

# PERIODIC ORBIT DISCOVERY ENHANCED BY PHYSICS-INFORMED NEURAL NETWORKS

John R. Martin\*, and Hanspeter Schaub†

No universal strategy exists for discovering periodic orbits in highly non-keplerian environments. A combination of analytical and numerical approaches have been proposed, but these approaches often rely on approximations of the gravitational potential or initial guesses that are sufficiently close to a true solution. In the case of small body settings, these problems are further exacerbated due to the computationally expensive gravity models needed to propagate spacecraft dynamics generated by the bodies' irregular geometries. This paper aims to extend the set of tools and algorithms used to search for periodic orbits in small body settings. An emphasis is placed on two novel contributions: First, reducing the computational overhead of this process through the use of the novel Physics-Informed Neural Network Gravity Model (PINN-GM); second, combining the differentiability of the PINN-GM with the Lagrange planetary equations (LPE) to conduct a shooting algorithm in orbital element space rather than cartesian space. Results show that the PINN-GM can achieve order-of-magnitude faster orbit search times than with the polyhedral model without sacrificing dynamical fidelity, and demonstrate how the PINN-GM enhanced orbital element shooting method can be used to assist mission designers in searching for desirable orbits that are otherwise challenging to determine using the past cartesian methods.

## INTRODUCTION

Periodic orbits are desirable because they allow spacecraft to remain in a fixed orbit without the need to expend fuel. Perhaps the most famous of such orbits can be found near the Earth-Sun Lagrange points where a neighborhood exists such that spacecraft motion remains bounded and periodic.<sup>1</sup> Notable spacecraft to leverage these periodic orbits in the three body problem include STEREO-A, STEREO-B, James Webb Space Telescope, among others. Despite the desirability of periodic orbits, their discovery is a challenging endeavour — particularly for highly non-keplerian environments.

The reason for this difficulty is two-fold. First, only in rare cases can periodic motion be discovered analytically. This is because closed form, analytic expression of the disturbing gravitational potential are required to compute such solutions. For simplified forms of the three-body problem, this is possible if one assumes that the gravitational accelerations exhibited by all bodies can be described by a single point mass potential model — a procedure which ultimately yields the aforementioned Lagrange points. In reality, however, this assumption is often not valid. Instead more

---

\*NSF Graduate Research Fellow, Ann and H.J. Smead Department of Aerospace Engineering Sciences, University of Colorado, Boulder, 431 UCB, Colorado Center for Astrodynamics Research, Boulder, CO, 80309.

†Glenn L. Murphy Chair in Engineering, Ann and H.J. Smead Department of Aerospace Engineering Sciences, University of Colorado, Boulder, 431 UCB, Colorado Center for Astrodynamics Research, Boulder, CO, 80309. AAS Fellow, AIAA Fellow.

complex gravity models are required, and dynamicists must turn to numerical shooting methods to solve the periodic orbit boundary value problem. Unfortunately, even the numerical shooting methods are not free of their own difficulties. In many cases, the numerical solvers will not converge if the initial guess deviates too far from the true solution. This disadvantage can be mitigated by testing many initial conditions until a solution is reached;<sup>2</sup> however this can be a major bottleneck if the gravity model used is computationally expensive. This problem can be further exacerbated by shooting methods which must use numerically-computed jacobians.

Small-body environments are among the settings for which these challenges are most apparent. Asteroids and comets often exhibit irregular geometries and corresponding gravity fields which produce complex, non-periodic motion. In these environments the gravitational potential cannot be adequately represented using a point mass gravity model, and instead, expensive polyhedral gravity models must be used to capture these unusual dynamics.<sup>3</sup> While the polyhedral gravity model offers a compelling analytic solution to the gravity modeling problem, this model can be extremely expensive to evaluate depending on the fidelity of the corresponding shape model used. This expense prevents the polyhedral model's practicality in the numerical methods used to solve the periodic orbit problem.

This paper seeks to investigate how a new gravity model can be leveraged to assist in the efficient discovery of periodic orbits in small-body settings. Specifically the Physics-Informed Neural Network Gravity Model (PINN-GM) has been shown to accurately represent the gravity fields of small-bodies by learning efficient basis functions rather than prescribing them.<sup>4</sup> This approach bypasses the computational inefficiencies of the polyhedral gravity model and retains differentiability for use of numerical methods thanks to automatic differentiation.<sup>5</sup> The efficient and differentiable form of the PINN-GM make it a compelling tool to be leveraged in periodic orbit discovery. Not only can traditional shooting methods be applied in small-body settings without risk of computational inefficiency, but there are also new opportunities to search for orbits in different element spaces. This paper investigates how the PINN-GM can assist in the discovery of such orbits and how they potentially expand the attracting basins around them such that fewer initial guesses are required to find a solution.

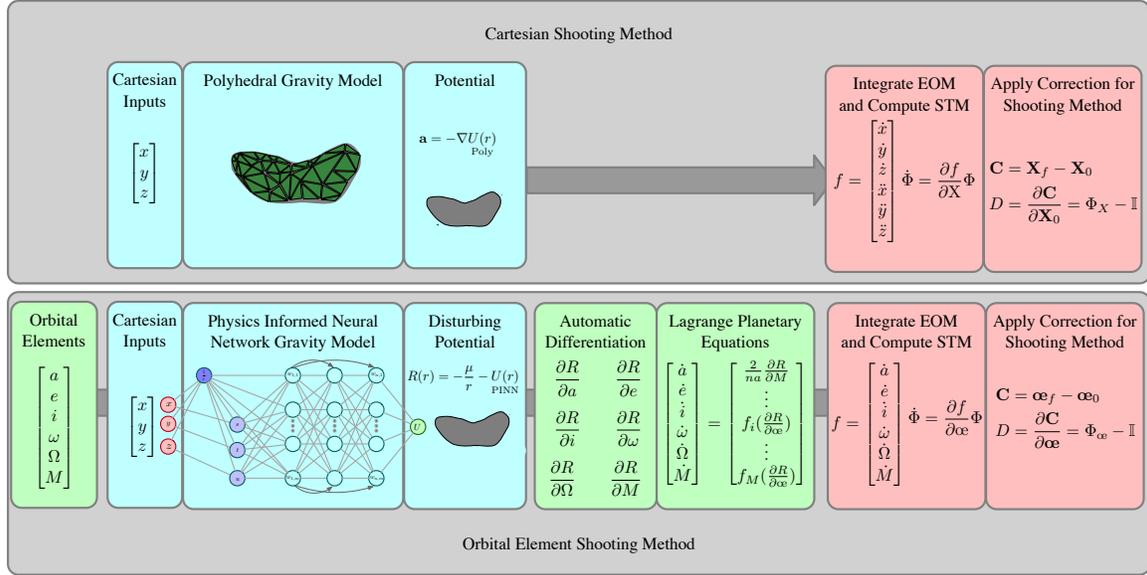
## **BACKGROUND**

Initial efforts to search for periodic orbits in small-body settings date back to the mid-90s. In 1995, Scheeres provided the first detailed characterization and analysis of the orbital dynamics about the asteroid 433-Eros.<sup>6</sup> These findings highlight how rotating asteroids stationary points can be used as the initial search space for periodic orbits, and the paper ultimately reveals three periodic orbit families about Eros: direct, near-equatorial orbits; retrograde, near-equatorial orbits; non-planar orbits with resonant nodal periods. In 1998, Scheeres also provided more algebraically involved derivations of bounded, frozen orbits about the asteroid 4179-Toutatis.<sup>7</sup> This strategy relies on period averaging the Lagrange planetary equations using a simplified low-degree spherical harmonic expansions of the potential. Additional analysis is conducted discussing how these discovered orbits in the simplified system can be iteratively corrected to produce increasingly stable and periodic motion in a full-fidelity gravitational potential.

