Numerics and validation of KAM tori

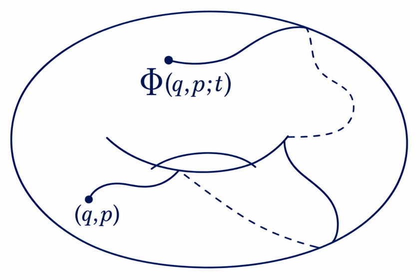

Jordi-Lluís Figueras,
Universitat Politècnica de Catalunya (UPC),
Barcelona, Spain.

Partial support from the Serra Húnter Senior Programme of the Generalitat de Catalunya and VR Grant 2024-04764.

XXVII International Workshop for Young Mathematicians
“Dynamical Systems” 20–26 Sept 2026
Jagiellonian University, Kraków

Thank you to the organizers, and especially to Paweł!

Returning to Kraków: a personal perspective

My first visit to Kraków was in **2011**, and I have been coming back quite regularly ever since.

What really impresses me on returning is **just how much Poland has grown** — especially compared with the growth I have seen in Spain.

I lived in **Sweden for 15 years**, so I have also watched these changes from a **Swedish perspective**.

That personal impression made me curious: what do the numbers show?

Poland has converged dramatically toward Sweden

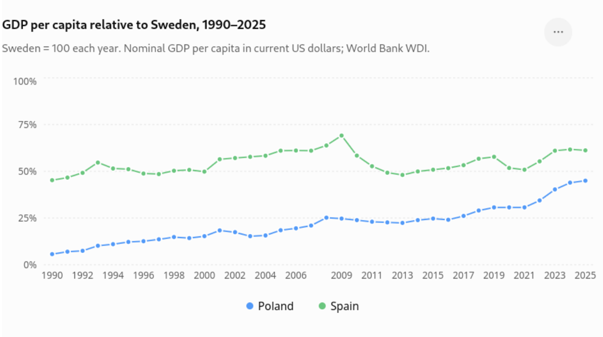

- **Sweden = 100** each year: a benchmark for relative economic performance.
- Poland: about **6% of Sweden's level in 1990**, rising to **45% in 2025** (nominal US dollars).
- Spain started much closer: **45% in 1990**, reaching around **61% in 2025**.

$$\text{Poland/Sweden: } 6\% \longrightarrow 45\% \qquad \text{Spain/Sweden: } 45\% \longrightarrow 61\%.$$

Poland has not just grown: it has **closed a large part of the gap with Western Europe**.

The Polish trajectory shows much stronger long-run convergence.

Poland is catching Spain much faster than Spain is catching Sweden

- Spain's position relative to Sweden has improved only moderately, fluctuating around roughly **50–60%** of the Swedish level.
- Poland, in contrast, has moved steadily upward:

$$\underbrace{15\%}_{2000} \longrightarrow \underbrace{24\%}_{2010} \longrightarrow \underbrace{31\%}_{2020} \longrightarrow \underbrace{45\%}_{2025} \quad \text{of Sweden's level.}$$

As a consequence, Poland has also moved much closer to Spain:

$$\frac{\text{Poland}}{\text{Spain}} \approx 0.13 \quad \text{in 1990,} \qquad \frac{\text{Poland}}{\text{Spain}} \approx 0.74 \quad \text{in 2025.}$$

This is consistent with Poland's broader structural transformation: **EU integration, industrial upgrading, foreign investment**, growth of business services, and a rapidly expanding **technology and software sector**.

Caveat: These figures use nominal GDP per capita in current US dollars, so exchange rates matter. PPP-adjusted GDP per capita shows an even stronger Polish convergence in terms of domestic purchasing power.

You are the future

To the Polish students: keep pushing! Your ideas, ambition, and hard work are helping Poland reach new heights.

To those from abroad: do not fall asleep — Poland is rocketing ahead! Let that energy inspire you to aim higher, too.

Learn from each other, challenge each other, and build the future together.

Future!?

(Illustration generated with ChatGPT)

Goals of this course

1. Recall what a **Hamiltonian system** is.
2. Understand what a **(Lagrangian) invariant torus** is and what role it plays in the dynamics.
3. Learn how to **numerically compute** these tori.
4. Learn how to **validate their existence** by combining mathematical theorems with computer-assisted proofs.

