

Toward an Exact Latent Encoding of the Gravitational Three-Body Problem

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Build-std v2.0+4758af4739 | 2026-07-10 12:22 | 99ed1ef5

Overview (non-technical)

The classical gravitational three-body problem has no closed-form trajectory formula comparable to Kepler’s conics. Sundman showed long ago that a convergent series exists away from triple collision, but it is useless in practice because it requires astronomically many terms for modest accuracy.

This working paper develops a **Latent** viewpoint: represent trajectories by Fourier data (and related generating-function objects) that satisfy an algebraic **Galerkin** reformulation of Newton’s equations. Near **exact periodic orbit families** — a measure-zero but structurally important slice of phase space — this yields fast, well-conditioned algebraic solves, and an empirical (fitted, not derived) scaling law linking representation size to a topological word length on the shape sphere. For general trajectories, the manuscript **argues** global coverage by chaining short time windows and using binary-collision regularization, under explicit analytic hypotheses in Section 11 (including a **uniform** quantitative windowing bound that is still **open**).

The paper does **not** claim a new conserved quantity, does not contradict Poincaré’s non-integrability result, does not assert that one explicit closed formula describes all orbit types, and does not claim the single-frequency closed form for generic (perturbed, non-periodic) initial data. “Encoding” is used in the sense of **existence of a finite representation to prescribed accuracy** under the stated hypotheses, not in the sense of a classical quadrature or a single formula for all initial data.

Limitations / open (read first). Three points scope everything below. **(1) Global existence is conditional.** The claim that *every* trajectory admits a finite Latent encoding rests on a **uniform** quantitative windowing bound (Extension B’, §11.2) that is **open** — proved pointwise (Painlevé) but not uniformly. Treat the global statement as *conditional*, not established. **(2) The grade topology law is an empirical fit, not a derived law.** §8 reports measured per-family linear fits $\text{grade} \approx \alpha |w|$ with $\alpha \in [1.55, 7.03]$ and strong within-family correlation. The mechanistic “derivation” of α from an analyticity radius is **not** self-consistent (an earlier draft’s constants were off by two–three orders of magnitude) and has been withdrawn; only the empirical fit is claimed. **(3) The single-frequency closed form is for exact periodic orbits only.** The boxed $x_i(t) = G_i(e^{i\omega t})$ and the atlas lookup represent *periodic* orbits and short-time neighborhoods. A generic perturbation of a periodic orbit is quasi-periodic/chaotic and is **not** captured by one frequency; atlas lookup then gives only a bounded-horizon approximation whose error grows per period.

Abstract

Conditional on the open uniform windowing bound (Extension B', §11.2), we argue that every trajectory of the planar gravitational three-body problem (excluding measure-zero triple collision) admits a finite Latent representation to arbitrary accuracy — an implicit, constructive *encoding* in Fourier space — via the global extension hypotheses of §11 (windowing regularity B', collision regularization C', and Saari-type triple-collision exclusion). The uniform bound is proved *pointwise* (Painlevé) but remains **open** in the uniform form the global claim requires; the global statement is therefore conditional, not established.

For an **exact periodic orbit**, the trajectory is determined by a generating function $G(z; \mathbf{v}_0)$ satisfying the Galerkin equation — Newton's law as an algebraic system in Fourier coordinates — and admits the single-frequency closed form $x_i(t; \mathbf{v}_0) = G_i(e^{i\omega(\mathbf{v}_0)t}; \mathbf{v}_0)$. This closed form is specific to periodic orbits and their short-time neighborhoods; a generic (perturbed, non-periodic) initial condition is quasi-periodic or chaotic and is not represented by one frequency. The **Rational Latent Theorem** frames the approximation chain: each truncation (Fourier, Taylor, Padé) targets an analytic object and is eliminable in the periodic case.

For arbitrary trajectories, two classical extensions are invoked to cover phase space: windowed step-chaining (Painlevé analyticity gives $\rho > 1$ per collision-free segment) and Levi-Civita regularization ($\rho_{\text{reg}} > 1$ through binary collisions). The logical skeleton is **Galerkin + Painlevé + Levi-Civita + Saari $\Rightarrow$ global Latent coverage**, but the passage from *pointwise* to *uniform* $\rho > 1$ across all windows is exactly the open gap of §11.2. Here “complete” means *representation existence* in the Latent sense of §6–11 (conditional on §11), not a single closed-form formula for all initial data.

Empirically, the information content of a periodic orbit tracks its topology: across 21 planar, equal-mass, zero-angular-momentum choreographies from 8 topological families, the Latent grade fits a per-family linear law grade $\approx \alpha |w|$ in braid-word length, with $\alpha \in [1.55, 7.03]$ (family-dependent) and within-family correlation $r = 0.968\text{--}1.000$. This is a measured fit, not a derived law, and is established only on this symmetric slice. A degree-3 / 18-mode truncation encodes the figure-eight orbit in 7,560 real numbers for the full atlas chart (a *neighborhood* of orbits; a single orbit needs $\approx 48\text{--}72$), where Sundman's single-trajectory series [1] would require $\sim 10^{10^6}$ terms.

The measured effective analyticity radius of the figure-eight is $\rho \approx 1.8$ (range 1.6–2.4 depending on the accuracy target; a raw per-mode amplitude fit gives $\rho \approx 1.2$). An earlier draft printed two mutually inconsistent values ($\rho \approx 4.8$ and $\rho \approx 1.004$) for this single coordinate-independent quantity; both were erroneous and have been corrected to the measured value throughout.

The mathematical core is developed below; §9 records a **Lean formalization target map** (module names, theorem hooks). **The cited Lean sources are not shipped in this repository snapshot** — until an artifact path is published, treat formal claims as specifications aligned with the informal proofs here. In that roadmap, Levi-Civita regularization and the *conditional* Painlevé bound (if minimum separation $\delta > 0$ then $\tau_{\text{sing}} > 0$) are the natural first formalization targets; a **uniform** lower bound on separation (hence on τ_{sing}) across an entire energy shell (Extension B', §11.2) and a full Mathlib proof of Saari-type triple-collision estimates remain **analytic inputs** for a end-to-end verifier, not “solved by citation” inside this manuscript.

1. Introduction

1.1 The Problem

The gravitational three-body problem asks: given three masses m_1, m_2, m_3 with initial positions $\mathbf{r}_i(0)$ and velocities $\dot{\mathbf{r}}_i(0)$, find $\mathbf{r}_i(t)$ for all t .

Newton’s equations are:

$$\ddot{\mathbf{r}}_i = - \sum_{j \neq i} m_j \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|^3}, \quad i = 1, 2, 3$$

For two bodies ($m_3 = 0$), Kepler (1609) gave the solution: conic sections with $r(\theta) = a(1 - e^2)/(1 + e \cos \theta)$. For three bodies, no comparably explicit formula has been found in 300 years.

1.2 What Is Known

Poincaré (1890): No additional analytic first integral $I(x, \dot{x}) = \text{const}$ exists beyond the 10 classical ones [2] (extending Bruns’ earlier algebraic result [22]). The problem cannot be reduced to quadratures — there is no analog of the vis-viva equation. Crucially, this does *not* forbid an exact implicit formula for the trajectory given initial conditions, nor a convergent series solution. It forbids a specific algebraic structure (a new conserved quantity), not every form of “solution.”

Sundman (1912): A globally convergent power series solution exists (away from triple collision) [1]. However, the convergence is impractically slow: $\sim 10^{10^6}$ terms for modest accuracy.

Picard–Lindelöf: The trajectory $\mathbf{r}(t; \mathbf{v}_0)$ exists and is unique on maximal intervals of definition while the vector field is Lipschitz (away from collisions); smooth dependence on parameters holds in the usual ODE sense, and real-analytic dependence follows where the vector field is real-analytic (standard ODE theory; see e.g. [12, 13] for classical analytic ODE background).

1.3 What This Paper Shows

We argue that **every trajectory of the planar three-body problem** (excluding measure-zero triple collision) admits a finite Latent representation to arbitrary accuracy **under the hypotheses of §11**. The construction has two layers.

Layer 1: The local machine (new). Near periodic orbit families, the trajectory is determined by a generating function $G(z; \mathbf{v}_0)$ satisfying an algebraic Galerkin equation, exploiting analyticity at four levels:

Level	What is analytic	Consequence	Computational form
Temporal	$x(t)$ in t	Fourier coefficients decay as ρ^{-k}	Fourier series
Parametric	$\text{IC} \rightarrow \Lambda$ in δv	Taylor coefficients decay as $\delta^{ \alpha }$	Polynomial map
Algebraic	EOM in Λ -space	Newton–Raphson converges quadratically	Algebraic system

Level	What is analytic	Consequence	Computational form
Continuous	$G(z) = \sum \Lambda_k z^k$	Singularity structure determines all Λ_k	Rational/algebraic closed form

Within the convergence basin of a periodic orbit family, this requires less iteration than Kepler’s solution of the two-body problem.

Layer 2: Global extension (new + classical). Two mechanisms cover all trajectories outside periodic basins:

Trajectory type	Mechanism	Why $\rho > 1$
Chaotic	Windowed step-chaining	Painlevé (1897): analytic away from collision $\Rightarrow \tau_{\text{sing}} > 0$ per window
Near-collision	Levi-Civita regularization	Smooth flow in $u = \sqrt{r}$ coordinates $\Rightarrow \rho_{\text{reg}} > 1$
Scattering	Decomposition	Bounded phase (windowing + regularization) + Kepler asymptotics
Triple collision	Excluded	Measure zero (Saari, 1977)

The combination is the global result (existence of a Latent cover): **Galerkin + Painlevé + Levi-Civita + Saari $\Rightarrow$ global Latent representation** (§11). The Galerkin equation is Newton’s law in Fourier coordinates; it is a reformulation, not a new equation. The novelty is structural: the reformulation yields exponential convergence, four-level analyticity exploitation, and a generating-function representation — combined with classical regularization theorems to cover the full phase space.

Layer 3: Complexity bound (new, empirical). The topological classification of periodic orbits by the free group F_2 (braid word w) gives a natural measure of orbit complexity: the word length $|w|$. Numerical experiments on 21 orbits from 8 topological families show that the Latent grade — the number of Fourier modes needed for ε -accuracy — scales linearly in $|w|$ within families ($r > 0.96$). This gives a practical characterization of “how many numbers” any given orbit needs, controlled by a computable topological invariant (§8).

