Variational Loss Landscapes for Periodic Orbits

Leo Yao
Department of Physics
Massachusetts Institute of Technology
Cambridge, MA 02139
leoy@mit.edu

Ziming Liu
Department of Physics
Massachusetts Institute of Technology
Cambridge, MA 02139
zmliu@mit.edu

Max Tegmark
Department of Physics
Massachusetts Institute of Technology
Cambridge, MA 02139
tegmark@mit.edu

Abstract

The discovery and classification of periodic orbits is fundamental to understanding chaotic dynamical systems, but existing algorithms typically search for individual orbits without considering underlying structure and connectivity. We consider the loss landscape of variational loops in phase space, devising a Hessian-based approach to numerically continue along periodic orbit families. Our method offers precise initializations of oscillations around unstable fixed points, an integrator-free variational continuation method, and efficient detection of orbit family intersections and subharmonic bifurcations. Leveraging autograd for computations, we present full continuations of periodic double pendulum oscillations from fixed points, demonstrate examples of orbit family intersections and bifurcations, and interpret branching orbits as combinations of perturbations in the periodic orbit structure.

1 Introduction

In chaotic dynamical systems, periodic orbits play a key role in studying long-term evolution. Periodic orbits are inputs to many statistical methods from classical chaos theory, such as cycle expansions (1, 5, 6, 4), and are also interpretable for humans due to the inherit simplicity of recurring trajectories.

The most basic method for finding orbits, the shooting method, integrates initial conditions and extracts close recurrences. Due to exponential divergence of trajectories in chaotic systems, obtaining recurrences requires exponentially precise initial guesses for longer orbits. Convergence becomes near impossible due to required initialization precision, or accumulating numerical integrator error. To avoid convergence issues, the variational method (8) is widely used. The initialization is a closed loop in phase space, not necessarily satisfying the equations of motion. The loop is then varied until satisfying evolution equations, resulting in a periodic orbit. As the loop is adapted locally, no integration of conditions occurs. This gives a much greater radius of convergence, and variations of the method modify loop representations (11) and optimization methods (2).

In Hamiltonian (energy-conserving) systems, periodic orbits are theoretically known to form continuous, single-parameter families (7, 12, 3). Individual families of orbits do not bifurcate, terminating only if a phase space coordinate or orbit period diverges. Bifurcations in the orbit spectrum are caused by crossings of families at original and higher period multiples. These connected, one-dimensional families of orbits are the perfect testbed for numerical continuation techniques. However, a majority of periodic orbit studies do not consider the underlying orbit structure when finding solutions, re-

Machine Learning and the Physical Sciences Workshop, NeurIPS 2024.

sorting instead to grid search. Older numerical continuation searches have involved integrated the Jacobian (13), a computationally expensive and numerically unstable task.

We consider the space of all possible variational loops, parametrizing a continuous loss with zeroes at periodic orbits. Using Hessian eigenvectors to analyze this loss landscape, we demonstrate automatic discovery of fixed point oscillations, continuation of orbit families, and detection of bifurcations in the orbit spectrum. Our orbit propagation method combines continuation-based initializations based on orbit structure with the greatly improved convergence of the variational method.

2 Background and Method

2.1 Variational Orbit Setup

We consider a general dynamical system $\frac{\mathrm{d}\vec{z}}{\mathrm{d}t} = \vec{f}(\vec{z})$ parametrized by some state vector $\vec{z}$ in a phase space $\mathcal{M}$. The time evolution of the system is given by integrating the phase space trajectory:

$$\vec{z}(t) \equiv f^t(\vec{z}_0) = \vec{z}_0 + \int_0^t \vec{f}(\vec{z}(t'))\mathrm{d}t' \tag{1}$$

A periodic orbit is defined by an initial condition $\vec{z}_0 \in \mathcal{M}$ and a time $T > 0$ such that $f^T(\vec{z}_0) = \vec{z}_0$.

To set up the variational method, we take a closed loop $\mathcal{L}$, defined as a period T and a trajectory $\vec{z}(t)$ in phase space, with $0 \leq t < T$, $\vec{z} \in \mathcal{M}$, and $\vec{z}(T) = \vec{z}(0)$. At each point $\vec{z}(t)$ on the trajectory, we quantify the deviation of the loop trajectory $\frac{\mathrm{d}\vec{z}(t)}{\mathrm{d}t}$ from physical evolution $f(\vec{z}(t))$, and average over the entire trajectory to obtain a loss function on the loop:

$$\ell(\mathcal{L}) = \frac{1}{T}\int_0^T \left|\frac{\mathrm{d}\vec{z}(t)}{\mathrm{d}t} - \vec{f}(\vec{z}(t))\right|^2 \mathrm{d}t \tag{2}$$

This loss function quantifies the overall deviation of the loop from a physical trajectory. Zeroes of ℓ correspond to loops which always locally match physical evolution, and therefore are periodic orbits. As the loss function is always nonnegative, at periodic orbits the loss ℓ and gradients $\nabla\ell$ both vanish.

