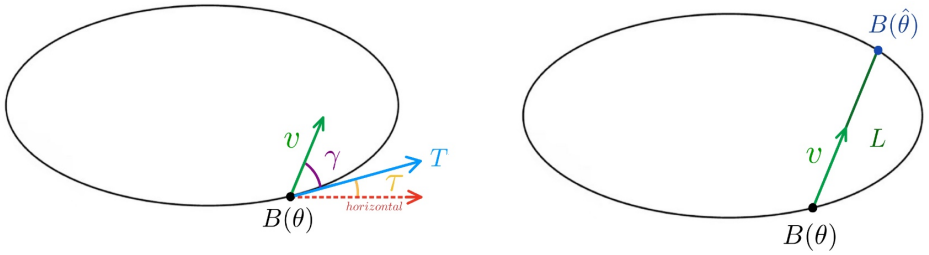

FIGURE 7. Notation for our numerical methods. Left: $\mathbf{v}, \gamma, T$, and τ. Right: The line L contains $B(\theta)$ and $B(\hat{\theta})$.

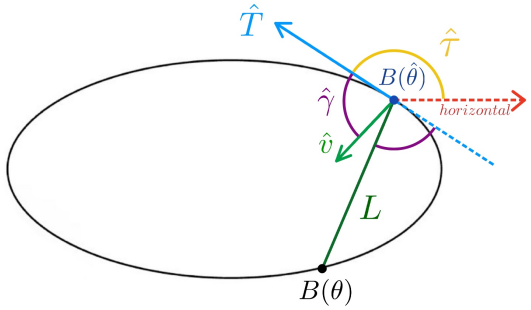

FIGURE 8. Computation of the $\hat{r}$ from $\mathbf{v}$ and $\hat{\theta}$.

3.3. **Iteration for general billiard tables.** Given θ and r, we use the following notation, depicted in Fig. 7 (left).

(1) $T = B'(\theta)$ the tangent vector to the boundary curve.

(2) γ is defined implicitly by $r = \cos\gamma$. We therefore define it directly by $\gamma = \arccos r$.

(3) τ is the angle between the positive horizontal direction and T. Thus τ is defined implicitly by $T_2/T_1 = \tan\tau$. To find τ, we use the quadrant-specific version of the arctan function (`atan2` in MATLAB) and set $\tau = \texttt{atan2}(T)$.

(4) $\mathbf{v}$ is defined implicitly by saying that γ is the angle between $\mathbf{v}$ and T. From the previous two definitions, we get

$$\mathbf{v} = \begin{pmatrix} \cos(\gamma+\tau) \\ \sin(\gamma+\tau) \end{pmatrix}.$$

Using the notation defined above, we are ready to find $\hat{\theta}$ using the fact that $B(\theta)$ and $B(\hat{\theta})$ lie on a line parallel to $\mathbf{v}$, given by

$$L(s) = B(\theta) + s\mathbf{v}, \tag{6}$$

where $s \in \mathbb{R}$ is a scalar. See right panel of Figure 7. Since the table is convex, this line has exactly two intersection points with the boundary, the first being the original point $B(\theta)$ at $s = 0$. The second point $B(\hat{\theta})$ occurs at $s > 0$. As long as we have a good guess near the desired solution, this intersection can be formulated as a zero finding problem

$$\mathbf{0} = h(s,\hat{\theta}) = L(s) - B(\hat{\theta}), \tag{7}$$

where the function $h : \mathbb{R}^2 \to \mathbb{R}^2$. Since L and B are both smooth, we solve using Newton's method.

We now explain how to find a good initial guess $(s_{\text{guess}}, \hat{\theta}_{\text{guess}})$ for our Newton's method to find a root for h. We saw in the last section that if f_e is the billiard map for an ellipse, then the iterate $f_e(\theta_e, r_e) = (\hat{\theta}_e, \hat{r}_e)$ can be computed in an explicit way, and this calculation also yields s_*. Therefore, for the map f_e ellipse associated with our table, we compute $f_e(\theta, r) = (\hat{\theta}_e, \hat{r}_e)$. We use $s_{\text{guess}} = s_*$, the s value for the ellipse given in (5) and $\theta_{\text{guess}} = \hat{\theta}_e$ for our guess for Newton's method. This is usually a good enough guess to allow us to find $(\hat{\theta}, \hat{r})$ for the perturbed table.

Though the method above works the vast majority of the time, it is not enough to guarantee convergence to the correct root in some small number of cases. We make two modifications. The first modification takes into account that when the ball only bounces a short distance, or more specifically $|r|$ is close to 1, then Newton's method is likely to find the spurious root that occurs at $s = 0$. In order to avoid this problem, and make sure that we converge to the correct root in these cases, we use the method of *deflation* on the function h to eliminate the incorrect zero to which we do not want Newton's method to converge. That is, since we know that h has a root at $(s, \hat{\theta}) = (0, \theta)$, we modify h as follows

$$\tilde{h}(s,\hat{\theta}) = \frac{h(s,\hat{\theta})}{\|(s-0, \hat{\theta}-\theta)\|}.$$

This has the same second root as h but avoids the spurious solution. For a more complete discussion on deflation see [BG71, PW71, Bee15, FBF15].

The second modification is a technical modification in making sure that the initial guess is well defined. It corrects for the fact that $\mathbf{v}$ for the perturbed table might point outside the associated ellipse table. See the discussion and Figure 16 in Appendix A.

After we have found $\hat{\theta}$, the last step is to find $\hat{r}$. Using similar notation as before, we use the notation below. See Fig. 8.

(1) $\hat{T} = B'(\hat{\theta})$ is the tangent vector.
(2) $\hat{\tau}$ is the angle between the tangent line and the positive horizontal axis.
(3) $\hat{\rho}$ is the angle between $\mathbf{v}$ and the positive horizontal axis.
(4) $\hat{\gamma} = \hat{\rho} - \hat{\tau}$. This is the angle between $\mathbf{v}$ and $\hat{T}$, which is up to sign the same as the angle between $\hat{T}$ and the reflected vector $\hat{\mathbf{v}}$. (We never need to compute $\hat{\mathbf{v}}$.)
(5) $\hat{r} = \cos\hat{\gamma}$.

We have found this method to work without failure for small perturbations of ellipses. In particular, since we initialize Newton's method with $(s_{\text{guess}}, \hat{\theta}_{\text{guess}}) = (s_*, \hat{\theta}_e)$, the boundary of a too greatly perturbed table may be insufficiently close to the guess to get convergence. In order to find an upper bound on the perturbation which results in Newton's method converging within a reasonable number of iterations, we studied the table with coefficients given in (3). This table was previously found to have lost convexity around $\epsilon \approx 9$ (see Figure 6), but Newton's method was found to converge reliably only for $\epsilon \in [0, 5]$. Convergence of Newton's was studied in depth for $\epsilon \in [5, 6.5]$. We calculated orbits for an initial location of (θ_0, r_0) with $\theta_0 = 0.3$ and 100 values of r_0 ranging from -0.9 to 0.9. These initial conditions were tested on the table given by (3) for each value of ϵ, where 23 values of ϵ were chosen, ranging from $\epsilon = 5$ to $\epsilon = 6.5$. We call Newton's non-convergent if for any bounce in the first 100 iterations of the orbit, Newton's method did not converge. Any ϵ which resulted in a non-convergent Newton's for any initial condition was said to be too large of a perturbation. We found that for $\epsilon \in [5.0, 5.778]$, the table was sufficiently close to elliptical for our method to converge, while for $\epsilon \geq 5.8434$, the table was non-convergent often enough to make the method impractical.

3.4. **Continuation – analytic billiard maps.** We extend the billiard map analytically to a complex domain via the following steps.

(1) Implicitly define $\mathbf{v}$ as a function of (θ, r): We rewrite our statements about $\tau, \gamma, \theta, T, r$, and $\mathbf{v}$ implicitly as a root finding problem.

$$g(\tau, \gamma, \theta, r) = \begin{pmatrix} \cos\gamma - r \\ \\ T_1 \sin\tau - T_2 \cos\tau \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \mathbf{v} = \begin{pmatrix} \cos(\gamma + \tau) \\ \sin(\gamma + \tau) \end{pmatrix}$$

Note that g is perfectly well defined for complex values of τ, γ, θ, and r. Indeed, g is complex analytic in these variables.

By implicit differentiation

$$\frac{\partial \mathbf{v}}{\partial(\theta, r)} = \frac{\partial \mathbf{v}}{\partial(\gamma, \tau)} \frac{\partial(\gamma, \tau)}{\partial(\theta, r)}$$

so that

$$\frac{\partial(\gamma, \tau)}{\partial(\theta, r)} = -\left(\frac{\partial g}{\partial(\gamma, \tau)}\right)^{-1} \frac{\partial g}{\partial(\theta, r)}$$

$$\frac{\partial \mathbf{v}}{\partial \gamma} = \frac{\partial \mathbf{v}}{\partial \tau} = \begin{pmatrix} -\sin(\gamma + \tau) \\ \cos(\gamma + \tau) \end{pmatrix}$$

$$\frac{\partial g}{\partial(\theta, r)} = \begin{pmatrix} -1 & 0 \\ 0 & \dfrac{\partial T_1}{\partial \theta}\sin\tau - \dfrac{\partial T_2}{\partial \theta}\cos\tau \end{pmatrix}$$

$$\frac{\partial g}{\partial(\gamma, \tau)} = \begin{pmatrix} -\sin\gamma & 0 \\ 0 & T_1\cos\tau + T_2\sin\tau \end{pmatrix}$$

The term $\partial g/\partial(\gamma,\tau)$, is non-singular under the following condition, which holds for any well defined billiard iterate.

(a) $\sin\gamma \neq 0$ meaning that $\gamma \neq 0, \pi$. These values correspond to the cases where the direction of bounce is tangent to the boundary – which is not a possibility for any convex billiard table.
(b) $T_1\cos\tau + T_2\sin\tau \neq 0$. If this was equal to zero and $g = \mathbf{0}$, then we can show that $T_1 = T_2 = 0$.

(2) Implicitly define $(s, \hat{\theta})$ as a function of (θ, r): We quickly restate the function h so that we can write down the derivative.

$$h(\theta, r, s, \hat{\theta}) = B(\theta) + s\mathbf{v} - B(\hat{\theta}) = \mathbf{0}$$

To compute the derivative, define $\hat{T} = B'(\hat{\theta})$ and $T = B'(\theta)$. Using implicit differentiation, we get

$$\frac{\partial(s,\hat{\theta})}{\partial(\theta, r)} = -\begin{pmatrix} \mathbf{v} & | & -\hat{T} \end{pmatrix}^{-1} \begin{pmatrix} s\dfrac{\partial \mathbf{v}}{\partial r} & | & T + s\dfrac{\partial \mathbf{v}}{\partial \theta} \end{pmatrix}$$

For the part in the inverse to be singular, we would need $\hat{T}_2/\hat{T}_1 = v_2/v_1$. This only would occur if $\mathbf{v}$ and $\hat{T}$ were parallel, which cannot happen for a convex billiard table.

(3) Implicitly define $\hat{r}$ as a function of (θ, r): This is formulated implicitly as follows.

$$k(\theta, r, \hat{\rho}, \hat{\tau}) = \begin{pmatrix} \hat{T}_1\sin\hat{\tau} - \hat{T}_2\cos\hat{\tau} \\ v_1\sin\hat{\rho} - v_2\cos\hat{\rho} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \text{ where } \hat{r} = \cos(\hat{\rho} - \hat{\tau}).$$

Note that this is written in terms of $\mathbf{v}$ and not $\hat{\mathbf{v}}$. Using implicit differentiation

$$\frac{\partial \hat{r}}{\partial(\theta, r)} = \frac{\partial \hat{r}}{\partial(\hat{\rho}, \hat{\tau})}\frac{\partial(\hat{\rho}, \hat{\tau})}{\partial(\theta, r)}$$

and

$$\frac{\partial(\hat{\rho}, \hat{\tau})}{\partial(\theta, r)} = -\left(\frac{\partial k}{\partial(\hat{\rho}, \hat{\tau})}\right)^{-1}\frac{\partial k}{\partial(\theta, r)}.$$

These partial derivatives are given by

$$\frac{\partial k}{\partial(\hat{\rho}, \hat{\tau})} = \begin{pmatrix} 0 & \hat{T}_1\cos\hat{\tau} + \hat{T}_2\sin\hat{\tau} \\ v_1\cos\hat{\rho} + v_2\sin\hat{\rho} & 0 \end{pmatrix},$$

and

$$\frac{\partial k}{\partial(\theta,r)} = \begin{pmatrix} \dfrac{\partial \hat{T}_1}{\partial r}\sin\hat{\tau} - \dfrac{\partial \hat{T}_2}{\partial r}\cos\hat{\tau} & \dfrac{\partial \hat{T}_1}{\partial \theta}\sin\hat{\tau} - \dfrac{\partial \hat{T}_2}{\partial \theta}\cos\hat{\tau} \\ \dfrac{\partial v_1}{\partial r}\sin\hat{\rho} - \dfrac{\partial v_2}{\partial r}\cos\hat{\rho} & \dfrac{\partial v_1}{\partial \theta}\sin\hat{\rho} - \dfrac{\partial v_2}{\partial \theta}\cos\hat{\rho} \end{pmatrix}, \text{ where } r = \cos\gamma.$$

(4) It only remains to show that $\partial k/\partial(\hat{\rho},\hat{\tau})$ is non-singular when $k = 0$. If the determinant of the Jacobian matrix is zero, then we can show that either $\mathbf{v}$ or $\hat{T}$ is the zero vector. By our assumptions on $\mathbf{v}$ and B, that cannot happen.

Unlike in the real-valued case, it is not feasible to use an ellipse as an initial guess. In fact, we can continue the ellipse into the complex plane along different paths to get different answers. Therefore to find the iterate of (θ, r) we require an initial guess $(\hat{\theta}_{\text{guess}}, \hat{r}_{\text{guess}})$, as well as guesses approximating the complex-valued angles $\tau, \gamma, \hat{\tau}, \hat{\gamma}$. These guesses come from the exact iterate values for a point (θ_p, r_p) which is close to (θ, r), so the analytic continuation is performed using numerical continuation.

