

---

# Active learning meets fractal decision boundaries: a cautionary tale from the Sitnikov three-body problem

---

**Nicolas Payot**

Département de Physique, Université de Montréal  
Mila - Quebec Artificial Intelligence Institute  
Ciela - Montreal Institute for Astrophysical Data Analysis and Machine Learning

**Mario Pasquato**

Département de Physique, Université de Montréal  
Mila - Quebec Artificial Intelligence Institute  
Ciela - Montreal Institute for Astrophysical Data Analysis and Machine Learning  
Dipartimento di Fisica e Astronomia, Università di Padova  
Vicolo dell'Osservatorio 5, I-35122, Padova, Italy

**Alessandro Alberto Trani**

Department of Earth Science and Astronomy,  
College of Arts and Sciences, The University of Tokyo,  
3-8-1 Komaba, Meguro-ku, Tokyo 153-8902, Japan Okinawa Institute of Science and Technology,  
1919-1 Tancha, Onna-son, Okinawa 904-0495, Japan

**Yashar Hezaveh**

Département de Physique, Université de Montréal  
Mila - Quebec Artificial Intelligence Institute  
Ciela - Montreal Institute for Astrophysical Data Analysis and Machine Learning

**Laurence Perreault-Levasseur**

Département de Physique, Université de Montréal  
Mila - Quebec Artificial Intelligence Institute  
Ciela - Montreal Institute for Astrophysical Data Analysis and Machine Learning

## Abstract

Chaotic systems such as the gravitational N-body problem are ubiquitous in astronomy. Machine learning (ML) is increasingly deployed to predict the evolution of such systems, e.g. with the goal of speeding up simulations. Strategies such as active Learning (AL) are a natural choice to optimize ML training. Here we showcase an AL failure when predicting the stability of the Sitnikov three-body problem, the simplest case of N-body problem displaying chaotic behavior. We link this failure to the fractal nature of our classification problem's decision boundary. This is a potential pitfall in optimizing large sets of N-body simulations via AL in the context of star cluster physics, galactic dynamics, or cosmology.

## 1 Introduction

The gravitational N-body problem lies at the core of computational astronomy. Since the work of Laplace and Lagrange in the 18th century, the stability of the Solar system has been a prominent issue, spurring research to this day. More recently, the gravitational N-body problem is being solved numerically to model systems ranging from icy fragments within the rings of Saturn [1] up to the cosmological scale [2]. Astronomers increasingly apply ML to predict properties of gravitational N-body and related systems, including their dynamical evolution in the chaotic regime [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13].

Predicting the evolution of a chaotic system presents challenges due to sensitive dependence on initial conditions [see e.g. 14], to the point that chaotic benchmarks were proposed for evaluating data-driven forecasting models [15]. Determining the limits of applicability of ML to these systems is thus crucial for astronomy. Here we focus on a narrower question: are techniques such as AL always helpful in training ML models on the gravitational N-body problem in the chaotic regime?

In AL, a model queries the data it deems most informative from an unlabeled pool, selectively requesting labeling with the goal of using fewer labeled samples. This is beneficial when labeled data is costly, e.g. when labelling requires running expensive simulations. AL is becoming increasingly more popular in astronomy, with an early focus on observational surveys [16]. We thus set out to test AL in the simplest N-body setting, a restricted version of the three-body problem known as the Sitnikov problem. This is arguably the simplest case of N-body problem capable of displaying chaotic behavior. The space of its initial conditions is two-dimensional, allowing easy visualization on the plane, and the boundary between stable and unstable systems is indeed fractal.

We train different deep learning models with the goal of predicting the stability of Sitnikov’s problem from its initial conditions, with and without AL. We find that our model’s performance is not improved, all else equal, by AL. We attribute this to the fact that the decision boundary of our problem is fractal, which results in our AL strategy wasting queries to probe it, resulting in a suboptimal sampling of feature space.