Notation conventions. $w \in F_2$ denotes a braid word (topological orbit type); $u = \sqrt{r}$ is the Levi-Civita regularization coordinate (§5). $G(z; \mathbf{v}_0)$ is the generating function (§4); the gravitational constant G appears only inside Newton’s equations and is set to 1 from §2 onward. ρ is the analyticity radius of the generating function; ρ_{reg} is its regularized counterpart near collisions.

2. The Latent of a Three-Body Orbit

2.1 Setup

We work with the planar three-body problem: 3 bodies in $\mathbb{R}^2$, giving 6 position coordinates $x = (x_1, y_1, x_2, y_2, x_3, y_3)$. The center-of-mass constraint $\sum m_i \mathbf{v}_i = 0$ removes 2 of 6 velocity degrees of freedom, leaving a 4-dimensional effective initial-condition space.

2.2 Definition

The **Latent** of a periodic orbit is its Fourier decomposition:

$$\mathbf{x}(t) = \sum_{k=0}^{\infty} [\Lambda_k \cos(k\omega t) + \Lambda'_k \sin(k\omega t)]$$

where $\Lambda_k, \Lambda'_k \in \mathbb{R}^6$ are the Fourier coefficient vectors and $\omega = 2\pi/T$ is the orbital frequency. By the Latent Theorem, if the orbit is real-analytic in time with analyticity parameter $\rho > 1$, the coefficients decay as $|\Lambda_k| \sim C\rho^{-k}$, and the orbit is determined to accuracy ε by $N = \Theta(\log(1/\varepsilon)/\log \rho)$ modes.

For the Chenciner–Montgomery figure-8 orbit with equal masses, the measured effective analyticity radius is $\rho \approx 1.8$ (range 1.6–2.4 depending on the target accuracy; obtained from the truncation grade $N(\varepsilon)$ via $N = \Theta(\log(1/\varepsilon)/\log \rho)$ — a raw per-mode amplitude fit gives a smaller $\rho \approx 1.2$). The rank is 4 (by SVD), and the truncated Latent at 18 modes is 72 real numbers (18 modes $\times$ 4 coordinates in Jacobi representation).

*(Correction note. An earlier draft of this manuscript reported $\rho \approx 4.8$ here and $\rho \approx 1.004$ in §8 for the *same* coordinate-independent quantity. Both were erroneous: $\rho \approx 4.8$ overstates the decay (it predicts ~ 6 modes at 10^{-4} where ≈ 13 are needed), and $\rho \approx 1.004$ understates it by two orders of magnitude. The measured effective value $\rho \approx 1.8$ is used consistently throughout.)**

2.3 The Parametric Map

The map from initial conditions to Latent is smooth:

$$F : \mathbb{R}^4 \rightarrow \mathcal{H}^{\otimes 2}, \quad \mathbf{z} \mapsto \Lambda(\mathbf{z})$$

where $\mathbf{z} = N \cdot (\mathbf{v} - \mathbf{v}_{\text{ref}}) \in \mathbb{R}^4$ is the CM-projected initial velocity perturbation. This map exists by Picard–Lindelöf and is real-analytic (because the gravitational force is real-analytic away from collisions).

3. The Solution Formula

3.1 Stage 1 — The Reference Latent (Galerkin Equation)

Substituting the Fourier ansatz $x_i(t) = \sum_k \Lambda_k^{(i)} \phi_k(t)$ into Newton’s equations and projecting onto each mode gives the **Galerkin equation**:

$$R_k(\Lambda, \omega) = -k^2 \omega^2 \Lambda_k^{(i)} + \mathcal{F}_k \left[\sum_{j \neq i} m_j \frac{x_j(t) - x_i(t)}{|x_j(t) - x_i(t)|^3} \right] = 0 \quad \forall k, i$$

Here $\mathcal{F}_k[\cdot]$ denotes the k -th Fourier-mode coefficient of the bracketed periodic function reconstructed from (Λ, ω) , and the sign matches $\ddot{x}_i = \sum_{j \neq i} m_j (x_j - x_i)/|x_j - x_i|^3$ (equivalently the Newton equations in §1.1).

This is a system of algebraic equations in the Latent coordinates Λ and the frequency ω . The right-hand side is an analytic function of Λ (because the gravitational force is analytic away from collisions). Newton–Raphson converges **quadratically**.

Computational validation (figure-8 orbit, $N = 18$ modes, 211 unknowns):

- Iteration 0: $\|R\| = 2.9 \times 10^{-4}$
- Iteration 1: $\|R\| = 3.0 \times 10^{-9}$ (quadratic)
- Iteration 2: $\|R\| = 1.0 \times 10^{-14}$ (machine precision)

No ODE was integrated. The reference orbit was found by solving an algebraic system.

3.2 Stage 2 — IC Dependence (Variational Equations)

The dependence of the Latent on initial conditions is given by the **variational equation** along the reference orbit:

$$\dot{J}(t) = \begin{pmatrix} 0 & I_6 \\ G(x_{\text{ref}}(t)) & 0 \end{pmatrix} J(t), \quad J(0) = \begin{pmatrix} 0 \\ I_6 \end{pmatrix}$$

where $G_{ab}(x) = \partial^2 \Phi / \partial x_a \partial x_b$ is the gravitational tidal matrix. The Fourier projection gives the IC $\rightarrow \Lambda$ Jacobian analytically:

$$\frac{\partial \Lambda_k^{(i)}}{\partial v_{0,j}} = \int_0^T J_{ij}^{(x)}(t) \phi_k(t) dt$$

Validation: Agreement with finite-difference Jacobian: 6.1×10^{-10} relative error (machine precision). SVD singular values (29.1, 15.7, 3.65, 2.28) — rank exactly 4. Higher-order variational equations give $\partial^n \Lambda / \partial \mathbf{v}_0^n$ to every order n .

The full IC dependence is:

$$\Lambda_k^{(i)}(\mathbf{v}_0) = \Lambda_{k,\text{ref}}^{(i)} + \sum_j \left(\int_0^T J_{ij} \phi_k dt \right) \Delta v_j + \sum_{j,l} \left(\int_0^T H_{ijl} \phi_k dt \right) \Delta v_j \Delta v_l + \dots$$

Each coefficient is a definite integral of a variational solution against a Fourier basis function — derived from Newton’s law, not fitted from samples.

3.3 Stage 3 — Trajectory (Fourier Synthesis)

Given the Latent $\Lambda(\mathbf{v}_0)$, the trajectory is:

$$x_i(t) = \sum_k \Lambda_k^{(i)} \cos(k\omega t) + \Lambda_k^{\prime(i)} \sin(k\omega t)$$

This is Fourier synthesis — explicit, non-iterative, instantaneous.

3.4 The Complete Formula

Combining all three stages:

$$x_i(t; \mathbf{v}_0, \mathbf{m}) = \sum_{k=0}^{\infty} \Lambda_k^{(i)}(\mathbf{v}_0, \mathbf{m}) \cdot \phi_k(t)$$

where Λ_k is determined by the Galerkin equation (Stage 1) with variational IC dependence (Stage 2). This formula uses a single frequency ω and is therefore exact for **periodic** orbits and their short-time neighborhoods; for a perturbed IC it is a bounded-horizon approximation (the perturbed orbit is generically quasi-periodic — see §6.2, §8.6.3). Within that scope it is more explicit than Kepler's in one respect: Kepler requires solving $M = E - e \sin E$ iteratively at **every** t ; the Latent formula evaluates a trigonometric sum — no iteration at any step after Stage 1.

	Kepler 2-body	Latent 3-body
Orbit shape	Explicit: $r(\theta)$ formula	Explicit: Fourier series
Position at time t	Iterative: Kepler's equation per t	Explicit: Fourier synthesis
IC dependence	Explicit: a, e from energy/momentum	Explicit: Taylor polynomial
Reference orbit	Not needed (universal conic)	Once: Galerkin algebraic solve

4. The Continuous Latent

4.1 The Generating Function

The Fourier series is discrete — a sum over integer mode indices. Its continuous counterpart is the **generating function**:

$$G(z) = \sum_{k=0}^{\infty} \Lambda_k z^k$$

Since $|\Lambda_k| \sim C\rho^{-k}$ with the measured $\rho \approx 1.8$ (§2.2), $G(z)$ is analytic in $|z| < \rho$. The trajectory is $x(t) = G(e^{i\omega t})$.

4.2 Singularity Structure

The singularities of $G(z)$ at $|z| = \rho$ correspond to **complex-time gravitational collisions**: times $t^* = t_{\text{real}} + i\tau$ where two bodies meet in the analytic continuation. The distance τ determines $\rho = e^{2\pi\tau/T}$.

If the singularities are poles, G is rational and the trajectory becomes a closed-form formula:

$$x_i(t) = \frac{P_i(e^{i\omega t})}{Q_i(e^{i\omega t})}$$

For the gravitational $|\mathbf{r}|^{-3}$ force, the singularities are generically branch points. The Padé approximant represents these as pole clusters, converging via Nuttall–Pommerenke.

4.3 Hierarchy of Representations

Representation	Object	Size (figure-8)	Sum/integral?
Fourier series (exact)	$\{\Lambda_k\}_{k=0}^{\infty}$	∞	Infinite discrete sum
Truncated Fourier	$\{\Lambda_k\}_{k=0}^N$	72 numbers	Finite sum (approximation)
Generating function	$G(z)$ analytic	1 function	Evaluation at $z = e^{i\omega t}$
Pole-residue (if meromorphic)	$\{(z_p, R_p)\}_{p=1}^M$	$\sim 2M$ complex numbers	Finite closed form

5. The Rational Latent Theorem

Theorem 8. *Let $G(z) = \sum_{k=0}^{\infty} c_k z^k$ be analytic in $|z| < \rho$ with $\rho > 1$. Then:*

(i) (*Latent Theorem bound.*) *The monomial-basis Latent of G has size $N_{\text{mono}} = \Theta(\log(1/\varepsilon)/\log \rho)$ for accuracy ε .*

(ii) (*Pole absorption.*) *If G extends meromorphically to $|z| < R$ with $R > \rho$ and has exactly M poles in $\rho \leq |z| < R$, then the pole-basis representation needs $N_{\text{pole}} = M^* + \Theta(\log(1/\varepsilon)/\log R) \ll N_{\text{mono}}$ terms.*

(iii) (*Exact termination.*) *If G is globally meromorphic (rational), the Padé $[L/M]$ at sufficient order recovers G exactly.*

(iv) (*Optimal basis.*) *The pole basis is the constructive realization of the optimal-basis bound (Theorem 3 in [6]): the Padé finds the poles, constructing the basis that maximizes ρ .*

Proof sketch. (i) is the Latent Theorem. (ii): subtract principal parts to remove poles, enlarging the analyticity disk from ρ to R . (iii): de Montessus de Ballore (1902). (iv): the monomial basis has effective parameter ρ ; after pole absorption, it rises to $R > \rho$. $\square$

For the 3-body problem: the figure-eight $G(z)$ has measured $\rho \approx 1.8$ (§2.2) with branch-point singularities. Padé with ~ 10 poles (~ 40 real numbers) captures G better than 72 Fourier coefficients, because pole clusters absorb the branch points and raise the effective analyticity parameter.

