Blur for the Three–Body Problem

Finite Families under Coarse Observation

Aleksandar Perišić

December 2025

Abstract

We treat the Newtonian three–body problem as a testbed for the general principle that *under blur there are only finitely many effective families of trajectories.* We introduce blur operators on both the phase space and the parameter space (masses, gravitational constant), and define *blur–families* of orbits: classes of solutions whose blurred trajectories remain uniformly close over a fixed time window. On any compact, non–collision energy region and finite time interval, we show that a finite set of blur–families suffices to approximate all trajectories at a given resolution. When parameters are blurred as well, we obtain finite families in the joint parameter–state space, and we describe *blur–stable dynamical regimes* where qualitative orbit types (bounded/escape) are constant on blur–classes.

Analytically, the results are straightforward applications of continuous–dependence estimates. Conceptually, they fit into a broader "blur vs. finiteness" picture that also appears in number–theoretic problems (e.g. Collatz certificates): one trades sharp, global knowledge for a finite number of coarse representatives. The three–body case provides a concrete, finite–dimensional illustration of how blur turns chaotic dynamics into finitely many observable histories.

Contents

1 Introduction

The Newtonian three–body problem is a classical laboratory for chaotic dynamics and structural complexity. Even in the simplest case of three point masses moving under mutual Newtonian attraction, the long–time behaviour of orbits exhibits extreme sensitivity to initial conditions and parameters.

From a blur perspective, this sensitivity is a feature, not a bug. Rather than asking for exact pointwise prediction, we ask:

> Given a blur scale r in phase space (and possibly in parameter space) and a tolerance ε on a finite time horizon $[0, T]$, how many distinct *blur–families* of trajectories does the system really have?

The answer is: finitely many. More precisely:

- fix a compact non–collision energy region and a time horizon T;
- fix blur scales r (state space) and σ (parameters), and a tolerance ε for the blurred trajectories;
- then there exists a *finite* set of representative trajectories such that every (parameter, initial condition) pair in the region is blur–close to one of these representatives on $[0, T]$.

This is almost tautological from the viewpoint of continuous dependence: the flow is continuous in initial data and parameters; on a compact set we have uniform continuity; therefore, finitely many sample points suffice at any fixed resolution.

The point of this note is to:

(i) formulate this in an explicit blur language, including blur on parameters;

(ii) state transparent finiteness results (blur–families and blur–stable regimes) that can be transplanted to more complicated systems;

(iii) isolate the pattern: *finite families under blur* as a kind of coarse certificate, echoing the finite certificates used in other contexts (e.g. Collatz Lyapunov proofs).

The results here are deliberately modest on the analytic side: nothing about global integrability, no regularity miracle. The contribution is a clean blur formalization in a very concrete, classical model.

2 The Newtonian three–body problem

2.1 Phase space and vector field

We work in three spatial dimensions; the planar case is a simplification but does not change the blur logic.

Let $q_i \in \mathbb{R}^3$ be the position and $p_i \in \mathbb{R}^3$ the momentum of the i–th body, $i = 1, 2, 3$. Denote

$$q = (q_1, q_2, q_3) \in \mathbb{R}^9, \qquad p = (p_1, p_2, p_3) \in \mathbb{R}^9, \qquad z = (q, p) \in \mathsf{X} := \mathbb{R}^{18}.$$

Let $m = (m_1, m_2, m_3) \in (0, \infty)^3$ be the masses and $G > 0$ the gravitational constant. The Hamiltonian is

$$H_{G,m}(q, p) = \sum_{i=1}^{3} \frac{|p_i|^2}{2m_i} - G \sum_{1 \le i < j \le 3} \frac{m_i m_j}{|q_i - q_j|}. \tag{1}$$

The equations of motion are Hamilton's equations

$$\dot{z} = F_{G,m}(z) = J\nabla H_{G,m}(z),$$

with J the standard symplectic matrix. For now we fix G and m and write $F = F_{G,m}$, $H = H_{G,m}$.

2.2 Non–collision energy shells

Let $\mathcal{C}$ be the collision set:

$$\mathcal{C} = \bigcup_{1\le i<j\le 3} \{z = (q,p) : q_i = q_j\}.$$

We restrict attention to a compact, non–collision region in phase space.