Note that the continuation version of the map can be continued to the case of $\hat{\theta}$ outside of $[0, 1)$, whereas the way that the real version was designed means that it will never be possible to find an iterate outside this range.

4. Parameterization Method for Hyperbolic Periodic Orbits of Planar Maps

In this section, after some initial setup, we present numerical methods for identifying dynamical structures for billiards maps.

4.1. **Multiple shooting.**

In order to study periodic orbits, we set up a multiple shooting map. We consider only the special case of planar maps, as this is the case used in the present work, and this restriction simplifies parts of the discussion. The material in this section and in Section 4.2 is standard, but it is essential for the remainder of the present work. We restate a number of well known results without proof in this context. We refer to [SN10, GMJ17, HLMJ23] for more complete discussion, including many additional references.

The following procedure referred to as *multiple shooting* for a periodic orbit of f. It converts a periodic orbit into a fixed point of the higher-dimensional map F in (8). Let $S \subset \mathbb{R}^2$ be an open and connected set and suppose that $f : S \to \mathbb{R}^2$ is a smooth map. For $k \in \mathbb{N}$, let f^k denote the composition of f with itself k times, and let $f^0 = \mathrm{Id}_{\mathbb{R}^2}$ denote the identity map. We say that $\hat{u} \in S$ is a period-K point for f if

$$f^K(\hat{u}) = \hat{u}.$$

We say that $\hat{u}$ has least period K if $f^k(\hat{u}) \neq \hat{u}$ for $0 < k < K$. Letting $u_1 = \hat{u}$, and $u_j = f(u_{j-1})$ for $1 < j \leq K$, we say that the ordered collection of points $\{u_1, \ldots, u_K\}$, each in S, is a periodic orbit for f, noting that $f(u_K) = u_1$.

Fix K, let $S^K \subset \mathbb{R}^{2K}$ denote the k-fold cartesian product of S. Define the multiple shooting map $F_K : S^K \to \mathbb{R}^{2K}$ by

$$F_K(u_1, u_2, \ldots, u_{K-1}, u_K) = (f(u_K), f(u_1), f(u_2), \ldots, f(u_{K-1})). \tag{8}$$

Note that $\mathbf{u} = (u_1, \ldots, u_K) \in S^K$ satisfies $F_K(\mathbf{u}) = \mathbf{u}$, if and only if $\{u_1, \ldots, u_K\}$ is a period-K orbit of f. That is, fixed points of F_K correspond to period-K orbits of f. Moreover, K is the least period if and only if $u_1 \neq \ldots \neq u_K$. We will suppress the subscript and write

$F_K = F$ when K is clear from context. The Jacobian matrix of F is the following $2K \times 2K$ matrix

$$\mathrm{DF}(\mathbf{u}) = \begin{pmatrix} 0_{2\times 2} & 0_{2\times 2} & \cdots & 0_{2\times 2} & \mathrm{Df}(u_K) \\ \mathrm{Df}(u_1) & 0_{2\times 2} & \cdots & 0_{2\times 2} & 0_{2\times 2} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0_{2\times 2} & 0_{2\times 2} & \cdots & \mathrm{Df}(u_{K-1}) & 0_{2\times 2} \end{pmatrix},$$

where $0_{2\times 2}$ denotes the 2×2 zero matrix and $\mathrm{Df}(u_i)$ is the 2×2 Jacobian of f evaluated at u_i. The virtue of the multiple-shooting framework is that it dispenses with compositions.

Fixed points of F are equivalent to zeros of the map $G : S^K \to \mathbb{R}^{2K}$ defined by

$$G(\mathbf{u}) = F(\mathbf{u}) - \mathbf{u},$$

and zeros of G can be computed using Newton's method. That is, if $\mathbf{u}^0 \in S^K$ has $\|G(\mathbf{u}^0)\|$ small enough, then the Newton sequence $\left\{\mathbf{u}^k\right\}$ defined by

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \boldsymbol{\Delta}^k,$$

where $\boldsymbol{\Delta}^k \in S^K$ solves the linear equation

$$DG(\mathbf{u}^k)\boldsymbol{\Delta}^k = -G(\mathbf{u}^k),$$

converges to $\mathbf{u}_*$ with $G(\mathbf{u}_*) = 0$. The components of $\mathbf{u}_*$ provide a periodic-K orbit for f. Here we have the explicit formula

$$DG(\mathbf{u}) = DF(\mathbf{u}) - \mathrm{Id}_{2K\times 2K},$$

where $\mathrm{Id}_{2K\times 2K}$ denotes the $2K \times 2K$ identity matrix.

4.2. **Stability for periodic orbits.** We state conditions for the stability of periodic orbits. As with the previous section, this material is standard but essential and thus is restated here. The stability of a periodic orbit is defined by applying fixed point stability (Hartman-Grobman theorem/stable manifold theorem) to the composition map f^K. Moreover, the eigenvalues of the matrices $Df^K(u_j)$ do not depend on $1 \le j \le K$. That is, each of these matrices has the same eigenvalues $\lambda_1, \lambda_2 \in \mathbb{C}$, and we refer to these as the multipliers of the periodic orbit. Note, however, that the associated eigenvectors are in general different for different j. The multipliers λ_1, λ_2 can be recovered from the Jacobian of the multiple shooting map. Indeed, using Proposition 1 of [SN10], Lemma 3.1 of [GMJ17] or Proposition 3.1 of [HLMJ23], we get the following lemma.

Lemma 4.1 (Multipliers via the multiple shooting eigenvalues)**.** *Let $\mathbf{u}_* = (u_1, \ldots, u_K)$ be a fixed point of the multiple shooting operator F associated with the map f, so that each u_j is a period-K under f.*

- *The eigenvalues of $DF(\mathbf{u}_*)$ are the K-th roots of the multipliers of the periodic orbit. More precisely, if (α, ξ) is an eigenvalue-eigenvector pair for $DF(\mathbf{u}_*)$ with $\xi = (\xi_1, \ldots, \xi_K)$, then (α^m, ξ_j) is an eigenpair for $Df^K(u_j)$, $1 \le j \le K$.*
- *If $\{u_1, \ldots, u_K\}$ is a period-K saddle orbit, then the multipliers $|\lambda_1| < 1 < |\lambda_2|$ are real, since the matrices $Df^K(u_j)$ are real, and each has a real eigenbasis.*
- *If (α, ξ) is an eigenpair with α and ξ complex then, since $DF(\mathbf{u}_*)$ is a real matrix, we have that $(\alpha, Real(\xi))$ and $(\alpha, imag(\xi))$ are both eigenpairs for $DF(\mathbf{u}_*)$.*

In the present work, the map f is orientation preserving, and the orbits we consider have two positive real multipliers (we did not find orbits with two negative multipliers, but we have not ruled them out). When there are two positive real multipliers, there is always a K-th root which is real, and this is the root we will always choose to work with. This leads to power series solutions with real coefficients, making it especially easy to recover the real image of the parameterization. Note however that, more generally, if $Df^K(u_j)$ has a negative multiplier, and if K is even, then there will be no real K-th root. In this case, recovering the real image of the perameterization is more delicate. This, and more general cases involving higher dimensional manifolds and maps, are discussed in detail in [GMJ17].

4.3. **Identifying orbit types – frequency and chaos.** If f is an area preserving map of the plane, then typical orbits of f are either periodic, quasiperiodic, or chaotic, and the method of weighted Birkhoff averages can be used to distinguish between these, as well as computing the frequency of quasiperiodic orbits. We briefly discuss this method, and refer to [DSSY17, DY18, DSSAY19, SM20, MS21] for more detailed discussion.

Consider an orbit of f $\{(\theta_k, r_k)\}_{k=0}^{\infty}$, where $(\theta_{k+1}, r_{k+1}) = f(\theta_k, r_k)$. The frequency (or rotation number) of the orbit is defined to be the average amount that the θ variable changes in a single iterate. More precisely, define

$$\text{frequency} = \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N-1} (\theta_{k+1} - \theta_k).$$

Under standard smoothness and ergodicity hypotheses, the limit of this average is an invariant of the orbit and does not depend on the initial condition/initial position along the orbit. In practice however this sum converges very slowly, with complexity $O(1/N)$.

The weighted Birkhoff average is a tool which accelerates the convergence, but only along nonchaotic orbits with Diaophantine frequency. This results in spectral convergence, faster than $O(1/N^k)$ for all $k \in \mathbb{N}$. The weighted Birkhoff average still converges slowly for chaotic orbits, and because of this it can be used as a sieve to distinguish chaotic from quasiperiodic orbits.

In order to distinguish chaotic orbits, given a length-N orbit segment, we compute a weighted Birkhoff average on iterates $0, \ldots, N/2$, and compare the answer to the same weighted Birkhoff average for iterates $N/2, \ldots, N$ (in practice it suffices to take $N = 1000$). Since a Birkhoff average depends only on the orbit, if the difference in the two answers is small, this an indication that the average is converging rapidly, and we can classify the orbit as nonchaotic; if the difference is large, the orbit is classified as chaotic. Applying this method to a large number of randomly selected points allows us to identify regions of chaos for billiard maps. These regions are colored gray in all phase plane figures.

Additionally, for nonchaotic orbits, the weighted-Birkhoff average allows us to give a highly accurate computation of the orbit frequency. This method was used to determine the color variation for all phase plane figures.

4.4. **Finding periodic orbits.** In searching for periodic orbits, our goal is to find ones that possess transverse homoclinic orbits. Therefore, we only seek periodic orbits within regions identified as chaotic.

Initial guesses for the multiple shooting Newton method for periodic orbits can be guessed by looking near elliptical islands. Indeed, since periodic orbits typically appear in hyperbolic

and elliptic pairs, and since the elliptic orbits are easy to see from the coloring of the phase space plot, we can often find initial data for hyperbolic periodic orbits by looking "between" the elliptic islands. Using this geometric intuition, it is possible for orbits of small period, such as say any period up to 5, to simply read good enough initial guesses directly off a phase plane plot colored by frequencies location of chaotic orbits, and to have these guesses converge under Newton's method without further effort. Finding periodic orbits of larger periods is a bit more ad hoc and involves several iterations of the same process, including zooming in on the colored phase plane figure to see more detail of the elliptical islands. With a bit of work we were able to find periodic orbits of periods up to 30. (We did not look for orbits beyond period 30.)

4.5. **Parameterization method for one-dimensional spectral submanifolds.** Suppose that $p_0 \in \mathbb{R}^d$ is a hyperbolic fixed point for a real analytic map F. In this case there exist $d_u, d_s \in \mathbb{N}$ with $d_u + d_s = d$ such that $DF(p_0)$ has d_s stable eigevalues $\lambda_1^s, \ldots, \lambda_{d_s}^s \in \mathbb{C}$, and d_u unstable eigenvalues $\lambda_1^u, \ldots, \lambda_{d_u}^u \in \mathbb{C}$. Assume that $DF(p_0)$ is diagonalizable, so that there are $\xi_1^s, \ldots, \xi_{d_s}^s \in \mathbb{R}^d$ associated stable, and $\xi_1^u, \ldots, \xi_{d_u}^u \in \mathbb{R}^d$ unstable eigenvectors. Since $DF(p_0)$ is diagonalizable, these are linearly independent and form a basis for $\mathbb{R}^n$. We are interested in the existence of one-dimensional invariant submanifolds, tangent at p_0 to one of these eigenvectors. Such manifolds are referred to as one-dimensional spectral submanifolds.

One-dimensional submanifolds of the stable/unstable manifold are typically not unique (in fact occur in continua). We select a unique spectral submanifold by imposing maximal regularity. The issue is that, while a typical one-dimensional stable/unstable manifold tangent to an eigenvector is only C^∞ at the fixed point p_0), there is a unique analytic one. The parameterization method is the right tool in this context. It works by looking for an analytic conjugacy to the linear dynamics generated by the eigenvalue. The following lemma provides a single equation which determines a fixed point, its eigendata, and a spectral submanifold. We include the elementary proof for the sake of completeness.

Lemma 4.2 (Parameterization Lemma)**.** *Let $U \subset \mathbb{R}^d$ be an open, connected set, and $F\colon U \to \mathbb{R}^d$ be a real analytic map. Suppose that $\lambda \in \mathbb{R}$, that $P\colon [-1,1] \to \mathbb{R}^d$ is real analytic on $(-1,1)$, that P is continuous on $[-1,1]$, and that the pair (λ, P) satisfy the conjugacy equation*

$$F(P(\sigma)) = P(\lambda\sigma), \qquad \sigma \in [-1,1]. \tag{9}$$

If $P'(0) \neq 0$ and $0 < |\lambda| < 1$, then:

- *$P(0)$ is a fixed point of F,*
- *$(\lambda, P'(0))$ is an eigenpair for $DF(P(0))$,*
- *$P([-1,1])$ is an analytic curve, tangent to the stable eigenvector associated with λ at $\sigma = 0$, and contained in the local stable manifold of $P(0)$.*

Proof. To see that the first point is true, simply evaluate the conjugacy of Equation (9) at $\sigma = 0$ to obtain

$$F(P(0)) = P(0).$$

To see the second point, differentiate Equation (9) with respect to σ to obtain

$$DF(P(\sigma))P'(\sigma) = \lambda P'(\lambda\sigma), \qquad \sigma \in (-1,1),$$

and evaluate at $\sigma = 0$ for

$$DF(P(0))P'(0) = \lambda P'(0),$$

with $P'(0) \neq 0$ by hypothesis.

Finally, to establish the third point note that the fixed point $P(0)$ is in the stable manifold of $P(0)$ by definition. Then choose $\sigma_0 \in [-1, 1]$, $\sigma_0 \neq 0$, and define

$$x_n = P(\lambda^n \sigma_0), \tag{10}$$

for each $n \geq 0$. Note that, since $0 < |\lambda| < 1$ is real, $P(\lambda^n \sigma_0)$ is well defined for each $n \geq 0$.

