

# Generative Design of Periodic Orbits in the Restricted Three-Body Problem

Alvaro Francisco Gil<sup>\*1</sup>, Walther Litteri<sup>2</sup>, Victor Rodriguez-Fernandez<sup>1</sup>, David Camacho<sup>1</sup>, Massimiliano Vasile<sup>2</sup>

<sup>1</sup>*School of Computer Systems Engineering, Universidad Politécnica de Madrid, Spain*

<sup>2</sup>*Aerospace Centre of Excellence, University of Strathclyde, Glasgow, United Kingdom*

**The Three-Body Problem has fascinated scientists for centuries and it has been crucial in the design of modern space missions. Recent developments in Generative Artificial Intelligence hold transformative promise for addressing this longstanding problem. This work investigates the use of Variational Autoencoder (VAE) and its internal representation to generate periodic orbits. We utilize a comprehensive dataset of periodic orbits in the Circular Restricted Three-Body Problem (CR3BP) to train deep-learning architectures that capture key orbital characteristics, and we set up physical evaluation metrics for the generated trajectories. Through this investigation, we seek to enhance the understanding of how Generative AI can improve space mission planning and astrodynamics research, leading to novel, data-driven approaches in the field.**

## 1 Introduction

Recent advancements in Artificial Intelligence (AI) have significantly impacted the space sector [1, 2]. AI techniques have been applied in tasks such as space traffic management [3, 4], space object characterization [5, 6], satellite pose estimation [7], natural language processing for space mission design [8], and spacecraft operations using large language models [9]. Research on large pre-trained models has also emerged in astronomy, with applications in both generative and discriminative tasks. Notable models include AstroCLIP [10], which uses cross-modal contrastive learning for astronomical images and spectra, and ASTROMER [11], a transformer-based model for creating representations of light curves in a self-supervised manner.

In astrodynamics, machine learning, especially deep learning, has seen significant growth. Models are used to learn guidance or control laws [12], and deep learning has been applied to design complex tra-

jectories in multi-body dynamics [13], the solution of two-point boundary value problems [14] and the classification of regular and chaotic motion [15]. Self-supervised learning techniques have shown effectiveness in tasks such as conjunction screening and maneuver detection, revealing potential for data-driven approaches in orbit analysis [16].

In the context of the Three-Body Problem (3BP), AI has been used to tackle complex challenges. Support vector machines classified trajectories in the circular restricted three-body problem [17], deep neural networks solved the chaotic three-body problem [18], and artificial neural networks predicted periodic orbits in three-body systems with arbitrary masses [19]. However, generative AI techniques have only recently been applied to astrodynamics [20, 21]. This paper presents some early results from the OrbitGPT project [22].

### 1.1 The Restricted Three-Body Problem

The Three-Body Problem (3BP), introduced by Newton in 1687 [23], involves predicting the motion of three celestial bodies interacting gravitationally. Euler and Lagrange made significant contributions in the 18th century, describing equilibrium points where a small, massless particle influenced by two larger bodies (primaries) can remain stationary. Euler identified three collinear points, while Lagrange added two more points, assuming the primaries move in circular motion about their center of mass [24]. These models are known as the Restricted Three-Body Problem (R3BP) and the Circular-Restricted Three-Body Problem (CR3BP).


The equations of motion (EOM) for the CR3BP are:

$$\ddot{x} = 2\dot{y} + U_x \quad (1a)$$

$$\ddot{y} = -2\dot{x} + U_y \quad (1b)$$

$$\ddot{z} = U_z \quad (1c)$$

<sup>\*</sup>Corresponding author. E-Mail: alvaro.francisco.gil@alumnos.upm.es



**Figure 1:** Geometry of the Circular-Restricted Three-Body Problem in a rotating reference frame.

The state vector  $\mathbf{X} = [x, y, z, \dot{x}, \dot{y}, \dot{z}]^T \in \mathbb{R}^6$  represents the position and velocity of the particle  $P$  relative to the center of mass in the non-inertial reference frame, as shown in Figure 1 [25].