Definition 2.1 (Non–collision compact region)**.** Fix energy level $E \in \mathbb{R}$, radius $R > 0$ and a collision–avoidance margin $\delta > 0$. Define

$$\Sigma_{E,R,\delta} = \left\{ z = (q,p) \in \mathbb{R}^{18} : H(z) = E,\ \max_i |q_i| \le R,\ \max_i |p_i| \le R,\ |q_i - q_j| \ge \delta \ \forall i < j \right\}.$$

Proposition 2.2 (Lipschitz vector field on $\Sigma_{E,R,\delta}$)**.** *On $\Sigma_{E,R,\delta}$ the vector field F is C^∞ and Lipschitz: there exists $L = L(E,R,\delta,G,m) > 0$ such that*

$$\|F(z) - F(z')\| \le L\,\|z - z'\| \qquad \forall z, z' \in \Sigma_{E,R,\delta}.$$

Sketch. On $\Sigma_{E,R,\delta}$ all mutual distances $|q_i - q_j|$ are bounded below by δ, so the potential $-Gm_im_j/|q_i - q_j|$ and its derivatives are bounded. The kinetic part is polynomial. Hence ∇H is smooth with bounded first derivatives on $\Sigma_{E,R,\delta}$, which implies the Lipschitz bound. □

By standard ODE theory, for each $z_0 \in \Sigma_{E,R,\delta}$ there exists a unique solution $z(t; z_0)$ defined on some interval around 0. Since $\Sigma_{E,R,\delta}$ is compact and F is Lipschitz on it, there is a uniform existence time $T_* = T_*(E,R,\delta,G,m)$ such that solutions starting in $\Sigma_{E,R,\delta}$ remain in $\Sigma_{E,R,\delta}$ for $|t| \le T_*$. For $T \le T_*$ we denote the flow by

$$\Phi_t : \Sigma_{E,R,\delta} \to \Sigma_{E,R,\delta}, \quad \Phi_t(z_0) = z(t; z_0).$$

Proposition 2.3 (Continuous dependence on initial data)**.** *For $|t| \le T$ with $T \le T_*$ and $z_0, z_0' \in \Sigma_{E,R,\delta}$,*

$$\|\Phi_t(z_0) - \Phi_t(z_0')\| \le e^{LT}\,\|z_0 - z_0'\|,$$

where L is the Lipschitz constant from Proposition 2.2.

Proof. Gronwall's inequality applied to $\dot{z} = F(z)$ and the Lipschitz estimate yields the standard bound $\|\Phi_t(z_0) - \Phi_t(z_0')\| \le e^{L|t|}\|z_0 - z_0'\|$. □

3 Blur on phase space and blur–families

We now introduce blur on the phase space and define blur–families of trajectories.

3.1 Phase–space blur

Since $\mathsf{X} = \mathbb{R}^{18}$ is finite–dimensional, many blur operators are available. For definiteness we use a Gaussian convolution blur on functions, and the associated action on trajectories.

Definition 3.1 (Phase–space blur operator)**.** For $r > 0$, let $\kappa_r : \mathsf{X} \to [0, \infty)$ be a Gaussian kernel

$$\kappa_r(z) = \frac{1}{(2\pi r^2)^9} \exp\Big(- \frac{\|z\|^2}{2r^2} \Big).$$

For a bounded continuous observable $\varphi : \mathsf{X} \to \mathbb{R}$ define its blurred version

$$(\mathsf{B}_r \varphi)(z) = \int_{\mathsf{X}} \kappa_r(z - z')\, \varphi(z')\, dz'.$$

For a trajectory $z(t)$ in X we define the *blurred trajectory* at scale r as the curve

$$t \mapsto (\mathsf{B}_r \delta_{z(t)}),$$

or more concretely via any fixed observable φ: we track $t \mapsto (\mathsf{B}_r\varphi)(z(t))$ instead of $t \mapsto \varphi(z(t))$.

For our finiteness results it is enough to work at the level of the state itself and view blur as a bounded linear map $\mathsf{B}_r : \mathsf{X} \to \mathsf{X}$ satisfying:

- $\|\mathsf{B}_r\| \le 1$ for all r;
- $\mathsf{B}_r \to \mathrm{Id}$ strongly as $r \downarrow 0$.

One may think of B_r as projecting onto a coarse grid of mesh r or onto low Fourier modes; the precise choice does not affect the logic.

Definition 3.2 (Blur distance on trajectories)**.** Fix $r > 0$ and a time horizon $T > 0$. For two trajectories $z(t), \tilde{z}(t)$ defined on $[0, T]$ we define their blurred distance

$$d_{r,T}(z, \tilde{z}) := \sup_{0 \le t \le T} \|\mathsf{B}_r z(t) - \mathsf{B}_r \tilde{z}(t)\|.$$

3.2 Blur–families on a compact energy shell

We now define blur–families on $\Sigma_{E,R,\delta}$.