Our systems: Hamiltonian mechanics

Newtonian mechanics (a simple pendulum)

Newtonian mechanics describes physical systems by **second-order ordinary differential equations** given by Newton's laws. In particular,

$$\underbrace{m\ddot{\boldsymbol{r}}}_{\text{mass}\times\text{acceleration}} = \underbrace{\sum \boldsymbol{F}}_{\text{sum of forces}}.$$

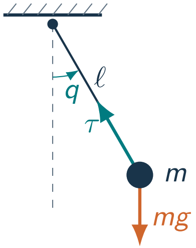

Fixed pivot, massless rod of length ℓ, no friction.

$$\boldsymbol{r}(q) = (\ell \sin q, -\ell \cos q).$$

$$m\ddot{\boldsymbol{r}} = -mg\,\boldsymbol{e}_y - \tau \frac{\boldsymbol{r}}{\ell}.$$

Here $\boldsymbol{e}_y$ points upward and τ is the rod tension.

Projecting onto the tangent eliminates the tension:

$$m\ell\ddot{q} = -mg \sin q \quad \Longrightarrow \quad \boxed{\ddot{q} + \frac{g}{\ell}\sin q = 0.}$$

Lagrangian mechanics (simple pendulum)

Instead of balancing forces and accelerations, we reinterpret conservative Newtonian mechanics as making the **action stationary** (a minimum in suitable cases).

For fixed endpoints and $L = T - V$, the action is

$$S[q] = \int_{t_0}^{t_1} L(q(t), \dot{q}(t), t)\, \mathrm{d}t.$$

For the simple pendulum,

$$S[q] = \int_{t_0}^{t_1} \left[\frac{1}{2} m\ell^2 \dot{q}^2 - mg\ell(1 - \cos q)\right] \mathrm{d}t.$$

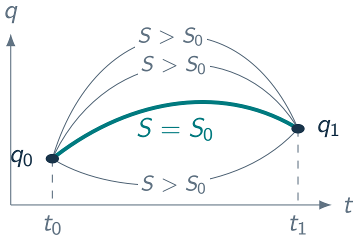

$$q(t_0) = q_0, \quad q(t_1) = q_1.$$

An extremal path satisfies the **Euler–Lagrange equation**:

$$\delta S[q] = 0 \quad \Longrightarrow \quad \frac{\mathrm{d}}{\mathrm{d}t} \partial_{\dot{q}} L - \partial_q L = 0 \quad \Longrightarrow \quad \boxed{m\ell^2 \ddot{q} + mg\ell \sin q = 0.}$$

Hamiltonian mechanics (simple pendulum)

We start with a physical fact: the **total energy** $H = T + V$ **is conserved**.

$$H = T + V = \frac{1}{2} m\ell^2 \dot{q}^2 + mg\ell(1 - \cos q).$$

Using the **conjugate momentum** $p = m\ell^2 \dot{q}$ we get

$$H(q,p) = \frac{p^2}{2m\ell^2} + mg\ell(1 - \cos q).$$

Newton's equation becomes a **first-order system**: **Hamilton's equations**.

$$\boxed{\begin{aligned} \dot{q} &= \quad \partial_p H = \frac{p}{m\ell^2}, \\ \dot{p} &= -\partial_q H = -mg\ell \sin q. \end{aligned}}$$

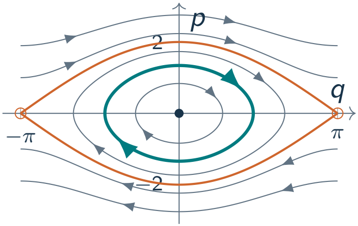

Energy contours $H(q,p) = E$; $m = \ell = g = 1$.
Separatrix: $H = 2$.

One scalar function H determines/generates the entire dynamics.