The state vector and its EOM are dimensionless, with the primaries at a fixed relative distance of 1 and rotating at an angular speed of 1. The potential function  $U$  is given by:

$$U(x, y, z) = \frac{1}{2}(x^2 + y^2) + \frac{1-\mu}{r_1} + \frac{\mu}{r_2} \quad (2)$$

where  $r_1$  and  $r_2$  are the distances between  $P$  and the two bodies. The components of its gradient,  $U_x, U_y, U_z$ , are the forcing terms in Eqs. 1. The gravitational parameter  $\mu = m_2/(m_1 + m_2)$  defines the system's behavior, with  $0 \leq \mu \leq 0.5$ . For the Earth-Moon system,  $\mu = 0.01215$ .

## 1.2 Periodic Orbits

Let  $\Phi(t, \mathbf{X}_0) : \mathbb{R} \times \mathbb{R}^6 \rightarrow \mathbb{R}^6$  be the evolution function for the dynamical system described by the equations of motion (EOM) in Eqs. 1. This function determines the state vector  $\mathbf{X}(t)$  at any time  $t$  based on the initial state  $\mathbf{X}_0$ :  $\mathbf{X}(t) = \Phi(t, \mathbf{X}_0)$ . A periodic orbit with period  $T$  is a trajectory where the state vector returns to its initial state after time  $T$ :  $\Phi(T, \mathbf{X}_0) = \mathbf{X}_0$ .

Periodic and quasi-periodic orbits in the Circular Restricted Three-Body Problem (CR3BP) have fascinated scientists since the 19th century, especially with the advent of numerical computation. In 1892, Poincaré proved that infinite periodic solutions exist for the Three-Body Problem and highlighted the potential for chaotic motion [26]. The discovery of unique families of periodic orbits in both two-dimensional and three-dimensional CR3BP began in the 1920s [27]. Within any family of periodic trajectories, characteristics such as period, energy, and stability indicators vary continuously.

Interest in these orbits surged with the dawn of space missions, offering cost-effective and scientifically valuable opportunities. For instance, the International Sun/Earth Explorer 3 (ISEE-3) mission, launched in 1978, was the first to utilize a Halo orbit around the Lagrange point  $L_1$  in the Sun-Earth system [28]. More recently, the James Webb Space Telescope, launched in December 2021, operates in a Halo orbit around the  $L_2$  point in the same system [29]. NASA's CAPSTONE mission, launched in June 2022, uses a more stable Near Rectilinear Halo Orbit (NRHO) in the cislunar environment [30].

The analysis of periodic orbits in the CR3BP is a vibrant research area in aerospace science, contributing to our fundamental understanding of this dynamical system and enabling complex and efficient space missions.

## 1.3 Generative AI

Generative AI has become a transformative force across scientific and industrial domains, advancing data generation and pattern synthesis. Models like Generative Adversarial Networks (GANs) [31], Variational Autoencoders (VAEs) [32], and diffusion models [33] have shown remarkable capabilities in producing realistic data, including images, speech, text, and time series. These models learn underlying data distributions to generate new, realistic instances.

Large pre-trained models, such as GPT-4 [34] and Llama2 [35], have revolutionized natural language processing (NLP) with their ability to understand and generate human-like text, enabling applications like conversational agents and automated content creation. In computer vision, models like DALL-E [36] and Stable Diffusion [37] have demonstrated the power of generative AI in creating photorealistic images from textual descriptions, transforming our interaction with visual data.

## 1.4 Variational Autoencoder

A Variational Autoencoder (VAE) is a generative model that combines variational inference and neural networks to generate new data points. VAEs are useful for tasks like image synthesis, anomaly detection, and data compression, as introduced by Kingma and Welling [32].