Definition 3.3 ((r, ε, T)–blur–family)**.** Let $K \subset \Sigma_{E,R,\delta}$ be compact. Fix a blur scale $r > 0$, tolerance $\varepsilon > 0$ and a time horizon $T \le T_*$. A family of reference initial data $\{z_j\}_{j=1}^N \subset K$ is called an (r, ε, T)*–blur–family* for K if for every $z_0 \in K$ there exists some j such that

$$d_{r,T}\big(\Phi_\cdot(z_0),\, \Phi_\cdot(z_j)\big) = \sup_{0 \le t \le T} \|\mathsf{B}_r \Phi_t(z_0) - \mathsf{B}_r \Phi_t(z_j)\| \le \varepsilon.$$

The corresponding blurred trajectories $t \mapsto \mathsf{B}_r \Phi_t(z_j)$ are called *blur–family representatives.*

In words: within the blur resolution r and tolerance ε, every trajectory starting in K is indistinguishable from one of the representative trajectories up to time T.

The main observation is that such families always exist and are finite.

Theorem 3.4 (Finite blur–families on a compact region)**.** *Let $\Sigma_{E,R,\delta}$ and $T \le T_*$ be as above, and let $K \subset \Sigma_{E,R,\delta}$ be compact. For any blur scale $r > 0$ and tolerance $\varepsilon > 0$ there exists a finite (r, ε, T)–blur–family for K.*

More precisely, one may choose reference points $\{z_j\}_{j=1}^N$ such that $K \subset \bigcup_{j=1}^N B(z_j, \delta_0)$ with

$$\delta_0 = \varepsilon\, e^{-LT},$$

where L is the Lipschitz constant of F on $\Sigma_{E,R,\delta}$, and N is the covering number of K by balls of radius δ_0.

Proof. Let L be the Lipschitz constant from Proposition 2.2. By Proposition 2.3,

$$\|\Phi_t(z_0) - \Phi_t(z_0')\| \le e^{LT} \|z_0 - z_0'\|$$

for all $0 \le t \le T$ and $z_0, z_0' \in K$. Since $\|\mathsf{B}_r\| \le 1$,

$$\|\mathsf{B}_r\Phi_t(z_0) - \mathsf{B}_r\Phi_t(z_0')\| \le \|\Phi_t(z_0) - \Phi_t(z_0')\| \le e^{LT} \|z_0 - z_0'\|.$$

Choose a finite δ_0–net $\{z_j\}_{j=1}^N$ in K with $\delta_0 = \varepsilon e^{-LT}$; such a net exists because K is compact. Then for any $z_0 \in K$ there is a j with $\|z_0 - z_j\| \le \delta_0$, and therefore

$$d_{r,T}(\Phi_\cdot(z_0), \Phi_\cdot(z_j)) \le e^{LT} \delta_0 = \varepsilon.$$

So $\{z_j\}$ is an (r, ε, T)–blur–family. The bound on N is immediate from the covering definition. □

The theorem confirms the intuitive picture: on a compact, non–collision energy region and finite time horizon, there are only finitely many blur–families at any fixed blur and tolerance. As we refine blur (smaller r or ε), the required number of families grows with the covering number of K and exponentially with the time window T (through e^{LT}).

Remark 3.5 (Dependence on blur scale)**.** Note that the bound for N in Theorem 3.4 does not depend explicitly on r: we did not use that B_r is approximating the identity as $r \to 0$. To incorporate that, one may let the tolerance $\varepsilon = \varepsilon(r)$ depend on r, shrinking as $r \to 0$, to reflect the fact that coarser blur allows larger tolerance in the sharp norm; the existence result then adapts with a rescaled $\delta_0(r)$.

4 Blurring parameters as well

The three–body flow depends not only on the initial condition but also on parameters (G, m_1, m_2, m_3). We now introduce blur on parameter space and obtain finiteness in the joint parameter–state space.

4.1 Parameter space and continuity

Let the parameter space be

$$\mathsf{P} = [G_{\min}, G_{\max}] \times [m_{1,\min}, m_{1,\max}] \times [m_{2,\min}, m_{2,\max}] \times [m_{3,\min}, m_{3,\max}],$$

with all bounds positive and finite. For $p = (G, m_1, m_2, m_3) \in \mathsf{P}$ we write F_p and Φ_t^p for the vector field and flow associated to that parameter tuple.