Now, observe that $\{x_n\}_{n=0}^{\infty}$ is an orbit of F. To see this, apply F to both sides of Equation (10), use the conjugacy of Equation (9), and then the definition x_n to obtain

$$\begin{aligned} F(x_n) &= F(P(\lambda^n \sigma_0)) \\ &= P(\lambda^{n+1} \sigma_0) \\ &= x_{n+1}. \end{aligned}$$

Repeated application of this identity gives

$$F^n(x_0) = x_n.$$

To see that the orbit $\{x_n\}_{n=0}^{\infty}$ is in the stable manifold of $P(0)$, note that combining the previous identity with the continuity of P and the functional equation gives

$$\begin{aligned} \lim_{n\to\infty} F^n(x_0) &= \lim_{n\to\infty} F(x_n) \\ &= \lim_{n\to\infty} P(\lambda^{n+1}\sigma_0) \\ &= P\left(\lim_{n\to\infty} \lambda^{n+1}\sigma_0\right) \\ &= P(0), \end{aligned}$$

again using that $0 < |\lambda| < 1$. Since $\sigma_0 \neq 0$ was otherwise arbitrary, this shows that the image of $[-1, 0) \cup (0, 1]$ under P is a subset of the stable manifold attached to $P(0)$ (and the image of $\sigma = 0$ was mentioned already above). Moreover, the image of P is by definition tangent to $P'(0)$, which is the stable eigenvector associated with the eigenvalue λ by the second bullet point. □

The Parameterization Lemma explains our interest in Equation (9), but leaves open questions of existence and uniqueness. Here we appeal to Theorem 4.1 of [CFdlL05], which deals precisely with the one dimensional analytic situation at hand (in fact, the theorem is a little more general than we need, as it allows for non-hyperbolicity, complex valued maps, and complex eigendata). More precisely, assume that the following non-resonance condition holds

$$\lambda^n \notin \mathrm{Spec}(DF(P(0))), \qquad n \geq 2. \tag{11}$$

From this assumption, Theorem 4.1 of [CFdlL05] concludes that Equation (9) has exactly one analytic solution for each choice of eigenvector $P'(0)$.

The choice of the scaling of the eigenvector determines the radius of convergence of P, and by using the rescaling argument from the proof of Theorem 4.1, one can arrange that P is analytic on a disk of radius $1 + \epsilon$ in $\mathbb{C}$. If F is real analytic then so is P, and this gives real analyticity on $(-1, 1)$ and continunity on $[-1, 1]$ as desired. We stress that the proof of Theorem 4.1 gives that the solution of Equation (9) is globally unique as soon as the length of the eigenvector $P'(0)$ is fixed.

We also note, as observed in Remark 4.2 of [CFdlL05], that the invertibility of $DF(P(0))$ implies that there is an $N \geq 2$ so that Equation (11) holds for all $n \geq N$. So, despite first appearances, the non-resonance condition imposes only a finite number of constraints on the eigenvalues of $DF(P(0))$, and Theorem 4.1 applies generically.

The discussion above motivates the following. Note that we use the usual trick of constraining the square of the norm of the eigenvector, rather than the norm, to obtain an equivalent but simpler formula for the derivative.

Lemma 4.3 (Functional equation, stable case)**.** *Fix $s > 0$ and define the operator $\Psi_s \colon \mathbb{R} \times C^\omega([-1,1],\mathbb{R}^d) \to \mathbb{R} \times C^\omega([-1,1],\mathbb{R}^d)$ by*

$$\Psi_s(\lambda, P(\sigma)) = \begin{pmatrix} \|P'(0)\|^2 - s \\ F(P(\sigma)) - P(\lambda\sigma) \end{pmatrix}$$

If (λ, P) is a zero of Ψ_s with $0 < |\lambda| < 1$, then P satisfies the hypotheses of Lemma 4.2. If the non-resonance condition of Equation (11) *is satisfied, then there exists an $\epsilon > 0$ so that for all $0 < s \leq \epsilon$, Ψ_s has a globally unique solution.*

Consideration of the unstable manifold leads to a similar zero finding problem. Since we assume that $P(0)$ is a hyperbolic fixed point, we have that F is a local diffeomorphism near $P(0)$. Let F^{-1} denote a local inverse with $F^{-1}(P(0)) = P(0)$, and suppose that $Q \colon [-1,1] \to \mathbb{R}^d$ is a real analytic function with

$$F^{-1}Q(\sigma) = Q(\mu\sigma) \qquad \sigma \in [-1,1], \tag{12}$$

with $Q'(0) \neq 0$ and $0 < |\mu| < 1$. Then by Lemma 4.2, $Q(0)$ is fixed for F^{-1}, $(\mu, Q'(0))$ are a stable eigenpair for $DF^{-1}(Q(0))$, and Q parameterizes a subset of the local stable manifold attached to $Q(0)$ and tangent to the eigenvector of μ. It is then a standard result from spectral theory that $(\lambda, Q'(0))$ with $\lambda = \mu^{-1}$ is an unstable eigenpair for $DF(Q(0))$. Moreover, Q parameterizes an arc in the unstable manifold attached to $Q(0)$ and tangent to the eigenvector $Q'(0)$ of λ. Summarizing, we have the following.

Lemma 4.4 (Functional equation, unstable case)**.** *Fix $u > 0$ and define the operator $\Psi_u \colon \mathbb{R} \times C^\omega([-1,1],\mathbb{R}^d)$*

$$\Psi_u(\mu, Q(\sigma)) = \begin{pmatrix} \|Q'(0)\|^2 - u \\ \\ F(Q(\mu\sigma)) - Q(\sigma) \end{pmatrix}$$

If (μ, Q) is a zero of Ψ_u with $0 < |\mu| < 1$, then $Q(0)$ is a fixed point of F, $(\mu^{-1}, Q'(0))$ is an unstable eigenpair for $DF(Q(0))$, and Q parameterizes a one-dimensional arc in the unstable manifold attached to $Q(0)$. This arc is tangent to the unstable eigenspace associated with μ^{-1}. If $\lambda = \mu^{-1}$ satisfies the non-resonance condition of Equation (11)*, then there exists an $\epsilon > 0$ so that for all $0 < u \leq \epsilon$, Ψ_u has a globally unique solution.*

The lemma is established by applying F^{-1} to the second component equation of Ψ_u and observing that this yields Equation (12).

In Section 4.7, we require the Frechét derivatives of Ψ_s and Ψ_u. The Frechét derivatives are written as follows, where $\lambda, \mu, h \in \mathbb{C}$ and $P, Q, H \in C^\omega(\mathbb{B}, \mathbb{R}^d)$.

$$
\begin{aligned}
D_\lambda \Psi_s(\lambda, P(\sigma))[h] &= \begin{pmatrix} 0 \\ -P'(\lambda\sigma)\sigma h \end{pmatrix}, \\
D_P \Psi_s(\lambda, P(\sigma))[H] &= \begin{pmatrix} 2\langle P'(0), H'(0)\rangle \\ DF(P(\sigma))H(\sigma) - H(\lambda\sigma) \end{pmatrix}, \\
D_\mu \Psi_u(\mu, P(\sigma))[h] &= \begin{pmatrix} 0 \\ DF(Q(\mu\sigma))Q'(\mu\sigma)\sigma h \end{pmatrix}, \\
D_P \Psi_u(\mu, P(\sigma))[H] &= \begin{pmatrix} 2\langle Q'(0), H'(0)\rangle \\ DF(Q(\mu\sigma))H(\mu\sigma) - H(\sigma) \end{pmatrix}
\end{aligned} \tag{13}
$$

Remark 4.1 (Generalizations)**.** The parameterization method generalizes to the case when there are resonances. However, one cannot then analytically conjugate to linear dynamics. Instead, one studies a conjugacy equation of the form

$$F(P(\sigma)) = P(K(\sigma)),$$

where K is a polynomial map whose degree and monomial terms are determined by the resonances. More than this, the parameterization method extends to higher dimensional spectral submanifolds, non-hyperbolic fixed points, C^k regularity, and even to other kinds of invariant objects, for both maps and vector fields defined on Banach spaces. Much more general treatment, and many additional references, are found in the paper [CFdlL03], and also to the book [HCF$^+$18]. See also Appendix B of [CFdlL05] for a scholarly discussion of the historical development of these ideas.

4.6. **Stable and unstable manifolds of a planar hyperbolic periodic orbit.** As an application of the ideas presented in the previous section, we discuss a parameterization method for stable/unstable manifolds of periodic orbits of maps. By restricting to saddle type orbits of planar maps, we obtain an especially complete picture. Note, the material in this section is adapted from [GMJ17], where generalizations to higher dimensional maps and manifolds are also found. We refer also to the paper [TMJ22] where extensions to implicitly defined maps given by polynomial relations are considered.

Let $S \subset \mathbb{R}^2$ be an open and connected set, and suppose that $f\colon S \to \mathbb{R}^2$ is a real analytic mapping. Let us also assume that f is orientation preserving (as is the case for the mathematical billiard maps considered in the present work). Let $F\colon S^K \to \mathbb{R}^{2K}$ be the multiple shooting map for f as defined by Equation (8). Let Ψ_s be the operator defined in Lemma 4.3 where $s > 0$ is fixed. Suppose that for a fixed value of λ with $0 < |\lambda| < 1$ and a fixed function $P\colon [-1, 1] \to \mathbb{R}^{2K}$ we have that $\Psi_s(\lambda, P) = \mathbf{0}$. Combining Lemma 4.1 from Section 4.1 with Lemma 4.2 from Section 4.5, we have the following:

- $P(0) \in \mathbb{R}^{2K}$ is a fixed point of the multiple shooting map F. Letting $P(0) = (p_1, \ldots, p_K)$ with $p_1, \ldots, p_K \in \mathbb{R}^2$, each p_j has period K for f. That is, the components of $P(0)$ provide a periodic orbit for f. If $p_k \neq p_j$ when $j \neq k$, then K is the least period.
- $\lambda^m = \alpha$ is a stable multiplier for the periodic orbit. Moreover, writing $P'(0) = (\xi_1, \ldots, \xi_K)$ with $\xi_1, \ldots, \xi_K \in \mathbb{R}^2$, we have that ξ_j is a stable eigenvector for $Df^K(p_j)$.

- Let $P(\sigma) = (P_1(\sigma), \ldots, P_K(\sigma))$ with each P_j taking values in $\mathbb{R}^2$. Then $P_j(\sigma)$ parameterizes a subset of the stable manifold attached to p_j, and is tangent at zero to the ξ_j.

Only the last point requires justification. Note that in the multiple shooting case, the functional equation

$$F(P(\sigma)) = P(\lambda\sigma),$$

can be expressed in terms of f as

$$\begin{aligned} f(P_K(\sigma)) &= P_1(\lambda\sigma) \\ f(P_1(\sigma)) &= P_2(\lambda\sigma) \\ f(P_2(\sigma)) &= P_3(\lambda\sigma) \\ &\vdots \\ f(P_{K-1}(\sigma)) &= P_K(\lambda\sigma) \end{aligned}$$

Applying f to the last equation gives

$$f^2(P_{K-1})(\sigma) = f(P_K(\lambda\sigma)),$$

combining this with the first equation leads to the following.

$$f^2(P_{K-1})(\sigma) = P_1(\lambda^2\sigma).$$

Continuing in this way, we obtain

$$f^K(P_1(\sigma)) = P_1(\lambda^K\sigma).$$

But this is the parameterization method for the map f^K.

Restarting the whole argument from the second to last equation leads to

$$f^K(P_2(\sigma)) = P_2(\lambda^K\sigma),$$

and working our way through each component equation in a similar way leads to

$$f^K(P_j(\sigma)) = P_j(\lambda^K\sigma).$$

for each $1 \leq j \leq K$.

Since we assume that the orbit is hyperbolic, and that the mapping f is a planar, real analytic, and orientation preserving, we have that the orbit has two real multipliers α, β with $0 < |\alpha| < 1$ and $|\beta| > 1$. Moreover, $\alpha\beta = 1$, so that they are either both positive, or both negative. In the present work we consider only the case when both multipliers are positive. In this case the eigenvalues of $DF(P(0))$ sort into two sets: the K complex K-th roots of α, and the K complex K-th roots of β. So, every eigenvalue of $DF(P(0))$ has either $|\lambda| = \alpha^{1/K}$ or $|\lambda| = \beta^{1/K}$. Since all the stable eigenvalue have the same (non-unit) modulus, no power of one can equal another, and the non-resonance conditions

$$\lambda^n \notin \mathrm{Spec}DF(P(0)),$$

holds for all $n \geq 2$, and similarly for the unstable eigenvalues.

So, the non-resonance conditions are automatically satisfied and the existence of the desired parameterization is assured. Moreover, since α and β are positive, each has a real K-th root, and associated real eigenvectors. Then the operators Ψ_s and Ψ_u defined in Equations (4.3)

and (4.4) have unique, real solutions as long as u and s remain in the range $0 < s, u < \epsilon$ for some $\epsilon > 0$. In practice it may not necessarily be the case that ϵ is small, and numerical heuristics are discussed below.

4.7. **Numerical computation of the local parameterizations.** Let

$$P(\sigma) = \sum_{n=0}^{\infty} p_n \sigma^n = \sum_{n=0}^{\infty} \begin{pmatrix} p_n^1 \\ \vdots \\ p_n^d \end{pmatrix} \sigma^n, \tag{14}$$

denote a formal power series with $p_n \in \mathbb{R}^d$. Now choose $N \in \mathbb{N}$ and write

$$P^N(\sigma) = \sum_{n=0}^{N} p_n \sigma^n, \tag{15}$$

to denote the truncation of P to order N. We then embed the coefficients of P^N in $\mathbb{R}^{d(N+1)}$ by forming the column vector (written as a row to save space)

$$\mathbf{p} = (p_0^1, \ldots, p_N^1, p_0^2, \ldots, p_N^2, \ldots, p_0^d, \ldots, p_N^d) \in \mathbb{R}^{d(N+1)}. \tag{16}$$

That is, for each $0 \leq n \leq N$, the numbers p_n^i with $1 \leq i \leq d$ are the components of $p_n \in \mathbb{R}^d$.