## 2 The Sitnikov problem

The Sitnikov three-body problem is a special case of the three-body problem wherein two masses orbit around their common barycenter while a third mass undergoes oscillations along the z axis. For eccentricities of the orbit of the first and the second body greater than 0, it exhibits chaotic behaviour, while the case with  $e = 0$  is integrable [17, 18, 19].

We studied the  $e = 0.5$  case and we selected  $m_1$  and  $m_2$  such that  $G(m_1 + m_2) = 1$  and  $m_1 = m_2$ . The motion of the third body in the Sitnikov problem is then governed by the following differential equations:

$$\frac{dz}{dt} = v_z \tag{1}$$

$$\frac{dv_z}{dt} = -\frac{G(m_1 + m_2)z}{(r(t)^2 + z^2)^{3/2}} \tag{2}$$

Here,  $z$  and  $v_z$  are the position and velocity of the third body on the z-axis, while  $m_i$  is the mass of the  $i^{th}$  body.  $r(t)$  represents the distance between either the first or the second body and their center of mass. The oscillations are independent of the mass of the oscillating body, which is considered massless [19].

This problem involves only two parameters: the initial position of the oscillating body, denoted as  $z_0$ , and its initial velocity along the z-axis, denoted as  $v_{z_0}$ . Changing the initial position of the two orbiting masses is equivalent to altering these two parameters. Consequently, all our simulations commence with the two massive bodies positioned at their respective apocenters.

We implemented our own Runge-Kutta-Fehlberg solver of order 7(8) (RK7(8)) to minimize the numerical error of our simulations [20]. We employed a variable time step to ensure that the error remained below  $\mathcal{O}(10^{-12})$  as it was demonstrated to preserve the analytical solution [21].

A grid of values was computed with 500 values of  $v_{z_0}$  for each of the 200  $z_0$  samples, resulting in a total of 100 000 simulations:

$$z_0 \in [-1, 1]$$

$$v_{z_0} \in [0, 3.5]$$

The simulations were conducted for 100 orbits of the massive bodies after which they were considered stable. If the oscillations of the massless particle were found to be unstable based on straight-forward criteria, the simulation was halted, and the initial conditions were classified as unstable. We define instability as an orbit in which the massless body does not pass through the center of mass again. The criteria used to ascertain instability are as follows:  $z > 10$  ;  $v_z > 0.5$  and  $a_z > -0.01$ .

Essentially, if the body is significantly distant from the center of mass and its acceleration is insufficient to complete another oscillation, it is classified as unstable. A similar criterion applies in cases where  $z < -10$ . These values were selected empirically and validated on a smaller subset of data to confirm their accuracy.

### 3 AL strategy

We began by implementing the classification uncertainty as the AL strategy function. Since the classifier predicts 1 if the simulation is stable or 0 if the simulation breaks up before the end of the simulated time, the maximum uncertainty for classification is 0.5. The classification uncertainty is defined as follows:

$$U(x) = 2|0.5 - P(\hat{x}|x)| \quad (3)$$

where  $x$  represents the simulation under study, and  $\hat{x}$  is the classifier’s prediction [22]. This formulation assigns an uncertainty value close to 0 for the most uncertain simulations and 1 for the most certain ones. While this strategy is straightforward and relatively computationally efficient - as all samples need to be passed through the network - it solely considers the current training model. Consequently, this method does not take into account the intrinsic data characteristics and largely overlooks the data space. Unfortunately, this approach of uncertainty sampling can potentially lead to overfitting when utilizing simulation data, which was observed in this work [23].