The algorithms responsible for the numerical search for periodic orbits date back even further. A particularly influential paper by Howell in 1984 introduces a numerical shooting method to identify periodic orbits in the Earth-Moon system.<sup>8</sup> This algorithm assisted in the discovery of the well known Halo orbit family. In 2003, Doedel et. al. provided a comprehensive description of the

constraints which can be used in conjunction with these shooting methods to find periodic orbits in conservative systems.<sup>9</sup> In 2009, Abad et. al. provided a detailed analytical approach to solving for periodic orbits for the asteroid 216-Kleopatra by using a Lie Transformation and Delaunay orbital elements with a zonal spherical harmonic gravitational potential assumption.<sup>10</sup> In 2012, Yu et. al. used a full polyhedral gravity model and a hierarchical grid search to identify 29 periodic orbits in approximately 12 days of compute time.<sup>2</sup> In 2003, Lan et. al. proposed a variational, cost-function minimization method to identify periodic solutions to high-dimensional problems rather than Poincare sections or traditional Newton-Raphson.<sup>11</sup> Additional useful references regarding the study of frozen and periodic orbits also include Refs. 1, 12, and 13.

To the best of the authors' knowledge, there have been no attempts to apply the shooting method in orbital elements space using the Lagrange planetary equations with a full-fidelity gravity model. Such effort is the primary contribution of this work.



**Figure 1:** Outline of the former cartesian shooting method (top) and the novel orbital element shooting method (bottom).

## METHODOLOGY

The search for periodic orbits is often constructed as a boundary value problem which seeks to minimize the difference between the initial state,  $X(t_0)$  or  $X_0$ , and the state after some period  $T$ ,  $X(t_f)$  or  $X_T$ :

$$X_f - X_0 = 0 \quad (1)$$

where  $X$  corresponds to the cartesian state vector  $[r, v, T]$  which evolves according to

$$\dot{X} = \begin{bmatrix} v \\ a \\ 0 \end{bmatrix} \quad (2)$$

For keplerian motion, this boundary value problem is naturally satisfied; however when non-keplerian forces are introduced like 3rd-body perturbations or non-symmetric gravity perturbations this is no

longer the case. This paper will focus exclusively on the non-symmetric gravity perturbations. To solve this problem, a traditional shooting method is introduced. This begins by taking the Taylor series expansion of the constraint in Equation 1 about the true solution and setting the result to zero.

$$C(\mathbf{X}_0^\star) = C(\mathbf{X}_0 + \delta\mathbf{X}_0) \quad (3)$$

$$0 = C(\mathbf{X}_0) + \frac{\partial C}{\partial \mathbf{X}_0} \delta\mathbf{X}_0 + \text{H.O.T.} \quad (4)$$

This expansion can then be solved for a value of  $\delta\mathbf{X}_0$  through a minimum norm solution:

$$\delta\mathbf{X}_0 = - \left( \left( \frac{\partial C}{\partial \mathbf{X}_0} \right) \left( \frac{\partial C}{\partial \mathbf{X}_0} \right)^T \right)^{-1} C(\mathbf{X}) \quad (5)$$

where

$$\frac{\partial C}{\partial \mathbf{X}_0} = \frac{\partial}{\partial \mathbf{X}_0} (\mathbf{X}_f - \mathbf{X}_0) \quad (6)$$

$$= \left[ \frac{\partial \mathbf{X}_f}{\partial \mathbf{X}_0} - \frac{\partial \mathbf{X}_0}{\partial \mathbf{X}_0}, \frac{\partial \mathbf{X}_f}{\partial T} \right] \quad (7)$$

$$= \left[ \Phi(t_f, t_0) - \mathbb{I}, \frac{\partial \mathbf{X}_f}{\partial T} \right] \quad (8)$$

and  $\Phi$  is the state transition matrix.

Due to the linearization of the original system, this process must be repeated until some termination criteria is reached such as  $|\delta\mathbf{X}_0| < \epsilon$ . This numerical procedure will converge for initial guesses of  $\mathbf{X}_0$  that are sufficiently close to the true periodic solution  $\mathbf{X}_0^\star$  assuming that a reliable gravity model exists which can provide accurate values of the acceleration vector,  $\mathbf{a}$ , in Equation 2.

Typically this procedure is performed in cartesian space as most gravity models are designed such that  $\mathbf{a}$  is a function of  $\mathbf{r}$ . One of the disadvantages of this choice is that cartesian coordinates span all of  $\mathbb{R}^6$  which makes the search space extremely large. Ideally there exists a different coordinate space which has smaller bounds that can be more readily searched. A natural candidate is the traditional orbital elements set for which only a single coordinate (the semi-major axis) can extend from  $(0, \infty)$  whereas the remaining coordinates  $e, i, \omega, \Omega, M$  remain bounded between  $[0, 1], [-\pi/2, \pi/2], [0, 2\pi], [0, 2\pi], [0, 2\pi]$ .

Unfortunately, traditional orbital elements cannot be easily leveraged in numerical shooting methods because the time-derivatives of these elements require a differentiable form of potential with respect to the coordinates themselves. These expressions are especially challenging to compute for high-fidelity gravity models. Explicitly, the dynamics of the osculating orbital elements are com-

puted using the Lagrange planetary equations:

$$\frac{\partial a}{\partial t} = \frac{2}{na} \frac{\partial R}{\partial M} \quad (9)$$

$$\frac{\partial e}{\partial t} = \frac{1-e^2}{na^2e} \frac{\partial R}{\partial M} - \frac{\sqrt{1-e^2}}{na^2e} \frac{\partial R}{\partial \omega} \quad (10)$$

$$\frac{\partial i}{\partial t} = \frac{\cos i}{na^2\sqrt{1-e^2}\sin i} \frac{\partial R}{\partial \omega} - \frac{1}{na^2\sqrt{1-e^2}\sin i} \frac{\partial R}{\partial \Omega} \quad (11)$$

$$\frac{\partial \omega}{\partial t} = -\frac{\cos i}{na^2\sqrt{1-e^2}\sin i} \frac{\partial R}{\partial i} + \frac{\sqrt{1-e^2}}{na^2e} \frac{\partial R}{\partial e} \quad (12)$$

$$\frac{\partial \Omega}{\partial t} = \frac{1}{na^2\sqrt{1-e^2}\sin i} \frac{\partial R}{\partial i} \quad (13)$$

$$\frac{\partial M}{\partial t} = n - \frac{1-e^2}{na^2e} \frac{\partial R}{\partial e} - \frac{2}{na} \frac{\partial R}{\partial a} \quad (14)$$

These equations rely on an differentiable form of the disturbing potential function,  $R$ , which is derived from the expression

$$R(\mathbf{r}) = -\frac{\mu}{r} - U(\mathbf{r}) \quad (15)$$

where  $U$  is the total gravitational potential and  $\mu$  is the gravitational parameter for the body in question.