Hamiltonian mechanics (simple pendulum)

Preserving Hamilton's equations

We value this structure: a single function H encodes the whole vector field,

$$X_H(q,p) = \begin{pmatrix} \partial_p H \\ -\partial_q H \end{pmatrix}.$$

Coordinates in which Hamilton's equations take this form for every Hamiltonian are called **canonical coordinates**.
We therefore restrict attention to smooth, invertible changes of coordinates $(Q,P) = \Phi(q,p)$ that **preserve the form of Hamilton's equations**:

$$\widetilde{H} = H \circ \Phi^{-1}, \qquad \boxed{\dot{Q} = \partial_P \widetilde{H}, \qquad \dot{P} = -\partial_Q \widetilde{H}.}$$

Time-independent changes that preserve this structure for every Hamiltonian are called **canonical transformations**.

The expression for H changes; the rule generating motion stays the same.

Hamiltonian mechanics (simple pendulum)

The geometric viewpoint: symplectic geometry

We study phase space through the transformations that preserve the Hamiltonian structure. For the pendulum, the **symplectic form** is its oriented area form:

$$\Omega = \mathrm{d}q \wedge \mathrm{d}p, \qquad \iota_{X_H}\Omega = \mathrm{d}H.$$

This encodes Hamilton's equations. Canonical transformations preserve Ω:

$$\boxed{\Phi^*\Omega = \Omega} \quad \Longleftrightarrow \quad \mathrm{d}Q \wedge \mathrm{d}P = \mathrm{d}q \wedge \mathrm{d}p.$$

These transformations form a group: compositions and inverses preserve Ω. The same geometric structure extends to n degrees of freedom:

$$\Omega = \sum_{j=1}^{n} \mathrm{d}q_j \wedge \mathrm{d}p_j.$$

Hamiltonian mechanics combines a scalar function H with symplectic geometry.

Three descriptions of the same motion

Formulation	Starting point	Pendulum equations
Newtonian	Forces and accelerations	$m\ddot{\boldsymbol{r}} = \sum \boldsymbol{F}$ $m\ell^2\ddot{q} = -mg\ell \sin q$
Lagrangian	$S[q] = \int_{t_0}^{t_1} (T - V)\,\mathrm{d}t$ Fixed endpoints	$\delta S[q] = 0$ (stationary action) $\frac{\mathrm{d}}{\mathrm{d}t}\partial_{\dot{q}} L = \partial_q L$
Hamiltonian	Conserved energy $H = T + V$ Momentum $p = m\ell^2\dot{q}$	$\dot{q} = \partial_p H, \quad \dot{p} = -\partial_q H$

All three give the same pendulum dynamics: $\boxed{\ddot{q} + (g/\ell)\sin q = 0.}$

Canonical changes preserve Hamiltonian form and $\Omega = \mathrm{d}q \wedge \mathrm{d}p$.

An invitation to explore mechanics

These few slides are only a **first glimpse**: they cannot do justice to the depth of Newtonian, Lagrangian and Hamiltonian mechanics.

Each viewpoint is **fascinating and rewarding to study**, revealing connections between physics, geometry and dynamical systems.

There is much more to discover: variational principles, symmetries and conservation laws, symplectic geometry, and celestial mechanics, among many other topics.

I encourage you to study these subjects in depth: read, work through examples, and explore the connections between the three viewpoints.

Astronomer Copernicus, or Conversations with God
Jan Matejko. Source: Wikimedia.

Dive into mechanics, and you may be as awestruck as Copernicus!

Suggested reading

Main references

- **V. I. Arnold**, *Mathematical Methods of Classical Mechanics*.
- **A. Giorgilli**, *Meccanica Celeste*.

Further reading (with a strong emphasis on celestial mechanics)

- **C. L. Siegel and J. K. Moser**, *Lectures on Celestial Mechanics*.
- **A. Morbidelli**, *Modern Celestial Mechanics: Aspects of Solar System Dynamics*.
- **R. Moeckel**, *Celestial Mechanics* (lecture notes).