Now, given $\mathbf{p} \in \mathbb{R}^{d(N+1)}$, we write

$$\mathbf{q} = \mathrm{DFT}[F](\mathbf{p}),$$

to signify that the column vector

$$\mathbf{q} = (q_0^1, \ldots, q_N^1, q_0^2, \ldots, q_N^2, \ldots, q_0^d, \ldots, q_N^d) \in \mathbb{R}^{d(N+1)},$$

is the vectorization of the series coefficients of

$$Q^N(\sigma) = \sum_{n=0}^{N} q_n \sigma^n \approx F(P^N(\sigma)), \qquad |\sigma| < 1,$$

where for each $1 \leq i \leq d$ we compute the series coefficients $\{q_0^i, \ldots, q_N^i\}$ by interpolating the components $F_i \circ P^N$, $1 \leq i \leq d$ of the composition via the DFT interpolation scheme discussed in Appendix B.4.

Similarly, for $1 \leq i, j \leq d$ define the coefficient sequences

$$\mathbf{b}^{ij}(\mathbf{p}) = (b_0^{ij}, \ldots, b_N^{ij}),$$

which result from applying the DFT algorithm discussed in Appendix B.4 to the composition maps $DF_{ij} \circ P$. That is, we let

$$\mathbf{b}^{ij}(\mathbf{p}) = \mathrm{DFT}\left[\partial_j F_i\right](P^N),$$

as these coefficients have

$$\sum_{n=0}^{N} b_n^{ij} \sigma^n \approx \partial_j F_i(P^N(\sigma)) \qquad |\sigma| < 1,$$

where F_i, $1 \leq i \leq d$ denote the component maps of F. We then define the d^2 many $(N+1) \times (N+1)$ matrices

$$B^{ij}(\mathbf{p}) = \mathrm{ConvMat}(\mathbf{b}^{ij}(\mathbf{p})), \tag{17}$$

where the convolution matrix is as defined in Equation (25) of Appendix B.5.

To vectorize the composition $P(\lambda\sigma)$ and its derivatives we first let Λ, Σ denote the $(N+1)\times(N+1)$ diagonal matrices

$$\Lambda = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & \lambda & 0 & \dots & 0 \\ 0 & 0 & \lambda^2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \lambda^N \end{pmatrix}, \qquad \text{and} \qquad \Sigma = \begin{pmatrix} 0 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & N \end{pmatrix},$$

and define $\mathbf{L}_\lambda$ and $\mathbf{N}_N$, the $d(N+1)\times d(N+1)$ diagonal matrices

$$\mathbf{L}_\lambda = \begin{pmatrix} \Lambda & \mathbf{0}_{\mathbb{R}^{N+1}} & \dots & \mathbf{0}_{\mathbb{R}^{N+1}} \\ \mathbf{0}_{\mathbb{R}^{N+1}} & \Lambda & \dots & \mathbf{0}_{\mathbb{R}^{N+1}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{\mathbb{R}^{N+1}} & \mathbf{0}_{\mathbb{R}^{N+1}} & \dots & \Lambda \end{pmatrix}, \qquad \text{and} \qquad \mathbf{N}_N = \begin{pmatrix} \Sigma & \mathbf{0}_{\mathbb{R}^{N+1}} & \dots & \mathbf{0}_{\mathbb{R}^{N+1}} \\ \mathbf{0}_{\mathbb{R}^{N+1}} & \Sigma & \dots & \mathbf{0}_{\mathbb{R}^{N+1}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{\mathbb{R}^{N+1}} & \mathbf{0}_{\mathbb{R}^{N+1}} & \dots & \Sigma \end{pmatrix},$$

and note that $\mathbf{L}_\lambda\mathbf{p}$ and $\mathbf{N}_N\mathbf{p}$ are, respectively, the vectorizations of the coefficients of $P^N(\lambda\sigma)$ and $\sigma(P^N)'(\sigma)$.

Truncating the domain and range of Ψ_s at order N and vectorizing the inputs and outputs leads to the map $\mathbf{\Psi}_s^N\colon \mathbb{R}^{d(N+1)+1} \to \mathbb{R}^{d(N+1)+1}$ given by

$$\mathbf{\Psi}_s^N(\lambda, \mathbf{p}) = \begin{pmatrix} (p_1^1)^2 + \ldots + (p_1^d)^2 - s \\ \\ \mathrm{DFT}[F](\mathbf{p}) - \mathbf{L}_\lambda\mathbf{p} \end{pmatrix}. \tag{18}$$

To vectorize the derivative, we recall the formulas for the action of the Frechét derivatives given in Equation (13). First, for $1 \le n \le N$ let e_n^* denote the $N+1$ dimensional row vector with a one in the n-th entry and zeros elsewhere. Keep in mind that we start with a zeroeth element. Therefore the row vector e_1^* is given by

$$e_1^* = (0, 1, 0, \dots, 0).$$

So, if

$$H^N(\theta) = \sum_{n=0}^{N} h_n \sigma^n,$$

with $h_n \in \mathbb{R}^d$, then the first component of the partial derivative $D_P\Psi_s(\lambda, P)[H]$ given in Equation (13) is

$$\begin{aligned} 2\left\langle (P^N)'(0), (H^N)'(0) \right\rangle &= 2\left(p_1^1 h_1^1 + \ldots + p_1^d h_1^d\right) \\ &= \begin{pmatrix} 2p_1^1 e_1^* & | & 2p_1^2 e_1^* & | & \dots & | & 2p_1^d e_1^* \end{pmatrix} \begin{pmatrix} h_0^1 \\ \vdots \\ h_N^1 \\ \vdots \\ h_0^d \\ \vdots \\ h_N^d \end{pmatrix}. \end{aligned}$$

Suppressing the action on H^N, the partial derivative is represented simply by the $1 \times d(N+1)$ row vector

$$n(\mathbf{p}) = \left(\begin{array}{ccccccc} 2p_1^1 e_1^* & | & 2p_1^2 e_1^* & | & \dots & | & 2p_1^d e_1^* \end{array} \right).$$

Similarily, the nontrivial component of the partial derivative $D_\lambda \Psi_s(\lambda, P)$ given in Equation (13) is the negative of $(P^N)'(\lambda\sigma)\sigma$, and this has vector representation as the $d(N+1)$ dimensional column vector

$$\mathbf{c}_\lambda(\mathbf{p}) = \mathbf{N}_N \mathbf{L}_\lambda(\mathbf{p}).$$

For the second component of partial derivative $D_P \Psi_s(\lambda, P)[H]$, following Equation (13), we vectorize $DF(P^N)$ as follows. Let $\mathbf{B}(\mathbf{p})$ denote the $d(N+1) \times d(N+1)$ matrix

$$\mathbf{B}(\mathbf{p}) = \left(\begin{array}{ccc} B^{00}(\mathbf{p}) & \dots & B^{0d}(\mathbf{p}) \\ \vdots & \ddots & \vdots \\ B^{d0}(\mathbf{p}) & \dots & B^{dd}(\mathbf{p}) \end{array} \right),$$

where for $1 \le i, j \le N$ the $B^{ij}(\mathbf{p})$ are the $(N+1) \times (N+1)$ convolution matrices defined in Equation (17). So, the derivative of $\mathbf{\Psi}_s^N$ has matrix representation

$$D\mathbf{\Psi}_s^N(\mathbf{p}) = \left(\begin{array}{cc} 0 & n(\mathbf{p}) \\ -\mathbf{c}_\lambda(\mathbf{p}) & \mathbf{B}(\mathbf{p}) - \mathbf{L}_\lambda(\mathbf{p}) \end{array} \right).$$

By a nearly identical argument, using the same notation, one checks that the vectorization of Ψ_u and its derivative are

$$\mathbf{\Psi}_u^N(\mu, \mathbf{p}) = \left(\begin{array}{c} (p_1^1)^2 + \ldots + (p_1^d)^2 - u \\ \\ \mathrm{DFT}[F](\mathbf{L}_\mu \mathbf{p}) - \mathbf{p} \end{array} \right), \tag{19}$$

and

$$D\mathbf{\Psi}_u^N(\mathbf{p}) = \left(\begin{array}{cc} 0 & n(\mathbf{p}) \\ \\ \mathbf{B}(\mathbf{L}_\mu \mathbf{p}) \mathbf{c}_\mu(\mathbf{p}) & \mathbf{B}(\mathbf{L}_\mu \mathbf{p}) \mathbf{L}_\mu - \mathbf{I} \end{array} \right),$$

where $\mathbf{I}$ is the $d(N+1) \times d(N+1)$ identity matrix.

Now consider the special case where $F = F_K$ is the multiple shooting map defined in Equation (8), with underlying planar map f, so that $d = 2K$. Let $\bar{\mathbf{p}}_0 \in \mathbb{R}^{2K}$ be an approximate fixed point of the multiple shooting map, in which case the planar components of $\bar{\mathbf{p}}_0$ give an approximate period K orbit for f. Suppose that these components are distinct so that K is the least period.

Assume that $(\bar{\lambda}, \bar{\mathbf{p}}_1)$ is a (numerically computed approximate) real stable (or unstable) eigenpair for $DF^K(\bar{\mathbf{p}}_0)$. Fix s (or u) a nonzero real number, and rescale $\bar{\mathbf{p}}_1$ so that $\|\bar{\mathbf{p}}_1\|^2 = s$ (or u). Let

$$P^0(\sigma) = \bar{\mathbf{p}}_0 + \sigma \bar{\mathbf{p}}_1. \tag{20}$$

If the data $\bar{\lambda}, \bar{\mathbf{p}}_0, \bar{\mathbf{p}}_1$ were exact, then P^0 would be a quadratically good approximate solution of $\Psi_{s,u} = 0$. That is, we would have

$$\|\Psi_s(P^0)\| \le C s^2.$$

So, assuming that $\bar{\mathbf{p}}_0$ and $\bar{\mathbf{p}}_1$ are good enough approximations, we expect to be able to find an $s, u > 0$ so that this initial error is small (limited only by the error in the initial data which is, presumably, on the order of machine epsilon).

Let $\mathfrak{p}^0$ be the vectorization of $P^0(\sigma)$, and define the Newton sequence

$$\mathfrak{p}^{m+1} = \mathfrak{p}^m + \boldsymbol{\Delta}^m,$$

where, for $m \geq 0$, $\boldsymbol{\Delta}^m$ solves the linear equation

$$D\boldsymbol{\Psi}^N_{s,u}(\mathfrak{p}^m)\boldsymbol{\Delta}_m = -\boldsymbol{\Psi}^N_{s,u}(\mathfrak{p}^m).$$

In practice we experiment with the scalings s, u, as we would like to choose the largest such values so that the Newton sequence converges. More sophisticated continuation schemes and optimization algorithms for choosing computational parameters for the parameterization method are discussed in [BLMJ16, HMJ21]

Remark 4.2 (Automatic reducibility)**.** Note that the Newton method just described is a "large matrix method". More precisely, if K is the period of the orbit, and N is the polynomial order of approximation, then the linear system determining the Newton step has size $(2K(N+1)+1) \times (2K(N+1)+1)$. While this presents no problem in the present work, the complexity can be greatly reduced by exploiting the symplectic structure of the billiard system to define a diagonal approximate inverse for the Newton matrix. This is referred to as "approximate reducibility" in the literature, and the interested reader will find much more on this topic, including many references to its use in the literature, in the book [HCF$^+$18]. Using approximate reducibility would be helpful for studying orbits with much higher period than those studied here.

Remark 4.3 (The method of jet transport)**.** We note that an alternative to the DFT based interpolation scheme described above would be to compute necessary compositions using the method of jet transport described in [GJJCZ25]. Indeed, the reference just cited uses jet transport to solve equations similar to the conjugacy equations considered in the present work, applying these ideas to Poincaré maps for ODEs in celestial mechanics. These implicit Poincaré maps are, if anything, more complicated than the billiard maps studied in the present work, so we assume that the method or jet transport could be applied to billiards as well. The only disadvantage of jet transport is that it requires special software like the TAYLOR package developed by the authors of the reference just cited. Our implementation relies only on standard numerical linear algebra routines like linear system solvers, eigenvalue/eigenvector solvers, and (if desired) the FFT. These algorithms are available in any standard numerical linear algebra package. Our codes for example are implemented in MATLAB.

A useful a-posteriori error indicator is obtained by measuring the defect in the conjugacy relation. So, define $E_{\text{conj}} \colon \mathbb{R} \times C^\omega(\mathbb{D}, \mathbb{R}^{2K}) \to \mathbb{R}$ by

$$E_{\text{conj}}(\lambda, P) = \sup_{\sigma \in \mathbb{D}} \|F(P(\sigma)) - P(\lambda\sigma)\|_{\mathbb{R}^{2K}}$$

Note that $E_{\text{conj}}(\bar{\lambda}, \bar{P}) = 0$ if and only if $(\bar{\lambda}, \bar{P})$ is the desired solution. A numerical proxy for the a-posteriori indicator is obtained by sampling the right hand side at $M \in \mathbb{N}$ points on the boundary of $\mathbb{D}$, and computing the maximum of this finite set. More precisely, by the maximum modulus principle we have that

$$E^M_{\text{conj}}(\lambda, P) = \max_{0 \leq m \leq M} \left\| F\left(P\left(e^{\frac{2\pi i m}{M+1}}\right)\right) - P^N\left(\lambda e^{\frac{2\pi i m}{M+1}}\right)\right\|_{\mathbb{R}^{2K}},$$

bounds $E_{\text{conj}}(\lambda, P)$, is computable in finitely many operations, and converges to the true a-posteriori indicator as $M \to \infty$. Note also that, in the case that $P = P^N$ is polynomial, it is efficiently evaluated at the desired roots of using the FFT.