In some simplified cases, the Lagrange planetary equations can be evaluated analytically as is shown by Ref. 7 and Ref. 10. By expressing the disturbing potential only as a function of low-degree spherical harmonic models, some first-order properties can be derived and candidate periodic or frozen orbits can be found. Unfortunately, this practice becomes intractable as the analytic representation of the disturbing potential function  $R$  increases in fidelity — preventing the use of high-degree spherical harmonics models or the polyhedral gravity model in the Lagrange planetary equations.

### Physic-Informed Neural Network Gravity Model

Fortunately in 2022, Martin and Schaub introduced the Physics-Informed Neural Network Gravity Model (PINN-GM) which simultaneously offers a computationally efficient, high-fidelity representation of the potential which can also be differentiated exactly with respect to arbitrary coordinate sets.<sup>4</sup>

The PINN-GM is a potential model learned directly from position and acceleration estimates of a given celestial body. From this data, a neural network learns both relevant and efficient basis functions to represent the gravitational potential of any arbitrary body by minimizing predictive errors and penalizing violations of underlying differential equations (Laplace’s Equation, conservative vector field properties, and gradients of potential fields). This construction allows for the PINN-GM to learn accurate and physically motivated models using smaller, noisier data sets.

A primary value of the PINN-GM applied to periodic orbit discovery is its ability to differentiated exactly with respect to any element set using automatic differentiation. This property bypasses the limitations of past gravity models such that orbital elements sets and their corresponding Lagrange planetary equations can be used with high-fidelity potential models within traditional shooting methods.

## Automatic Differentiation

Automatic differentiation is method to compute the exact derivative of an algorithm with respect to any input. This is done by constructing a computational graph and using either a forward or backward form of chain rule such as:

$$\begin{aligned}\frac{\partial y}{\partial x} &= \frac{\partial y}{\partial w_{n-1}} \frac{\partial w_{n-1}}{\partial x} \\ &= \frac{\partial y}{\partial w_{n-1}} \left( \frac{\partial w_{n-1}}{\partial w_{n-2}} \frac{\partial w_{n-2}}{\partial x} \right) \\ &= \dots\end{aligned}$$

where  $x$  is the input to some arbitrary algorithm,  $w_i$  are the sequence of intermediate calculations performed to produce the final output,  $y$ .

Because all algorithms are constructed from elementary functions with known derivatives, the partials of each intermediate expressions can always be computed alongside the original calculation. This property ensures that partial of the output  $y$  with respect to any input  $x$  can be computed automatically. Automatic differentiation is best known for its application within deep learning, where it commonly applied in the stochastic gradient descent algorithms used to train neural networks.<sup>14</sup>

Explicitly, neural networks are trained though the the iterative minimization of some loss function. A simple and common loss function is the minimization of a mean squared error function

$$\mathcal{J}(\Theta) = \frac{1}{N_f} \sum_{i=1}^{N_f} |(y_i - \hat{y}_i(x_i|\Theta))|^2, \quad (16)$$

where  $y_i$  is the true output,  $\hat{y}_i(x_i|\Theta)$  is the output predicted by the artificial neural network. The network is parameterized by the vector of trainable parameters  $\Theta$  which includes the weights,  $\mathbf{w}$ , and biases,  $\mathbf{b}$ . Finally,  $N_f$  corresponds to the total number of training data used to update the network. The neural network is trained to minimize Eq. (16) by iteratively updating its weights and biases:

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \Theta} (\mathcal{J}(\mathbf{w})); \quad \mathbf{b}^* = \arg \min_{\mathbf{b} \in \Theta} (\mathcal{J}(\mathbf{b})), \quad (17)$$

This is performed via a gradient descent algorithm like Adam or SGD:<sup>15,16</sup>

$$\Theta^{m+1} = \Theta^m - \eta \nabla_{\Theta^m} \mathcal{J}^m(\Theta), \quad (18)$$

where  $\eta$  is the learning rate and  $m$  is the training iteration.

Equation 18 is where automatic differentiation is applied. The gradient of the prescribed loss function with respect to the weights and biases of a network can be computed automatically and applied in the descent to move the network's parameters in a direction that will minimize that loss function.

This paper extends the application of automatic differentiation beyond training neural networks. Instead, automatic differentiation is applied to a pre-trained PINN-GM which produces a high-fidelity prediction of the disturbing potential function,  $\hat{R}$ , of a celestial body. Consequently, the PINN-GM can be automatically differentiated with respect to arbitrary input coordinates, be that a cartesian set or orbital element set. This allows for the use of the Lagrange planetary equations for periodic orbit discovery in a manner that was not possible before. This paper aims to explore the utility of this novel approach.

## PROBLEM STATEMENT

The aim of this paper is two-fold. First, the paper seeks to quantify the advantages of leveraging the PINN-GM in the periodic orbit problem — specifically seeking to quantify computational advantages and characterize model accuracy. Second, this paper investigates the extent to which it is helpful to frame the periodic orbit discovery problem using an orbital element description as opposed to the traditional cartesian description.

### Characterization of the PINN-GM

To begin, an experiment is proposed which seeks to characterize the speed and accuracy of the PINN-GM compared to that of the polyhedral gravity model. This will ensure the PINN-GM can be exchanged with its more expensive polyhedral model counterpart without loss of dynamical fidelity. The PINN-GM in question is pre-trained on an 8,000 vertex polyhedral gravity model of the asteroid Eros in a manner closely related to Ref. 17. The particular PINN-GM used produces an average acceleration error of 0.5% and a maximum acceleration error of 3.5% on a 100,000 point test set. The training details of the network are omitted from this paper for the purpose of brevity.

**Table 1:** Initial Orbital Element Distribution

Parameter	Value	Parameter	Value
Semi-major axis	$\mathcal{U}[3R_{\text{Eros}}, 7R_{\text{Eros}}]$	Argument of periapsis	$\mathcal{U}[0, 2\pi]$
Eccentricity	$\mathcal{U}[0.1, 0.3]$	Longitude of the Ascending Node	$\mathcal{U}[0, 2\pi]$
Inclination	$\mathcal{U}[-\frac{\pi}{2}, \frac{\pi}{2}]$	Mean Anomaly	$\mathcal{U}[0, 2\pi]$