Newton, Lagrange and Hamilton discussing mechanics. (Illustration generated with ChatGPT)

More degrees of freedom: coupled pendula and their dynamics

Coupled pendula: the mechanical model

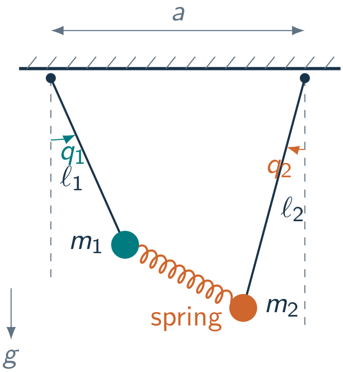

Two independent angles: $q = (q_1, q_2) \in \mathbb{T}^2$.

Fixed pivots $\boldsymbol{a}_1 = (0,0)$, $\boldsymbol{a}_2 = (a, 0)$.

$$\boldsymbol{r}_j(q_j) = \boldsymbol{a}_j + \ell_j \begin{pmatrix} \sin q_j \\ -\cos q_j \end{pmatrix}.$$

- Masses m_1, m_2; massless rods of lengths ℓ_1, ℓ_2.
- A spring joins the bobs, with stiffness εk, $\varepsilon \geq 0$.
- Uniform gravity; no friction.

Two degrees of freedom: each pendulum has its own angle and momentum.

Coupled pendula: the coupling potential

The spring length depends on both angles:

$$d(q) = |\boldsymbol{r}_2(q_2) - \boldsymbol{r}_1(q_1)|,$$
$$d(q)^2 = (a + \ell_2 \sin q_2 - \ell_1 \sin q_1)^2 + (\ell_1 \cos q_1 - \ell_2 \cos q_2)^2.$$

Choose the natural length so that the spring is unstretched at $q_1 = q_2 = 0$:

$$d_0 = d(0,0) = \sqrt{a^2 + (\ell_1 - \ell_2)^2}.$$

Hooke's law gives the coupling energy

$$\boxed{U(q) = \frac{k}{2}\left(d(q) - d_0\right)^2, \qquad V_{\text{spring}}(q) = \varepsilon U(q).}$$

We use $\ell_1 \neq \ell_2$ and work where $d(q) > 0$.

$\varepsilon = 0$: **independent pendula.** $0 < \varepsilon \ll 1$: **weak coupling.**

The Hamiltonian of the coupled pendula

Introduce the canonical momenta

$$p_j = m_j \ell_j^2 \dot{q}_j, \qquad j = 1, 2.$$

The **Hamiltonian is the total energy**: kinetic, gravitational, and elastic.

$$\boxed{H_\varepsilon(q,p) = \sum_{j=1}^{2} \left[\frac{p_j^2}{2 m_j \ell_j^2} + m_j g \ell_j (1 - \cos q_j) \right] + \varepsilon \frac{k}{2} \big(d(q) - d_0 \big)^2.}$$

Equivalently, separating the two one-pendulum Hamiltonians,

$$H_\varepsilon(q,p) = H_1(q_1,p_1) + H_2(q_2,p_2) + \varepsilon U(q).$$

The phase space is four-dimensional, restricted to $d(q) > 0$:

$$(q,p) = (q_1, q_2, p_1, p_2) \in T^*\mathbb{T}^2 \simeq \mathbb{T}^2 \times \mathbb{R}^2.$$

A weakly coupled system is a perturbation of two independent pendula.

Coupled pendula: Hamilton's equations

Differentiate H_ε to obtain **four first-order ODEs**:

$$\boxed{\begin{aligned} \dot q_j &= \quad \partial_{p_j} H_\varepsilon = \frac{p_j}{m_j \ell_j^2}, \\ \dot p_j &= -\partial_{q_j} H_\varepsilon = -m_j g \ell_j \sin q_j - \varepsilon \partial_{q_j} U(q), \qquad j = 1, 2. \end{aligned}}$$

The coupling enters through the momentum equations, with

$$\partial_{q_j} U(q) = k\big(d(q) - d_0\big)\partial_{q_j} d(q).$$

The individual pendulum energies can change:

$$\frac{\mathrm{d}H_j}{\mathrm{d}t} = -\varepsilon \dot q_j \partial_{q_j} U, \qquad \frac{\mathrm{d}H_\varepsilon}{\mathrm{d}t} = 0.$$

The pendula exchange energy through the spring; total energy is conserved.