The VAE consists of an encoder and a decoder, forming an autoencoder architecture with stochastic elements and probabilistic principles. The encoder transforms input data  $x$  into a latent representation  $z$ , producing parameters (mean  $\mu(x)$  and standard deviation  $\sigma(x)$ ) for the probability distribution  $q_\phi(z|x)$ :

$$q_\phi(z|x) = \mathcal{N}(z; \mu(x), \sigma(x))$$

Here,  $\phi$  represents the encoder’s parameters.

VAEs sample  $z$  from the distribution parameterized by  $\mu(x)$  and  $\sigma(x)$ :

$$z = \mu(x) + \sigma(x) \cdot \epsilon$$

where  $\epsilon$  is a random variable from a standard normal distribution  $\mathcal{N}(0, 1)$ .

The decoder reconstructs the input data  $x$  from  $z$ , maximizing the likelihood of the data given the latent variable:

$$p_{\theta}(x|z)$$

Here,  $\theta$  represents the decoder’s parameters.

The VAE is trained to maximize the Evidence Lower Bound (ELBO), which consists of the reconstruction loss and the KL-divergence:

$$\mathcal{L}(\phi, \theta; x) = \mathbb{E}_{q_{\phi}(z|x)}[\log p_{\theta}(x|z)] - \text{KL}(q_{\phi}(z|x)||p(z))$$

where  $\text{KL}(q_{\phi}(z|x)||p(z))$  is the Kullback-Leibler divergence between the encoder’s output distribution and the prior distribution  $p(z)$ .

### 1.5 Research Motivation

This paper is part of the OrbitGPT project [22], which aims to apply generative AI to astrodynamics. The primary goal is to develop a large orbit model (LOM) that generates orbital trajectories with desired features, reducing the need for conventional design or orbit determination algorithms. This approach could revolutionize space mission design by generating new types of orbits, minimizing mission analysis costs, capturing past orbital knowledge, and training specialized models through synthetic data generation.

## 2 Methodology

### 2.1 Dataset

For this study, we use an extensive dataset precomputed by NASA, comprising 44,112 periodic initial conditions in the Earth-Moon CR3BP, classified into 40 families of orbits [38]. The dataset includes various types of orbits, such as those located at the libration points (e.g., planar Lyapunov, Axial, Halo, and Vertical Orbits) and those developing around the entire system (e.g., Butterfly, Dragonfly, planar Distant Retrograde and Prograde Orbits, and Long Period Orbits).


The initial dataset from NASA consists only of the initial conditions of the orbits, along with their periods, stabilities, and Jacobi constants. To obtain the full vector of positions for each orbit, the initial conditions are integrated over time for one period using 100 nodes (N=100) with Matlab’s ODE113 solver

[39], with both absolute and relative tolerances set to  $1 \times 10^{-13}$ .

The dataset is structured in a three-dimensional numpy array with a shape of `data.shape = (num_orbits, 7, num_time_points)`, where `num_orbits` is 44,112, indexing each distinct orbit. The second dimension (7) contains seven scalar values for each orbit at every time point: position components (`posX`, `posY`, `posZ`), velocity components (`velX`, `velY`, `velZ`), and time. The third dimension, `num_time_points`, represents the number of time points at which data for each orbit is recorded, initially set to 100 nodes. The dataset is standardized using the min-max method for each scalar to facilitate efficient data handling and analysis.

### 2.2 Model

The model trained on the orbit data is a Variational Autoencoder (VAE), as depicted in Figure 2. This VAE utilizes a convolutional neural network (CNN) architecture specifically designed for time series data, implemented using the TSGM library [40].