Subsequently, another strategy was implemented: the ranked batch-mode sampling, where space exploration is prioritized, and uncertainty sampling is applied subsequently [24]. The scoring function is defined as:

$$\text{score}(x) = \alpha(1 - S(x, Y)) + (1 - \alpha)U(x) \quad (4)$$

$$\alpha = \max\left(\frac{N_X}{N_Y + N_X}c, 1\right) \quad (5)$$

Here,  $Y$  represents the pool of labeled data or the labeled dataset,  $X$  is the set of unlabeled data, and  $N_i$  denotes the number of samples in a dataset. The parameter  $\alpha$  acts as a weighted control that decides whether the sampling process emphasizes space exploration or relies on uncertainty sampling. Initially,  $\alpha$  is set to a higher value, but as more data is labeled using the AL strategy, the unlabeled dataset decreases in size, causing  $\alpha$  to decrease. When  $\alpha$  approaches 1, the strategy prioritizes unexplored regions of the space for selecting new samples, while  $\alpha$  close to 0 signifies reliance on model uncertainty. The scalar parameter  $c$  is accessible to the user and determines the rate at which the transition between the two strategies occurs, making it a new hyper-parameter. The similarity function  $S(x, Y)$  measures the similarity between  $x$  - the sample under study - and  $y$  - the labeled dataset. The value returned by  $S(x, Y)$  should fall between 0 and 1, with 1 indicating that the sample  $x$  lies in a completely unexplored portion of the space, and 0 denoting that it coincides with another sample. For instance, cosine similarity can be employed as a suitable choice [24]. The sample with the minimum score is then selected for labeling, added to the labeled dataset, and removed from the unlabeled dataset. This process is repeated for the required number of queries.

### 4 Results

The neural network employed is a straightforward fully-connected network comprising seven layers. We conducted tests with this model on all our simulations to confirm its capability to effectively learn the spatial information [see figure 1]. The neural network generated a single output value: 1 if it predicted that the initial conditions would lead to a stable system or 0 if not. Since most of the data could be easily classified as stable or unstable, the dataset was pruned to eliminate these trivial cases. Out of the original 100 000 simulations, only 46 502 fell within the selected criteria. For AL, 400 000 potential initial conditions were generated from a uniform distribution and selected using the same criteria as for the dataset. The remaining 278 980 initial conditions were designated for AL sampling. When AL was not employed, the same number of simulations as would be seen in total by the AL strategy were selected from the dataset.

We utilized a learning rate of 0.001 and set  $c$  to 0.5, to facilitate exploration of the parameter space, with a primary focus on modeling uncertainty. We examined a range of values for  $c$  from 0 to 1, determining that  $c = 0.5$  yielded the most favorable results. All models underwent a training process that spanned 400 epochs, utilizing a NVIDIA A100. In the context of our AL approach,