6. The Exact Solution for Periodic Orbits

This section constructs the exact solution for an **exact periodic orbit**. The single-frequency closed form and “eliminable approximation” chain below are specific to periodic orbits and their short-time neighborhoods; generic (non-periodic) initial data are handled — conditionally — by the global extension of §11, and atlas lookup on perturbed ICs is only a bounded-horizon approximation (§6.2, §8.6.3).

6.1 Every Approximation Is Eliminable (Periodic Case)

At each step of the solution chain, we used a finite approximation. But each targets a smooth object with an exact Latent:

Step	Approximation used	Exact object	Why exact
Reference orbit	Fourier truncation ($N = 18$)	$G(z)$	Galerkin equation $R(\Lambda, \omega) = 0$
IC dependence	Taylor truncation (degree 3)	Full variational series	Variational ODEs to all orders
Trajectory	Padé [L/M]	$G(z)$ itself	Padé converges to G

Replace each approximation with its exact counterpart:

1. **Galerkin equation** $R_k(\Lambda, \omega) = 0$ for all k — Newton’s law in Latent coordinates. Defines $G(z)$ implicitly as an **infinite** coupled algebraic system (countably many mode equations); exact in the sense of no truncation of the *statement*, not a single finite matrix equation (§10.5).
2. **Variational equations** $\dot{J} = A(t)J$ — define IC dependence to all orders. Linear ODEs, exact.
3. **Fourier synthesis** $x(t) = G(e^{i\omega t})$ — evaluates G directly. No sum computed.

6.2 The Exact Solution (Periodic Orbits)

Scope of this box. The single-frequency closed form below holds for an **exact periodic orbit** — one $\mathbf{v}_0$ on a periodic family — and its short-time neighborhood. It is *not* a formula for generic initial data: a generic perturbation of a periodic orbit is quasi-periodic (a KAM torus) or chaotic, hence not T -periodic and not represented by a single frequency ω . Such orbits require either a multi-frequency Latent (a finite sum $\sum_m G^{(m)}(e^{i\omega_m t})$, cf. §8.6, §11.5) or the windowed construction of §11. See §7.3/§8.6.3 for the (bounded-horizon) accuracy of atlas lookup on perturbed ICs.

For an exact periodic orbit, after substitution zero approximations remain:

$$x_i(t; \mathbf{v}_0) = G_i(e^{i\omega(\mathbf{v}_0)t}; \mathbf{v}_0) \quad (\mathbf{v}_0 \text{ periodic})$$

where $G_i(z; \mathbf{v}_0)$ is the unique analytic function satisfying the Galerkin equation with variational IC dependence.

- **Exact (on the periodic family):** no truncation, no iteration at evaluation time, no approximation.
- **Encoding:** the defining object is either the full Galerkin system $\{R_k = 0\}_{k=0}^{\infty}$ (countably infinite constraints) or an equivalent functional equation for $G(z)$ — not a single finite list of scalar equations (§10.5).
- **Constructive:** on any finite truncation, Newton–Raphson converges quadratically under standard nondegeneracy assumptions; variational equations are linear ODEs.
- **More explicit than Kepler:** Kepler’s equation must be solved at every t ; the Latent evaluates $G(e^{i\omega t})$ directly.

6.3 Hierarchy of Exactness

Object	Exact?	Finite?	Form
Newton’s ODE	Yes	Yes (the equation)	Requires integration
Galerkin equation $R(\Lambda, \omega) = 0$	Yes	Infinite (countable); finite per truncation	Algebraic mode-by-mode — no time integration
Generating function $G(z)$	Yes	1 function	Defined by Galerkin equation
Padé $[L/M]$ rational	No (converges)	$L + M + 1$ coefficients	Rational approximation
Fourier truncation N modes	No (converges)	$2N + 1$ coefficients	Trigonometric sum
Taylor truncation degree p	No (converges)	$\binom{p+4}{4}$ coefficients	Polynomial evaluation

The Galerkin equation and the generating function are exact; everything below them is a convergent evaluation scheme. The relationship is the same as between $x^2 = 2$ and $\sqrt{2} \approx 1.41421 \dots$ — the equation IS the exact solution; the decimal expansion is one way to evaluate it.

6.4 Comparison with Kepler

Criterion	Kepler (2-body)	Latent (3-body)
Orbit shape	$r(\theta) = a(1 - e^2)/(1 + e \cos \theta)$ — explicit	$G(z)$ satisfying $R(\Lambda, \omega) = 0$ — implicit
Position at time t	Solve $M = E - e \sin E$ iteratively	Evaluate $G(e^{i\omega t})$ directly
IC dependence	a, e from energy/momentum — explicit	Variational series — explicit
Convergence	Kepler iteration: linear	Newton–Raphson on Galerkin: quadratic
Poincaré obstruction	None	No closed-form integral — but implicit formula exists

The Kepler problem is considered “solved” even though computing position at time t requires

iteratively solving a transcendental equation. The 3-body Latent solution is solved in the same sense — and requires **less** iteration.

7. Numerical Validation

7.1 Finite Truncations

When the exact formula is evaluated at finite truncation (degree 3, 18 modes), we obtain 7,560 real numbers with the following performance:

δ (IC scale)	Degree p	Object size	Test error	Speedup
10^{-4}	3	7,560	1.44e-3	28×
10^{-3}	3	7,560	1.99e-3	28×
5×10^{-3}	3	7,560	4.91e-3	28×
10^{-2}	3	7,560	5.39e-3	28×

7.2 Convergence Rates

Both truncation errors decay exponentially:

- Fourier: $O(\rho^{-N})$ with the measured $\rho \approx 1.8$ (§2.2)
- Taylor: $O(\delta^{p+1})$ within convergence radius

For comparison: Sundman’s single-trajectory series requires $\sim 10^{10^6}$ terms for 10^{-3} accuracy. The Latent atlas chart for the figure-eight (a *neighborhood* of orbits) uses 7.6×10^3 numbers; a single figure-eight orbit needs ≈ 48 –72 Fourier coefficients at that accuracy.

7.3 Global Atlas

The local solution extends to multiple orbit families via the **Global Latent Atlas**:

Family	Period T	Object size	Local error
Figure-8	6.3259	7,560	1.1%
Lagrange equilateral	8.2691	7,560	1.1%
Broucke A2	5.2247	7,560	1.1%
Hierarchical triple	20.6487	7,560	29% [†]
Total atlas		30,240	

[†] The hierarchical triple has a much longer period ($T \approx 20.6$ vs. ~ 5 –8 for the other families), meaning 18 Fourier modes under-resolve it. Increasing to 60 modes reduces the error to $< 1\%$ — consistent with the grade–topology prediction of §8.

30,240 real numbers encode the global solution across 4 orbit families at the baseline truncation. Adding families from the literature catalogs (§8.8) extends coverage; increasing mode count handles longer-period orbits.

8. Latent Grade and Topological Complexity

The preceding sections established — under the same §11 hypotheses — that every three-body orbit admits a finite Latent representation. This section answers the quantitative question: **how large is the Latent?** The answer connects the Latent framework to the topological classification of periodic orbits — and provides the precise sense in which the three-body problem is “solved.”

8.1 The Topological Classification of Periodic Orbits

The **shape sphere** (Montgomery [15]) maps the triangle formed by three bodies to a point on S^2 . Removing the three collision points (where two bodies coincide) gives a punctured sphere $S^2 \setminus \{p_{12}, p_{13}, p_{23}\}$. Its fundamental group is the **free group on two generators**:

$$\pi_1(S^2 \setminus \{3 \text{ points}\}) \cong F_2 = \langle a, b \rangle$$

Each periodic orbit traces a closed curve on this punctured sphere, defining a conjugacy class in F_2 — a **word** $w = w(a, b, a^{-1}, b^{-1})$. The word length $|w|$ counts the number of **syzygies** (collinear passages) per period: moments when all three bodies align. This classification is complete: Montgomery and Moeckel [16] proved that every free homotopy class is realized by a periodic solution.

The number of topological orbit types is **countably infinite** — fundamentally different from central configurations, which are finite. The free group F_2 has infinitely many distinct conjugacy classes: $a, ab, aab, abba, \dots$ Each defines a different topological family of periodic orbits.

A remarkable empirical law (Dmitrašinović and Šuvakov [17]) connects topology to dynamics: the period T of a periodic orbit scales linearly with the word length:

$$T \approx T_0 \cdot |w|$$

where T_0 is a family-dependent fundamental period unit. This “quasi-Kepler’s third law” has been verified across hundreds of orbit families [18].

8.2 The Grade of a Latent Representation

Definition. The **Latent grade** of a periodic orbit at accuracy ε is the minimum number of Fourier modes $N(\varepsilon)$ such that the truncated Fourier sum (in Jacobi coordinates) reconstructs the orbit with max error $< \varepsilon$:

$$\text{grade}(\varepsilon) = \min \left\{ N : \max_t \left\| \mathbf{x}(t) - \sum_{k=0}^N [\Lambda_k \cos(k\omega t) + \Lambda'_k \sin(k\omega t)] \right\| < \varepsilon \right\}$$

The grade measures the **information content** of the orbit: how many real numbers are needed to encode it to a given accuracy. For the figure-eight orbit (§2.2), $\text{grade}(10^{-3}) = 12 \text{ modes} \times 4 \text{ Jacobi coordinates} = 48 \text{ real numbers}$. For a simpler orbit (Lagrange equilateral), the grade is lower. For a more complex orbit (moth, yarn), it is higher.

The question is: **what determines the grade?**

8.3 Experimental Design

We tested the relationship between Latent grade and topological word length across 21 periodic orbits from 8 topological families, using the Šuvakov–Dmitrašinović catalog [19] (Table 1). The orbits span word lengths from $|w| = 6$ (figure-eight) to $|w| = 92$ (yin-yang family, high index).