**Figure 2:** Schematic representation of an orbit Variational Autoencoder architecture (VAE).

The encoder consists of five Conv1D layers, each followed by dropout layers to enhance the robustness of the model and prevent overfitting. After passing through the convolutional layers, the data is flattened and fed into two fully connected (dense) layers. The first dense layer has 512 neurons with ReLU activation, reducing the dimensionality of the feature space while retaining essential information. The second dense layer further reduces the feature space to 64 dimensions. This step ensures that the high-dimensional input data is compressed into a more manageable form before reaching the latent space. The output from these dense layers is then used to compute the mean ( $\mu$ ) and log-variance ( $\log(\sigma^2)$ ) of the latent variable distribution. These parameters define the latent space where the encoder maps the input orbit data.

The decoder follows an inverse architecture to the encoder. It first passes through dense layers with 64 and 512 neurons using ReLU activation, then a final dense layer reshapes the data for the five Conv1DTranspose layers. These layers, interspersed with dropout layers for robustness, reconstruct the input sequences, ending with a Conv1DTranspose layer with a sigmoid activation to produce the final output sequences matching the input dimensionality.

### 2.3 Convergence

The generative model described above may not produce fully physical solutions, but the generated trajectories can serve as initial guesses for trajectory optimization algorithms to generate actual physical periodic orbits.

A Multiple-Shooting (MS) algorithm iteratively adjusts a discretised trajectory  $\{\mathbf{X}_i^0\}_{i=1}^N$  to satisfy a set of constraints:

$$\mathbf{F}(\bar{\mathbf{X}}) = \mathbf{0} \tag{3}$$

Here,  $\bar{\mathbf{X}} \in \mathbb{R}^{N \times 6 + (N-1)}$  includes  $N$  states  $\{\mathbf{X}_i^0\}_{i=1}^N$  and  $N - 1$  time intervals  $dt_i = t_{i+1} - t_i, i = 1, \dots, N - 1$ . The constraint vector  $\mathbf{F} \in \mathbb{R}^{N \times 6}$  enforces the respect of both the dynamical equations as in Eq. 4, and the periodicity condition reported in Eq. 5.

$$\mathbf{F}_i = \mathbf{X}_i^f - \mathbf{X}_{i+1}^0, i = 1, \dots, N - 1 \tag{4}$$


$$\mathbf{F}_N = \mathbf{X}_1^0 - \mathbf{X}_N^0 \tag{5}$$

The notation  $\mathbf{X}_i^f$  stands for the state vector obtained through the numerical integration of Eq. 1 from the initial condition  $\mathbf{X}_i^0$  over the time interval  $dt_i$ . The algorithm then uses a Newton-Raphson method for iterative correction:

$$\bar{\mathbf{X}}^j = \bar{\mathbf{X}}^{j-1} - \mathbf{DF}^{-1} \mathbf{F}(\bar{\mathbf{X}}^{j-1}), j = 1, \dots, N_{max} \tag{6}$$

where  $\bar{\mathbf{X}}^j$  is the solution at iteration  $j$ , and  $\mathbf{DF}$  is the Jacobian of the constraint vector, at iteration  $j - 1$ .

Convergence is achieved when the constraint vector meets an assigned tolerance. A good initial guess  $\bar{\mathbf{X}}^0$  is crucial for convergence. To ensure significant guesses, we set a maximum of 20 iterations ( $N_{max} = 20$ ) and used 10% of the states equally spaced in the orbit as prompts. To visualize the effect of the refinement algorithm, we plotted an example of a generated orbit and its refinement in Figure 3.

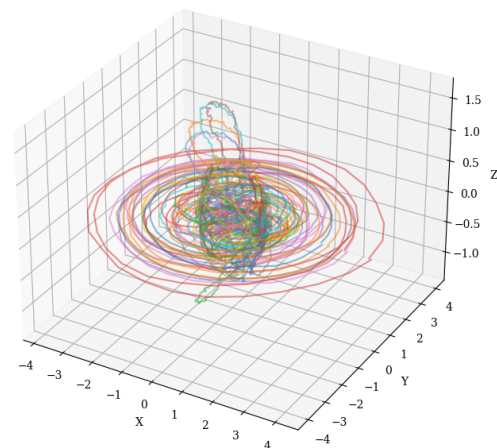


**Figure 3:** Example of a refined orbit. The blue trajectory is the initial output of the model, whereas the orange one is the trajectory refined.