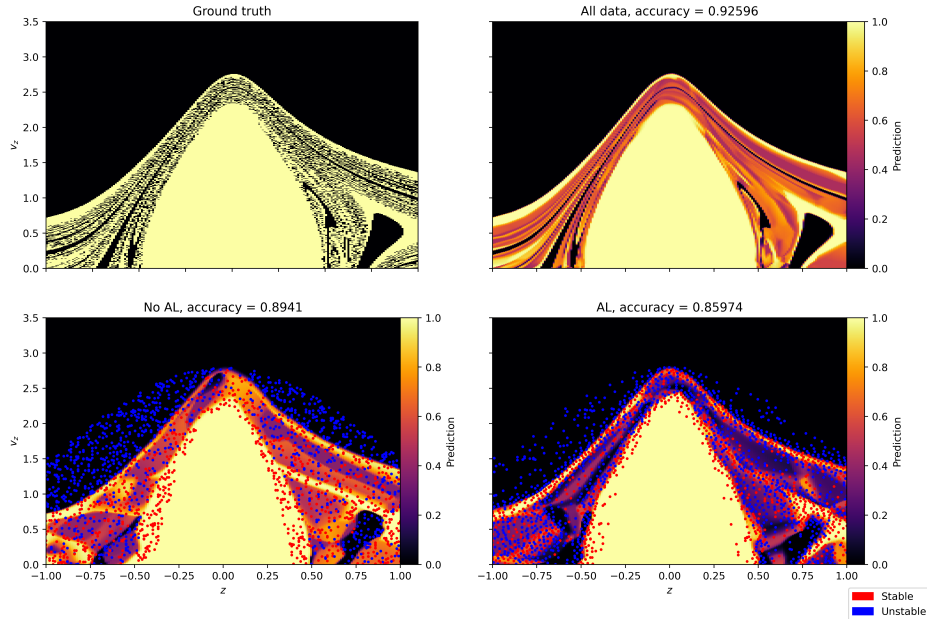


Figure 1: Top row: Ground truth and a model trained on all the data, achieving an accuracy of 92.596%. The ground truth is limited by the number of simulations performed. Bottom row: A model trained without AL, which achieved an accuracy of 89.410%, and a model trained with AL, with an accuracy of 85.974%. The blue and red points respectively represent unstable and stable simulations used during training.

the neural network began with an initial dataset of 500 simulations. As the training progressed, we incrementally added 10 additional simulations from the initial condition dataset during each epoch, specifically from epochs 50 to 200, accumulating a total of 1500 new simulations.

Incorporating AL significantly reduces the number of simulations sampled in the trivial cases, as expected. It places a focus on boundary cases where predicting the stability of a simulation is more challenging. However, a drawback arises due to the fractal-shaped boundaries in the Sitnikov problem, as documented in prior research [25, 26]. The model never truly converges to a satisfactory result because in the chaotic regime arbitrarily small perturbations in initial conditions can lead to different outcomes.

Increasing  $c$  pushes the model to explore the space more extensively. However, this exploration eventually essentially resembles random sampling. As the model begins to prioritize uncertainty sampling, it zooms in on these intricate and ill-defined boundaries. This excessive focus on narrow boundary regions results in the neglect of other parts of the space, leading to poor generalization. This behavior is depicted in Figure 2, which illustrates the impact of different values of  $c$  on the model's ability to generalize and its performance.

We also investigated the model's performance when trained on the complete dataset gathered via the AL process, as opposed to employing random sampling. This approach was intended to ascertain if the simulations selected through AL were inherently the source of the poor performance. The outcomes revealed that the model trained on the AL-derived dataset demonstrated performance closely aligned with that of the model trained using AL itself [see figure 3].

## 5 Conclusions

We used the Sitnikov restricted three-body problem to compare the performance achieved by a neural network trained to predict stability using an AL approach versus random sampling. We found that

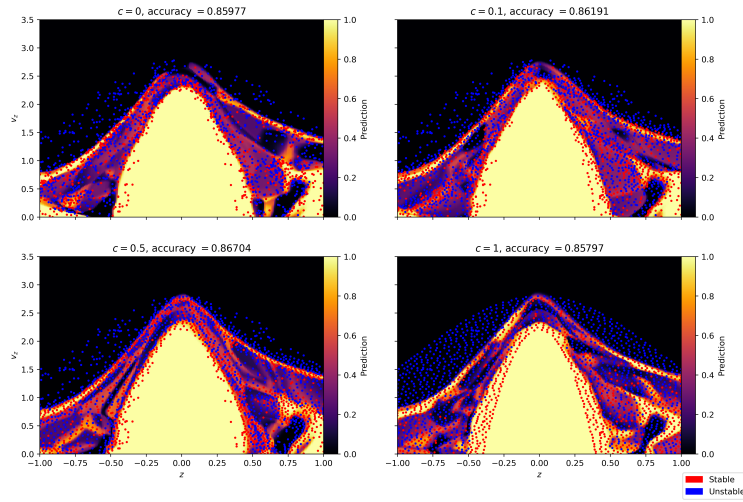


Figure 2: Models trained using the same parameters as previously stated, but with different values for  $c$  ranging from 0 to 1. The model achieving the best accuracy is the one trained with  $c = 0.5$ . The blue and red points respectively represent unstable and stable simulations used during training.