4.8. **Growing the local manifolds.** To compute larger portions of the stable and unstable manifolds, we employ a fundamental domain approach that leverages the periodic structure of the system. Let $\mathcal{W}^s$ and $\mathcal{W}^u$ denote the global stable and unstable manifolds, respectively.

We work within a fundamental domain $\mathcal{D} \subset \mathbb{R}^2$, a bounded region such that every point on the global manifold can be obtained by applying forward or backward iterates of the map F to points in $\mathcal{D}$.

Given the local parameterizations $P^s(\theta)$ and $P^u(\theta)$ of the stable and unstable manifolds, the global manifolds are constructed by first selecting parameterization points θ_j such that $P^{s,u}(\theta_j) \in \mathcal{D}$. Then for each selected point $x_j = P^{s,u}(\theta_j)$, where $j = 1, \ldots, J$, compute:

$$\mathcal{W}^u_{\text{global}} = \bigcup_{j=1}^{J} \bigcup_{n=0}^{M} F^n(x_j^u) \tag{21}$$

$$\mathcal{W}^s_{\text{global}} = \bigcup_{j=1}^{J} \bigcup_{n=0}^{M} F^{-n}(x_j^s)$$

where M is the maximum number of iterates. This approach ensures that the computed global manifolds cover the relevant phase space region while avoiding redundant computation outside the fundamental domain.

5. Results

In this section, we give our results. As mentioned in the introduction, the importance of our findings are twofold: the novelty of the use of the method, and also the implications for billiard maps. Correspondingly, we divide this section into two parts.

5.1. **Results for the parameterization method.** To compute the local stable and unstable manifolds, we solve the functional equations from Section 4.5:

$$F(P(z)) = P(\lambda z), \qquad \text{and} \qquad F(P(\mu z)) = P(z),$$

with $|\lambda| < 1$ for the stable case and $|\mu| = 1/|\lambda| > 1$ for the unstable case. These equations are solved numerically via Newton iteration formulated in the space of spectral coefficients, as described in Sections B and 4.7. The two crucial sets of parameters that we vary in our numerics are the truncation order of the polynomial approximation N, introduced in Equation (15) and the scaling coefficients. The scaling coefficients are labeled s in Lemma 4.3 and u in Lemma 4.4, but in this section and in our Tables of numerical parameters, we refer to these parameters by $s = \text{scale}_s$ and $u = \text{scale}_u$. The resulting local manifolds are shown in the left panels of Figures 9–11 for the period-10 orbit for Table D, the period-3 orbit for Table C, and the period-2 orbit for Table B (cf. Tbl. 1), respectively. The corresponding global manifolds are obtained by iterating the local parameterizations under F. The right panels of the same figures illustrate these global structures, together with the truncation order N of the polynomial approximation $P^N(z)$ and the number of iterates used.

The novelty of the present computation lies in performing the Newton iteration directly in the spectral domain, allowing both stable and unstable parameterizations to be obtained with

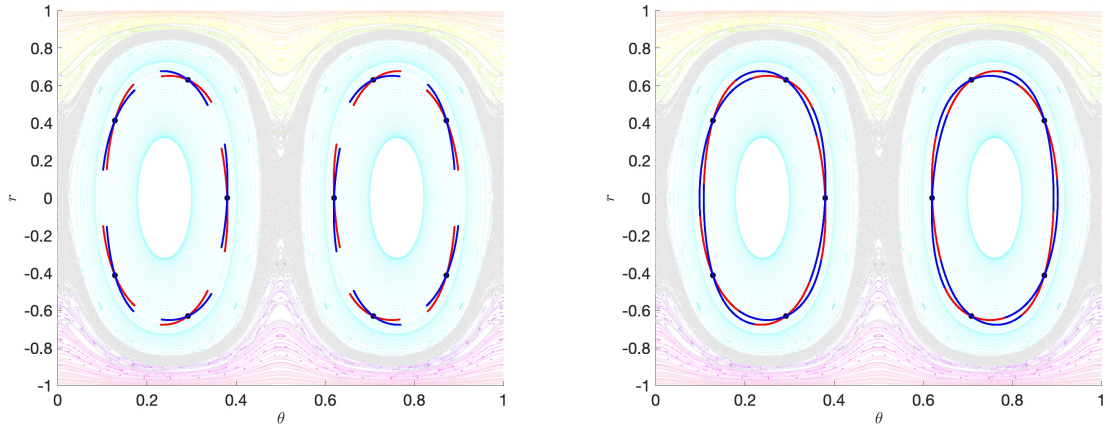

FIGURE 9. Local (left) and global (right) stable (blue) and unstable (red) manifolds for a period-10 orbit for Table D (cf. Tbl. 1). The global manifolds were computed with 20 iterates of the fundamental domain and $N = 60$.

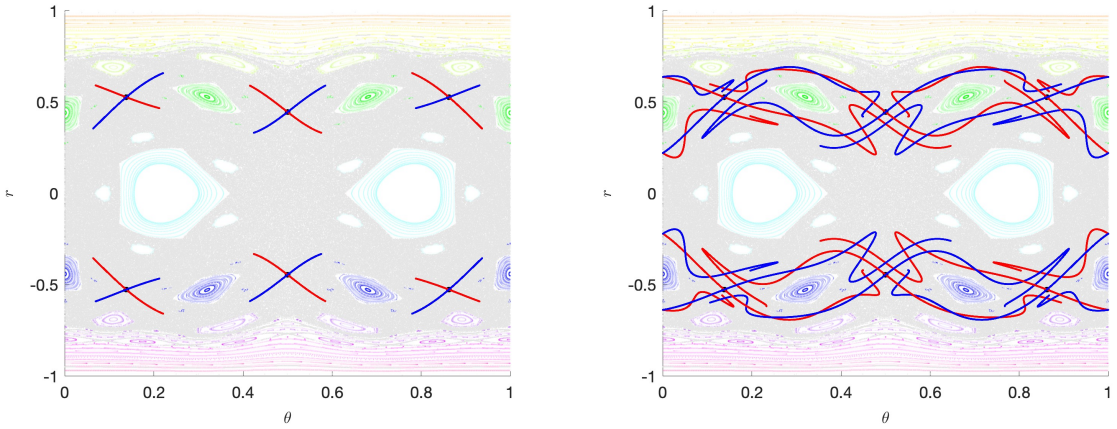

FIGURE 10. Local (left) and global (right) stable (blue) and unstable (red) manifolds for a period-3 orbit for Table C (cf. Tbl. 1). The global manifolds were computed with 6 iterates of the fundamental domain and $N = 60$.

high-order accuracy and without explicit recursion on the coefficients. In contrast to traditional recursive or pointwise approaches, this formulation combines spectral evaluation and FFT-based transforms to achieve rapid convergence, ensuring both robustness and precision in the manifold computation.

To verify the spectral accuracy of the computed parameterizations, we analyze the decay of the coefficients a_n of P^N obtained from the Newton iteration. For both the stable and unstable manifolds, the magnitude of each coefficient is evaluated as $\|a_n\|_\infty$, and its logarithmic scale $\log_{10}(\|a_n\|_\infty)$ is plotted against the mode index n. The resulting coefficient profiles, shown in the center and right panels of Figure 12 for a period-5 orbit of Table B (cf. Tbl. 1),

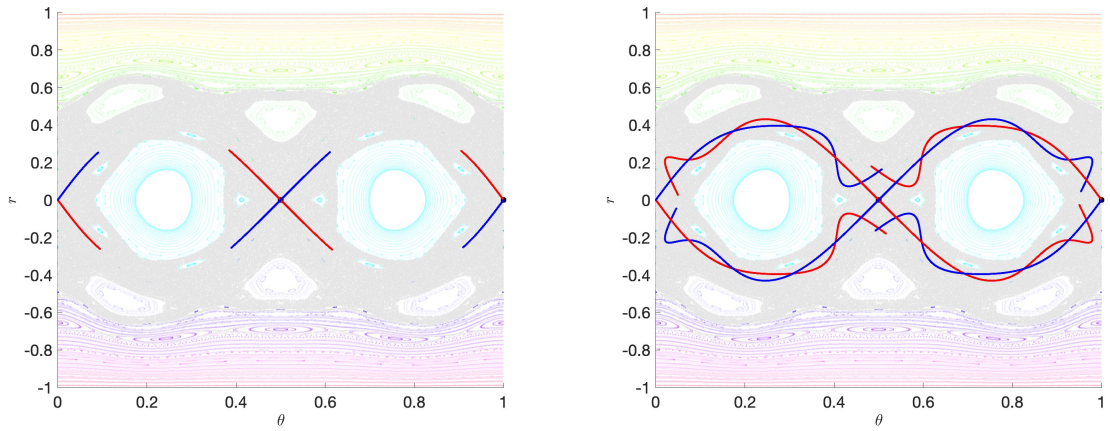

FIGURE 11. Local (left) and global (right) stable (blue) and unstable (red) manifolds for a period-2 orbit for Table B (cf. Tbl. 1). The global manifolds were computed with 4 iterates of the fundamental domain and $N = 60$.

exhibit a rapid exponential decay of the coefficients, reaching values of order 10^{-16} by the truncation index N. This confirms that the spectral representation $P^N(z) = \sum_{n=0}^{N} a_n z^n$ provides a highly accurate analytic approximation of the local manifolds within the domain of convergence of the parameterization.

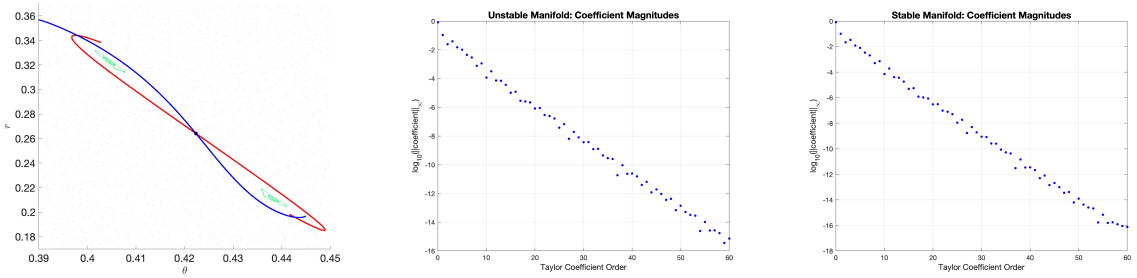

FIGURE 12. (Left) Computed global stable and unstable manifolds for one point of the period-5 orbit of Table B (cf. Tbl. 1). (Center, right) Logarithmic decay of the infinity norm of Taylor coefficients $\|\mathbf{a}_n\|_\infty$ for the local unstable and stable manifolds, respectively. Computations performed using $N = 60$ and 5 iterates of the fundamental domain.

Another key advantage of our approach is that it does not stop at the linear tangent eigenvector chart, but by increasing the truncation order N of the expansion of $P^N(z)$, one can capture the nonlinear deformation of the invariant manifold at larger amplitude. In particular, while the leading term a_1 corresponds to the linear eigenvector scaling, the higher coefficients a_n with $n \geq 2$ encode the curvature, folding and higher-order geometry of the manifold. By systematically computing and monitoring these coefficients, our implementation allows the eigenvector to be scaled to increasingly large magnitudes and hence captures local nonlinear behavior of the manifold near the periodic orbit, see Figures 13-14 for a visual

illustration of this behavior. In both figures, increasing N captures more nonlinear geometry of the manifolds, with scale parameters scale_u and scale_s (see Appendix A.2 for definitions) adjusted accordingly to control the extent of the parameterization from the periodic orbit.

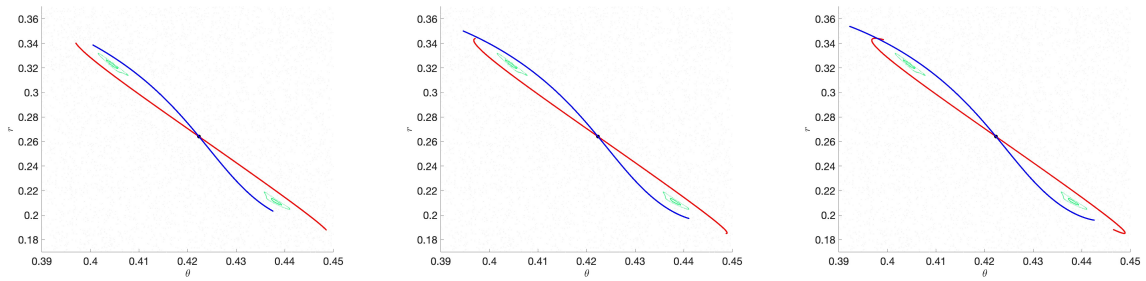

FIGURE 13. Computed local stable and unstable manifolds for the same period-5 orbit from Figure 12 with varying polynomial order N. (Left) $N = 60$, $\text{scale}_u = 0.21$, $\text{scale}_s = 0.2$. (Center) $N = 100$, $\text{scale}_u = 0.26$, $\text{scale}_s = 0.26$. (Right) $N = 180$, $\text{scale}_u = 0.32$, $\text{scale}_s = 0.293$; transverse intersection occur even for the local manifolds.

To assess the robustness of our numerical methods, we systematically increase the perturbation amplitude applied to the coefficients in Table C (cf. Tbl. 1). Specifically, we multiply all coefficients by a factor of 3.5, representing an extreme perturbation far beyond the nominal parameter regime. In this strongly perturbed case, the phase space reveals extensive chaotic behavior: iterates densely populate the entire accessible domain in the (θ, r) plane, forming a nearly uniform cloud of points that obscures much of the underlying structure. The chaotic sea extends throughout the region $0 \leq \theta \leq 1$ and $-0.5 \leq r \leq 0.5$, with particularly dense regions near $r \approx \pm 0.5$ where iterates exhibit wild oscillations. Despite this extreme perturbation—a 250% increase over the nominal values—we successfully compute the stable and unstable manifolds of the hyperbolic periodic-2 orbits located near $(\theta, r) \approx (0.25, 0)$ and $(0.75, 0)$. The successful resolution of these invariant manifolds under such severe perturbation, where the system dynamics have deteriorated into an almost completely chaotic state, demonstrates the remarkable effectiveness of our computational approach even as it operates near the boundary of numerical feasibility.