## 3 Results

### 3.1 Generation

For this experiment, the training dataset sequence length, and consequently the timesteps of the generation, were set to 100. A VAE with 2 latent dimensions was trained on the dataset over several epochs. Although we believe that the loss may not be entirely meaningful for performance due to the refinement algorithm used at the end of the pipeline, the final model loss was 8.6, with a reconstruction loss of 2.8 and a KL loss of 5.8. After the training, 100 new orbits were generated by sampling from the latent space (Figure 4).




**Figure 4:** Generation of 100 synthetic orbits.

By examining the results, it is evident that the model has learned and represented orbits that closely resemble actual orbits in essence and shape. However,

the generated orbits exhibit an inherent fuzziness not present in real data, resulting in the orbits appearing as fuzzy lines rather than distinct ones. This fuzziness is further corroborated by the physical error check, which shows an average error of 34 for each node, whereas a physical orbit should exhibit an error in the order of magnitude of  $1 \times 10^{-13}$  (absolute tolerance of the ODE solver). Another important aspect is that the model consistently learned that periodic orbits need to close, meaning that the final value of the position needs to be close to the initial value.

### 3.2 Latent Space

We explored the latent space without the need for dimensionality reduction, as the latent dimensions were already set to two. This exploration is shown in Figure 5. To properly represent each data point in the latent space, we plotted the mean of each distribution, as it is the most representative point. As expected, orbits of the same class are observed grouping together in the latent space, forming filament-like structures. This clustering occurs despite the absence of labels during training, indicating that the model has learned an internal representation of the orbit families and groups orbits with similar characteristics. To further corroborate this insight, we applied Gaussian Mixture Models to cluster the latent space and then applied certain metrics to quantify the extent of the clustering. For a clustering of 40 classes, the performance yielded a Normalized Mutual Information (NMI) score of 0.78 and an accuracy rate of 0.56.




**Figure 5:** Visualization of the latent space from the experiment colored by orbital family labels, with plots showing the average distribution of the Jacobi constant, period, and stability.

Furthermore, we had a list of three features—Jacobi constant, period, and stability—for every orbit in the training dataset. We decided to plot the average of these features for each line in the latent space, both vertically and horizontally, effectively creating a representation of the features’ distribution in the latent space (Figure 5). We found that the period increases incrementally from up to down along the vertical axis. This suggests that the model has encoded the period in an unsupervised manner along this axis, reminiscent of the insight from the seminal VAE paper by Kingma and Welling [32], where moving along one dimension of the latent space caused an incremental change in the generated faces’ orientation.

Regarding the other features, no clear correlation was found. We observed high stability in the middle of the distribution, which might indicate another type of encoding. The Jacobi constant appeared to be equally distributed throughout the entire latent space, showing no specific pattern.

### 3.3 Refinement

After the generation process, the 100 synthetic orbits were refined using the convergence algorithm described in Section 2.3. We found that 46 out of the 100 generated orbits were sufficiently accurate guesses for the refinement algorithm to successfully converge, achieving a ratio of convergence of 0.46 and an average of 10.1 iterations for convergence. The refined orbits, computed from the ones shown in Figure 4, are displayed in Figure 6. Interestingly, all of the final refined orbits were new, meaning none of them were present in the training data.