Method. For each orbit:

1. Integrate one period with DOP853 ($\text{rtol} = 10^{-13}$, $\text{atol} = 10^{-14}$, 16,384 time points).
2. Count syzygies: sign changes in the signed triangle area $S(t) = (\mathbf{r}_2 - \mathbf{r}_1) \times (\mathbf{r}_3 - \mathbf{r}_1)$ give $|w|$ directly.
3. Transform to Jacobi coordinates $(\rho, \mathbf{R})$ and extract Fourier coefficients via FFT (up to 2000 modes).
4. Compute the grade $N(\varepsilon)$ at $\varepsilon \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ from the cumulative tail energy bound.
5. Estimate the analyticity radius ρ from the Fourier amplitude decay rate.

All initial conditions use the standardized convention of [19]: bodies at $(-1, 0)$, $(0, 0)$, $(1, 0)$ with equal masses $m_1 = m_2 = m_3 = 1$ and zero angular momentum. Velocities are parameterized by (p_1, p_2) .

8.4 Results

All experiments use $n_{\max} = 2000$ Fourier modes to avoid ceiling effects.

Table 2. Per-family linear fit $\text{grade}(\varepsilon) = \alpha \cdot |w| + \beta$ at $\varepsilon = 10^{-2}$ (the cleanest regime, where no orbit hits the mode ceiling).

Family	n orbits	$ w $ range	Slope α	Correlation r
I (butterfly)	4	14–74	5.51	0.990
II (dragonfly)	3	24–66	4.78	0.982
III (yin-yang)	3	20–92	5.64	0.997
IVa (moth)	3	16–48	7.03	0.998
IVb (butterfly III)	2	24–40	6.00	1.000
V (figure-eight)	4	6–66	1.55	0.968
VI (yarn)	2	16–44	6.68	1.000

Every family shows strong to near-perfect linear correlation ($r \geq 0.968$). The within-family slopes α range from 1.55 (figure-eight) to 7.03 (moth). The slope is a **measured, family-dependent constant of proportionality**; we do not derive it from an analyticity radius (see §8.5 for why an earlier attempted derivation was inconsistent and has been withdrawn). Qualitatively, families whose orbits avoid close binary encounters (larger effective analyticity radius) need fewer modes per syzygy and so have smaller α , but the present data do not pin the per-family ρ precisely enough to convert this into a quantitative prediction of α .

Table 3. Global regression across all 21 orbits and per-family at multiple accuracy levels.

Scope	ε	Slope α	Intercept	Correlation r
Global (21 orbits)	10^{-2}	5.17	−2.8	0.856
Global	10^{-3}	18.55	128.5	0.725

Scope	ε	Slope α	Intercept	Correlation r
Seq. I (butterfly)	10^{-3}	20.11	201.1	0.995
Seq. IVa (moth)	10^{-3}	37.32	−398.9	0.996

The global correlation ($r = 0.856$) is weaker than within-family correlations because different families have different analyticity radii. Within each family, the linearity is near-perfect.

The quasi-Kepler law is confirmed: $T = 0.78|w| + 4.2$ ($r = 0.836$) across all 21 orbits, consistent with Dmitrašinović and Šuvakov [17].

8.5 Interpretation: Information Content of Three-Body Orbits

The result $\text{grade}(\varepsilon) = O(|w|)$ has a precise physical interpretation. Each syzygy (collinear passage) corresponds to a close interaction between two bodies. During a close approach:

1. The motion speeds up (Kepler’s second law — conservation of angular momentum).
2. The trajectory develops **high-frequency content** (fast oscillation near the passage).
3. This creates a **near-singularity** in the complex- t plane, reducing the local analyticity radius.

Each such event contributes $O(1)$ additional Fourier modes to the orbit’s spectral content. The total mode count is therefore $O(|w|)$: the information content of a three-body orbit is, empirically, **proportional to its topological complexity**. This is the content of the measured law $\text{grade}(\varepsilon) \approx \alpha(\varepsilon)|w|$ (Table 2), with a family-dependent slope α .

On the constant of proportionality (withdrawn derivation). A mechanistic derivation would like to write, for an orbit with per-syzygy analyticity radius ρ_0 ,

$$\text{grade}(\varepsilon) \stackrel{?}{\approx} \frac{|w| \cdot \log(1/\varepsilon)}{\log \rho_0}, \quad \text{i.e. } \alpha \stackrel{?}{=} \frac{\log(1/\varepsilon)}{\log \rho_0}.$$

This is not self-consistent with the measured constants and is withdrawn. Reproducing the measured slopes $\alpha \in [1.55, 7.03]$ at $\varepsilon = 10^{-2}$ through $\alpha = \log(1/\varepsilon)/\log \rho_0$ would require $\rho_0 \in [1.9, 19.5]$; conversely, the small per-time-point decay rates ($\rho_0 \approx 1.002\text{--}1.005$) that an earlier draft reported here would predict $\alpha \approx 900\text{--}2300$, i.e. $10^2\text{--}10^3\times$ larger than observed. In other words, the small ρ_0 values were a *misabeled* quantity (a per-sample or per-time-point decay rate, not the per-mode analyticity radius ρ of §2.2/§4.1), and the “derivation” bridging them to α was arithmetically impossible. We therefore report the slope α **only as an empirical, family-dependent fit** (Table 2) and make no claim to derive it from a per-syzygy radius. The measured effective per-mode radius of the figure-eight is $\rho \approx 1.8$ (§2.2), consistent with its low grade (≈ 5 modes at 10^{-2} , ≈ 13 at 10^{-4}) and with the small slope $\alpha = 1.55$; a first-principles formula for α from orbit geometry is left open.

8.6 What “Solving” Means for a Chaotic System

8.6.1 Why Newton Could Give a Single Formula — and We Cannot

Newton’s formula for the two-body orbit,

$$r(\theta) = \frac{a(1 - e^2)}{1 + e \cos \theta},$$

describes a single geometric shape: a conic section. Ellipse, parabola, and hyperbola are all parametrizations of this one formula. Every two-body orbit, regardless of initial conditions, is topologically the same — a curve of genus zero in the orbital plane. The two-body problem is *topologically trivial*: there is only one orbit shape.

The three-body problem is fundamentally different because it is *topologically rich*. The shape sphere $S^2 \setminus \{\Delta_1, \Delta_2, \Delta_3\}$ has fundamental group F_2 , the free group on two generators, which has countably infinitely many conjugacy classes. Each class corresponds to a qualitatively distinct orbit shape: the figure-eight traces a different curve from the butterfly, which traces a different curve from the moth, the dragonfly, the yarn, the yin-yang — and these are just the first few words in an infinite alphabet.

	Two-body	Three-body
Distinct orbit shapes	1 (conic section)	Countably infinite (F_2 free group)
Topological types	None — all orbits are topologically equivalent	Figure-eight, butterfly, moth, dragonfly, yarn, yin-yang, ...
Representation	1 formula, 5 numbers	1 encoding system, $O(\ w\)$ numbers per type

A single closed-form formula *cannot* describe all three-body orbits because the orbits are not parametric deformations of one shape — they are genuinely, qualitatively distinct curves. This is not a failure of ingenuity; it is a topological obstruction. The space of three-body orbits is infinitely richer than the space of two-body orbits, and any complete description must reflect this richness.

8.6.2 The Latent as Universal Encoding

The Latent framework resolves this obstruction by providing a *universal encoding* rather than a universal formula. The Fourier coefficient set $\{\Lambda_k\}$ can represent any periodic orbit of any topological type. What changes between types is not the encoding scheme but the *size* of the encoding — how many coefficients are needed:

	Two-body	Three-body
What describes the orbit	5 orbital elements	$N(\varepsilon)$ Fourier coefficients (the Latent)
How many numbers	5 (exact, universal)	$O(w)$ (to accuracy ε , per family)
What determines the count	Nothing — always 5	Topological word length $ w $
Information content	Constant	Proportional to topological complexity

Newton solved the two-body problem by showing that every orbit is a conic, described by 5 parameters. Our result shows that every three-body orbit has a family-specific finite description, and its size is controlled by topology:

Every periodic three-body orbit has a finite Latent representation that reconstructs the orbit to any desired accuracy $\varepsilon > 0$ (the constant depends logarithmically on $1/\varepsilon$). Empirically — measured on planar, equal-mass, zero-angular-momentum choreographies — the representation size fits grade $\approx \alpha |w|$ in the braid-word length, with a family-dependent slope α .

Scope. The *existence* of a finite Latent for a periodic orbit is general (it is Fourier truncation of an analytic periodic function). The *linear-in- $|w|$ size law* is an empirical fit established only on the symmetric slice above; we do not claim it, or the values of α , for unequal masses or nonzero angular momentum, where the S_3 symmetry that keeps these choreographies well-conditioned is broken (see §8.7). In this scoped sense the framework extends “solved” from integrable to chaotic systems: a finite representation exists, and its size is *empirically* controlled by a computable topological invariant.

8.6.3 From Initial Conditions to Orbits Without Simulation

The reduced IC-space. The full planar three-body problem has 12 phase-space dimensions (3 bodies $\times$ 2D position $\times$ 2 for position and velocity). Symmetries and conservation laws reduce this to an effective 4-dimensional search space for periodic orbits:

Reduction	Dimensions removed	Remaining
Center of mass at origin	−2 (position)	10
Total momentum = 0	−2 (velocity)	8
Energy fixed (scale)	−1	7
Angular momentum fixed	−1	6
Rotational symmetry SO(2)	−1	5
Phase along periodic orbit	−1	4

This is the 4-dimensional IC-space $\mathcal{I} \subset \mathbb{R}^4$ in which Šuvakov–Dmitrašinović and Li–Liao search for periodic orbits, and in which the Latent atlases live.

The two-body baseline. For two bodies, the map from IC to orbit is global and explicit: $(E, L) \mapsto (a, e)$ via conservation laws. No integration is needed — a single formula covers the entire IC-space.

The Latent atlas construction. For three bodies, the variational equations (§3.2) provide an analogous — though local — map. The construction proceeds in three stages:

Stage 1 (Reference orbit catalog). Take a catalog of periodic orbits — for instance, the Li–Liao catalog of 10,059 orbits [18] or the Šuvakov–Dmitrašinović families [19]. Each reference orbit γ_j has a known Latent $\Lambda_{k,j}$ obtained by solving the Galerkin algebraic system (§3.1), not by integration.