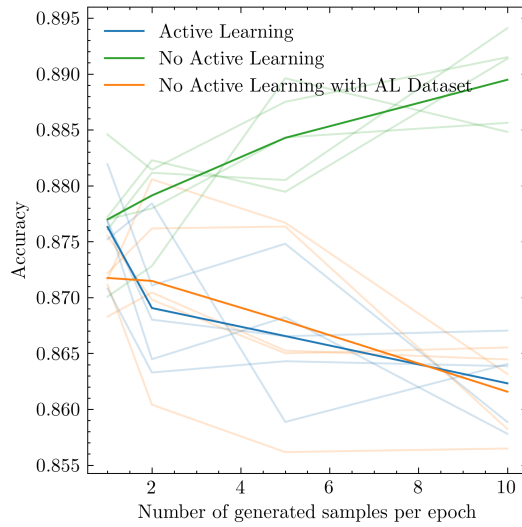


Figure 3: Comparison of model accuracies trained with AL and without AL, against the number of generated samples per epoch. The 'No Active Learning with AL Dataset' curve represents a model trained on a dataset selected through AL but without the AL strategy during training. The accuracy values represent averages over multiple training runs with different random seeds to ensure reliability of the performance measure.

AL does not result in improved performance. We attribute this failure to the fractal nature of the decision boundary for our problem. Real life applications of AL to optimize sets of gravitational N-body simulations are likely to face the same issue. Our results are therefore valuable as a cautionary tale for astronomers planning such an undertaking.

## Acknowledgments and Disclosure of Funding

M. P. acknowledges financial support from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 896248. A.A.T. acknowledges support from the European Union’s Horizon 2020 and Horizon Europe research and innovation programs under the Marie Skłodowska-Curie grant agreements no. 847523 and 101103134.

This work is supported by the Simons Collaboration on “Learning the Universe”. The Flatiron Institute is supported by the Simons Foundation. The work is in part supported by computational resources provided by Calcul Quebec and the Digital Research Alliance of Canada. Y.H. and L.P. acknowledge support from the Canada Research Chairs Program, the National Sciences and Engineering Council of Canada through grants RGPIN-2020- 05073 and 05102, and the Fonds de recherche du Québec through grants 2022-NC-301305 and 300397. P.L acknowledges support from the Simons Foundation.

## References

- [1] Hiroshi Daisaka, Hidekazu Tanaka, and Shigeru Ida. Viscosity in a Dense Planetary Ring with Self-Gravitating Particles. *Icarus*, 154(2):296–312, December 2001.
- [2] Liang Wang, Aaron A. Dutton, Gregory S. Stinson, Andrea V. Macciò, Camilla Penzo, Xi Kang, Ben W. Keller, and James Wadsley. NIHAO project - I. Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations. *Monthly Notices of The Royal Astronomical Society*, 454(1):83–94, November 2015.
- [3] T. A. F. Comerford and R. G. Izzard. Estimating the outcomes of common envelope evolution in triple stellar systems. *Monthly Notices of The Royal Astronomical Society*, 498(2):2957–2967, October 2020.
- [4] Philip G. Breen, Christopher N. Foley, Tjarda Boekholt, and Simon Portegies Zwart. Newton versus the machine: solving the chaotic three-body problem using deep neural networks. *Monthly Notices of The Royal Astronomical Society*, 494(2):2465–2470, May 2020.
- [5] Maxwell X. Cai, Simon Portegies Zwart, and Damian Podareanu. Neural Symplectic Integrator with Hamiltonian Inductive Bias for the Gravitational  $N$ -body Problem. *arXiv e-prints*, page arXiv:2111.15631, November 2021.
- [6] Pablo Lemos, Niall Jeffrey, Miles Cranmer, Shirley Ho, and Peter Battaglia. Rediscovering orbital mechanics with machine learning. *arXiv e-prints*, page arXiv:2202.02306, February 2022.
- [7] Hongwei Yang, Jiumei Yan, and Shuang Li. Fast computation of the Jovian-moon three-body flyby map based on artificial neural networks. *Acta Astronautica*, 193:710–720, April 2022.
- [8] M. I. Ikhsan and M. I. Arifyanto. Exploring multi-planet system wasp-148 using n-body simulation and deep learning. In *Journal of Physics Conference Series*, volume 2243 of *Journal of Physics Conference Series*, page 012010, June 2022.
- [9] Jiumei Yan, Hongwei Yang, and Shuang Li. ANN-based method for fast optimization of Jovian-moon gravity-assisted trajectories in CR3BP. *Advances in Space Research*, 69(7):2865–2882, April 2022.
- [10] Alessandra Celletti, Catalin Gales, Victor Rodriguez-Fernandez, and Massimiliano Vasile. Classification of regular and chaotic motions in Hamiltonian systems with deep learning. *Scientific Reports*, 12:1890, January 2022.
- [11] Xin Li, Jian Li, Zhihong Jeff Xia, and Nikolaos Georgakarakos. Machine-learning prediction for mean motion resonance behaviour - The planar case. *Monthly Notices of The Royal Astronomical Society*, 511(2):2218–2228, April 2022.