**Figure 6:** Successful refinement of the generated orbits.

## 4 Discussion

The paper presented what can be considered a first example of Generative Astrodynamics. The results in the paper demonstrate how families of periodic orbits can be encoded into a two-dimensional latent space of features and decoded back into families of approximated periodic orbits. By sampling the latent space we demonstrated the generation of approximated periodic trajectories that converge to physical ones after a local refinement. The results in this paper are evidence of the transformative potential of generative AI in astrodynamics. One particularly promising area is the use of generative models for orbit discovery. We have demonstrated the feasibility of creating a pipeline that generates physical trajectories rather than merely optimizing existing ones. This achievement represents a significant advancement not documented in the literature at the inception of this project. By exploring other AI architectures and expanding our datasets to include other planetary systems, we aim to revolutionize space mission design, minimize mission analysis costs, capture and transfer past orbital knowledge, and enable the training of specialized models through synthetic data generation.

## 5 Acknowledgements

This research was partially funded by the European Space Agency (ESA) project OrbitGPT, grant number 40001435.

## References

1. Furano, G. *et al.* Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities. *IEEE Aerospace and Electronic Systems Magazine* **35**, 44–56 (2020).
2. Hassan, K. *et al.* Application of Artificial Intelligence in Aerospace Engineering and Its Future Directions: A Systematic Quantitative Literature Review. *Archives of Computational Methods in Engineering*, 1–56 (2024).
3. Vasile, M., Rodríguez-Fernández, V., Serra, R., Camacho, D. & Riccardi, A. *Artificial intelligence in support to space traffic management* English. in *68th International Astronautical Congress, IAC 2017* (2017), 3822–3831. ISBN: 9781510855373.
4. Sánchez Fernández-Mellado, L. & Vasile, M. On the use of Machine Learning and Evidence Theory to improve collision risk management. *Acta Astronautica* **181**, 694–706. ISSN: 0094-5765. <https://www.sciencedirect.com/science/article/pii/S0094576520304914> (2021).
5. Furfaro, R., Linares, R. & Reddy, V. *Space Objects Classification and Characterization via Deep Learning and Light Curves: Applications to Space Traffic Management in Space Traffic Management Conference* (2019).
6. Vasile, M. *et al.* Intelligent characterisation of space objects with hyperspectral imaging. *Acta Astronautica* **203**, 510–534. ISSN: 0094-5765. <https://www.sciencedirect.com/science/article/pii/S0094576522006464> (2023).
7. Kisantal, M. *et al.* Satellite pose estimation challenge: Dataset, competition design, and results. *IEEE Transactions on Aerospace and Electronic Systems* (2020).
8. Berquand, A., Moshfeghi, Y. & Riccardi, A. *Space mission design ontology: extraction of domain-specific entities and concepts similarity analysis in AIAA Scitech 2020 Forum* (2020).
9. Rodríguez-Fernández, V. *et al.* *Language Models are Spacecraft Operators* 2022.
10. Lanusse, F. *et al.* *AstroCLIP: Cross-modal contrastive learning for astronomical images and spectra* 2023.
11. Donoso-Oliva, C. *et al.* *ASTROMER: A transformer-based model for creating representations of light curves in a self-supervised manner* 2022.
12. Izzo, D. *et al.* *Learning guidance and control laws with deep learning* 2019.
13. Bonasera, G. *et al.* *Deep learning applied to the design of complex trajectories in multi-body dynamics* 2022.
14. Yang, B., Li, S., Feng, J. & Vasile, M. Fast Solver for J2-Perturbed Lambert Problem Using Deep Neural Network. *Journal of Guidance, Control, and Dynamics* **45**, 875–884. eprint: <https://doi.org/10.2514/1.G006091>. <https://doi.org/10.2514/1.G006091> (2022).
15. Celletti, A., Gales, C., Rodríguez-Fernández, V. & MassimilianoVasile. Classification of regular and chaotic motions in Hamiltonian systems with deep learning. *Sci Rep* **1890** (2022).
16. Stevenson, E., Rodríguez-Fernández, V., Urrutxua, H., Morand, V. & Camacho, D. *Self-supervised Machine Learning Based Approach to Orbit Modelling Applied to Space Traffic Management* 2023. arXiv: 2312.06854.
17. Li, A. Classifying trajectories in the circular restricted three-body problem using support vector machines. *Journal of Computational Physics* (2015).
18. Breen, P. G., Foley, C. N., Boekholt, T. & Zwart, S. P. Newton versus the machine: solving the chaotic three-body problem using deep neural networks. *Monthly Notices of the Royal Astronomical Society* **494**, 2465–2470 (2020).
19. Liao, S., Li, X. & Yang, Y. Three-body problem: From Newton to supercomputer plus machine learning. *New Astronomy* (2022).
20. Wilson, C. & Vasile, M. *Generation and Classification of Critical Points in Uncertain N-Body Problems via Machine Learning in IAC2024, paper IAC-24,C1,9,8* (October 2014, Milan, Italy).
21. Engineering and Physical Sciences Research Council. *Generative Mapping and Control of Stationary Points in Complex Dynamical Systems* 2022. <https://gow.epsrc.ukri.org/NGBOViewGrant.aspx?GrantRef=EP/X018288/1>.
22. European Space Agency. *OrbitGPT: ESA Activity 4000143521* 2024. <https://activities.esa.int/4000143521>.
23. Newton, I. *Philosophiae Naturalis Principia Mathematica* (S. Pepys, 1687).
24. Lagrange, J.-L. *Essai sur le probleme des trois corps. Prix de l'académie royale des Sciences de Paris* **9**, 292 (1772).
25. Ross, S., Koon, W., Lo, M. & Marsden, J. *Dynamical Systems, the Three-Body Problem, and Space Mission Design* ISBN: 978-0-615-24095-4 (Oct. 2022).
26. Poincaré, H. *Les méthodes nouvelles de la mécanique céleste* (Gauthier-Villars et fils, 1892).