Coupled pendula: interactive simulation

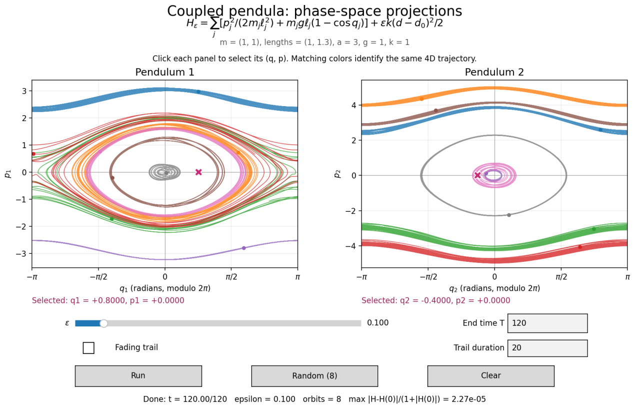

▸ Run the coupled-pendula simulation

(Launch slides with `python3 present.py`.)

One Hamiltonian, many kinds of motion

For positive coupling $\varepsilon > 0$, the coupled pendula can display a rich variety of orbits.

- **Periodic:** the motion repeats after a fixed time.
- **Chaotic:** non-repeating motion with sensitive dependence on the initial condition.
- **Quasiperiodic:** several incommensurate frequencies combine into ordered, non-repeating motion.

The behaviour depends on the **initial condition**, the energy, and the coupling strength.

Different kinds of motion can coexist in the same Hamiltonian system.

Periodic orbits: returning and repeating

Write the full phase-space state as $z(t) = (q_1(t), q_2(t), p_1(t), p_2(t))$.

A nonconstant orbit is **periodic** if there is a time $T > 0$ such that

$$\boxed{z(t+T) = z(t) \qquad \text{for every } t.}$$

The smallest such positive time is its **period**.

- Start at any point on the orbit: after one period, you return to that same point.
- The motion then repeats, again and again.
- Both angles (q_j) **and** both momenta (p_j) must return (angles are understood modulo 2π).

A periodic orbit is a closed curve in phase space.

Chaotic orbits: sensitivity to initial conditions

A chaotic trajectory does **not repeat periodically** and exhibits **sensitive dependence on its initial condition**.

- Suppose we want to start at z_0, but instead start at a very nearby point $\widetilde{z}_0$.
- Even a tiny initial error can grow: the two trajectories can eventually follow very different paths.
- This limits long-time prediction, even though the equations are deterministic and the total energy is conserved.

Non-repetition alone is not enough to identify chaos.

Our focus: quasiperiodic orbits and invariant tori

In this course we are interested in **quasiperiodic orbits** and the **invariant tori** on which they lie.

Our goal is to **compute these tori numerically** and **validate their existence**.

Before computing these tori, let us first understand quasiperiodic motion.

The uncoupled case

The uncoupled case: two independent pendula

When $\varepsilon = 0$ our Hamiltonian is of the form

$$H_0 = H_1(q_1, p_1) + H_2(q_2, p_2),$$

and the two individual energies are conserved.

$$H_1 = E_1, \quad H_2 = E_2.$$

Choose a nonconstant periodic motion for each pendulum, away from the separatrices.

- Pendulum 1 travels around a closed phase-space curve C_1 with period T_1.
- Pendulum 2 travels around a closed phase-space curve C_2 with period T_2.
- The full state records where **both** pendula are on their curves.

Will the two pendula return to their starting states at the same time?