- [12] Florian Lalande and Alessandro Alberto Trani. Predicting the Stability of Hierarchical Triple Systems with Convolutional Neural Networks. *Astrophysical Journal*, 938(1):18, October 2022.
- [13] Xin Li, Jian Li, Zhihong Jeff Xia, and Nikolaos Georgakarakos. Large-step neural network for learning the symplectic evolution from partitioned data. *Monthly Notices of The Royal Astronomical Society*, 524(1):1374–1385, September 2023.
- [14] Tianli Hu and Shijun Liao. On the risks of using double precision in numerical simulations of spatio-temporal chaos. *Journal of Computational Physics*, 418:109629, October 2020.
- [15] William Gilpin. Chaos as an interpretable benchmark for forecasting and data-driven modelling. *arXiv e-prints*, page arXiv:2110.05266, October 2021.
- [16] M. Leoni, E. E. O. Ishida, J. Peloton, and A. Möller. Fink: Early supernovae Ia classification using active learning. *Astronomy and Astrophysics*, 663:A13, July 2022.
- [17] W. D. MacMillan. An integrable case in the restricted problem of three bodies. *The Astronomical Journal*, 27:11, May 1911.
- [18] K. Sitnikov. The Existence of Oscillatory Motions in the Three-Body Problem. *Soviet Physics Doklady*, 5:647, January 1961.
- [19] Johannes Hagel and Christoph Lhotka. A High Order Perturbation Analysis of the Sitnikov Problem. *Celestial Mech Dyn Astr*, 93(1-4):201–228, September 2005.
- [20] Erwin Fehlberg. *Classical Fifth-, Sixth-, Seventh-, and Eighth-order Runge-Kutta Formulas with Stepsize Control*. National Aeronautics and Space Administration, 1968.
- [21] Rudolf Dvorak and Christoph Lhotka. *Celestial Dynamics: Chaoticity and Dynamics of Celestial Systems*. Wiley, 1 edition, April 2013.
- [22] Burr Settles. Active learning literature survey. In *Active Learning Literature Survey*, 2009.
- [23] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G. Hauptmann. Multi-Class Active Learning by Uncertainty Sampling with Diversity Maximization. *Int J Comput Vis*, 113(2):113–127, June 2015.
- [24] Thiago N.C. Cardoso, Rodrigo M. Silva, Sérgio Canuto, Mirella M. Moro, and Marcos A. Gonçalves. Ranked batch-mode active learning. *Information Sciences*, 379:313–337, February 2017.
- [25] T. Kovács and B. Érdi. The structure of the extended phase space of the Sitnikov problem. *Astron. Nachr.*, 328(8):801–804, October 2007.
- [26] R. Dvorak. Numerical Results to the Sitnikov-Problem. *Celestial Mechanics and Dynamical Astronomy*, 56:71–80, May 1993. ADS Bibcode: 1993CeMDA..56...71D.
- [27] Alessandro A. Trani, Mario Spera, Nathan W. C. Leigh, and Michiko S. Fujii. The keplerian three-body encounter. II. comparisons with isolated encounters and impact on gravitational wave merger timescales. *ApJ*, 885(2):135, 2019.
- [28] Alessandro A. Trani, Michiko S. Fujii, and Mario Spera. The keplerian three-body encounter. i. insights on the origin of the s-stars and the g-objects in the galactic center. *ApJ*, 875(1):42, 2019.
- [29] Florian Lalande and Alessandro Alberto Trani. Predicting the stability of hierarchical triple systems with convolutional neural networks. *ApJ*, 938(1):18, 2022.