27. Zhang, R. A review of periodic orbits in the circular restricted three-body problem. *Journal of Systems Engineering and Electronics* **33**, 612–646 (2022).
28. Farquhar, R. W. The flight of ISEE-3/ICE: origins, mission history, and a legacy. *The Journal of the Astronautical Sciences* **49**, 23–73 (2001).
29. Gardner, J. P. *et al.* The James Webb Space Telescope. *Space Science Reviews* **123**, 485–606 (2006).
30. Thompson, M. R. *et al.* Navigation design of the CAPSTONE mission near NRHO insertion in AAS/AIAA Astrodynamics Specialist Conference (2021).
31. Goodfellow, I. e. a. Generative adversarial nets in *Advances in Neural Information Processing Systems* **27** (2014).
32. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
33. Ho, J., Jain, A. & Abbeel, P. Denoising diffusion probabilistic models in *Advances in Neural Information Processing Systems* **33** (2020), 6840–6851.
34. OpenAI, Achiam, J., Adler, S., Agarwal, S. & et al, L. A. *GPT-4 Technical Report* 2024. arXiv: 2303.08774 [cs.CL].
35. Touvron, H., Martin, L. & et al, K. S. *Llama 2: Open Foundation and Fine-Tuned Chat Models* 2023. arXiv: 2307.09288.
36. Ramesh, A. *et al.* Zero-Shot Text-to-Image Generation. *arXiv preprint arXiv:2102.12092* (2021).
37. Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. *arXiv preprint arXiv:2112.10752* (2022).
38. NASA, Jet Propulsion Laboratory. *Three-Body Periodic Orbits* [https://ssd.jpl.nasa.gov/tools/periodic\\_orbits.html](https://ssd.jpl.nasa.gov/tools/periodic_orbits.html).
39. MathWorks. *Solve nonstiff differential equations — variable order method - MATLAB ode113* 2023. <https://www.mathworks.com/help/matlab/ref/ode113.html>.
40. Nikitin, A., Iannucci, L. & Kaski, S. *TSGM: A Flexible Framework for Generative Modeling of Synthetic Time Series* 2023. arXiv: 2305.11567.