Stage 2 (Variational Taylor polynomial). For each reference orbit γ_j , compute the variational equations to obtain the Taylor polynomial map:

$$\Lambda_k(\mathbf{v}_0) = \Lambda_{k,j} + J_{k,j} \cdot \Delta \mathbf{v} + \frac{1}{2} H_{k,j} : \Delta \mathbf{v}^{\otimes 2} + \dots$$

This polynomial maps any nearby initial condition $\mathbf{v}_0 = \mathbf{v}_j + \Delta \mathbf{v}$ in $\mathcal{I}$ directly to its Fourier coefficients — without integrating Newton’s equations. The convergence basin $B_j \subset \mathcal{I}$ of this polynomial is a *Latent atlas chart*.

Stage 3 (Lookup and synthesis). Given an arbitrary initial condition $\mathbf{v}_0 \in \mathcal{I}$:

1. Identify the atlas chart: find the reference orbit γ_j whose basin B_j contains $\mathbf{v}_0$.
2. Evaluate the Taylor polynomial: $\Lambda_k(\mathbf{v}_0) = \Lambda_{k,j} + J_{k,j} \cdot \Delta \mathbf{v} + \dots$
3. Reconstruct the trajectory: $\mathbf{r}(t) = \sum_k \Lambda_k(\mathbf{v}_0) e^{ik\omega t}$.

No ODE integration appears in Steps 2 or 3. The computation is polynomial evaluation (Step 2) followed by trigonometric synthesis (Step 3).

Horizon of validity (important). Step 3 uses the single reference frequency $\omega = \omega(\gamma_j)$, so the synthesized $\mathbf{r}(t)$ is T_j -periodic. Unless $\mathbf{v}_0$ is itself periodic, the true perturbed orbit is *not* periodic (it lies on a KAM torus or a chaotic set), so atlas lookup is a **bounded-horizon approximation** whose error grows secularly with t : numerically, the period-closure defect of a velocity perturbation of size δ scales as $O(\delta)$ per period and accumulates (e.g. $\delta = 10^{-4}$ gives closure error $\sim 2 \times 10^{-3}$ over $1T$ and $\sim 1 \times 10^{-2}$ over $5T$; $\delta = 10^{-2}$ gives ~ 0.25 over $1T$). For horizons beyond $\sim$ one period, or for larger δ , either a multi-frequency Latent or the windowed chaining of §11 is required.

The GPS analogy. The atlas approach is analogous to satellite navigation: computing the satellite ephemerides (the atlas) is a hard, one-time numerical problem. But once the ephemeris tables are computed, the GPS receiver “looks up” its position from the tables — a fast, algebraic operation. Similarly, computing the reference orbit catalog and its variational polynomials is the hard pre-computation; looking up any specific orbit from initial conditions is a polynomial evaluation.

Atlas geometry. The 4D IC-space $\mathcal{I}$ is partitioned into convergence basins $\{B_j\}$, one per orbit family. The boundaries between basins are fractal — a direct manifestation of chaos. The grade-topology theorem (§8.4) controls the atlas chart *size*: for an orbit of topological type $w \in F_2$, the atlas chart stores $O(|w|)$ Fourier coefficients and $O(|w|)$ variational derivatives. Simple orbits (figure-eight, $|w| = 6$) have compact charts; complex orbits (yin-yang, $|w| = 92$) have larger charts.

Figure 1. *Fractal atlas boundaries in the 4D initial-condition space.* A 2D slice of the reduced IC-space (p_1, p_2) at the standard Šuvakov–Dmitrašinović positions, computed at 1000×1000 resolution (10^6 orbits, each integrated for $T = 15$). **Left:** syzygy count (= topological word length $|w|$) — dark regions are simple orbits, bright regions are topologically complex, white is escape or near-collision. Known orbit families are marked with stars. **Right:** minimum pairwise separation — dark regions are near-collision zones. The fractal basin boundaries are visible at all scales, directly demonstrating the chaotic structure of the three-body IC-space that the Latent atlas must navigate.

Figure 2. *Compression ladder for the figure-eight choreography.* The same orbit reconstructed from an increasing number of Fourier modes (the Latent grade). At grade 3 (24 numbers), the shape is recognizable but distorted. At grade 8 (64 numbers), the error is below 4×10^{-3} . At grade 20 (160 numbers), the orbit is visually indistinguishable from the DOP853 ground truth at 10^{-14} tolerance. This is the Latent representation in action: a finite set of numbers that encodes the orbit to any desired accuracy.

Figure 3. *Atlas lookup speed comparison (one-period horizon).* For 100 perturbed initial conditions near the figure-eight orbit, the atlas polynomial evaluation produces a one-period trajectory in 0.17 ms (mean), compared to 34 ms for DOP853 numerical integration — a **199** $\times$ speedup. The quoted mean position error 4.2×10^{-3} (first-order Taylor) is measured **over a single period**; it matches the $O(\delta)$ one-period closure defect for perturbations of size $\delta \sim 10^{-4}$ and therefore does **not** represent whole-orbit recovery for non-periodic perturbed ICs — the error accumulates over successive periods

(see §8.6.3, Horizon of validity). Higher-order polynomials reduce the *one-period* error; extending the horizon requires multi-frequency or windowed representations.

Existence versus efficiency. An important distinction separates two claims about the Latent atlas:

	Existence of solution	Efficiency via atlas
Claim	Every trajectory has a finite Latent representation	Every IC can be resolved by a single atlas lookup
Status	Argued (§11: $B' + C' +$ classical inputs; quantitative B' uniform bound still open, §11.2)	Open (Conjecture A, §11.5)
What it gives	Universal coverage, but via many windowed segments	One polynomial evaluation per IC

The *existence* of a global Latent cover **given B' and C'** does not require atlas completeness. The argument is:

1. **Painlevé (1897)**: away from collisions, the solution is analytic in a strip of width $\rho > 1$ around the real time axis. Every collision-free time window therefore admits a finite Latent.
2. **Levi-Civita (1920)**: binary collisions are regularizable — the transformed flow is analytic through the collision, giving $\rho_{\text{reg}} > 1$ even at collision passages.
3. **Saari (1977)**: simultaneous triple collision is measure zero in phase space.

Combining these: for any initial condition (except a measure-zero set), the trajectory is a sequence of time windows, each with a finite Latent. The representation always exists — the atlas is not needed.

What the atlas provides is *efficiency*. If $\mathbf{v}_0$ falls in a basin B_j , the entire orbit is recovered from one polynomial evaluation — no windowing, no chaining, no per-segment computation. This is the difference between reading a map (atlas lookup) and surveying the terrain yourself (windowed integration).

The open question is whether the basins $\{B_j\}$ cover a full-measure subset of $\mathcal{I}$. Three classical results suggest they should:

1. **Poincaré recurrence**: on a bounded energy surface, almost every orbit returns arbitrarily close to its starting point, shadowing periodic orbits.
2. **KAM theory**: positive-measure sets of quasi-periodic tori persist under perturbation, and these have multi-frequency Latent representations.
3. **Recurrence / shadowing heuristics**: bounded-energy dynamics often revisit neighborhoods of periodic skeletons; turning this into full-measure basin coverage is exactly Conjecture A (§11.5).

But the basin boundaries are fractal, and proving that the union of countably many fractal-bounded regions has full measure is a deep open problem in Hamiltonian ergodic theory (Conjecture A, §11.5). If the conjecture holds, the atlas is a “complete map” — the three-body problem reduces to table lookup for almost every initial condition. If not, the windowed chaining fallback (§11.2–11.3) fills the gaps unconditionally.

8.7 Relation to Prior Work: Quantification Over Existence

The topological classification of periodic three-body orbits by the free group F_2 is due to Montgomery [15]. The realization theorem — every conjugacy class in F_2 is realized by a periodic or collision orbit — is due to Moeckel and Montgomery [16]. The quasi-Kepler law $T \sim |w|$ connecting period to word length was discovered empirically by Dmitrašinović and Šuvakov [17]. These are existence and classification results: they tell us *which* orbits exist and *how to label them*.

Our contribution is the next question: **given that these orbits exist, how large is their description?**

The relationship is analogous to two other milestone pairs in celestial mechanics:

1. **Euler and Lagrange (1767–1772)** [26, 27] found all 5 central configurations for $N = 3$. **Smale’s 6th Problem (1998)** [21] asked: is the number of central configurations always finite for general N ? For $N = 4$, finiteness was proved by Hampton and Moeckel [20]. The existence was classical; the finiteness question came a century later.
2. **Sundman (1912)** proved that a convergent power series solution exists for the three-body problem. **Our result** shows that the practical description needs far fewer numbers: ~ 48 real numbers for the figure-eight orbit ($|w| = 6$, $\varepsilon = 10^{-3}$), ~ 530 for a high-index yin-yang orbit ($|w| = 92$), and in general $O(|w|)$ for any orbit — where Sundman’s series requires $\sim 10^{10^6}$ terms.

In each case, the second result is a **quantification over an existence theorem**. Moeckel–Montgomery proved that every braid type is realized. **This paper does not supply a theorem-level linear bound for all periodic orbits:** within fixed topological families, §8 gives strong numerical evidence that Latent grade scales approximately linearly in $|w|$; a uniform rigorous bound is left open. The classification is old; the empirical complexity law and the Latent computational pipeline are new.

8.8 The Orbit Catalog and the 10,000-Orbit Frontier

The topological classification generates a countable infinity of orbit families, indexed by conjugacy classes in F_2 . The experimental catalog has grown rapidly:

Year	Authors	Orbits	Method	Classification
1767–1772	Euler, Lagrange	5	Analytic (central configurations)	None
2000	Chenciner, Montgomery	1	Variational (figure-eight)	Shape sphere
2013	Šuvakov, Dmitrašinović [19]	13	Numerical search	F_2 free group
2015	Dmitrašinović, Hudomal <i>et al.</i>	~ 300	Extended numerical search	Sequences I–VI
2025	Li, Liao [18]	10,059	High-precision computation (3D)	Mass configuration

The Li–Liao catalog [18] extends the classification to **three-dimensional** orbits with unequal masses ($m_1 = m_2 = 1$, $m_3 \in \{0.1, 0.2, \dots, 2.0\}$), discovering 21 choreographies, 273 “piano-trio” orbits, and $\sim 20\%$ linear stability. Their classification is by mass symmetry (equal / two-equal / general), not by topological word.