Two librating pendula: the curves C_1 and C_2

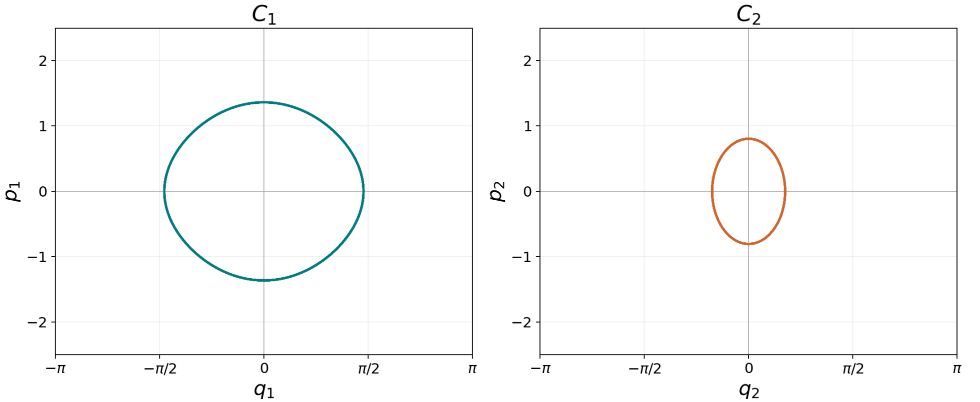

The periods of C_1 and C_2 are T_1 and T_2, respectively.

Commensurate periods: periodic motion

The periods are **commensurate** if their ratio is rational:

$$\frac{T_1}{T_2} = \frac{m}{n}, \qquad m, n \in \mathbb{N}.$$

After a common time, both pendula have completed a whole number of cycles:

$$\boxed{T = nT_1 = mT_2 \quad \Longrightarrow \quad z(t+T) = z(t).}$$

For example, if $T_1/T_2 = 2/3$, pendulum 1 completes three cycles while pendulum 2 completes two.

Commensurate individual periods give a periodic orbit of the full system.

Incommensurate periods: quasiperiodic motion

The periods are **incommensurate** if their ratio is irrational:

$$\boxed{\frac{T_1}{T_2} \notin \mathbb{Q}.}$$

There is no positive time that is a whole number of **both** periods.

- Each pendulum repeats its own motion, but the full state **never repeats exactly**.
- The two phases nevertheless return arbitrarily close to their initial values simultaneously.
- This is **quasiperiodic motion**, with two frequencies $\omega_j = 2\pi / T_j$.

Non-repeating motion can be ordered: quasiperiodicity is not chaos.

Both kinds of motion live on two-dimensional tori

Allowing both initial phases to vary gives the product of the two closed curves:

$$\boxed{C_1 \times C_2 \simeq S^1 \times S^1 = \mathbb{T}^2.}$$

This is a **two-dimensional invariant torus** in the **four-dimensional phase space** (q_1, q_2, p_1, p_2).

- **Commensurate periods:** the torus is filled by a family of periodic orbits. Each individual orbit is a closed curve.

- **Incommensurate periods:** each individual orbit is dense on the torus; its closure is the whole torus.

The same torus geometry supports periodic or quasiperiodic motion.

Maximal tori for two degrees of freedom

⚠

For readers already familiar with mechanics or symplectic geometry

These product tori are **Lagrangian**: the symplectic form vanishes when restricted to their tangent directions,

$$\Omega|_{C_1 \times C_2} = 0, \qquad \Omega = \mathrm{d}q_1 \wedge \mathrm{d}p_1 + \mathrm{d}q_2 \wedge \mathrm{d}p_2.$$

In a $2n$-dimensional symplectic phase space, an isotropic submanifold has dimension at most n.

In our example of two coupled pendula there are two degrees of freedom:

$$\text{dim(phase space)} = 4, \qquad \boxed{\text{dim(Lagrangian torus)} = 2.}$$

Thus these are **maximal-dimensional isotropic tori**.

Our goal: compute and validate quasiperiodic tori when $\varepsilon > 0$.

Perturbing the periodic tori

A torus filled with periodic orbits is resonant

At $\varepsilon = 0$, commensurate periods satisfy

$$nT_1 = mT_2, \qquad m, n \in \mathbb{N}.$$

With $\omega_j = 2\pi/T_j$, this is equivalent to

$$\boxed{m\omega_1 - n\omega_2 = 0.}$$

This integer relation is called a **resonance**.

The two pendula return simultaneously: their **relative timing repeats**.

Any choice of initial phases gives periodic motion, so the torus is filled with a **continuous family of closed orbits**.

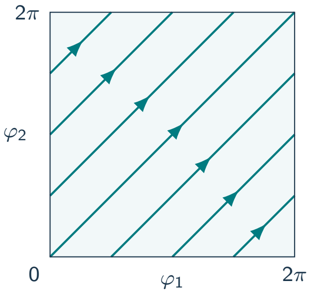

A 1 : 1 example in phase coordinates. Opposite edges of the square are identified: the lines continue into closed orbits.