5.2. **Results for billiards.** Here we concentrate on observations in terms of what our manifolds indicate about the billiard tables. In Appendix A, we have listed all of the periodic orbits found for the five tables. For example, for Billiard Tables A and B, we found a total of seven periodic orbits and manifolds, with periods 2, 3, 5, 10, and 30, with a similar scope in all other cases. In Figures 2–4, we give an overview of all of our computations of invariant manifolds of Tables A–E listed in (2) and Tbl. 1.

In all of the manifolds we computed, either already in the local manifold, or with a small number of iterates, we see transverse intersection of the stable and unstable manifolds of the periodic orbits, implying chaos. The manifolds are trapped inside the connected chaotic region containing the original periodic orbit, wrapping around to avoid the islands, such as the period-2 manifolds wrapping around the two main islands in Figure 12 (left), the period-5 manifolds wrapping around a small pair of islands in Figure 14, and the period-10 manifolds

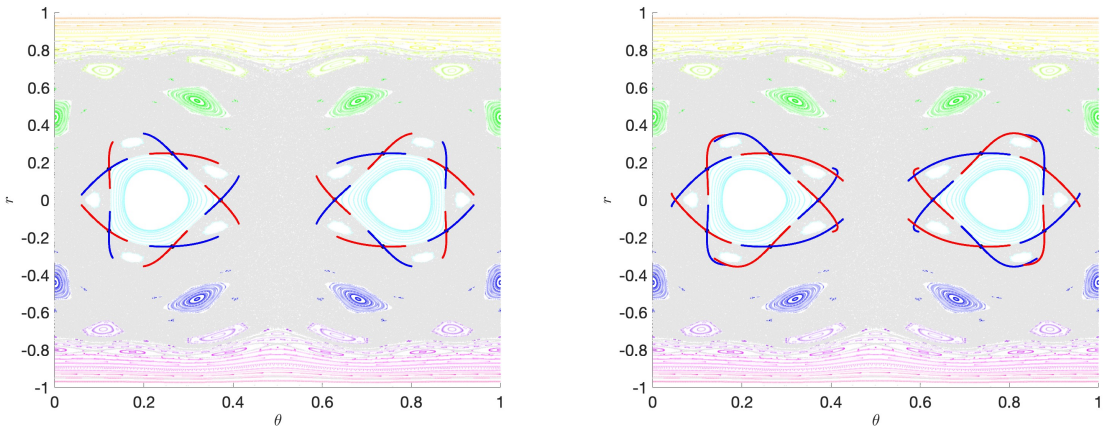

FIGURE 14. Computed local stable and unstable manifolds for the period-10 orbit for Table C (cf. Tbl. 1) with varying polynomial order N. (Left) $N = 60$, $\text{scale}_u = 0.34$, $\text{scale}_s = 0.34$. (Right) $N = 200$, $\text{scale}_u = 0.51$, $\text{scale}_s = 0.51$; transverse intersections occur even for the local manifolds.

wrapping around two star-shape pieces consisting of twelve islands. The larger eccentricity cases have the same behavior, but all the islands are squeezed into a smaller space, meaning that it is harder to observe.

The resulting manifolds and homoclinic tangle beautifully align with the chaotic regions computed using the weighted Birkhoff average. Each set of manifolds and their intersections fills up the entire connected chaotic sea that contains it, giving good agreement between these two methods. This is already somewhat clear in Figures 2-4, although we had kept the number of iterations rather small to avoid overcrowding. It is much more obvious in Figure 15, where we show only one set of manifolds and allow the number of iterations of the local manifolds to be ten.

We are certainly not the first to observe that these maps bear a striking resemblance to the behavior of the Chirikov standard map. In the perturbed case (positive parameter for the Chirikov standard map), they both have elliptic islands, along with small chaotic regions, which are filled with homoclinic tangles of hyperbolic periodic orbits. For a strictly convex billiard table, if the boundary map B is parameterized by arclength, then the billiard map is symplectic [Mei92]. Though we do not assume arclength parameterization, since the frequency is independent of the parameterization, we can conclude from this that the frequency of the rotational tori increases as r increases, as we can observe from the figures of the phase space.

We note that there are a few key differences. The Chirikov standard map has large elliptic period-one islands, whereas here the lowest period possible is two. For the Chirikov standard map, there is a famous parameter value for which all rotational orbits break up [SM20]. In the billiards case, no perturbation breaks up all rotational orbits. In particular, for every table with a C^6 smooth strictly convex boundary, there exist "whispering gallery" orbits [Laz73, Dou82]. Finally, the Chirikov standard map is entire, whereas billiard maps can only be continued to a finite thickness strip of the complex plane. Of particular interest is how well

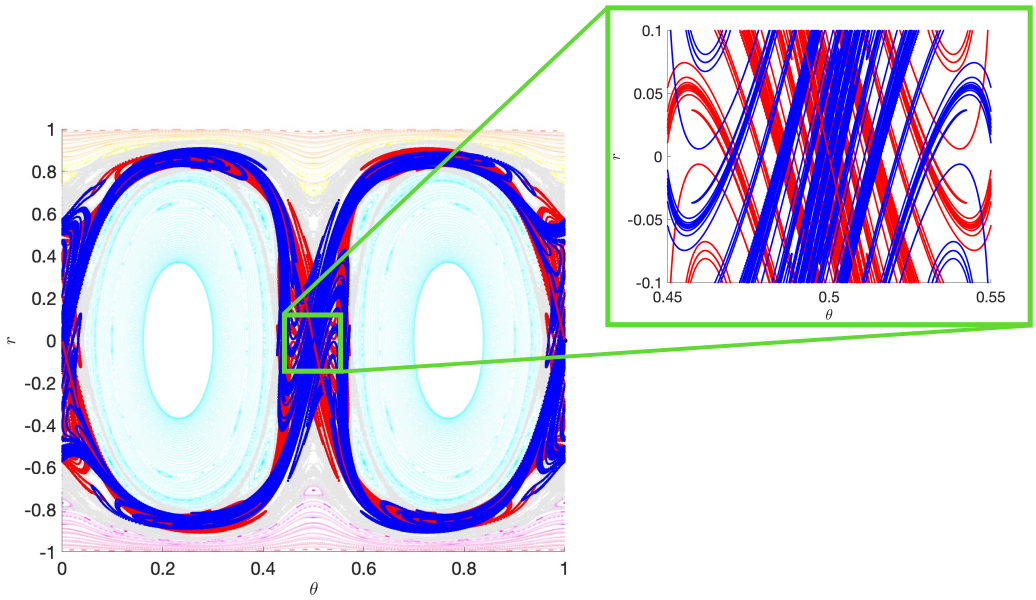

FIGURE 15. (Left) Computed global stable and unstable manifolds for a period-2 orbit for Table E (cf. Tbl. 1), $N = 40$, $\text{scale}_u = 0.52$, $\text{scale}_s = 0.48$ and 10 iterations of the fundamental domain. (Right) Zoomed in plot at one of the period-2 points $(\theta, r) = (0.5, 0)$ to clearly discern chaotic behavior in that particular region.

we can see homoclinic tangle within the chaotic region, such as the Cantor set type behavior of the manifolds, and how well it illustrates in this context the results of the λ-lemma.

The λ–lemma describes a fundamental geometric mechanism by which images of disks (or arcs) transverse to a stable (resp. unstable) manifold are stretched and asymptotically aligned with the unstable (resp. stable) manifold under forward (resp. backward) iteration of the dynamics; informally, a small transverse disk is "carried along" and accumulates onto the invariant manifold while its tangent directions converge to the tangent bundle of that manifold. This statement is a classical ingredient in hyperbolic dynamics and its ramifications for homoclinic and heteroclinic behavior are well documented (see, [Rob98, CFJ20]).

Numerically, the mechanism predicted by the λ–lemma is captured here in two complementary ways. First, the high-order parameterizations $P^N(z)$ provide accurate local charts (disks) transverse to the periodic orbit together with reliable approximations of their tangent bundles: the leading coefficient a_1 recovers the tangent eigenvector while the higher coefficients $a_{n\geq 2}$ encode the local nonlinear deformation. Second, by iterating the computed local chart under the map F (or its inverse) and monitoring both pointwise images and the evolution of tangent information, we observe the characteristic stretching, alignment and accumulation described by the λ–lemma. Practically, the method's spectral accuracy (exponential coefficient decay to machine precision) and the Newton method applied ensure that these iterated images remain within the radius of analyticity of the parameterization and that the observed alignment is not a numerical artifact. Thus, by combining high-order spectral charts with controlled iteration in the fundamental domain, our implementation reproduces

the geometric content of the λ–lemma in a numerically verifiable way. A visualization of this phenomena can be seen in Figure 15 for a period-2 orbit of Table E (cf. Tbl. 1).

6. Conclusions and future plans

In this paper, we have developed methods for computing billiard iteration for billiard tables with smooth convex boundaries. We have then a variety of powerful dynamical tools to such maps. We used the weighted-Birkhoff average to distinguish between chaos and quasiperiodic orbits, at the same time as computing the frequency of the quasiperiodic orbits. We then used the method of multiple-shooting to find periodic orbits for the billiard map. While this method would find arbitrary periodic orbits, we selected the hyperbolic periodic orbits. We then used the parameterization method to compute the local stable and unstable manifolds for the periodic orbits. This use of the parameterization method is novel, in that the billiard map is fully implicit. Therefore, in order to find the coefficients Taylor series approximation of the manifolds, we continued our billiard map analytically into the complex plane, allowing us to compute a Fourier series with the same coefficients as the Taylor series. Finally, we are able to recover global manifolds, up to a chosen number of iterations.

In addition to the computational novelty of this setting, the application to billiards is itself of great interest. Previous results on homoclinic tangles for billiard maps have concentrated on the singular limit as the perturbation approaches the ellipse, showing that there are exponentially small angles of intersection of the manifolds. In contrast, we are interested in homoclinic tangles for billiard tables which are non-limiting perturbations from the ellipse.

One of the advantages of the parameterization method for computing manifolds is that in addition to giving high accuracy approximations, the method can be used for computer assisted proofs. Not just local manifolds, but also a finite number of iterations of these local manifolds can be rigorously validated. In future work, we plan to validate our manifolds. Since the existence of transverse homoclinic points implies chaos, such a validation will lead to a rigorous validation of chaos for perturbed elliptic billiard maps.

The numerical algorithms described in this work have been implemented in MATLAB and are available as open-source code at `https://github.com/efleurantin2103/Billiards`, enabling reproduction of all computational results presented here.

Acknowledgements

SC was partially supported by the GMU Math Industrial Immersion Program. EF was partially supported by the National Science Foundation under Grant No. DMS-2137947. JMJ was partially supported by National Science Foundation grant DMS-2307987. ES was partially supported by the Simons Foundation under Award 636383.

Appendix A. Computational Parameters and Technical Details

This appendix provides the parameters used to compute the global stable and unstable manifolds. It also gives a technical modification for our numerical method for finding an initial guess for f.

A.1. **Modification to find our initial guess.** We now describe our second technical modification to our method for computing the intitial guess in order to perform Newton's method to find the billiard iterate $f(\theta, r)$. To make sure that the initial guess is always well defined, we consider $\mathbf{v}_{-1}$, the previous direction of $\mathbf{v}$, and use the $\mathbf{v}_e$ that is the reflection of $\mathbf{v}_{-1}$ on the elliptical table.

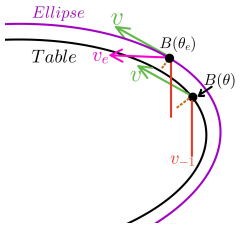

FIGURE 16. The reflection vector v for our table might not point into the interior of the associated elliptical table. To make sure it is always well defined, we use v_e in finding the guess $(\hat{\theta}_e, \hat{r}_e)$.

A.2. We now present the parameters and values of the periodic orbits seen in Figures 2–4 in a series of data Tables, one for each period. More specifically, each data Table corresponds to orbits of a specific period and references the billiard tables in the main text where these orbits are analyzed. The parameter r represents the radial coordinate, θ denotes the parameter describing the point on the boundary of the billiard table, and scale_u and scale_s are the scaling factors applied to the unstable and stable eigenvectors, respectively, in the parameterization method. (These were labeled s and u respectively in the theoretical discussion in Section 4.5.) The variable N indicates the order of the polynomial approximation, while iterations refer to the number of iterates M used in the computational scheme. These parameters were carefully chosen to ensure convergence and accuracy of the computed invariant manifolds associated with each periodic orbit.

Period 2	r	θ	scale_u	scale_s	Iterations	N
Table A	0.0000	0.5000	0.48	0.46	6	60
Table B	0.0	0.5	0.45	0.43	5	60
Table C	0.0	0.5	0.32	0.32	6	60
Table D	0.0	0.5	0.12	0.1	2	10
Table E	0.0	0.5	0.12	0.12	2	10
Period 3	r	θ	scale_u	scale_s	Iterations	N
Table A	0.5333	0.8534	0.25	0.25	10	60
Table A	-0.5333	0.8534	0.25	0.25	10	60
Table B	0.5328	0.3348	0.27	0.27	7	60
Table B	-0.5328	0.3348	0.27	0.27	7	60
Table C	0.4454	0.5	0.28	0.28	6	60
Table C	-0.4454	0.5	0.28	0.28	6	60
Period 4	r	θ	scale_u	scale_s	Iterations	N
Table C	0.7078	0	0.2	0.2	8	60
Table C	-0.7334	0.2152	0.2	0.2	8	60
Table D	0.5	0	0.114	0.114	8	35
Table D	-0.8923	0.2216	0.114	0.114	8	35
Table E	0.89	0.22	0.14	0.115	3	25
Table E	-0.89	0.22	0.14	0.115	3	25

TABLE 2. Numerical parameters for the orbits of periods 2, 3, and 4. The value N is the order of the polynomial approximation given in (14), Iterations is the number M of iterates in (21), and scale_s and scale_u are the scalings of the eigenvectors in the parameterization method defined in Definitions 4.3 and 4.4.