Proposition 4.1 (Uniform Lipschitz and continuous dependence on p)**.** *Let $K \subset \Sigma_{E,R,\delta}$ be compact and let P be as above, with all m_i bounded away from 0 and infinity and G in a compact interval. Then:*

(a) There exists $L_K > 0$ such that for all $p \in \mathsf{P}$,

$$\|F_p(z) - F_p(z')\| \le L_K \|z - z'\| \qquad \forall z, z' \in K.$$

(b) For each $T \le T_*$, the flow $(p, z_0) \mapsto \Phi_t^p(z_0)$ is uniformly continuous on $\mathsf{P} \times K$ for $0 \le t \le T$.

Sketch. On $K \times \mathsf{P}$, the masses and G are bounded and bounded away from 0, and the positions are collision–separated by δ. The Hamiltonian and its derivatives up to second order are bounded uniformly in (p, z), so the Lipschitz constant in z can be chosen uniformly in p. For (b), F_p depends smoothly on p, so the solution depends continuously on both p and z_0; on a compact domain $\mathsf{P} \times K$ and finite time interval, this implies uniform continuity. □

4.2 Blur on parameter space

We now define blur on parameter space analogously.

Definition 4.2 (Parameter blur)**.** Let $\mathsf{B}_\sigma^{(\mathsf{P})} : \mathsf{P} \to \mathsf{P}$ be a family of blur operators at scales $\sigma > 0$, for instance given by convolution with a smooth compactly supported kernel on the box P (extended by reflection at the boundary), or by projection to a coarse grid of mesh σ. We require:

- $\|\mathsf{B}_\sigma^{(\mathsf{P})}\| \le 1$;
- $\mathsf{B}_\sigma^{(\mathsf{P})} \to \mathrm{Id}$ as $\sigma \downarrow 0$.

We consider the joint blur operator on parameter–state pairs $(p, z) \in \mathsf{P} \times \mathsf{X}$:

$$\mathsf{B}_{\sigma,r}(p, z) := (\mathsf{B}_\sigma^{(\mathsf{P})} p,\ \mathsf{B}_r z).$$

4.3 Joint blur–families in parameter and state

We now define blur–families in the joint space $\mathsf{P} \times K$.

Definition 4.3 ($(\sigma, r, \varepsilon, T)$–joint blur–family)**.** Fix compact $K \subset \Sigma_{E,R,\delta}$, parameter box P, blur scales $\sigma > 0$, $r > 0$, tolerance $\varepsilon > 0$ and $T \le T_*$. A finite collection of representatives

$$\{(p_j, z_j)\}_{j=1}^N \subset \mathsf{P} \times K$$

is a $(\sigma, r, \varepsilon, T)$*–joint blur–family* if for every $(p, z_0) \in \mathsf{P} \times K$ there exists j such that

$$\sup_{0 \le t \le T} \left\|\mathsf{B}_{\sigma,r}(p, \Phi_t^p(z_0)) - \mathsf{B}_{\sigma,r}(p_j, \Phi_t^{p_j}(z_j))\right\| \le \varepsilon.$$

In words: under the observer's blur on parameters and phase space, every realised three–body configuration over $[0, T]$ is indistinguishable from one of the representative parameter–trajectory pairs.

Theorem 4.4 (Finite joint blur–families)**.** *Under the assumptions of Proposition 4.1, for any blur scales $\sigma > 0$, $r > 0$, tolerance $\varepsilon > 0$ and time horizon $T \le T_*$, there exists a finite $(\sigma, r, \varepsilon, T)$–joint blur–family for $\mathsf{P} \times K$.*

More concretely, there exists a finite δ–net $\{(p_j, z_j)\}_{j=1}^N$ in $\mathsf{P} \times K$ such that the above property holds, where $\delta = \delta(\varepsilon)$ is determined by the uniform continuity modulus of $(p, z_0) \mapsto (\mathsf{B}_{\sigma,r} p, \mathsf{B}_{\sigma,r} \Phi_t^p(z_0))$ on $\mathsf{P} \times K$ and $[0, T]$.

Proof. By Proposition 4.1, the map

$$(p, z_0, t) \mapsto (\mathsf{B}_\sigma^{(\mathsf{P})} p,\ \mathsf{B}_r \Phi_t^p(z_0))$$

is continuous on the compact set $\mathsf{P} \times K \times [0, T]$, hence uniformly continuous. Thus there exists a modulus $\omega(\cdot)$ such that

$$\|(p, z_0) - (p', z_0')\| \le \delta \implies \sup_{0 \le t \le T} \left\|\mathsf{B}_{\sigma,r}(p, \Phi_t^p(z_0)) - \mathsf{B}_{\sigma,r}(p', \Phi_t^{p'}(z_0'))\right\| \le \omega(\delta).$$

Pick $\delta > 0$ with $\omega(\delta) \le \varepsilon$, and choose a finite δ–net $\{(p_j, z_j)\}_{j=1}^N$ in $\mathsf{P} \times K$ (possible by compactness). Then for any (p, z_0) there exists j with $\|(p, z_0) - (p_j, z_j)\| \le \delta$, which gives the desired inequality. □

The theorem expresses precisely the finite–families–under–blur intuition in the joint parameter–state space.