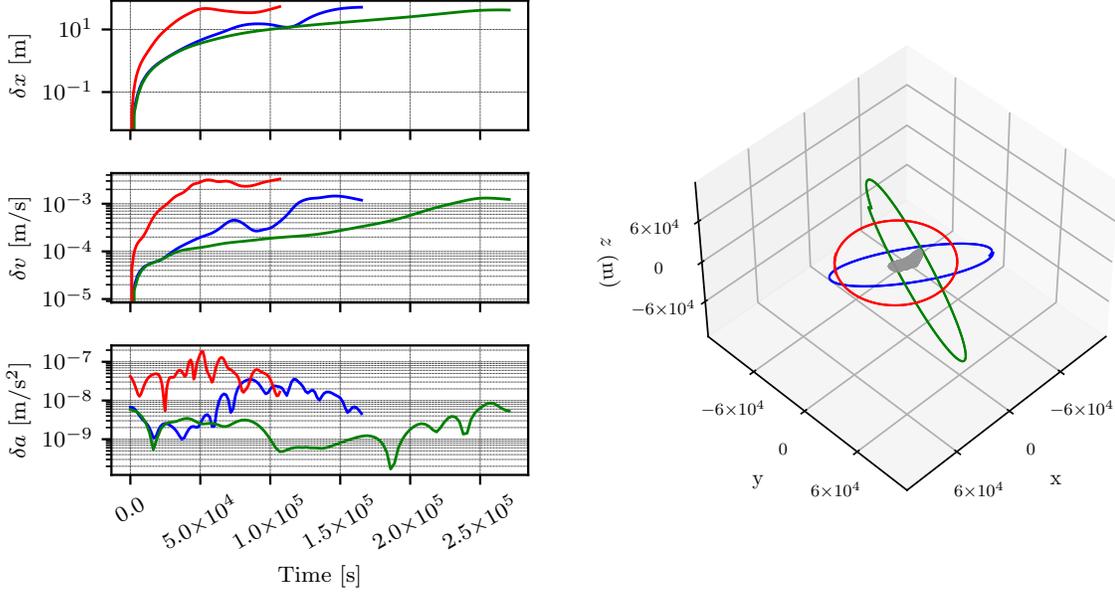
The experiment begins by uniformly sampling three initial conditions about the asteroid 433-Eros according to the orbital element distributions listed in Table 1. These initial conditions are converted to cartesian coordinates and then propagated for a keplerian orbit period,  $T$ , defined as:

$$T = 2\pi\sqrt{\frac{a^3}{\mu}} \quad (19)$$

The initial conditions are propagated using an adaptive Runge-Kutta integrator three times. First, the conditions are propagated with the PINN-GM, second, with a polyhedral gravity model of approximately 8,000 vertices, and third, with a polyhedral gravity model of 200,000 vertices. Note that the 8,000 vertex shape model will be considered ground truth for the purpose of this and future experiments as it was used to train the PINN-GM. The 200,000 vertex model is provided only as reference to demonstrate the computational expense of high-fidelity polyhedral models.

Once propagated, various accuracy and speed metrics between the 8k polyhedral gravity model orbit and the PINN-GM orbit are reported in Figure 2. Figure 2a plots the discrepancy between the position, velocity, and accelerations of the PINN-GM orbit versus the 8k polyhedral orbit. For all three orbits, the PINN-GM model remains below 0.1 km error in position from the true orbit, or  $<0.1\%$  error relative to the magnitude of the orbits’ semi-major axes. The discrepancy between the accelerations produced by the PINN-GM versus the 8k polyhedral model never exceeds  $2\text{E-}7 \text{ m/s}^2$ . These metrics demonstrate that the PINN-GM is a viable representation to replace the true polyhedral model of the asteroid, such that it can be used in the periodic orbit discovery framework.

The computational speed metrics of the PINN-GM are also provided in Figure 2c. The PINN-GM integrates these orbits between approximately 6x and 20x faster than the 8k polyhedral orbit, and



(a) State error of PINN vs 8k Polyhedral

(b) Integrated PINN and Polyhedral Orbits

Orbit	$\delta r$ [m]	$\delta v$ [m/s]	$a_0$ [m]	$T_{\text{Poly 200k}}$	$T_{\text{Poly 8k}}$	$T_{\text{PINN}}$	Speedup ( $T_{\text{Poly 8k}}/T_{\text{PINN}}$ )
Blue	51.404	0.001	67578	351.805	15.163	2.580	5.875x
Green	41.743	0.001	93937	363.103	14.294	0.738	19.353x
Red	53.831	0.003	50523	331.250	14.729	0.840	17.523x

(c) Metrics comparing the PINN and Polyhedral Integrated Orbits

**Figure 2:** Metrics comparing the orbits generated by the PINN-GM and the polyhedral gravity models.

between 135x and 490x faster than the 200k polyhedral orbit. This speed benefit alone makes the PINN-GM a compelling tool in periodic orbit discovery as these trajectories need to be integrated many times given the iterative nature of the shooting method. In Ref. 2, the search for 29 periodic orbit families about the asteroid Kleopatra took 12 days using a high-fidelity polyhedral model. These findings suggest that the same search conducted using a PINN-GM could be conducted on the order of 30 minutes.

### PINN-GM Cartesian Shooting Method

Accepting the PINN-GM as a representative model of the true potential, it can then be used to search for periodic orbits using the shooting method discussed earlier. In this second proposed experiment, the same three initial conditions tested before will be used as initial conditions to then begin the search for periodic orbits using a cartesian shooting method. To compute the state transition matrix used in Equation 5, automatic differentiation is again leveraged to compute the jacobian

$$A = \frac{\partial \dot{\mathbf{X}}}{\partial \mathbf{X}} \quad (20)$$

to then propagate the STM through

$$\dot{\Phi}(t) = A\Phi(t) \quad (21)$$

where  $\Phi(0) = \mathcal{I}$ .

The shooting algorithm's termination is triggered if any of the following criteria are met:

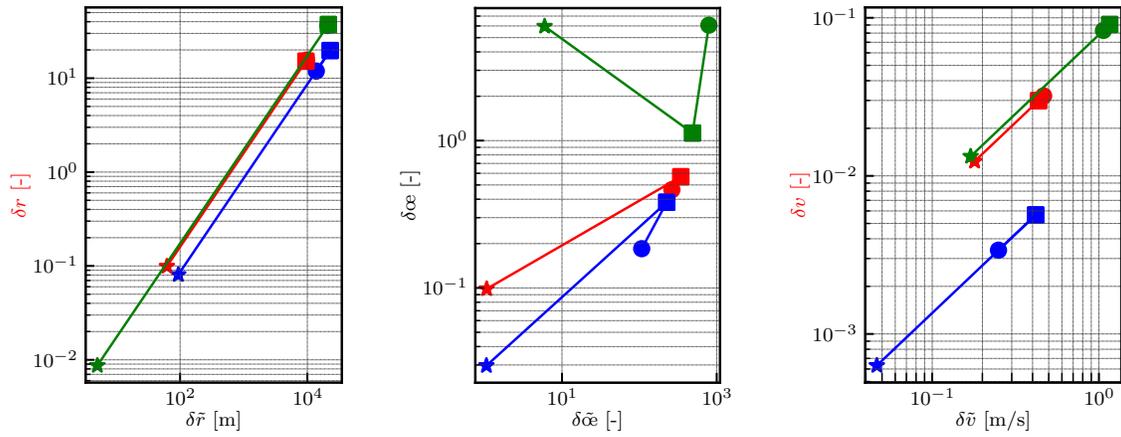
1. The norm of the correction vector,  $|\delta\mathbf{X}_0|$ , is less than  $\epsilon_x(\epsilon_x + |\mathbf{X}_0|)$  where  $\epsilon_x = 10^{-8}$ .
2. The normalized change in the cost function,  $\delta F/F$ , is less than  $\epsilon_f$  where  $\epsilon_f = 10^{-8}$ .
3. The number of iterations exceed 50.

The experiment proceeds by running the shooting methods in two modes: coarse and fine. In the coarse shooting method, higher Runge-Kutta integration error is tolerated than during that of the fine shooting method. This allows for rapid integration of the state and state transition matrix, albeit at lower accuracy. Once the coarse shooting method is completed and the initial conditions are shifted closer to a solution, the fine shooting method is executed starting from the output found with the coarse method. The lower error tolerance of the Runge-Kutta method ultimately assists the algorithm in converging accurately on the local minimum.