2.2 Hessian Analysis of Loss Landscapes

The loss function ℓ on the space of all possible loops $\mathscr{P}$ defines a loss landscape. The Hessian $\mathbf{H}_\ell$ describes the local curvature of this loss landscape around a loop $\mathcal{L} \in \mathscr{P}$. For a periodic orbit $\ell = 0$, and the Hessian $\mathbf{H}_l(\mathcal{L}, \mathscr{P})$ must be positive semi-definite. The minimum eigenvalues in an eigendecomposition give minimum curvature magnitudes, and associated directions.

A zero eigenvalue gives a perturbation to the loop $\mathcal{L}$ maintaining zero loss and a periodic orbit. Each zero eigenvalue corresponds to a unique direction, creating a subspace of possible perturbations. The dimension of the connected periodic orbit space at $\mathcal{L}$ is given by the number of zero eigenvalues of the Hessian $\mathbf{H}_l(\mathcal{L}, \mathscr{P})$.

Taking a finite step in loop space along a flat direction gives an adjacent connected orbit. The new initialization is an entire loop, meaning the subsequent periodic orbit optimization benefits from the much larger radius of convergence of the variational method. This gives a continuation method for orbits in the loss landscape of periodic loops with variational initialization and convergence.

2.3 Applications to Orbit Discovery

By propagating through the loss landscape, we can take a methodical approach to finding periodic orbits that leverages the underlying structure of solutions (Figure 1).

Hessians find fixed point oscillations To test for periodic perturbations, we can initialize loops at the fixed point with various periods T, and take the Hessian. A drop of the minimal eigenvalue to zero indicates a period T with a periodic perturbation given by the corresponding eigenvector. By sweeping values of T, all possible small oscillations around fixed points, which are starting points for periodic orbit branches, can be systematically discovered.

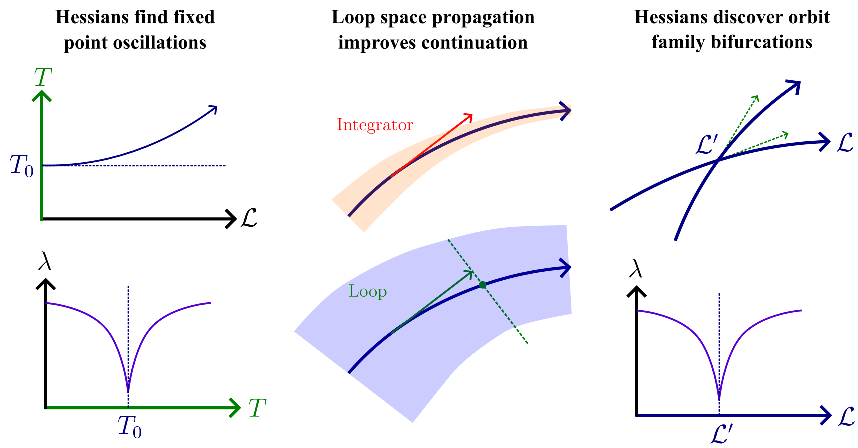

Figure 1: On left, a branch of orbits $\mathcal{L}$ emerging from a fixed point at period T_0 (above) is found by sweeping period T at the fixed point, and detecting a drop in the minimal eigenvalue λ (below). At center, the increased radius of convergence (shaded) of the variational method allows for exploration further along orbit families, with an orthogonal loop space constraint. On right, a crossing of orbit families at $\mathcal{L}'$ (above) is found by detecting a drop in the off-branch eigenvalue λ (below).

Loop space propagation improves continuation The larger radius of convergence in loop space allows for larger steps to be taken and longer orbits to be continued, allowing exploration further along orbit families even when integrator methods fail due to precision issues. The eigenvalue direction gives an orthogonal orbit constraint in loop space, better constraining orbit optimization dynamics compared to just an initial condition in phase space.