The Latent framework provides a **quantitative bridge** between these catalogs:

1. **Each orbit in the catalog has a definite Latent grade.** Given an orbit’s initial conditions and period, the procedure of §8.3 yields $\text{grade}(\varepsilon)$ in seconds.
2. **The grade predicts storage cost (within the measured class).** An orbit with $|w| = k$ syzygies requires $\sim 4\alpha(\varepsilon)k$ real numbers (4 Jacobi coordinates $\times \alpha$ modes per syzygy) *for the planar, equal-mass, $L = 0$ orbits on which α was measured*. Extrapolating the same α to the Li–Liao catalog of 10,059 orbits would give $\sim 10^6$ real numbers total at $\varepsilon = 10^{-2}$ — but this is a **conjectural extrapolation**, not a measured figure: the Li–Liao orbits are three-dimensional and unequal-mass, a regime where neither α nor the linear law has been validated (§8.7).
3. **The Latent orders orbits by complexity.** Simple orbits (figure-eight, $|w| = 6$, grade ≈ 8) and complex orbits (yin-yang high-index, $|w| = 92$, grade ≈ 532) live on the same linear scale, enabling systematic comparison across families.

The 3D extension is structurally natural but not verified here: in 3 dimensions each body has 3 coordinates instead of 2, and the Jacobi representation has 6 signals instead of 4. One would *expect* a fit of the same form $\text{grade}(\varepsilon) \approx \alpha(\varepsilon) \cdot |w|$ — the word length $|w|$ still counts collinear passages on the (now 3D) shape sphere — but the slope α and even the linearity remain to be measured for 3D, unequal-mass orbits.

8.9 Connection to the Global Extension

For non-periodic (chaotic) trajectories, the windowed step-chaining of §11 produces a sequence of local Latents, one per time window. The grade of each local Latent is determined by the number of syzygies within that window. The total information content of a chaotic trajectory over time interval $[0, T_{\text{final}}]$ is:

$$\text{total Latent size} = \sum_{\text{windows}} \text{grade}_n(\varepsilon) = O\left(\frac{T_{\text{final}}}{\tau} \cdot \overline{|w|}_\tau\right)$$

where τ is the window length and $\overline{|w|}_\tau$ is the average number of syzygies per window. This connects the local result (grade $\sim$ word length) to the global result (total representation size $\sim$ total topological complexity).

9. Lean 4 formalization (specification)

Repository note. The paths and theorem names below describe an intended **Mathlib-aligned formalization** of the Latent / Padé / N -body chain in this paper. **They are not present as .lean files under this repository snapshot** (as of 2026-04-10). Until a tagged artifact ships, readers should treat this section as a **verification roadmap** tied to the informal proofs in §§4–8 and §11 — not as a claim of lake build reproducibility from this checkout.

The primary formal targets are the Rational Latent Theorem (Theorem 8, §5), the “eliminable approximation” chain (§6), and the global extension scaffolds (§11). A natural entry file name (when the artifact exists) would be `RationalLatentTheorem.lean` inside a `PadéResummation` (or successor) package.

9.1 Intended theorem map (illustrative names)

Paper section	Planned Lean hook	Intended mathematical content
§4.1	<code>latent_geometric_decay</code> ,	Analyticity of G from
	<code>latent_tail_bound</code>	coefficient decay
§4.2	<code>basis_change_preserves_latent</code>	Change of basis (Fourier $\leftrightarrow$ Padé) preserves rank notions
§5 Thm 8	<code>rational_latent_theorem</code> (bundle)	Size bound + pole absorption + Padé exactness for rational G
§6	<code>all_approximations_elimidable</code> ,	Eliminable approximations;
	<code>threebody_exact_latent_solution</code>	3-body chain as specification
§8	<code>grade_topology_bound</code> ,	Grade-topology scaling;
	braid-realization lemmas	topological prerequisites
§11	<code>global_latent_solution</code> ,	Classification $\Rightarrow$ finite Latent
	windowing lemmas	given $B' + C'$ -type inputs

9.2 What would be verified first

1. **Latent tail bound** from geometric decay ($\rho > 1$).
2. **Padé / pole bookkeeping** (Theorem 8, classical de Montessus layer).
3. **Conditional Painlevé window** ($\delta > 0 \Rightarrow \tau_{\text{sing}} > 0$) and **Levi-Civita** regularization of binary collisions — matching §11.3 and `ClassicalTheorems`-style packaging *when implemented*.
4. **Global extension** only after the quantitative B' gap (§11.2) is either proved or explicitly axiomatized with audited Mathlib statements.

9.3 Honesty boundary (paper prover)

Any future formal proof must separate: **(i)** algebraic / Padé / Latent-compression lemmas; **(ii)** ODE analyticity on a *strip* given **pointwise** lower bounds on pairwise distance; **(iii)** **uniform** such bounds over families of windows or over an energy surface — the step (iii) is exactly where §11.2 flags open analysis. The `TrajectoryType` / hypothesis-parameter pattern in §11.7 is the correct abstraction: **do not** claim the prover derives $\rho > 1$ for a physical orbit without supplying the missing quantitative input.

9.4 Status summary (this checkout)

Layer	In this repo?	Note
Informal proofs in §§1–8, 11	Yes (paper.md)	Authoritative for claims until Lean artifact ships

Layer	In this repo?	Note
Cited kernel/LeanProofs/PadeResummation/*.lean	No	Paths are targets , not dependencies
Mathlib proof of Saari-type triple-collision measure	No	Cite classical literature; do not confuse citation with formal proof

10. Discussion

10.1 What Is New

1. **The exact solution exists as a finite implicit formula** — the Galerkin equation in Latent coordinates. This is not a numerical method; it is the mathematical structure of the solution.
2. **Four levels of analyticity** — temporal, parametric, algebraic, continuous — each exploited constructively. Previous work used at most one or two.
3. **The Rational Latent Theorem** (Theorem 8) — shows that Padé convergence is a consequence of the Latent Theorem applied to the generating function. The theory predicts its own approximation efficiency.
4. **Every approximation is eliminable** — the exact objects behind each truncation are identified, and the substitution chain from approximate to exact is explicit.
5. **Latent complexity is topologically controlled** (§8) — the number of Fourier modes needed to represent a periodic orbit to accuracy ε scales as $O(|w|)$, where $|w|$ is the topological word length in F_2 . This is the first quantitative link between the Latent framework and the topological classification of three-body orbits.

10.2 In What Sense Is the Three-Body Problem “Solved”?

The three-body problem is addressed in the same *implicit-algebraic* sense that the two-body problem is addressed by Kepler’s equation (§6.4). The narrative has three layers: (1) near periodic orbits, the Galerkin Latent with exponential convergence (§2–7); (2) a topological complexity bound — the Latent grade scales as $O(|w|)$ empirically within families (§8); and (3) for general trajectories, **conditional** global coverage via windowing + Levi-Civita + Saari-type exclusions (§11). Classical theorems (Painlevé, Levi-Civita, Saari) are used as **analytic inputs**; **Conjecture A** (§11.5) is optional for existence but relevant for atlas efficiency; the **quantitative uniform windowing bound** in Extension B’ (§11.2) is a genuine remaining analytic gap, not a mere formality.

The key conceptual point is §8.6.1: the two-body problem is topologically trivial (one orbit shape), while the three-body problem is topologically rich (countably infinite shapes in F_2). A single formula cannot describe all three-body orbits — but a universal *encoding system* with topology-controlled size can (§8.6.2).

10.3 What Is the Galerkin Equation, Really?

An important conceptual clarification: the Galerkin equation $R_k(\Lambda, \omega) = 0$ for all k is Newton’s law $\ddot{x} = F(x)$ projected onto Fourier modes. It is the **same** equation in different coordinates — not a new equation. Newton’s ODE was already an implicit exact characterization of the trajectory.

What the Galerkin reformulation provides is not new *information* but new *structure*:

1. **Algebraic rather than differential.** The Galerkin equation is an algebraic system; Newton’s ODE requires integration. This is a qualitative improvement in computational form.
2. **Exponential convergence guarantees.** The Latent Theorem bounds the truncation error, giving a priori control over accuracy.
3. **Separation of concerns.** The reference orbit (Stage 1), IC dependence (Stage 2), and trajectory evaluation (Stage 3) are cleanly separated, each with its own convergence guarantee.
4. **Basis-free exactness.** The generating function $G(z)$ is the same object regardless of representation (Fourier, Padé, pole-residue). Different bases are different coordinate systems for the same mathematical object.

The analogy is: the wave equation $\partial^2 u / \partial t^2 = c^2 \nabla^2 u$ and its Fourier-mode decomposition $\ddot{a}_k = -c^2 k^2 a_k$ describe the same physics. The mode decomposition does not “solve” the wave equation in a fundamentally new sense — but it makes the solution structure transparent and computation tractable.

10.4 Relation to Sundman

Sundman (1912) already proved that a convergent power series solution exists for the three-body problem (away from triple collision). In a formal sense, Sundman “solved” the problem. The reason this is not considered a satisfactory solution is practical: the convergence is so slow ($\sim 10^{10^6}$ terms for modest accuracy) that the series is useless for computation.

The Latent solution achieves exponential convergence — a single figure-eight orbit needs ≈ 48 – 72 numbers for 10^{-3} accuracy (its atlas chart, encoding a whole neighborhood, is $\sim 7.6 \times 10^3$) — by exploiting analyticity in both time and initial-condition parameters, structure that Sundman’s single-variable series ignores. (The comparison is per-object: a single trajectory here versus Sundman’s per-trajectory term count; the atlas-chart figure encodes many orbits and should not be compared to a single Sundman series.) The novelty is therefore not the *existence* of a solution (which Sundman established) but the *rate* of convergence and the *algebraic structure* of the representation.

10.5 Scope and Limitations

- **Triple collision excluded.** Simultaneous collision of all three bodies ($r_{12} = r_{13} = r_{23} = 0$) is not regularizable and is excluded. This is measure zero in phase space (Saari, 1977).
- **Efficiency varies by regime.** Near periodic orbits, the Latent has exponential convergence ($\rho \gg 1$, few modes). In chaotic regimes, windowed step-chaining gives $\rho > 1$ per segment **when** the windowing hypotheses of §11 hold, but the mode count per window may be large. Global *existence* of a Latent cover is claimed **under §11**, not pointwise unconditionally in the strongest naive reading.
- **Planar.** All computations are 2D. The method generalizes to 3D with Kustaanheimo–Stiefel regularization [25] (replacing Levi-Civita) and no conceptual change.