## 6 Appendix

---

### Algorithm 1 Score pseudo code

---

```
function SCORE( $x, Y, c, u$ )  $\triangleright x$  is the total data set,  $Y$  is the labeled or training data set,  $u$  is the uncertainty  
 $\alpha \leftarrow \frac{\text{len}(x)}{\text{len}(x)+\text{len}(Y)} * c$   
 $\text{term}_1 \leftarrow \alpha * (1 - S(x, Y))$   
 $\text{term}_2 \leftarrow (1 - \alpha) * u$   
return  $\text{term}_1 + \text{term}_2$   
end function
```

---

---

### Algorithm 2 Uncertainty pseudo code

---

```
function U( $x$ )  $\triangleright$  See Algorithm 3 for usage  
  subdiv  $\leftarrow$  64  $\triangleright$  Same as batch size  
  res  $\leftarrow$  []  $\triangleright$  Will contain the uncertainties  
  for  $i$  from 0 to  $(\text{len}(x) // \text{subdiv} + 1)$  do  
    end  $\leftarrow$   $\min(\text{len}(x), (i + 1) * \text{subdiv})$   
    predictions  $\leftarrow$  model( $x[i * \text{subdiv} : \text{end}]$ )  
    predictions  $\leftarrow$   $2 * \text{abs}(0.5 - \text{predictions})$   
    res  $\leftarrow$  Concatenate(res, predictions)  
  end for  
  return res  
end function
```

---

---

### Algorithm 3 AL query pseudo code

---

```
function QUERY( $n, x, Y, c$ )  $\triangleright n$  is the amount of samples to query,  $x$  is the total data set,  $Y$  is the labeled or training data set  
  toLabel  $\leftarrow$  []  
   $u \leftarrow u(x)$   
  for  $i$  from 0 to  $n$  do  
    scores  $\leftarrow$  SCORE( $x, Y, c, u$ )  
    idx  $\leftarrow$  argmin(scores)  
    toLabel  $\leftarrow$  Concatenate(toLabel,  $x[\text{idx}]$ )  
     $Y \leftarrow$  Concatenate( $Y, x[\text{idx}]$ )  
     $x \leftarrow$  Concatenate( $x[:\text{idx}], x[\text{idx}+1:]$ )  
     $u \leftarrow$  Concatenate( $u[:\text{idx}], u[\text{idx}+1:]$ )  
  end for  
  return toLabel  
end function
```

---

After 3, the new samples need to be labeled using an oracle. In our case, we run the simulations to see if they are stable or not.



A neural network’s capability to effectively classify three-body problems using time series data was demonstrated in [12]. Building on this, we employed a simple neural network to the initial conditions of a generalized three-body simulation to predict stability, utilizing the same active learning scheme. Our investigation focused on a scenario where a star orbits a binary system. The simulations were carried out within the TSUNAMI framework [27, 28], exploring a wide range of initial conditions given in Table 1. For simplicity, all masses were assumed to be equal.