Hessians discover orbit family bifurcations If two orbit families cross, there will be a two-dimensional null subspace at the bifurcation point, and a second eigenvalue will drop to zero along the orbit family. Subharmonic bifurcations can be detected by multiply-winding the original loop to represent multiple orbits, introducing fractional frequency components into the parametrization and allowing perturbations over multiple periods to be detected.

3 Experiment

3.1 Implementation

We use the ideal double pendulum as a test system, with equal point masses and equal length massless arms ($m_1 = m_2 = l_1 = l_2 = g = 1$). A system state is specified by arm angles θ_1, θ_2 and angular velocities $\dot{\theta}_1, \dot{\theta}_2$, giving a four-dimensional phase space. A loop $\mathcal{L}$ is represented by a period T and a closed trajectory $\vec{z}(t)$, which we parametrize using a Fourier decomposition in time with a finite frequency cutoff K:

$$\vec{z}(t) = \vec{a}_0 + \sum_{k=1}^{K} \vec{a}_k \cos\left(\frac{2k\pi t}{T}\right) + \vec{b}_k \sin\left(\frac{2k\pi t}{T}\right), \quad 0 \le t \le T. \tag{3}$$

We implement computations in PyTorch (9) with `float64` precision, combining the Rprop algorithm (10) and linear extrapolation to handle varying magnitudes of parameters and gradients. We optimize to a minimum of the integrated error, and take a convergence condition of $\sqrt{\ell} < 10^{-10}$. We propagate along an orbit family in steps of size $0.5°$ in initial conditions space, doubling the frequency cutoff if the convergence condition for ℓ is not reached. Orbit symmetries are used to reduce the parameter count and Hessian size. For spatially symmetric double pendulum oscillations, where the two halves of the orbit are mirrored, we evolve only components satisfying the symmetry constraint, which both enforces a phase condition and avoids accidentally propagating to perturbed orbits.

All computation was conducted on CPUs; converging an orbit took between seconds and a few hours, and memory requirements were primarily due to Hessian size. We started with $K = 16$ Fourier terms per state parameter for small oscillations, rising to $K = 1024$ for the longest converged orbits.

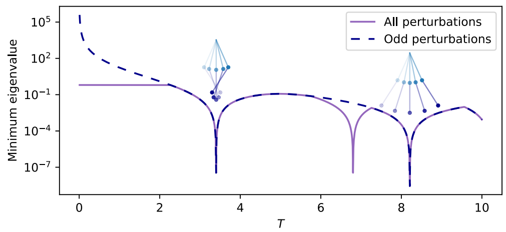

Figure 2: Minimum eigenvalue for period T with both masses down ($\theta_1 = 0$, $\theta_2 = 0$); eigenvalue minimums give out of phase and in-phase normal modes and correct corresponding periods T.

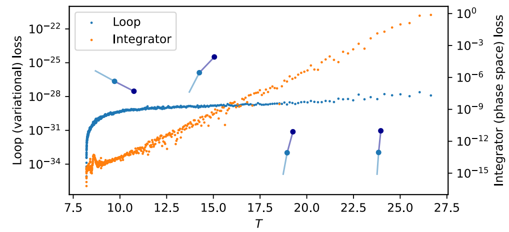

Figure 3: Loop and integrator losses for in-phase oscillations; variational convergence succeeds even for longer periods T as masses approach vertical, while integrator error diverges.

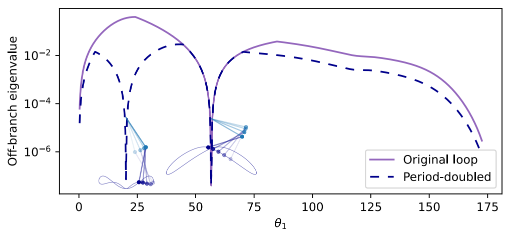

Figure 4: Minimum off-branch eigenvalue for out of phase oscillations; additional eigenvalue minimums give same period and period-doubling bifurcations leading to new orbit families.

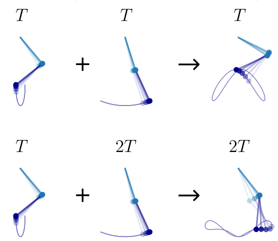

Figure 5: Bifurcating orbits are created by combinations of perturbations at same and higher multiples of period, which can be interpreted by following propagated branches.