*Non-dimensionalization* An important consideration to the shooting method is how the state and equations of motion are non-dimensionalized. For the cartesian shooting method, two intuitive options exist: First, normalize distance by  $l^* = |\mathbf{r}_0|$  and normalize time by the initial period  $t^* = 2\pi\sqrt{a_0^3/\mu}$ . This choice is useful for achieving numerical stability; however, it can bias the results of the optimization process towards the velocity coordinates. Specifically, in certain circumstances the velocity coordinates non-dimensionalize to values that are considerably larger than that of the position coordinates. The shooting method will therefore prioritize solutions which provide closer matches between the initial and final velocity at the expense of the position vector. This behavior leads to false solutions which are not periodic. The second non-dimensionalization option is to continue to scaling distance by scale  $l^* = |\mathbf{r}_0|$ , but instead scale the velocity such that  $t^* = l^*/|\mathbf{v}|$ . This will ensure that the position and velocity are of similar magnitudes in the constraint vector. Unfortunately this will yield extremely large values of the period variable  $T$  which can trigger termination criteria (1) prematurely. In attempts to avoid each of these failure cases, the non-dimensionalization adopted for the cartesian shooting methods moving forward is to normalize distance by  $l^* = |\mathbf{r}_0|/100$  and normalize time by  $t^* = l^* * 10^4/|\mathbf{v}|$ .

*Results* Figure 3 shows a collection of results for the cartesian coarse and fine shooting methods run on the three randomly selected initial conditions. On the first row (Figures 3a, 3b, 3c), the magnitude of the state error between  $\mathbf{X}_f$  and  $\mathbf{X}_0$  in both dimensionalized and non-dimensionalized space is shown. This error is shown for the first and final iterations of the coarse shooting method (circle markers to box markers) and for the first and final iterations of the fine shooting method (box markers to star markers). The figure should be interpreted as follows: The cartesian shooting method seeks to minimize the non-dimensionalized cartesian coordinates  $\mathbf{r}$  and  $\mathbf{v}$  (as opposed to the dimensionalized  $\tilde{\mathbf{r}}$  and  $\tilde{\mathbf{v}}$ ). Figures 3a and 3c demonstrate that this is being accomplished given the shift from high values of  $\delta r$  and  $\delta v$  on the y-axis to low values. Interestingly, this desirable shifting of solutions to smaller values of  $\delta r$  and  $\delta v$  does not guarantee that the solutions found the same initial and final orbital elements (note the green curve in Figure 3b).

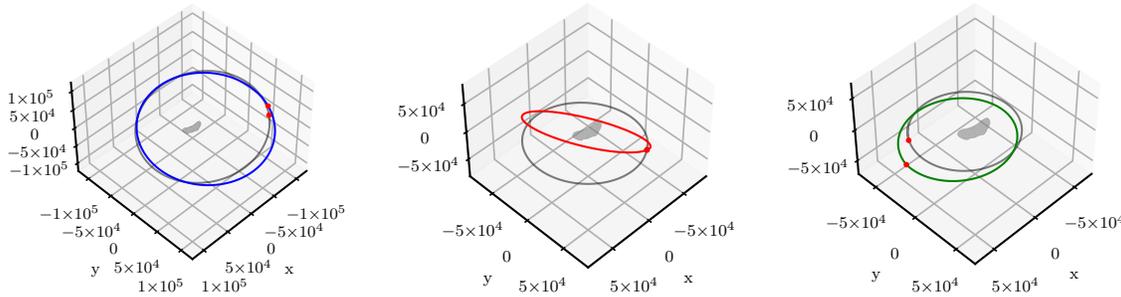
This inconsistency between cartesian state error and orbital element state error between  $\mathbf{X}_f$  and  $\mathbf{X}_0$  manifests an interesting long-term behavior. If the solutions found via the cartesian shooting method are propagated for a single orbit, the difference between their initial coordinate and final



(a) Position Error

(b) Orbital Element Error

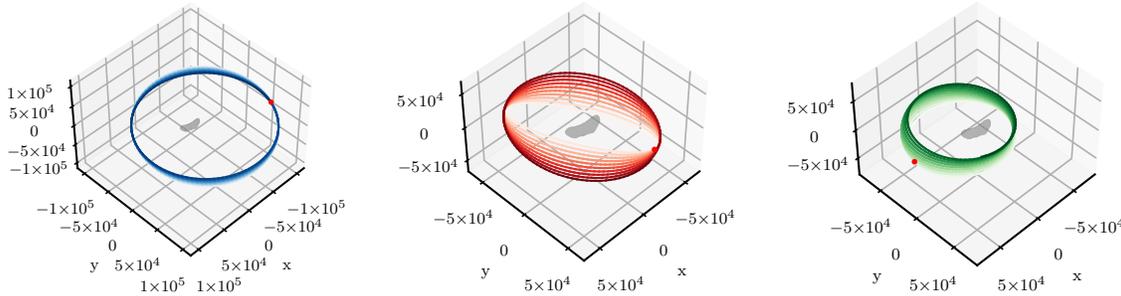
(c) Velocity Error



(d) Solution 1, One Period

(e) Solution 2, One Period

(f) Solution 3, One Period



(g) Solution 1, Ten Periods

(h) Solution 2, Ten Periods

(i) Solution 3, Ten Periods

**Figure 3:** Results from the Cartesian Shooting Method. Top row: Solution error in dimensionalized coordinates (tilde) and non-dimensionalized coordinates (no tilde). Middle row: Starting orbit (gray) and the discovered solutions (color) propagated for one orbit. Bottom row: Solutions found propagated for 10 orbits.

coordinate will be very small (Figures 3d, 3e, 3f); however, over longer periods of time — for example, 10 periods — the orbit will begin to deviate from the initial orbit (Figures 3g, 3h, 3i) and transition into a considerably different orbit than before. This unstable drift away from periodicity in orbital element space provides motivation for why it may be advantageous to solve for periodic

orbits in an orbital element description instead.

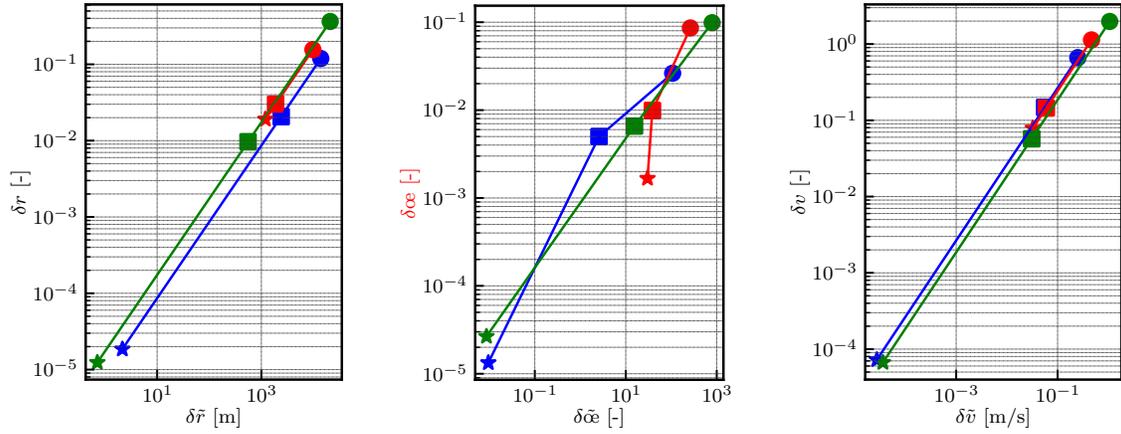
### PINN-GM Orbit Element Shooting Method