Category	Parameter	Symbol	Value
Binary	Semimajor axis	$a_1$	1
	Eccentricity	$e_1$	$U[0, 1]$
	Argument of periapsis	$\omega_1$	$U[-\pi, \pi]$
	Longitude of the ascending node	$\Omega_1$	$U[0, 2\pi]$
	Inclination	$i_1$	$U[0, \pi]$
Orbiting star	Semimajor axis	$a_2$	$U[0, 3]$
	Eccentricity	$e_2$	$U[0, 1]$
	Argument of periapsis	$\omega_2$	$U[-\pi, \pi]$

Table 1: Orbital elements sampled for the generalized three-body simulations.  $U[a, b]$  denotes a uniform distribution between the limits  $a$  and  $b$ .

The distribution of the second semimajor axis,  $a_2$ , was chosen to balance the number of stable and unstable simulations. Orbital elements were sampled using the Latin hypercube sampling method, initially between 0 and 1, then scaled to the appropriate values. The inclination of the orbiting star and its longitude of the ascending node were set to 0, as they merely result in a rotation of the plane in which the simulation is conducted, providing no additional information. The true anomaly of both the binary and the orbiting star was initialized at 0.

Systems were simulated for a duration expressed in units of the binary’s first complete orbit time divided by  $2\pi$ , denoted as  $P_1$ .

$$P_1 = \sqrt{\frac{a_1^3}{G(m_1 + m_2)}} = \sqrt{\frac{1}{2m_1 G}} \quad (6)$$

Simulations were retained only if the system remained intact for at least  $500P_1$ . They were then continued until  $10,000P_1$  or a breakup occurred. Breakup was detected if the ratio of the two semimajor axes deviated from the initial conditions by more than 15%, a criterion also utilized in other classification tasks [29]. Orbital elements were recorded every  $\pi P_1$  for the first  $500P_1$ , provided no breakup occurred before this threshold. Simulations were deemed unstable if  $e_1$  or  $e_2$  fell outside the inclusive range of 0 to 1 exclusive.

A dataset comprising 48,262 stable and 44,533 unstable simulations was assembled to perform the active learning task. Unlike in the Sitnikov problem, applying the previously described AL scheme to this dataset did not exhibit the same drawbacks. We hypothesize that this difference arises from the smaller fractal and chaotic regions within the space of the generalized three-body problem. Consequently, the AL scheme is less likely to become ensnared in these complex areas. Additionally, the exploratory component of our AL strategy should aid in avoiding getting trapped in such smaller challenging regions, thereby not perturbing the overall effectiveness of the learning process.

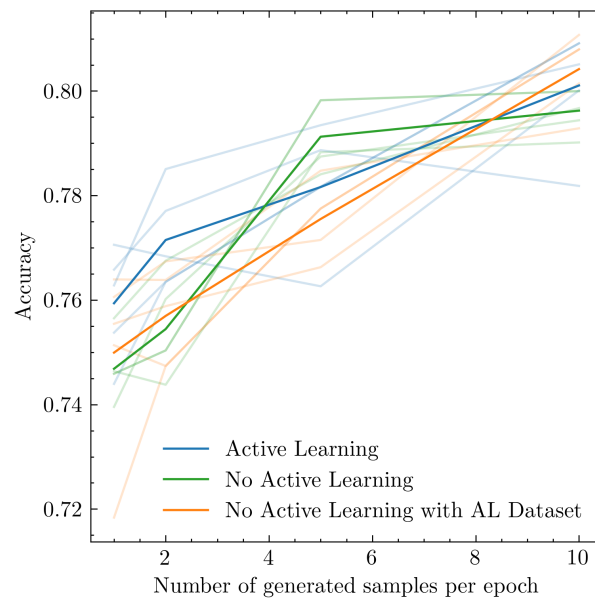


Figure 4: Comparison of model accuracies using Active Learning (AL) and without AL against the number of samples generated per epoch in the generalized three-body problem. The 'No Active Learning with AL Dataset' curve shows the accuracy of a model trained on an AL-selected dataset but without implementing AL during training. Accuracy values are averaged over multiple runs with varied random seeds to guarantee performance measure reliability.