5 Blur–stable dynamical regimes

The finiteness theorems above concern trajectories as curves. Often one is more interested in qualitative *orbit types*: bounded motion, escape, collision, and refinements thereof (e.g. binary formation, ejection of one body).

We sketch how blur–families interact with coarse classification maps.

5.1 Coarse orbit types

Fix a classification time horizon $T_{\max} > 0$. Define a coarse orbit type map

$$\mathsf{C} : \mathsf{P} \times K \to \{1, 2, 3\},$$

where, say,

- $\mathsf{C}(p, z_0) = 1$ if the orbit remains in a bounded region (no escape, no collision) up to time $T_{\max}$;
- $\mathsf{C}(p, z_0) = 2$ if a collision occurs before $T_{\max}$;
- $\mathsf{C}(p, z_0) = 3$ if some body escapes beyond a large fixed distance R_{esc} before $T_{\max}$ without collision.

The precise definitions are not crucial; the key point is that C is a map from (p, z_0) to a finite set of labels.

In general C may be highly discontinuous on $\mathsf{P} \times K$ due to chaotic boundaries between regimes. Nevertheless, there are regions where it is locally constant.

Definition 5.1 (Blur–stable regime)**.** Let $(p_*, z_*) \in \mathsf{P} \times K$. We say that (p_*, z_*) lies in a *blur–stable regime at scales* (σ, r) if there exists $\varepsilon > 0$ such that

$$\mathsf{C}(p, z_0) = \mathsf{C}(p_*, z_*)$$

for all (p, z_0) with

$$\|\mathsf{B}_{\sigma,r}(p, z_0) - \mathsf{B}_{\sigma,r}(p_*, z_*)\| \leq \varepsilon.$$

In other words, within the blur–class around (p_*, z_*) the coarse orbit type is constant; the regime is stable under the observer's blur.

Proposition 5.2 (Local blur–stability under structural stability)**.** *Suppose* C *is locally constant in a (sharp) neighbourhood* $U \times V$ *of* (p_*, z_*) *in* $\mathsf{P} \times K$*. Then there exist* $\sigma_0, r_0 > 0$ *and* $\varepsilon > 0$ *such that* (p_*, z_*) *is blur–stable at all blur scales* (σ, r) *with* $0 < \sigma \leq \sigma_0$*,* $0 < r \leq r_0$*, with the same label* $\mathsf{C}(p_*, z_*)$*.*

Proof. If C is locally constant near (p_*, z_*), there exists $\eta > 0$ such that $\mathsf{C}(p, z_0) = \mathsf{C}(p_*, z_*)$ whenever $\|(p, z_0) - (p_*, z_*)\| \leq \eta$. Choose $\sigma_0, r_0 > 0$ so small that $\|\mathsf{B}_{\sigma,r}(p, z_0) - (p_*, z_*)\| \leq \eta$ whenever $\|(p, z_0) - (p_*, z_*)\| \leq \eta/2$ and $0 < \sigma \leq \sigma_0$, $0 < r \leq r_0$; this is possible since $\mathsf{B}_{\sigma,r} \to \mathrm{Id}$ as $\sigma, r \to 0$. Then take $\varepsilon = \eta/2$; the blur–ball of radius ε around $\mathsf{B}_{\sigma,r}(p_*, z_*)$ in blurred norm sits inside the sharp ball of radius η around (p_*, z_*), where C is constant. □

Combined with Theorem 4.4, this implies that on any compact region where C is piecewise locally constant, the number of distinct blur–stable orbit types at a given blur resolution is finite. Each blur–family representative in the joint space can be labelled by its orbit type, and the blur–stable regimes correspond to clusters of representatives with the same label.

While this is still completely in the realm of classical continuity and structural stability, the blur language turns it into a clean statement: *at any finite resolution, the three–body system has finitely many blur–stable fates.*

Conceptual takeaway: why blur is not optional

From a distance, blur may look like a repackaging of things we already know: error bars, coarse grids, numerical tolerances. The point of the present analysis is that, for genuinely history–dependent systems such as the three–body problem, blur is not just convenient, it is mathematically unavoidable.