A second experiment is proposed which investigates this claim. The search for periodic orbits using the PINN-GM is repeated. However, rather than representing the problem in cartesian coordinates, the problem is reformulated in orbital element space (i.e.  $(r_x, r_y, r_z, v_x, v_y, v_z) \rightarrow (a, e, i, \omega, \Omega, M)$ ). By framing the periodic orbit problem in element space, three advantages are acquired. First, the search space for periodic solutions is considerably reduced. Four of the six elements are naturally bounded by angle wrapping ( $i, \omega, \Omega, M$ ) and the eccentricity must remain less than 1 for the orbit to remain bounded. Only the semi-major axis is allowed to scale to infinity, whereas the solutions in cartesian coordinates can span anywhere in  $(-\infty, \infty)$  for each coordinate. The second advantage of using orbital elements is that it gives mission designers greater control over the types of solutions reached by the shooting method. A shooting method that relies on cartesian coordinates may reach a minimum norm solution for which the state difference is small, but the difference in orbital elements is large. By solving the periodic orbit problem in element space, the shooting method guarantees that the spacecraft will remain in a very similar orbit regardless of the cartesian state error. The third advantage of this approach is that certain orbital elements can be held fixed during the minimization process. For example, if a particular semi-major axis value and inclination are required for a particular mission phase, they can be removed from the set of decision variables but kept as part of the constraint vector. All solutions must thereby maintain the prescribed values of  $a$  and  $i$  and seek periodicity using the other unconstrained elements. Enforcing constraints of this nature is not possible with a cartesian shooting method.

*Non-dimensionalization* Similar to the cartesian shooting method, careful attention must be paid to the non-dimensionalization of the orbital element state vector, jacobian, and constraint vector. Given the varying domains of the different coordinates within the orbital element vector, the non-dimensionalization must be performed as follows: Time is non-dimensionalized using the keplerian period (Equation 19). Distance is non-dimensionalized using the value of the initial semi-major axis. Angles  $i, \omega, \Omega$  and  $M$  are all scaled by  $2\pi$ . In addition, the  $\delta i, \delta \omega, \delta \Omega$ , and  $\delta M$  coordinates in the constraint vector must be the minimum signed angle between the initial and final coordinate as to not disproportionately bias the update (i.e.  $\delta M = 1.99\pi \rightarrow \delta M = -0.01\pi$ ).

*Results* For the experiment, the same three initial conditions are used but are instead propagated and corrected in orbital element space. All elements remain unconstrained in the minimization, and the corresponding improvement in orbital element state error (both dimensionalized and non-dimensionalized) are provided in Figure 4. The experiment also makes use of the coarse and fine shooting methods, where the square corresponds to the final iteration of the coarse method and the star corresponds to the final iteration of the fine method.

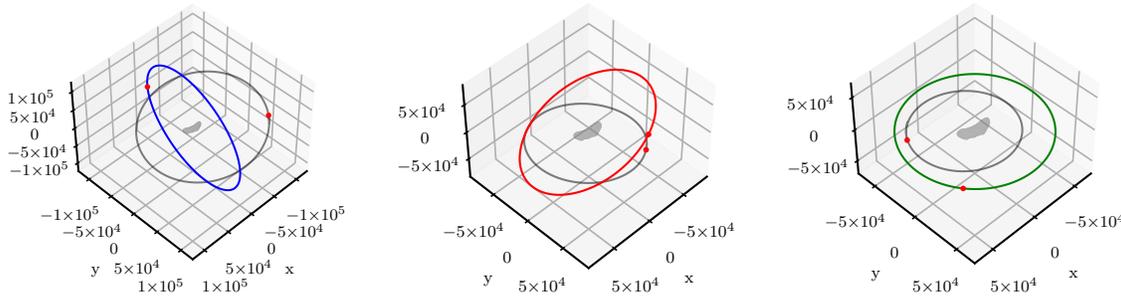
In this experiment, all values of  $\delta r$ ,  $\delta v$ , and  $\delta \epsilon$  decrease, rather than only the cartesian state error as seen Figures 4a, 4b, and 4c. Not only are the solutions found periodic over a single period (Figures 4d, 4e, 4f), but they also remain periodic even after 10 orbits (Figures 4g, 4h, 4i). These findings suggest that the solutions found using the orbital element shooting method are more likely to maintain the desired element set, even if not perfectly periodic. This is best exhibited by the red orbit whose dimensionalized position difference between  $r_0$  and  $r_f$  is in excess of 1 km, yet its orbit elements remain nearly identical over the full 10 periods.



(a) Position Error

(b) Orbital Element Error

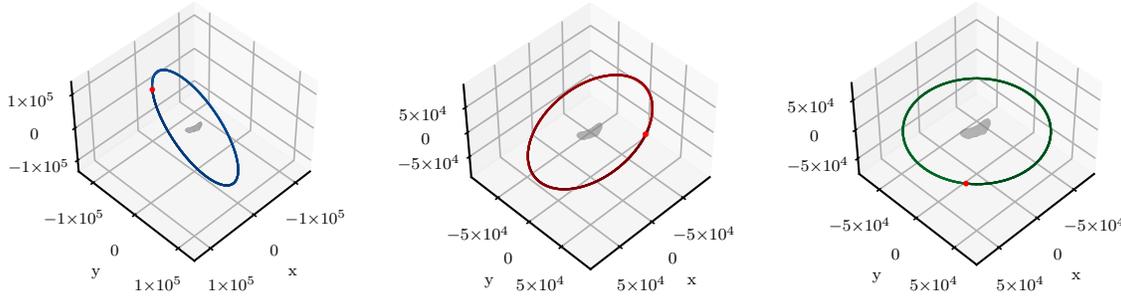
(c) Velocity Error



(d) Solution 1, One Period

(e) Solution 2, One Period

(f) Solution 3, One Period



(g) Solution 1, Ten Periods

(h) Solution 2, Ten Periods

(i) Solution e, Ten Periods

**Figure 4:** Results from the Orbital Elements Shooting Method. Top row: Solution error after one orbit in dimensionalized coordinates (tilde) and non-dimensionalized coordinates (no tilde). Middle row: Starting orbit (gray) and the discovered solutions (color) propagated for one orbit. Bottom row: Solutions found propagated for 10 orbits.

### PINN-GM Constrained Orbital Element Shooting Method

Another advantage of the orbital element shooting method comes from the ability to constrain particular elements of the solution orbit. For example, in some mission settings a specific semi-

major axis and inclination may be required for a mission phase. There are no obvious ways to embed these mission requirements into a cartesian shooting method; however, in the orbital element shooting method, forcing these mission constraints is trivial. By removing the semi-major axis and inclination from the decision variables being optimized over ( $\delta\mathbf{ae}' = \delta\mathbf{ae}_0 / \{\delta a_0, \delta i_0\}$ ), these constraints can remain enforced and the periodicity requirements must be managed by the elements that remain.