3.2 Results

For double pendulum fixed points, we conducted sweeps of periods T from 0 to 10, with a step size of $\Delta T = 0.01$. From eigenvalue minimums, we successfully recovered small oscillations of $T = 3.39$ and $T = 8.21$ around the stable fixed point (both masses down, Figure 2), and $T = 5.28$ around both unstable saddle points (one mass down, one mass up), matching known normal modes.

Orbit families were successfully propagated from each small oscillation to orbits approaching a period divergence, with masses within 5° of vertical. Loop loss remained below the convergence condition, while once-around phase space integrator error scaled exponentially with orbit period T (Figure 3). Even using `scipy.integrate` at double precision, longer period orbits are destroyed by accumulating error, demonstrating the advantage of the variational method for more complex orbits.

Bifurcation detection via minimum off-branch eigenvalue was successful for both original period perturbations and period doubling (Figure 4). By twice-winding the original loop, subharmonic perturbations appear as additional eigenvalue minimums. Eigenvector directions give perturbed loop initializations from bifurcation points, which become new starting points for propagating orbit families.

By systematically building up a collection of branching orbit families, we can interpret more complicated orbits in terms of simpler ones. Along orbit families, higher energy periodic solutions can be understood as extensions of lower energy orbits. By following bifurcation perturbations, we can intuitively break down a complicated motion into sums of simpler oscillations (Figure 5).

Propagating perturbations was successful across the phase space, with many branches also approaching period divergences. A variety of rich behavior in the periodic orbit spectrum is revealed, such as orbit families forming closed loops and additional bifurcation points. A plot of all converged periodic orbits is shown in Figure 6.

4 Conclusion

We have introduced a loss landscapes perspective to searching periodic orbits. By probing loss curvatures using the Hessian, we can systematically explore the structure of orbit families, determining

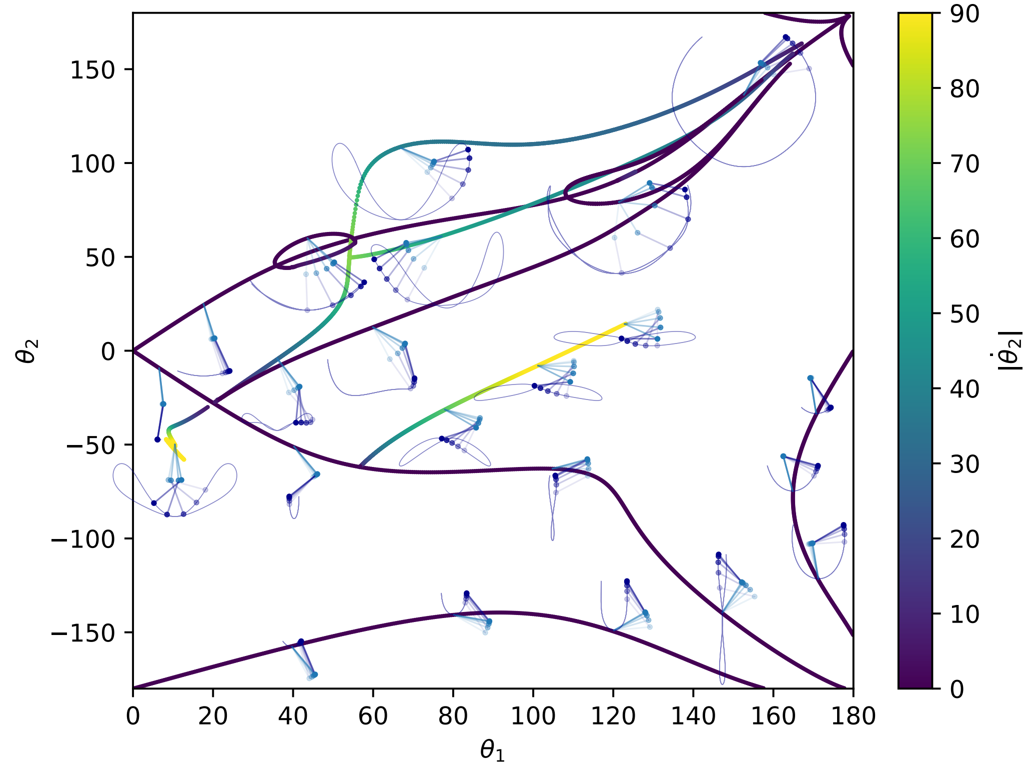

Figure 6: Converged double pendulum periodic orbits in the Poincaré section with $\dot{\theta}_1 = 0$; bifurcating branches intersect and extend off continuous orbit families between fixed points.

connected classes of orbits and bifurcations. We obtain a numerical continuation method in loop space, with variational initialization and convergence properties.