A deterministic three–body system with exact parameters and exact initial data has a unique history. In practice, however, the parameters and initial states live in a small but irreducible uncertainty region: measurement noise, modelling error, unresolved forces, and numerical truncation. Because the dynamics is strongly sensitive to small perturbations, this uncertainty region does not remain small when transported along the flow; it stretches, folds, and threads itself through phase space in a way that makes *sharp* histories meaningless beyond a short horizon.

The blur viewpoint says: do not fight this, build it into the definition of the object you are studying. Fix from the start a *precise blur scale* in parameter and state space—that is, a region of uncertainty that you accept as constitutive. Instead of one exact trajectory, consider the entire cloud of trajectories compatible with that blur. Our analysis shows that, at any fixed blur scale and any fixed finite time horizon, this cloud breaks into only finitely many *blurred families of histories*: equivalence classes of solutions that are indistinguishable within the accepted error for all observables we are willing to pay for. Inside each family, the system still has genuine microscopic freedom, but that freedom no longer propagates into new macroscopic scenarios; it only changes which representative of the same blurred history we see.

In this sense, blur trades microscopic uniqueness for macroscopic finiteness. The infinite history of the system is not a *single fragile curve* in phase space but a compatible sequence of blurred histories, one for each time window and each resolution. Once the blur scale is honestly declared, the usual "pictures" of the evolution (configurations, energy exchanges, qualitative regimes) become stable objects: they are read as invariants of these blurred families rather than as shadows of an unattainable exact trajectory. Far from being a weakness, this is precisely the level at which the three–body problem is observable in the real world.

Principle 5.3 (Conservation of blur budget)**.** Fix a system with flow Φ^t on a state space X, and a blur scale r that encodes an accepted initial uncertainty (a small, honest chunk of missing information in parameter–state space). Suppose this uncertainty is modelled by a blur operator B_r that is transported covariantly along the dynamics,

$$\Phi^t_* \circ \mathsf{B}_r \;\approx\; \mathsf{B}_r \circ \Phi^t_*,$$

so that we always look at the system through the *same* informational lens. Then the initial error does not blow up in the sense of readability: the missing information stays of order r for all times. What grows is the *internal complexity* inside each blurred cell, not the blur budget itself.

In more concrete terms: an initial error that is treated as a genuine reduction of information inside a proper information field does not, by itself, make the system unreadable. What fails is the Platonic ideal of zero error. If we insist on tracking a single sharp trajectory, then even a tiny unknown fragment explodes under a chaotic flow and the description becomes useless. But if we *ride together with the error*—fix from the start a legitimate blur scale and propagate it with the dynamics—then the informational gap remains bounded: 0.1 bits of missing information remain 0.1 bits of missing information, only rearranged.

The usual mistake is to take the first few time steps, where the system is almost sharp, as the reference standard, and then complain when that level of precision is lost as the unknown fragment starts to affect everything else. From the blur point of view, the correct move is the opposite: accept the uncertainty at $t = 0$ as structural, design the dynamics and observables so that blur is pushed exactly where blur belongs, and then read the system consistently at that scale. The price for reducing blur is that more and more distinct histories appear (eventually too many for any computational method), but for any fixed, tame blur scale we retain a finite collection of macroscopic outcomes, each no more than the accepted error away from the "true" history. Classical, pointlike determinism is recovered only as the singular limit where the blur scale is driven to zero, and it is precisely in that limit that the description becomes unstable.

6 Blur with an Observer: Finite Observable Histories

So far we have blurred the three–body problem at the level of phase space and parameters, without explicitly modelling *who* is looking. In a Blurrichevsky spirit, we now insert an observer and reinterpret the finite blur–families as finite sets of *observable histories*.

6.1 Observer, measurement map, and observation metric

Let O denote an observer equipped with an instrument model. We encode this as a measurement map

$$M_O : \mathsf{P} \times \mathsf{X} \longrightarrow Y,$$

where Y is a (finite–dimensional) data space; typical examples are:

- Y a space of apparent positions and radial velocities (astrometric + spectroscopic data);
- Y a space of light curves or time–dependent fluxes (photometric data);
- Y a vector of derived quantities (pairwise distances, angles).

The map $(p, z) \mapsto M_O(p, z)$ includes the observer's location, line of sight, projection effects, and instrument response.