APPENDIX B. NUMERICAL MANIPULATION OF POWER SERIES

Consider a power series

$$P(\sigma) = \sum_{n=0}^{\infty} a_n \sigma^n,$$

with $a_n \in \mathbb{R}$ (or $\mathbb{C}$), and let $g\colon \mathbb{C} \to \mathbb{C}$ be an analytic function. We seek the power series coefficients of

$$(g \circ P)(\sigma) = \sum_{n=0}^{\infty} b_n \sigma^n.$$

These are given explicitly by the Faa da Bruno formula (multivariate generalization of the Leibnitz rule). This however leads to exponential complexity, and we consider alternatives.

Period 5	r	θ	scale_u	scale_s	Iterations	N
Table A	0.2585	0.4308	0.29	0.29	11	60
Table A	-0.2585	0.4308	0.29	0.29	11	60
Table B	0.2641	0.4223	0.21	0.2	10	60
Table B	-0.2641	0.4223	0.21	0.2	10	60
Table C	0.7695	0.5	0.17	0.17	25	60
Table C	-0.7695	0.5	0.17	0.17	25	60
Table D	0.75	0.44	0.125	0.123	10	30
Table D	-0.75	0.44	0.125	0.123	10	30
Table E	0.7	0.46	0.125	0.123	7	25
Table E	-0.7	0.46	0.125	0.123	7	25
Period 10	r	θ	scale_u	scale_s	Iterations	N
Table A	-0.3283	0.701	0.39	0.39	10	60
Table B	-0.1746	0.3767	0.46	0.48	4	60
Table C	0.1653	0.122	0.34	0.34	8	60
Table D	0	0.3802	0.7	0.7	20	60
Table E	-0.4106	0.1275	0.65	0.65	20	60
Period 30	r	θ	scale_u	scale_s	Iterations	N
Table A	0.1667	0.5103	0.15	0.15	20	60
Table A	-0.1667	0.4897	0.15	0.15	20	60
Table B	0.4877	0.5	0.15	0.15	30	60
Table B	-0.4877	0.5	0.15	0.15	30	60

TABLE 3. Numerical parameters for the orbits of periods 5, 10, and 30.

B.1. Polynomial case. A common situation is that g is polynomial, in which case explicit recursive formulas for the coefficients of integer powers of P are worked out using the Cauchy product. For example we have that

$$P(\sigma)^2 = \sum_{n=0}^{\infty} \left(\sum_{k=0}^{n} a_{n-k} a_k \right) \sigma^n, \tag{22}$$

and that

$$P(\sigma)^3 = \sum_{n=0}^{\infty} \left(\sum_{k=0}^{n} \sum_{j=0}^{k} a_{n-k} a_{k-j} a_j \right) \sigma^n.$$

In these cases

$$b_n = \sum_{k=0}^{n} a_{n-k} a_k, \qquad \text{and} \qquad b_n = \sum_{k=0}^{n} \sum_{j=0}^{k} a_{n-k} a_{k-j} a_j,$$

respectively. Formulas for higher order powers generalize in an obvious way. Therefore Cauchy products are sufficient for working out the power series coefficients for $g \circ P$ in the polynomial case.

B.2. D-finite functions: differential algebra. If g is non-polynomial, but elementary (sometimes referred to as D-finite), then the coefficients of $g \circ P$ are often found by combining the fact that the chain rule turns composition into multiplication with the Cauchy product. For example, suppose $g(z) = e^z$. We write

$$Q(\sigma) = \sum_{n=0}^{\infty} b_n \sigma^n = g(P(\sigma)) = e^{P(\sigma)},$$

to denote the unknown power series of the composition and have that

$$Q'(\sigma) = \sum_{n=0}^{\infty} (n+1) b_{n+1} \sigma^n = g'(P(\sigma)) P'(\sigma) = e^{P(\sigma)} P'(\sigma),$$

which is

$$Q'(\sigma) = Q(\sigma) P'(\sigma).$$

Expressing this last expression in terms of power series and calling on the Cauchy product, we have that

$$\begin{aligned} \sum_{n=0}^{\infty} (n+1) b_{n+1} \sigma^n &= \left(\sum_{n=0}^{\infty} b_n \sigma^n \right) \left(\sum_{n=0}^{\infty} (n+1) a_{n+1} \sigma^n \right) \\ &= \sum_{n=0}^{\infty} \left((k+1) b_{n-k} a_{k+1} \right) \sigma^n. \end{aligned}$$

Matching like powers of σ^n, yields the recursive expression for the power series coefficients of $g \circ P$ given by

$$b_{n+1} = \frac{1}{n+1} \sum_{k=0}^{n} (k+1) b_{n-k} a_{k+1}, \qquad n \geq 0,$$

with

$$b_0 = e^{a_0}.$$

Using this differential-algebraic framework, we find that the power series coefficients of $g \circ P$ are no more difficult to compute in this transcendental case than they are in the quadratic case.

B.3. Remarks on the literature. This idea illustrated in the previous section is sometimes referred to as automatic differentiation for power series, polynomial recast, or polynomial embedding, and it is sufficient for many problems with reasonably complicated nonlinearities. An important general reference is Chapter 4.6 of the second volume of Knuth's the Art of Computer Programming [Knu98]. See also [JZ05] for explicit recursion relations for many elementary functions. Indeed, the reference just cited combines automatic differentiation for power series with a state-of-the-art parser as part of a powerful C++ software library called TAYLOR, and suggests that the boundary between simple and complex nonlinearities is fluid, when approached with the right tools. We refer also to the papers [FGJ25, GJJCZ25] where more recent enhancements and applications of the TAYLOR package are discussed. Further extension of these ideas to multivariable power series is discussed at length in Chapter 2.3 of [HCF$^+$18].

This said, there may be (depending on one's tolerance for complexity) problems where explicit formulas for g are either too much to hope for, or are more complicated than one wishes consider (especially if one is unwilling to adopt special purpose software). For example, Poincaré sections for ODEs and billiard maps studied in the present work approach this boundary. For such maps, we have ready access only to numerical procedures that compute the values g (and its derivatives) at points. In this case, interpolation is a useful tool for determining the power series coefficients of $f \circ P$, and when the functions are analytic, the interpolation can be carried out via the discrete Fourier transform (DFT). An excellent reference for this material is the book of [BH95].

B.4. Composition of power series via the DFT. Suppose that $P\colon \mathbb{R} \to \mathbb{R}^d$, and $g\colon \mathbb{R}^d \to \mathbb{R}$ are real analytic, $g(P(\sigma)) = Q(\sigma)$, whose power series we approximate with the polynomial interpolant

$$Q^N(\sigma) = \sum_{n=0}^{N} \bar{b}_n \sigma^n.$$

The coefficients $\bar{b}_n \approx b_n$, $0 \le n \le N$ of the interpolant are determined by the values of Q at any $(N+1)$ distinct points, and we choose values equidistributed on the unit circle in $\mathbb{C}$ with the first equal to one.

More precisely, define the $N+1$ grid points

$$\sigma_k = e^{2\pi i k/(N+1)}, \quad 0 \le k \le N.$$

and evaluate the composition on this mesh to obtain

$$Q^N(\sigma_k) = \sum_{n=0}^{N} \bar{b}_n \sigma_k^n \approx g(P(\sigma_k))$$

(23)

for $0 \le k \le N$. Letting $\omega = e^{\frac{2\pi i}{N+1}}$, the b_n are approximated by $\bar{b}_n$ solving the linear system

$$\begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega_1 & \omega_1^2 & \cdots & \omega_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N & \omega_N^2 & \cdots & \omega_N^N \end{pmatrix} \begin{pmatrix} \bar{b}_0 \\ \bar{b}_1 \\ \bar{b}_2 \\ \vdots \\ \bar{b}_N \end{pmatrix} = \begin{pmatrix} g(P(\sigma_0)) \\ g(P(\sigma_1)) \\ g(P(\sigma_2)) \\ \vdots \\ g(P(\sigma_N)) \end{pmatrix}$$

Remark B.1 (Analytic continuation)**.** In this the step, we see that computing $f(P(\sigma_k))$ requires evaluating the real analytic function P on complex inputs. This assumes that we are able to analytically continue P to the unit disk in $\mathbb{C}$. Moreover, since the analytic continuation of P is complex valued, we have to evaluate the real analytic function g for complex values as well. Again, this may require analytic continuation of g.

Let

$$V = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega_1 & \omega_1^2 & \cdots & \omega_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N & \omega_N^2 & \cdots & \omega_N^N \end{pmatrix},$$

Observe that V is invertible with

$$V^{-1} = \frac{1}{N+1} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega_1^{-1} & \omega_1^{-2} & \cdots & \omega_1^{-N} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N^{-1} & \omega_N^{-2} & \cdots & \omega_N^{-N} \end{pmatrix}.$$

These are the well known DFT matrices. Since we have an explicit inverse, solving the linear system determining the coefficients of the composition is reduced to matrix-vector multiplication. Moreover, when N is large, this matrix-vector multiplication is accelerated using the Fast Fourier Transform (FFT).

When P itself is a polynomial (as will be the case in our Newton procedure) we write

$$P(\sigma) = \sum_{n=0}^{N} a_n \sigma^n,$$

and have that evaluation of P at the interpolation points σ_k is expressed as

$$P(\sigma_k) = \sum_{n=0}^{N} a_n \sigma_k^n.$$

That is

$$\begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega_1 & \omega_1^2 & \cdots & \omega_1^N \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_N & \omega_N^2 & \cdots & \omega_N^N \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_N \end{pmatrix} = \begin{pmatrix} (P(q_0)) \\ (P(q_1)) \\ (P(q_2)) \\ \vdots \\ (P(q_N)) \end{pmatrix}$$

Let $p_n = P(q_k)$ for $0 \le k \le N$ and write

$$\mathbf{p} = \begin{pmatrix} p_1 \\ \vdots \\ p_N \end{pmatrix}, \quad \mathbf{a} = \begin{pmatrix} a_1 \\ \vdots \\ a_N \end{pmatrix} \quad \text{and} \quad \mathbf{b} = \begin{pmatrix} \bar{b}_1 \\ \vdots \\ \bar{b}_N \end{pmatrix}.$$

Given the coefficients $\mathbf{a}$ of a polynomial P, the procedure for approximating the series coefficients of the composition is as follows:

- Compute $\mathbf{p}$ via
$$V\mathbf{a} = \mathbf{p}.$$
This step is actually known as the DFT.
- Compute the numbers $f_k = f(p_k)$, $0 \le k \le N$. Let
$$\mathbf{f} = \begin{pmatrix} \bar{f}_1 \\ \vdots \\ \bar{f}_N \end{pmatrix}.$$
- Compute
$$\mathbf{b} = V^{-1}\mathbf{f}.$$
This step is actually known as the inverse discrete Fourier transform (IDFT).

The components of $\mathbf{b}$ are the approximate Taylor coefficients of the composition $f \circ P$.

B.5. Multiplication matrix for power series. In the context of a Newton procedure, we often have to solve equations like

$$Q(\sigma)P(\sigma) = g(\sigma), \tag{24}$$

where Q and g are given truncated power series: that is, polynomials. (The equation is an oversimplification, but it will make the necessary point). While this equation can be solved by recursion, it is also amiable to numerical linear algebra.

To see this, let

$$\mathbf{a} = \begin{pmatrix} a_0 \\ \vdots \\ a_N \end{pmatrix}, \quad \text{and} \quad \mathbf{g} = \begin{pmatrix} g_0 \\ \vdots \\ g_N \end{pmatrix},$$

denote the vectors of known polynomial coefficients, and

$$\mathbf{b} = \begin{pmatrix} b_0 \\ \vdots \\ b_N \end{pmatrix},$$

denote the unknown coefficients of P (the quotient of g and Q).

Since this equation is linear in P, there is a matrix representing the action of Q on P. Applying the formula of Equation (22) to the standard basis vectors leads a representation of the action of the Cauchy product on the coefficient vector for P given by

$$Q \cdot P = \begin{pmatrix} a_0 & 0 & 0 & \dots & 0 \\ a_1 & a_0 & 0 & \dots & 0 \\ a_2 & a_1 & a_0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_N & a_{N-1} & a_{N-2} & \vdots & a_0 \end{pmatrix} \begin{pmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_N \end{pmatrix}$$

Let

$$\text{CauchyMat}(\mathbf{a}) = \begin{pmatrix} a_0 & 0 & 0 & \dots & 0 \\ a_1 & a_0 & 0 & \dots & 0 \\ a_2 & a_1 & a_0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_N & a_{N-1} & a_{N-2} & \vdots & a_0 \end{pmatrix}. \tag{25}$$

We can solve for the unknown coefficients $\mathbf{b}$ of $P(\sigma)$ in Equation (24) by solving the (lower triangular) linear system

$$\text{CauchyMat}(\mathbf{a})\mathbf{b} = \mathbf{g}.$$

References

[ADSK16] Artur Avila, Jacopo De Simoi, and Vadim Kaloshin. An integrable deformation of an ellipse of small eccentricity is an ellipse. *Annals of Mathematics*, 184(2):527–558, Sep 2016.

[AV21] Keegan D Anderson and Charles M Villet. Computational study of the dynamics of an asymmetric wedge billiard. *International Journal of Bifurcation and Chaos*, 31(2):2130006, 2021.

[BB24] Luca Baracco and Olga Bernardi. Totally integrable symplectic billiards are ellipses. *Advances in Mathematics*, 454:109873, Oct 2024.

[BDBT23] Vivina L Barutello, Irene De Blasi, and Susanna Terracini. Chaotic dynamics in refraction galactic billiards. *Nonlinearity*, 36(8):4209–4246, Jun 2023.