- **The “infinite system” subtlety.** The Galerkin equation $R_k = 0$ for all k is an infinite system of equations (one per Fourier mode). Each individual equation is finite, but there are infinitely many. The “finite equation” claim refers to the generating-function form: $G(z)$ satisfies a single functional equation derived from Newton’s law. Truncation to N modes is the computational realization — analogous to discretizing a PDE.
- **Not a new first integral.** The Galerkin equation does not provide a new conserved quantity. It is a reformulation, not a reduction of the dynamics. Poincaré’s impossibility result is not circumvented — it is simply not relevant to what we construct.
- **Classical dependence.** The global extension relies on Painlevé’s analyticity theorem (1897), Levi-Civita’s regularization (1920), and Saari’s measure-zero result (1977). These are proved theorems, not conjectures — but the Latent solution is only as strong as its classical foundations.

10.6 Relation to Prior Work

- **Sundman (1912):** Convergent series, impractically slow. We achieve exponentially faster convergence by exploiting analyticity in both time and parameters.
- **Poincaré (1890):** Forbids closed-form integrals. Does not forbid implicit formulas.
- **Chenciner & Montgomery (2000):** Discovered the figure-8 orbit. We provide its exact Latent representation.
- **Galerkin/spectral methods:** Known in fluid dynamics and engineering. Applied here to the 3-body problem as an exact algebraic characterization, not merely a numerical method.
- **KAM theory (Kolmogorov–Arnold–Moser) [23, 24]:** Guarantees persistence of quasi-periodic orbits under perturbation. The Latent solution provides the explicit quantitative representation of orbits whose existence KAM guarantees qualitatively.

11. The Global Extension

The local construction of Sections 2–7 yields the exact Latent solution near periodic orbit families. This section extends the result to all trajectories. A key structural insight simplifies the problem: only **two** classical ingredients are needed beyond the Galerkin framework. Scattering follows from windowing + Kepler, and atlas completeness is an efficiency optimization, not a logical prerequisite.

11.1 The Dependency Structure

The four gaps identified in the literature are not independent:

$$\begin{array}{ccc}
 & & \text{C' (Regularization)} \\
 \text{Local result} & \rightarrow & \text{GLOBAL SOLUTION} \\
 & & \text{B' (Windowing)} \quad \rightarrow \\
 & & \text{D (Scattering)} = \text{B' + C' + Kepler} \\
 & & \text{A (Atlas)} = \text{efficiency, NOT required for existence}
 \end{array}$$

B' + C' is sufficient. If every collision-free trajectory segment has $\rho > 1$ (B'), and collision passages are regularized with $\rho_{\text{reg}} > 1$ (C'), then every trajectory of every type has a finite Latent representation. Scattering trajectories decompose into bounded interaction (B' + C') and asymptotic Kepler phases (already solved). Atlas completeness determines *how efficiently* the solution is represented, but is not required for its *existence*.

11.2 Extension B' — Quantitative Painlevé (Chaotic Windowing)

Extension B'. *There exists $c = c(E, \mathbf{m}, \epsilon) > 0$ such that for all points on the energy surface with $r_{ij} \geq \epsilon$, the complex-time singularity distance satisfies $\tau_{\text{sing}} \geq c$.*

Painlevé's theorem already guarantees that the 3-body flow is real-analytic away from collisions. This means $\tau_{\text{sing}} > 0$ for every collision-free point — so B' is *qualitatively* true. What B' adds is the **uniform** lower bound.

For a window of length τ , the analyticity parameter is:

$$\rho(\tau) = e^{2\pi\tau_{\text{sing}}/\tau}$$

Once $\tau_{\text{sing}} > 0$ and $\tau > 0$ are given, the inequality $\rho(\tau) = e^{2\pi\tau_{\text{sing}}/\tau} > 1$ is immediate. The uniform bound $\tau_{\text{sing}} \geq c > 0$ makes this quantitative: $\rho(\tau) \geq e^{2\pi c/\tau} > 1$ uniformly — and **that** uniform bound is what §11.2 flags as open.

The strategy is **windowed step-chaining** — the same approach used in the Padé paper [7]. Each window has its own local Latent:

$$x(t) = G_{[n]}(e^{i\omega_n(t-t_n)}) \quad \text{for } t \in [t_n, t_{n+1}]$$

Chaining K windows with per-window error budget ε/K gives total error $\leq \varepsilon$ by the triangle inequality (standard step-chaining bookkeeping).

What remains to prove: The uniform bound $\tau_{\text{sing}} \geq c > 0$. This is related to Painlevé's conjecture that non-collision singularities do not exist in the 3-body problem (widely believed true for $n = 3$; proved for $n \geq 5$ by Xia 1992). If no non-collision singularity exists, then $\tau_{\text{sing}} \rightarrow 0$ only as $r_{ij} \rightarrow 0$, and the bound follows from the collision exclusion $r_{ij} \geq \epsilon$.

Difficulty: moderate. The qualitative statement is classical. The quantitative bound requires complex-analysis estimates on the singularity structure.

11.3 Extension C' — Collision Regularization

Extension C'. *After Levi-Civita regularization ($w^2 = z$), the Galerkin equation in regularized coordinates has $\rho_{\text{reg}} > 1$ through binary collision passages.*

The Levi-Civita transformation maps the $r \rightarrow 0$ collision singularity to a smooth flow:

$$w'' = -\frac{E}{2}w + \frac{|w|^2}{2}\nabla_w U_{\text{other}}(w)$$

This is an analytic ODE even at $|w| = 0$ (binary collision). Since the regularized flow is analytic, it has exponentially decaying Fourier coefficients with $\rho_{\text{reg}} > 1$.

Heuristically, $\rho_{\text{orig}} \leq 1$ at a binary-collision singularity while regularization can restore $\rho_{\text{reg}} > 1$; making this quantitative for full 3-body passages is the program of §11.3 and §11.8.

Triple collision ($r_{12} = r_{13} = r_{23} = 0$) is not regularizable — but is measure-zero (Saari, 1977) and is excluded.

What remains to prove: Explicit computation of ρ_{reg} for near-collision passages. A numerical demonstration (Galerkin in Levi-Civita coordinates through a known near-collision orbit) is the immediate next step.

Difficulty: low. Levi-Civita regularization is well-understood (Siegel & Moser, 1971). The regularized flow is known to be analytic (Sundman, 1912). The computation is standard spectral methods.

11.4 Scattering: Not an Independent Extension

Escape/capture trajectories decompose into three phases:

1. **Approach:** far from interaction $\rightarrow$ approximately Keplerian (solved)
2. **Interaction:** bounded, finite-time close encounter $\rightarrow$ B' applies
3. **Departure:** far from interaction $\rightarrow$ approximately Keplerian (solved)

If B' and C' hold, the interaction phase has $\rho > 1$ (windowing + regularization). The approach and departure phases have $\rho \gg 1$ (nearly Kepler). Scattering is therefore not an independent extension: it is the same $\rho > 1$ story applied phase-by-phase.

11.5 Atlas Completeness: Efficiency, Not Existence

Conjecture A (*For energy $E < 0$, periodic orbit convergence basins cover a full-measure subset of Σ_E*) determines whether the **atlas** representation (one Latent per orbit family) is globally efficient. Heuristic support includes:

1. **Poincaré recurrence** (bounded components): typical orbits revisit neighborhoods infinitely often.
2. **KAM theory** [23, 24]: positive-measure families of quasi-periodic tori persist under perturbation and admit multi-frequency Latent representations.
3. **Shadowing intuition:** periodic orbits are dense in reasonable model systems, but **Hamiltonian** 3-body dynamics does *not* automatically inherit a hyperbolic “Smale horseshoe” package — Conjecture A should be read as an **open** Hamiltonian question, not as a quotation of a classical theorem.

But A is NOT required for the global solution to exist. With B' + C', the windowed step-chaining already covers everything. A makes the representation compact; B' + C' makes it possible.

Difficulty: very hard. Related to open problems in Hamiltonian ergodic theory.

11.6 The Minimal Global Solution

Given B' and C', the global Latent solution for any trajectory (except measure-zero triple collision) is:

1. **Classify** trajectory segments as regular, chaotic, near-collision, or asymptotic
2. **Regular segments:** direct Galerkin Latent (Section 6)

3. **Chaotic segments:** windowed step-chaining with $\rho \geq e^{2\pi c/\tau} > 1$ (B')
4. **Near-collision passages:** Levi-Civita regularized Galerkin with $\rho_{\text{reg}} > 1$ (C')
5. **Asymptotic phases:** Kepler (solved)
6. **Compose** all segments by step-chaining

Conditional global lemma (informal). Given B' and C' as in §11.1–11.3, every trajectory type listed there admits a finite Latent to accuracy ε by windowing + regularization + Kepler tails. The **remaining mathematical work** is to supply the **uniform** B' input (§11.2) and to justify Saari-type exclusions at the level of a journal proof (not only by citation).

11.7 Formalization hooks for the global extension (see §9)

The table below lists **illustrative Lean names** for packaging the §11 dependency graph. It is **not** a build report: the corresponding sources are **not included** in this repository snapshot.

Extension	Illustrative lemma name	Intended content
B' (Windowed ρ)	windowed_rho_gt_one	$\tau_{\text{sing}} > 0, \tau > 0 \Rightarrow \rho > 1$
B' (Uniform bound)	quantitative_painleve_bound	Uniform $\tau_{\text{sing}} \geq c \Rightarrow$ uniform $\rho > 1$
B' (Per-segment Latent)	windowed_latent_per_segment	N modes suffice per window
C' (Regularized Latent)	regularized_latent_exists	Finite-mode representation after Levi-Civita
Global (Main)	global_latent_solution	Hypothesis-parameterized completion

Prover hygiene. Any future formal proof must not pretend that carrying $\rho > 1$ in a type family substitutes for proving **uniform** lower bounds on complex-time regularity across windows — the gap in §11.2 is exactly about exchanging **pointwise** non-collision for **uniform** estimates.

Two analytic inputs remain prominent in this manuscript:

1. **Saari-type triple-collision exclusions** for the planar problem (see [10] and the discussion in §10.5; verify that the cited Saari paper matches the exact statement you need, or cite a more targeted source).
2. **Uniform separation / quantitative B'** (§11.2): supported numerically in §11.8, but not closed as a theorem here.

11.8 Numerical Verification

We verify both extensions numerically on the Pythagorean three-body problem (Burrau [8]: masses 3, 4, 5 at vertices of a 3-4-5 right triangle, zero initial velocities). This is a canonical chaotic test case with close encounters down to $r_{\min} \approx 10^{-2}$.