By combining the Rprop optimizer and linear extrapolation, we are able to converge orbits of varying complexity across phase space. The success of this optimization routine demonstrates the applicability of machine learning optimizers to problems in physical systems, allowing to take advantage of developed tools and methods from machine learning such as parametrizations. However, Hessian probing of loss landscapes and convergence techniques can be decoupled, and our propagation method is also compatible with other variational methods in the literature.

Methodically searching orbits allows more complicated motions to be interpreted as combinations of simple oscillations, giving a more interpretable method for analyzing periodic solutions. We look forward to appllying our method of systematically propagating through the orbit spectrum to a wide range of dynamical systems, and for further advances in periodic orbit search methods based on machine learning principles.

Acknowledgments and Disclosure of Funding

Z.L. and M.T. are supported by IAIFI through NSF grant PHY-2019786.

References

[1] D. Auerbach, P. Cvitanović, J.-P. Eckmann, G. Gunaratne, and I. Procaccia. Exploring chaotic motion through periodic orbits. *Phys. Rev. Lett.*, 58:2387–2389, Jun 1987. doi: 10.1103/PhysRevLett.58.2387. URL `https://link.aps.org/doi/10.1103/PhysRevLett.58.2387`.

[2] S. Azimi, O. Ashtari, and T. M. Schneider. Constructing periodic orbits of high-dimensional chaotic systems by an adjoint-based variational method. *Phys. Rev. E*, 105:014217, Jan 2022. doi: 10.1103/PhysRevE.105.014217. URL `https://link.aps.org/doi/10.1103/PhysRevE.105.014217`.

[3] G. D. Birkhoff. Sur le problème restreint des trois corps (second mémoire). *Annali della Scuola Normale Superiore di Pisa - Classe di Scienze*, 5(1):9–50, 1936. URL `http://eudml.org/doc/82917`.

[4] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay. *Chaos: Classical and Quantum*. Niels Bohr Inst., Copenhagen, 2016. URL `http://ChaosBook.org/`.

[5] P. Cvitanović. Invariant measurement of strange sets in terms of cycles. *Phys. Rev. Lett.*, 61:2729–2732, Dec 1988. doi: 10.1103/PhysRevLett.61.2729. URL `https://link.aps.org/doi/10.1103/PhysRevLett.61.2729`.

[6] M. C. Gutzwiller. *Chaos in Classical and Quantum Mechanics*. Springer New York, 1990. ISBN 9781461209836. doi: 10.1007/978-1-4612-0983-6. URL `http://dx.doi.org/10.1007/978-1-4612-0983-6`.

[7] M. Hénon. *Generating Families in the Restricted Three-Body Problem*. Springer Berlin Heidelberg, 1997. ISBN 9783540638025. doi: 10.1007/3-540-69650-4. URL `http://dx.doi.org/10.1007/3-540-69650-4`.

[8] Y. Lan and P. Cvitanović. Variational method for finding periodic orbits in a general flow. *Phys. Rev. E*, 69:016217, Jan 2004. doi: 10.1103/PhysRevE.69.016217. URL `https://link.aps.org/doi/10.1103/PhysRevE.69.016217`.

[9] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

[10] M. Riedmiller and H. Braun. A direct adaptive method for faster backpropagation learning: The rprop algorithm. In *IEEE international conference on neural networks*, pages 586–591. IEEE, 1993.

[11] D. Wang, P. Wang, and Y. Lan. Accelerated variational approach for searching cycles. *Phys. Rev. E*, 98:042204, Oct 2018. doi: 10.1103/PhysRevE.98.042204. URL `https://link.aps.org/doi/10.1103/PhysRevE.98.042204`.

[12] A. Wintner. Grundlagen einer genealogie der periodischen bahnen im restringierten dreikörperproblem: Erste mitteilung beweis des e. strömgrenschen dynamischen abschlußprinzips der periodischen bahngruppen. *Mathematische Zeitschrift*, 34(1):321–349, Dec. 1932. ISSN 1432-1823. doi: 10.1007/bf01180594. URL `http://dx.doi.org/10.1007/BF01180594`.

[13] C. Wulff and A. Schebesch. Numerical continuation of symmetric periodic orbits. *SIAM Journal on Applied Dynamical Systems*, 5(3):435–475, Jan. 2006. ISSN 1536-0040. doi: 10.1137/050637170. URL `http://dx.doi.org/10.1137/050637170`.