[Bee15] Casper Beentjes. *Computing Bifurcation Diagrams with Deflation*. PhD thesis, University of Oxford, 2015.

[BFB+11] C.C. Bordeianu, D. Felea, C. Beşliu, Al. Jipa, and I.V. Grossu. Chaotic dynamics in classical nuclear billiards. *Communications in Nonlinear Science and Numerical Simulation*, 16(1):324–340, Jan 2011.

[BFG+22] Misha Bialy, Corentin Fierobe, Alexey Glutsyuk, Mark Levi, Alexander Plakhov, and Serge Tabachnikov. Open problems on billiards and geometric optics. *Arnold Mathematical Journal*, 8(3–4):411–422, Jan 2022.

[BG71] Kenneth M. Brow and William B. Gearhart. Deflation techniques for the calculation of further solutions of a nonlinear system. *Numerische Mathematik*, 16(4):334–342, Jan 1971.

[BH95] William L. Briggs and Van Emden Henson. *The DFT*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1995. An owner's manual for the discrete Fourier transform.

[Bir27] George Birkhoff. Dynamical systems. *Colloquium Publications*, Dec 1927.

[Bis24] Patrick Bishop. *Computation of Dynamical Structures for Dynamical Billiards and Parameter Drift Equations*. PhD thesis, George Mason University, 2024.

[BK96] N. Berglund and H. Kunz. Integrability and ergodicity of classical billiards in a magnetic field. *Journal of Statistical Physics*, 83(1–2):81–126, Apr 1996.

[BLMJ16] Maxime Breden, Jean-Philippe Lessard, and Jason D. Mireles James. Computation of maximal local (un)stable manifold patches by the parameterization method. *Indag. Math. (N.S.)*, 27(1):340–367, 2016.

[BM22] Misha Bialy and Andrey E. Mironov. The Birkhoff-Poritsky conjecture for centrally-symmetric billiard tables. *Annals of Mathematics*, 196(1), Jul 2022.

[Bun75] L. A. Bunimovich. On ergodic properties of certain billiards. *Functional Analysis and Its Applications*, 8(3):254–255, 1975.

[Bun79] L. A. Bunimovich. On the ergodic properties of nowhere dispersing billiards. *Communications in Mathematical Physics*, 65(3):295–312, Oct 1979.

[Bun19] L. A. Bunimovich. Physical versus mathematical billiards: From regular dynamics to chaos and back. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(9), Sep 2019.

[CF21] Kevin E. M. Church and Clément Fortin. Computer-assisted methods for analyzing periodic orbits in vibrating gravitational billiards. *International Journal of Bifurcation and Chaos*, 31(8):2130021, 2021.

[CFdlL03] Xavier Cabré, Ernest Fontich, and Rafael de la Llave. The parameterization method for invariant manifolds I: Manifolds associated to non-resonant subspaces. *Indiana University Mathematics Journal*, 52(2):283–328, 2003.

[CFdlL05] Xavier Cabré, Ernest Fontich, and Rafael de la Llave. The parameterization method for invariant manifolds III: Overview and applications. *Journal of Differential Equations*, 218(2):444–515, Nov 2005.

[CFJ20] Maciej J Capiński, Emmanuel Fleurantin, and JD James. Computer assisted proofs of two-dimensional attracting invariant tori for odes. *Discrete & Continuous Dynamical Systems: Series A*, 40(12), 2020.

[CM06] Nikolai Chernov and Roberto Markarian. *Chaotic billiards*. American Mathematical Society, 2006.

[Dat17] George Datseris. DynamicalBilliards.jl. *The Journal of Open Source Software*, 2(19), 2017.

[dCHSL22] Diogo Ricardo da Costa, Matheus Hansen, Mário Roberto Silva, and Edson D. Leonel. Tangent method and some dynamical properties of an oval-like billiard. *International Journal of Bifurcation and Chaos*, 32(04), Mar 2022.

[DHF19] George Datseris, Lukas Hupe, and Ragnar Fleischmann. Estimating Lyapunov exponents in billiards. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(9), Sep 2019.

[DHN03] T Damour, M Henneaux, and H Nicolai. Cosmological billiards. *Classical and Quantum Gravity*, 20(9), Apr 2003.

[Dou82] R. Douady. *Applications du théorème des tores invariantes*. 1982.

[DRR96] Amadeu Delshams and Rafael Ramírez-Ros. On Birkhoff's conjecture about convex billiards. *Proceedings of nd Catalan Days on Applied Mathematics*, 10 1996.

[DSSAY19] Suddhasattwa Das, Yoshitaka Saiki, Evelyn Sander, and James A. Yorke. Solving the Babylonian problem of quasiperiodic rotation rates. *Discrete & Continuous Dynamical Systems - S*, 12(8):2279–2305, 2019.

[DSSY17] Suddhasattwa Das, Yoshitaka Saiki, Evelyn Sander, and James A Yorke. Quantitative quasiperiodicity. *Nonlinearity*, 30(11):4111–4140, 2017.

[DY18] Suddhasattwa Das and James A Yorke. Super convergence of ergodic averages for quasiperiodic orbits. *Nonlinearity*, 31(2):491–501, 2018.

[FBF15] P. E. Farrell, Birkisson, and S. W. Funke. Deflation techniques for finding distinct solutions of nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 37(4), Jan 2015.

[FGJ25] Philip P. Forrier, Joan Gimeno, and Àngel Jorba. A note on the local behavior of the Taylor method for stiff ODEs. *Appl. Math. Comput.*, 496:Paper No. 129344, 18, 2025.

[FJ20] Emmanuel Fleurantin and JD Mireles James. Resonant tori, transport barriers, and chaos in a vector field with a Neimark–Sacker bifurcation. *Communications in Nonlinear Science and Numerical Simulation*, 85:105226, 2020.

[Fra03] John Franks. Rotation numbers and instability sets. *Bulletin of the American Mathematical Society*, 40(3):263—279, 2003.

[GJJCZ25] Joan Gimeno, Àngel Jorba, Marc Jorba-Cuscó, and Maorong Zou. Explicit numerical computation of normal forms for Poincaré maps. *Commun. Nonlinear Sci. Numer. Simul.*, 149:Paper No. 108913, 25, 2025.

[GMJ17] J. L. Gonzalez and J. D. Mireles James. High-order parameterization of stable/unstable manifolds for long periodic orbits of maps. *SIAM J. Appl. Dyn. Syst.*, 16(3):1748–1795, 2017.

[Gut86] Eugene Gutkin. Billiards in polygons. *Physica D: Nonlinear Phenomena*, 19(3):311–333, Apr 1986.

[Gut12] Eugene Gutkin. Billiard dynamics: An updated survey with the emphasis on open problems. *Chaos*, 22:026116, 2012.

[HCF+18] A. Haro, M. Canadell, J. Figueras, A. Luque, and J. Mondelo. *Parameterization method for invariant manifolds: From rigorous results to effective computations*. Springer, 2018.

[HLMJ23] Olivier Hénot, Jean-Philippe Lessard, and Jason D. Mireles James. Numerical computation of transverse homoclinic orbits for periodic solutions of delay differential equations. *SIAM J. Appl. Dyn. Syst.*, 22(4):3093–3129, 2023.

[HMJ21] Wouter Hetebrij and J. D. Mireles James. Critical homoclinics in a restricted four-body problem: numerical continuation and center manifold computations. *Celestial Mech. Dynam. Astronom.*, 133(2):Paper No. 5, 48, 2021.

[HW13] Markus Himmelstrand and Victor Wilén. A survey of dynamical billiards, 2013.

[JZ05] Àngel Jorba and Maorong Zou. A software package for the numerical integration of ODEs by means of high-order Taylor methods. *Experiment. Math.*, 14(1):99–117, 2005.

[KAdlL21] Bhanu Kumar, Rodney L. Anderson, and Rafael de la Llave. High-order resonant orbit manifold expansions for mission design in the planar circular restricted 3-body problem. *Commun. Nonlinear Sci. Numer. Simul.*, 97:Paper No. 105691, 15, 2021.

[KAdlL22] Bhanu Kumar, Rodney L. Anderson, and Rafael de la Llave. Rapid and accurate methods for computing whiskered tori and their manifolds in periodically perturbed planar circular restricted 3-body problems. *Celestial Mech. Dynam. Astronom.*, 134(1):Paper No. 3, 38, 2022.

[KAdlL25] Bhanu Kumar, Rodney L. Anderson, and Rafael de la Llave. Rapid GPU-assisted search and parameterization-based refinement and continuation of connections between tori in periodically perturbed planar circular restricted 3-body problems. *SIAM J. Appl. Dyn. Syst.*, 24(1):219–258, 2025.

[Kni98] Oliver Knill. On nonconvex caustics of convex billiards. *Elemente der Mathematik*, 53(3):89–106, 1998.

[Knu98] Donald E. Knuth. *The art of computer programming. Vol. 2*. Addison-Wesley, Reading, MA, third edition, 1998. Seminumerical algorithms.

[KS18a] Vadim Kaloshin and Alfonso Sorrentino. On the integrability of Birkhoff billiards. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2131):20170419, Sep 2018.

[KS18b] Vadim Kaloshin and Alfonso Sorrentino. On the local Birkhoff conjecture for convex billiards. *Annals of Mathematics*, 188(1), Jul 2018.

[KS22] Vadim Kaloshin and Alfonso Sorrentino. Inverse problems and rigidity questions in billiard dynamics. *Ergodic Theory and Dynamical Systems*, 42(3):1023–1056, 2022.

[Laz73] V.F. Lazutkin. The existence of caustics for a billiard problem in a convex domain. *T. Math. USSR Izvestija*, 7, 1973.

[LP04] Steven Lansel and Mason A. Porter. A graphical user interface to simulate classical billiard systems, May 2004.

[LT07] Mark Levi and Serge Tabachnikov. The Poncelet grid and billiards in ellipses. *The American Mathematical Monthly*, 114(10):895–908, 2007.

[Mei92] J. D. Meiss. Symplectic maps, variational principles, and transport. *Reviews of Modern Physics*, 64(3):795–848, Jul 1992.

[MRRTS16] P. Martín, R. Ramírez-Ros, and A. Tamarit-Sariol. Exponentially small asymptotic formulas for the length spectrum in some billiard tables. *Experimental Mathematics*, 25(4):416–440, 2016.

[MS21] J.D. Meiss and E. Sander. Birkhoff averages and the breakdown of invariant tori in volume-preserving maps. *Physica D: Nonlinear Phenomena*, 428:133048, 2021.

[Poi93a] Henri Poincaré. *New methods of celestial mechanics. Vol. 1*, volume 13 of *History of Modern Physics and Astronomy*. American Institute of Physics, New York, 1993. Periodic and asymptotic solutions, Translated from the French, Revised reprint of the 1967 English translation, With endnotes by V. I. Arnold, Edited and with an introduction by Daniel L. Goroff.

[Poi93b] Henri Poincaré. *New methods of celestial mechanics. Vol. 2*, volume 13 of *History of Modern Physics and Astronomy*. American Institute of Physics, New York, 1993. Approximations by series, Translated from the French, Revised reprint of the 1967 English translation, With endnotes by V. M. Alekseev, Edited and with an introduction by Daniel L. Goroff.

[Poi93c] Henri Poincaré. *New methods of celestial mechanics. Vol. 3*, volume 13 of *History of Modern Physics and Astronomy*. American Institute of Physics, New York, 1993. Integral invariants and asymptotic properties of certain solutions, Translated from the French, Revised reprint of the 1967 English translation, With endnotes by G. A. Merman, Edited and with an introduction by Daniel L. Goroff.

[PRK22] Michal Pnueli and Vered Rom-Kedar. Near tangent dynamics in a class of hamiltonian impact systems. *SIAM Journal on Applied Dynamical Systems*, 21(3):2000–2046, Jul 2022.

[PTZ23] Michael Plum, Airi Takeuchi, and Lei Zhao. Boltzmann's billiard systems: Computation of the billiard mapping and some numerical results. (arXiv:2311.09699), 2023.

[PW71] G. Peters and J. H. Wilkinson. Practical problems arising in the solution of polynomial equations. *IMA Journal of Applied Mathematics*, 8(1):16–35, 1971.

[RGK21] Dan Reznik, Ronaldo Garcia, and Jair Koiller. Fifty new invariants of n-periodics in the elliptic billiard. *Arnold Mathematical Journal*, 7:341–355, 2021.

[Rob98] Clark Robinson. *Dynamical systems: stability, symbolic dynamics, and chaos*. CRC press, 1998.

[Sch22] Richard Schwartz. Survey lecture on billiards. In *Proceedings of the International Congress of Mathematicians*, pages 2392–2429, 2022.

[SLR16] J. Solanpää, P.J.J. Luukko, and E. Räsänen. Bill2d — a software package for classical two-dimensional Hamiltonian systems. *Computer Physics Communications*, 199:133–138, Feb 2016.

[SM20] E. Sander and J.D. Meiss. Birkhoff averages and rotational invariant circles for area-preserving maps. *Physica D: Nonlinear Phenomena*, 411:132569, October 2020.

[Sma65] Stephen Smale. Diffeomorphisms with many periodic points. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)*, pages 63–80. Princeton Univ. Press, Princeton, NJ, 1965.

[SN10] Juan Sánchez and Marta Net. On the multiple shooting continuation of periodic orbits by Newton-Krylov methods. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 20(1):43–61, 2010.

[Tab09] Serge Tabachnikov. *Geometry and Billiards Serge Tabachnikov*. American Mathematical Soc, 2009.

[TMJ22] Archana Neupane Timsina and J.D. Mireles James. Parameterized stable/unstable manifolds for periodic solutions of implicitly defined dynamical systems. *Chaos, Solitons & Fractals*, 161:112345, August 2022.

[Tur16] Michael Turaev. *Numerical Experiments in Billiards*. PhD thesis, Uppsala University, 2016.

Department of Mathematical Sciences, George Mason University, Fairfax, VA 22030

Department of Mathematics and Statistics, Florida Atlantic University, Boca Raton, FL 33431