We equip Y with a norm $\|\cdot\|_Y$ (e.g. Euclidean on a finite vector of observables) and define the *observation metric* on $\mathsf{P} \times \mathsf{X}$ by

$$d_O\big((p, z), (p', z')\big) := \big\|M_O(p, z) - M_O(p', z')\big\|_Y.$$

Thus two parameter–state pairs are close for O if and only if they produce nearly the same raw data at a single instant.

We also allow blur at the level of the data: for a blur scale $\rho > 0$ on Y we let

$$\mathsf{B}_\rho^{(Y)} : Y \to Y$$

be a linear operator with $\|\mathsf{B}_\rho^{(Y)}\| \le 1$ and $\mathsf{B}_\rho^{(Y)} \to \mathrm{Id}_Y$ as $\rho \downarrow 0$ (e.g. a convolution with a narrow kernel, or a projection to a coarse grid in data space). Composing with our earlier blur on parameters and state we obtain a *total observed blur*

$$\mathcal{B}^O_{\sigma,r,\rho}(p, z) := \mathsf{B}_\rho^{(Y)}\Big(M_O\big(\mathsf{B}_\sigma^{(\mathsf{P})} p,\ \mathsf{B}_r z\big)\Big).$$

6.2 Observable blur–distance and finite observable histories

For a fixed observer O, blur scales (σ, r, ρ), and time horizon $T > 0$, we define the *observable blur–distance* between two parameter–trajectory pairs $(p, \Phi^p_\cdot(z_0))$ and $(p', \Phi^{p'}_\cdot(z_0'))$ as

$$d^O_{\sigma,r,\rho;T}\big((p, z_0), (p', z_0')\big) := \sup_{0 \le t \le T} \Big\|\mathcal{B}^O_{\sigma,r,\rho}(p, \Phi^p_t(z_0)) - \mathcal{B}^O_{\sigma,r,\rho}(p', \Phi^{p'}_t(z_0'))\Big\|_Y.$$

Definition 6.1 (($(\sigma, r, \rho, \varepsilon, T)$–observable blur–family)**.** Let $K \subset \Sigma_{E,R,\delta}$ be compact and P a compact parameter box as before. A finite set of representatives

$$\{(p_j, z_j)\}_{j=1}^{N} \subset \mathsf{P} \times K$$

is called an $(\sigma, r, \rho, \varepsilon, T)$*–observable blur–family for the observer* O if for every $(p, z_0) \in \mathsf{P} \times K$ there exists j such that

$$d^{O}_{\sigma,r,\rho;T}\big((p, z_0), (p_j, z_j)\big) \le \varepsilon.$$

The time series

$$t \longmapsto \mathcal{B}^{O}_{\sigma,r,\rho}(p_j, \Phi^{p_j}_t(z_j)) \ \in Y$$

are then called the *observable blur–family representatives* for O on $[0, T]$.

In words: from the point of view of O and at the given blur scales, *every* three–body configuration in $\mathsf{P} \times K$ produces a data stream that is indistinguishable (within ε) from one of these finitely many representative data streams.

Proposition 6.2 (Finite observable histories for a fixed observer)**.** *Under the assumptions of Theorem 4.4 and with O and M_O as above, for any blur scales $\sigma, r, \rho > 0$, tolerance $\varepsilon > 0$, and time horizon $T \le T_*$ there exists a finite $(\sigma, r, \rho, \varepsilon, T)$–observable blur–family for O.*

Equivalently: at that resolution, O can see only finitely many distinct blurred histories of the three–body system on $[0, T]$.

Proof. The map

$$(p, z_0, t) \longmapsto \mathcal{B}^{O}_{\sigma,r,\rho}(p, \Phi^{p}_t(z_0))$$

is continuous on the compact set $\mathsf{P} \times K \times [0, T]$, hence uniformly continuous. The proof of Theorem 4.4 applies verbatim with Y in place of $\mathsf{P} \times \mathsf{X}$, and we obtain a finite ε–net in the sup–norm on $[0, T]$ of blurred data streams. Each net point can be realized by some (p_j, z_j), giving the desired observable blur–family. □

From the observer's perspective, this is the natural strengthening of Theorem 4.4: instead of finitely many *abstract* blur–families in parameter–state space, we have finitely many *observable* histories at any fixed resolution.

6.3 Insolvability and the role of irreducible blur

The classical "insolvability" of the three–body problem is often phrased as follows: exact long–time prediction is impossible in practice because of extreme sensitivity to initial conditions and parameters. The blur/observer formalism clarifies that this is not just a technical nuisance but a structural consequence of the irreducible blur in any real observation.