Verification tables withdrawn (provisional). Earlier drafts of this section tabulated specific per-window ρ values for B' and per-impact-parameter ρ values for C'. Those tables were found to be **internally inconsistent with the paper's own** $\rho(\tau) = e^{2\pi\tau_{\text{sing}}/\tau}$ **law (§11.2) and physically implausible**, and have been removed rather than reported as trustworthy evidence:

- **B'**: the tabulated $\rho_{\min}$ was essentially flat in τ (≈ 1.027 – 1.028 for $\tau = 0.5, 1, 2, 4$), whereas §11.2's own law requires strong τ -dependence — calibrating τ_{sing} from $\rho(4.0) = 1.028$ forces $\rho(0.5) \approx 1.25$, not 1.027 . The reported per-window minimum separations were also identical across different window sizes, which cannot hold for many small windows.
- **C'**: the tabulated ρ was identical to three significant figures across four impact parameters spanning $500\times$ in b , while $r_{\min}$ barely moved — physically implausible for hyperbolic two-body encounters, where $r_{\min}$ should track b .

These signatures indicate a broken or templated measurement pipeline, not a verified result. Regenerating both tables with a documented, reproducible pipeline — one that exhibits ρ varying with τ (B') and $r_{\min}$ tracking b (C'), with honest per-window separations — is required before any quantitative “100% of windows $\rho > 1$ ” or “ $2.4\times$ mode saving” claim can be made. Pending that regeneration, §11.8 provides **no** numerical verification.

What remains defensible qualitatively. Away from collisions the flow is analytic (Painlevé), so each collision-free window has $\rho > 1$ *pointwise*; and Levi-Civita regularization is a classical, analytic transformation of binary collisions, so a mode reduction near close encounters is expected on general grounds. Neither statement is a substitute for the withdrawn quantitative tables, and neither closes the **uniform** bound of §11.2.

11.9 Research Program

Priority	Extension	Status
1	C' (Collision Reg.)	Classical / informal. Numerical mode-saving table withdrawn (§11.8, inconsistent pipeline); regeneration pending.
2	B' (Quantitative Painlevé)	Conditional chain ($\delta > 0 \Rightarrow \tau_{\text{sing}} > 0 \Rightarrow \rho > 1$) is standard ODE lore (pointwise). Per-window verification table withdrawn (§11.8). Uniform $\delta / \tau_{\text{sing}}$: open (§11.2).
3	Saari (Triple collision measure zero)	Classical literature [10]; match theorem statement to 3-body claim before treating as black box
4	A (Atlas efficiency)	Open (long-term, not required for existence)

12. Conclusion

Under the global hypotheses of §11, the planar gravitational three-body problem admits an exact Latent *encoding* for trajectories outside the usual triple-collision pathology. The construction combines one new ingredient — the Galerkin–Latent framework — with classical analytic inputs: Painlevé-type complex-time regularity away from collisions, Levi-Civita’s binary regularization, and Saari-type measure-zero exclusions (verify statement-level match with [10]).

The solution is not a number, not a series, and not an approximation. It is an **equation** — Newton’s law in Latent coordinates — and the equation IS the solution, in the same sense that $x^2 = 2$ is the exact definition of $\sqrt{2}$. Every numerical evaluation (Fourier, Taylor, Padé) is a way to compute digits of this exact object.

The result. For initial conditions in the regime covered by §11 ($B' + C' +$ classical exclusions), the trajectory $x_i(t)$ has a finite Latent representation to arbitrary accuracy ε :

- **Exact periodic orbits:** the generating function $G_i(z; \mathbf{v}_0)$ satisfying the Galerkin equation, with the single-frequency closed form $x_i(t) = G_i(e^{i\omega t})$. Exponential convergence at the measured $\rho \approx 1.8$, less iteration than Kepler; a single figure-eight orbit needs ≈ 48 –72 coefficients (its atlas chart, a neighborhood, $\sim 7.6 \times 10^3$) versus Sundman’s $\sim 10^{10^6}$ terms.
- **Chaotic segments:** windowed step-chaining. Painlevé guarantees $\tau_{\text{sing}} > 0$ per collision-free segment, so $\rho > 1$ *pointwise* for every finite window; the **uniform** bound over all windows (§11.2) is open. (The earlier per-window verification table is withdrawn — §11.8.)
- **Near-collision passages:** Levi-Civita regularization transforms the binary singularity into smooth flow with $\rho_{\text{reg}} > 1$, a classical fact; the earlier quantitative mode-reduction table is withdrawn pending a reproducible pipeline (§11.8).
- **Scattering trajectories:** decompose into bounded interaction (windowing + regularization) and asymptotic Kepler phases (already solved).


Formal verification. This repository snapshot **does not ship** the Lean modules named in earlier drafts. §9 and §11.7 record **targets only**. The informal chain is: Levi-Civita removes binary singularities in regularized coordinates; Painlevé-type estimates give $\tau_{\text{sing}} > 0$ on **collision-free** strips; hence $\rho > 1$ **once** a positive lower bound on separations is available per window. Closing the **uniform** B' gap (§11.2) and aligning Saari citations with the exact 3-body statement remain part of a submission-grade proof, not a footnote.

What is new. The Galerkin equation is Newton’s law in Fourier coordinates; what is new is not the equation itself but the structural insight that this reformulation yields exponential convergence, four-level analyticity exploitation, and an exact generating-function representation *for periodic orbits* — and that, **conditional on the uniform windowing bound of §11.2**, combining it with classical regularization theorems would cover the entire phase space. The three-body problem has no closed-form first integral (Poincaré, 1890); for exact periodic orbits it has a finite, exact, implicit representation, and — conditional on §11 — a finite encoding for every trajectory. The gap between “we prove this for periodic orbits” and “this holds for every trajectory” is precisely the open uniform bound B' (§11.2), not a settled result.

Efficiency. Where §11 applies, the Latent cover exists, but it is **most efficient** near periodic orbits. In chaotic regimes, windowed chaining produces a finite representation but may require many segments. Atlas completeness (§11.5) remains an open **efficiency** question.

Poincaré showed that the 3-body problem has no closed-form first integral. He did not show it has

no implicit characterization. The Galerkin system is such a characterization — **as an encoding**, not as a single closed formula for all F_2 types — with exponential precision **when** $\rho \gg 1$.



Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Cursor and its integrated AI models (Claude and GPT families) in order to assist with manuscript drafting, structural editing, and language polishing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.



References

- Sundman, K. F (1912). Mémoire sur le problème des trois corps. *Acta Mathematica*, 105-179. DOI: 10.1007/bf02422379
- Poincaré, H (1890). Sur le problème des trois corps et les équations de la dynamique. *Acta Mathematica*, 1-270.
- Chenciner, A. and R. Montgomery (2000). A remarkable periodic solution of the three-body problem in the case of equal masses. *Annals of Mathematics*, 881-901. DOI: 10.2307/2661357
- de Montessus de Ballore, R (1902). Sur les fractions continues algébriques. *Bull. Soc. Math. France*, 28-36. DOI: 10.1007/bf03014011
- Baker, G. A. and P. Graves-Morris (1996). Padé Approximants. *Padé Approximants*.
- Nagy, T. (2026). The Latent: Finite Sufficient Representations of Smooth Systems. *Zenodo*. DOI: 10.5281/zenodo.19101209
- Nagy, T. (2026). Practical Padé Representations of the Gravitational Three-Body Problem. *Zenodo*. DOI: 10.5281/zenodo.19101253
- Burrau, C (1913). Numerische Berechnung eines Spezialfalles des Dreikörperproblems. *Astronomische Nachrichten*, 113-118. DOI: 10.1002/asna.19131950602
- Siegel, C. L. and J. K. Moser (1971). Lectures on Celestial Mechanics. *Lectures on Celestial Mechanics*. DOI: 10.1007/978-3-642-87284-6
- Saari, D. G (1977). A global existence theorem for the four-body problem of Newtonian mechanics. *J. Differential Equations*, 80-111. DOI: 10.1016/0022-0396(77)90100-0
- Xia, Z (1992). The existence of noncollision singularities in Newtonian systems. *Annals of Mathematics*, 411-468. DOI: 10.2307/2946572
- Painlevé, P (1897). Leçons sur la théorie analytique des équations différentielles. *Leçons sur la théorie analytique des équations différentielles*.
- Hille, E (1976). Ordinary Differential Equations in the Complex Domain. *Ordinary Differential Equations in the Complex Domain*. DOI: 10.1090/simon/002.2/03
- Levi-Civita, T (1920). Sur la régularisation du problème des trois corps. *Acta Mathematica*, 99-144. DOI: 10.1007/bf02404404

- Montgomery, R (2005). Fitting hyperbolic pants to a three-body problem. *Ergodic Theory and Dynamical Systems*, 25(3), 921-947.
- Moeckel, R. and R. Montgomery (2015). Realizing all reduced syzygy sequences in the planar three-body problem. *Nonlinearity*, 28(6), 1919-1935.
- Dmitrašinović, V. and M. Šuvakov (2015). Topological dependence of Kepler's third law for planar periodic three-body orbits with vanishing angular momentum. *Physics Letters A*, 34-35.
- Li, X. and S. Liao (2025). Collisionless periodic orbits in the free-fall three-body problem. *New Astronomy*.
- Šuvakov, M. and V. Dmitrašinović (2013). Three classes of Newtonian three-body planar periodic orbits. *Physical Review Letters*, 110(11).
- Hampton, M. and R. Moeckel (2006). Finiteness of relative equilibria of the four-body problem. *Inventiones Mathematicae*, 163(2), 289-312.
- Smale, S (1998). Mathematical problems for the next century. *The Mathematical Intelligencer*, 20(2), 7-15.
- Bruns, H (1887). Über die Integrale des Vielkörper-Problems. *Acta Mathematica*, 25-96.
- Kolmogorov, A. N (1954). On conservation of conditionally periodic motions for a small change in Hamilton's function. *Doklady Akademii Nauk SSSR*, 98(4), 527-530.
- Arnold, V. I (1963). Proof of a theorem of A. N. Kolmogorov on the invariance of quasi-periodic motions under small perturbations. *Russian Mathematical Surveys*, 18(5). DOI: 10.1007/978-3-642-01742-1_21
- Kustaanheimo, P. and E. Stiefel (1965). Perturbation theory of Kepler motion based on spinor regularization. *Journal für die reine und angewandte Mathematik*, 204-219. DOI: 10.1515/crll.1965.218.204
- Euler, L (1767). De motu rectilineo trium corporum se mutuo attrahentium. *Novi Commentarii Academiae Scientiarum Petropolitanae*, 144-151.
- Lagrange, Joseph Louis (1772). Étude sur le problème des trois corps. *Prix de l'Académie des Sciences de Paris*.