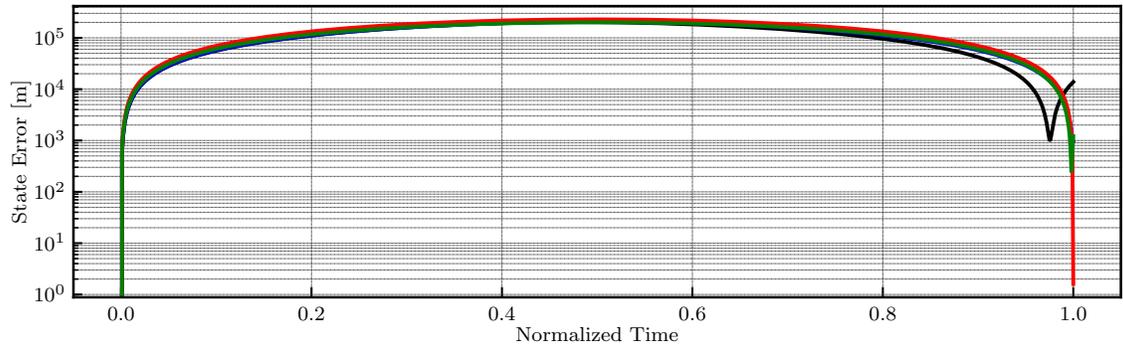
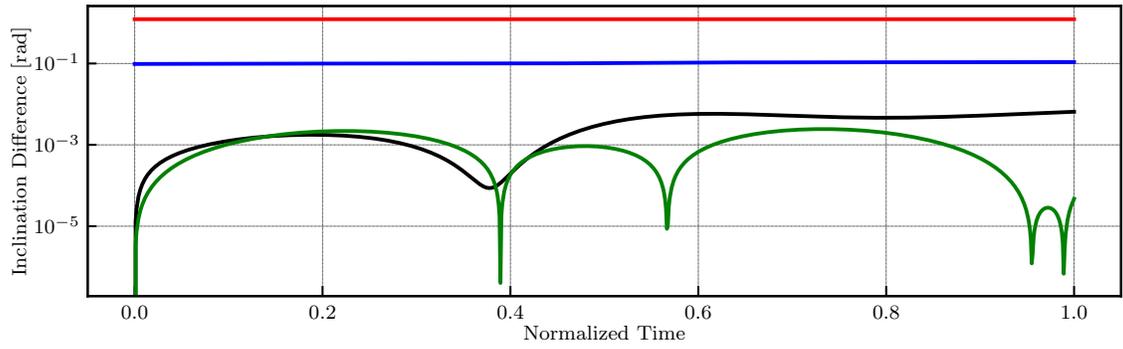
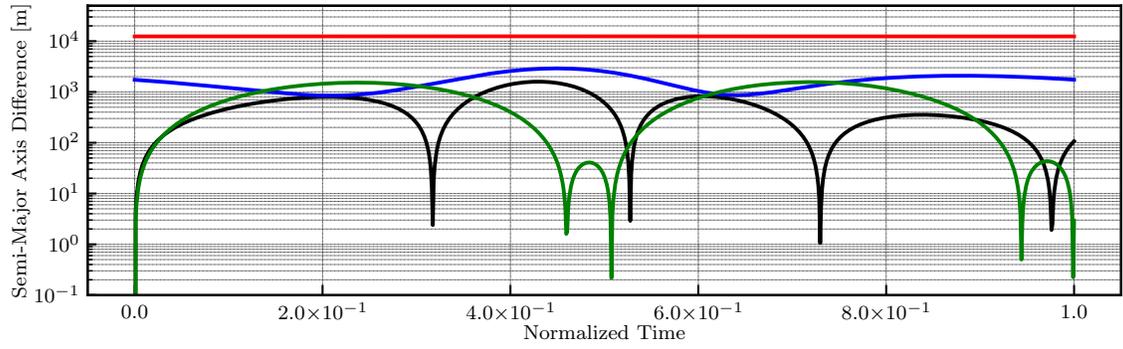
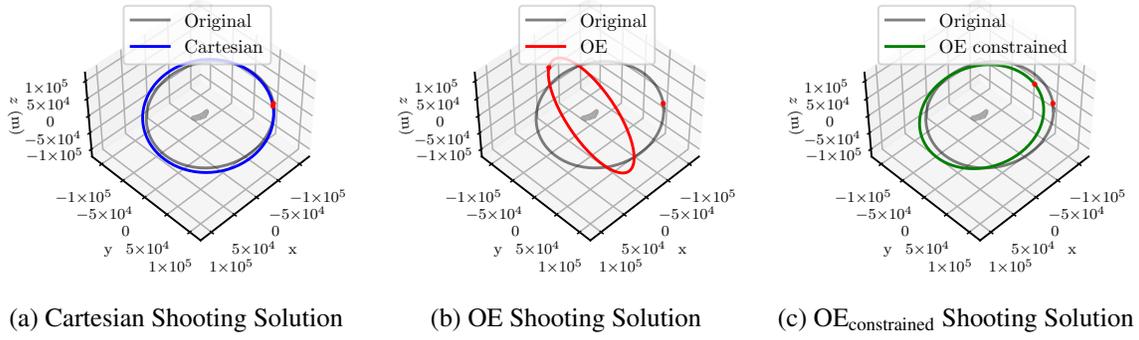
To demonstrate this behavior, an experiment is proposed which samples a single random initial condition and solves for a periodic orbit using three methods. First, a periodic orbit solution is sought using a cartesian shooting method. Second, a periodic orbit is sought with the orbital element shooting method but with no variable constraints. Third and finally, a solution is sought with the orbital element shooting method, but it removes the semi-major axis and inclination from the decision variables of the system. The corresponding solutions found using each method are listed alongside the original initial conditions in Table 2. The propagated orbits are plotted in Figures 5a, 5b, and 5c and solution percent deviations from the desired semi-major axis and inclination as well as the percent deviation from the initial cartesian state are plotted in Figure 5d, 5e, and 5f respectively.

**Table 2:** Initial Conditions and Solutions

Scenario	a	e	i	$\omega$	$\Omega$	M	T
Initial Condition	102324.3326	0.1357	2.8000	-0.8699	-1.4111	-2.9530	307843.3143
Cartesian	104061.6453	0.0990	2.7026	-0.4607	-1.1100	-3.0906	307843.3145
Orbital Elements	114733.6974	0.0100	1.5712	-2.9123	-1.7343	3.1062	364109.1254
OE Constrained	102324.3326	0.0001	2.8000	-1.2218	-1.7581	2.9023	297946.7411

Note how the constrained orbital element shooting method is the only algorithm that converges on an orbit that maintains the specified semi-major axis ( $a^*$ ) and inclination ( $i^*$ ) posed by the original mission requirements. The cartesian shooting method and unconstrained orbital element shooting method both find solutions, but neither solution converges on an orbit which shares values of  $a$  or  $i$  remotely close to the prescribed values. This is best seen by comparing the final y-values along each curve in Figures 5d and 5e.