Even if the underlying Newtonian system is perfectly deterministic, an observer O only ever has access to $M_O(p, z)$ up to some finite blur scale (σ, r, ρ):

- instrument noise and finite resolution impose a lower bound on ρ in data space;
- limited knowledge of masses and the coupling constant imposes a lower bound on σ in parameter space;
- finite spatial and velocity resolution impose a lower bound on r in phase space.

At those nonzero blur scales, Proposition 6.2 says that there are only finitely many observable histories on any finite time window. Inverting the data means choosing one branch among these finitely many.

Where, then, does unpredictability hide? Not in the finite catalogue itself, but in the *switching* between its elements. The map that assigns to each microscopic state (p, z_0) an observable branch index j is typically organised along fractal boundaries in parameter–state space. Arbitrarily small changes in (p, z_0) can move the system across such a boundary and thus select a different branch, even though the underlying blur scales (σ, r, ρ) are the same. If, in addition, the observer effectively changes their blur (for instance by reading different time windows at slightly different resolutions), the sequence of branch indices can look statistically random: the system appears to jump erratically between coarse states, not because new states are being created, but because the selection mechanism is extremely sensitive. This is precisely where the "small change in the beginning, huge difference at the end" lives. The practical cure is not to chase zero blur, but to keep blur *tame and consistent*: choose a reasonable, slightly conservative blur scale from the start and stick to it, so that successive observations stay in the same blur–class rather than wandering across fractal boundaries.

The paradox appears only if one insists on sending (σ, r, ρ) all the way to 0. In that singular limit, the number of observable branches explodes, reflecting the well–known fact that infinitesimal uncertainties in initial conditions and parameters lead to macroscopically different evolutions. From the blur point of view, this explosion is precisely what marks the limit as unphysical: it corresponds to an observer with infinite information budget and infinite sensitivity, which no real experiment supplies.

Thus the blur/observer picture separates two statements:

- *Mathematically*, the three–body flow is well–posed and uniquely determined by exact initial data and parameters.
- *Operationally*, any physically meaningful description lives at nonzero blur scales, and at each such scale there are only finitely many observable histories. The "insolvability" lies in the impossibility of taking the blur to zero without losing stability of the description, not in any defect of the dynamics itself.

From this angle, blur is not a cosmetic layer on top of an already understood problem; it is the natural level at which the three–body problem becomes *readable*. The finite observable blur–families are the correct objects to speak about when we discuss what an observer can know about a chaotic gravitational system.

7 Discussion and outlook

The results in this note are intentionally elementary on the analysis side: the finiteness of blur–families and blur–stable regimes derives from three ingredients:

(i) local Lipschitz continuity of the vector field on non–collision compact regions;

(ii) continuous dependence of solutions on initial data and parameters;

(iii) compactness of the region under consideration.

What is new is the systematic framing in terms of *blur*:

- phase–space blur B_r as the observer's finite resolution;
- parameter blur $\mathsf{B}_\sigma^{(\mathsf{P})}$ as finite knowledge of masses and coupling;
- blur–families as the minimal finite set of effective trajectories at that resolution;
- blur–stable regimes as coarse orbit types that do not change under blur.

This mirrors the pattern in other contexts where *finite certificates under blur* appear. For instance:

- In Lyapunov proofs for Collatz–type maps, one constructs a finite certificate (a function on residues modulo 2^k) that controls the drift for all larger scales; reducing blur (increasing k) refines the certificate but keeps it finite.
- Here, for the three–body problem, one obtains finite blur–families in the joint parameter–state space; reducing blur refines the families but does not change the basic theorem: *finiteness at each scale*.

A natural next step (beyond the scope of this note) is to formulate a general meta–theorem for dynamical systems:

> Given a well–posed dynamical system with suitable continuity in initial data and parameters, any compact, non–singular region admits only finitely many blur–families at each fixed blur scale in state and parameter space. Moreover, any coarse property that is locally structurally stable induces finitely many blur–stable regimes at that scale.

The three–body case then becomes a worked example of this meta–theory in a nontrivial but still classical setting.

Acknowledgments. This note is part of a broader programme on blur as a mathematical object and as an organizing principle: from dynamical systems to number theory and cosmology. Any coherence it has comes from that underlying view; any gaps are the author's.

References

[1] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, Springer, 1978.

[2] C. Simo, Dynamical properties of the restricted three–body problem, in *Hamiltonian Systems with Three or More Degrees of Freedom*, Kluwer, 1999.

[3] A. Perišić, *A Category of Blur and the Grand Lemma*, preprint, 2025.

[4] A. Perišić, *Blur as a Universal Principle*, preprint, 2025.