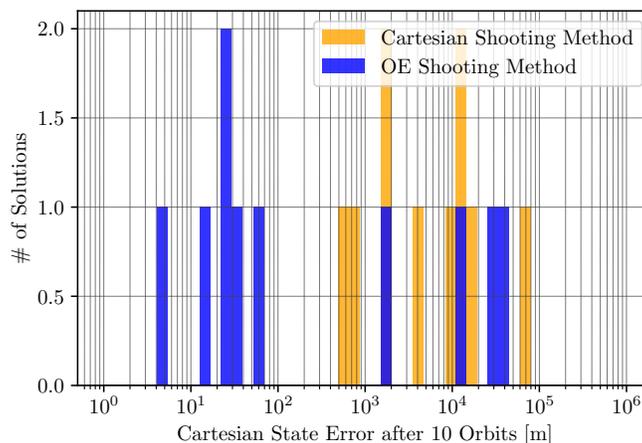
An important observation about the constrained orbital element solution (green) is that its cartesian state error is larger than that of the unconstrained orbital element solution (red) — see the final values in Figure 5f. This is to suggest that there are trade-offs being made by constraining particular elements during the minimization. In principle, a more periodic solution may exist if these constraints are relaxed. Fortunately, such formulation is possible through trust region minimization constraints like those proposed in Ref. 18 and available through the popular `scipy.optimize` Python package. Trust region minimization allows for softer constraints to be applied to the orbital element shooting method where solutions are to be bounded by some prescribed limits. This optimization bounding can be used in conjunction with the orbital element shooting method to assist in finding periodic solution within a wider range of mission parameters rather than enforcing a single value. Characterizing the performance of these algorithms, among others, is left for future work.



**Figure 5:** Orbit solutions found by cartesian shooting method (blue), unconstrained OE shooting method (red), and constrained OE shooting method (green) compared to the original orbit (gray/black).

## Additional Initial Conditions

While the findings presented thus far are encouraging, there are insufficient samples to yield statistically significant conclusions contrasting the two shooting methods. In an attempt to close this gap, an additional ten initial conditions are tested by both the cartesian and the orbital element shooting methods. Once a solution is found or the solver terminates, the solutions are propagated for ten orbits and the resulting position error for each solution is saved and plotted in Figure 6.



**Figure 6:** Average Percent Error of Solutions using the Cartesian LPE and OE LPE

Figure 6 provides further evidence of the advantages posed by the orbital element shooting method over that of the traditional cartesian approach. Of the ten initial guesses, the orbital element shooting method consistently discovers orbits that are closer to periodic than its cartesian counterpart with an average position error on the order of tens of meters versus tens of kilometers in the case of the cartesian solver. This result further supports the idea that by searching for periodic orbits in a orbital element space, more robust solutions are found using fewer initial guesses.

## CONCLUSIONS

This paper investigates the utility of the physics-informed neural network gravity model (PINN-GM) when applied to the periodic orbit discovery problem. Unlike past representations, the PINN-GM can be differentiated with respect to both cartesian coordinates or traditional orbital elements all the while being orders-of-magnitudes faster to evaluate. This allows for traditional shooting methods for boundary value problems to leverage the Lagrange planetary equations and to search for solutions which are periodic or near periodic in orbit element space rather than cartesian space. Preliminary results of this approach demonstrate how the orbital element shooting method discovers orbits that are better at remaining fixed over long periods of time than that of the cartesian shooting method. Future work will investigate if additional orbital element sets yield advantages to this problem, in addition to exploring how collocation or multiple shooting algorithms can be combined with the PINN-GM and the Lagrange planetary equations for faster orbit discovery.

## ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 2040434.

## REFERENCES

- [1] D. J. Scheeres, *Orbital Motion in Strongly Perturbed Environments*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, 10.1007/978-3-642-03256-1.
- [2] Y. Yu and H. Baoyin, “Generating Families of 3D Periodic Orbits about Asteroids: Generating 3D Periodic Orbits about Asteroids,” *Monthly Notices of the Royal Astronomical Society*, Vol. 427, Nov. 2012, pp. 872–881, 10.1111/j.1365-2966.2012.21963.x.
- [3] R. Werner and D. Scheeres, “Exterior Gravitation of a Polyhedron Derived and Compared with Harmonic and Mascon Gravitation Representations of Asteroid 4769 Castalia,” *Celestial Mechanics and Dynamical Astronomy*, Vol. 65, No. 3, 1997, pp. 313–344, 10.1007/BF00053511.
- [4] J. Martin and H. Schaub, “Physics-Informed Neural Networks for Gravity Field Modeling of the Earth and Moon,” *Celestial Mechanics and Dynamical Astronomy*, Vol. 134, Apr. 2022, 10.1007/s10569-022-10069-5.
- [5] A. Güne, G. Baydin, B. A. Pearlmutter, and J. M. Siskind, “Automatic Differentiation in Machine Learning: A Survey,” *Journal of Machine Learning Research*, Vol. 18, 2018, pp. 1–43.
- [6] D. J. Scheeres, “Analysis of Orbital Motion around 433 Eros,” *The Journal of the Astronautical Sciences*, Vol. 43, No. 4, 1995, pp. 427–452.
- [7] D. J. Scheeres, S. J. Ostro, R. S. Hudson, E. M. Dejong, and S. Suzuki, “Dynamics of Orbits Close to Asteroid 4179 Toutatis,” *Icarus*, Vol. 132, No. 1, 1998, pp. 53–79, 10.1006/icar.1997.5870.
- [8] K. C. Howell, “Three-Dimensional, Periodic, ‘Halo’ Orbits,” *Celestial Mechanics and Dynamical Astronomy*, 1984, pp. 53–71.
- [9] E. J. Doedel, R. C. Paffenroth, H. B. Keller, D. J. Dichmann, J. Galán-Vioque, and A. Vanderbauwhede, “Computation of Periodic Solutions of Conservative Systems with Application to the 3-Body Problem,” *International Journal of Bifurcation and Chaos in Applied Sciences and Engineering*, Vol. 13, No. 6, 2003, pp. 1353–1381, 10.1142/S0218127403007291.
- [10] A. Abad, A. Elipse, and E. Tresaco, “Analytical Model to Find Frozen Orbits for a Lunar Orbiter,” *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 3, 2009, pp. 888–898, 10.2514/1.38350.
- [11] Y. Lan and P. Cvitanovi’c, “Variational Method for Finding Periodic Orbits in a General Flow,” 2003.
- [12] D. J. Scheeres, “Orbit Mechanics about Asteroids and Comets,” *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 3, 2012, pp. 987–997, 10.2514/1.57247.
- [13] N. Baresi, *Spacecraft Formation Flight on Quasi-periodic Invariant Tori*. PhD thesis, University of Colorado Boulder, 2017.
- [14] Y. Lecun, Y. Bengio, and G. Hinton, “Deep Learning,” *Nature*, Vol. 521, No. 7553, 2015, pp. 436–444, 10.1038/nature14539.
- [15] S. Ruder, “An Overview of Gradient Descent Optimization Algorithms,” *arXiv*, Sept. 2016, pp. 1–14.
- [16] L. Bottou, *Stochastic Gradient Descent Tricks*, Vol. 7700. Springer, Jan. 2012, 10.1007/978-3-642-35289-8-25.
- [17] J. R. Martin and H. Schaub, “Applications of Physics-Informed Neural Networks for Gravity Field Modeling of Small Bodies,” *Celestial Mechanics and Dynamical Astronomy*, 2022.
- [18] M. A. Branch, T. F. Coleman, and Y. Li, “A Subspace, Interior, and Conjugate Gradient Method for Large-Scale Bound-Constrained Minimization Problems,” *SIAM Journal on Scientific Computing*, Vol. 21, Jan. 1999, pp. 1–23, 10.1137/S1064827595289108.