The resonance is already present before the coupling is switched on.

Resonant coupling: repeated pushes

For $\varepsilon > 0$ each pendulum now pushes or pulls the other through the spring.

- At resonance, the interaction repeats with the same relative timing in the uncoupled motion.
- Energy transfers **can reinforce one another** instead of averaging out.
- Think of pushing a swing at the right moments: here the "pusher" is another moving pendulum.

$$\text{Pendulum 1} \longleftrightarrow \text{Spring} \longleftrightarrow \text{Pendulum 2}$$

$$\frac{\mathrm{d}H_j}{\mathrm{d}t} = -\varepsilon \dot{q}_j \partial_{q_j} U, \qquad \boxed{H_1 + H_2 + \varepsilon U = \text{constant.}}$$

The direction and strength of the transfer depend on the **relative phase**.

A weak coupling can have an accumulating effect at resonance.

Jordi-Lluís Figueras

A swing: well-timed pushes add energy

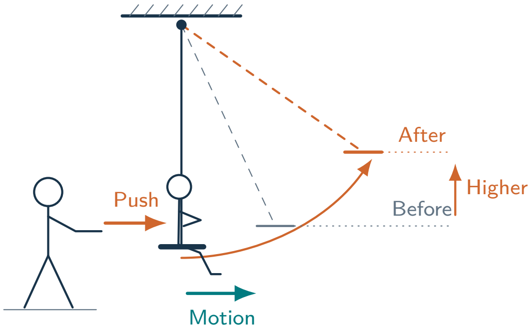

Push in the direction of motion:

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \boldsymbol{F}_{\mathrm{push}} \cdot \boldsymbol{v} > 0.$$

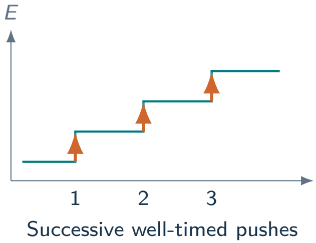

Repeat on each forward pass.

In-phase pushes add energy and increase the swing amplitude.

Here the person supplies energy; the coupled pendula exchange energy through the spring.

Energy transfer changes the timing

For a **nonlinear pendulum**, the period depends on its energy.

For librations, more energy means a larger amplitude and a longer period.

Energy exchange generally upsets the original simultaneous return of the two pendula.

Under generic weak coupling, most members of the original family **do not continue as nearby periodic orbits**.

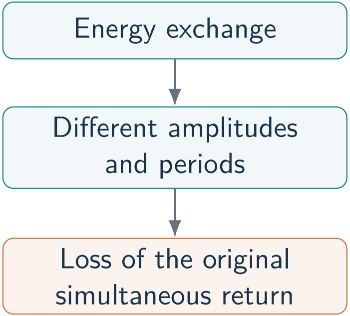

Energy transfer can later reverse; neither pendulum gains energy indefinitely.

The continuous family of closed orbits generally breaks up.

Which periodic orbits survive?

At certain **special relative timings**, the interaction can be compatible with a repeating motion.

- Energy can pass back and forth, but the **full state returns** after a cycle.
- Any surviving orbits are **deformed**: their shapes and periods adjust to the coupling.
- At a fixed energy, generic weak coupling leaves at most **isolated periodic orbits** from the original family.

Any surviving periodic motions are isolated.

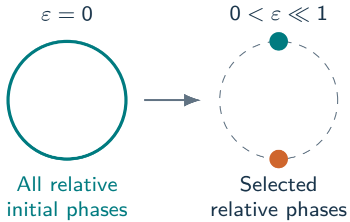

Each circle represents the choices of relative phase.
Schematic example with two survivors; their number depends on the coupling.

Proving that resonant tori break

Preparing the proof

Our goal is to **prove that resonant tori break under generic perturbations**.

Before doing so, we will make a **canonical change of variables** to write the Hamiltonian in a form better suited to the proof.

First choose suitable coordinates; then analyse the effect of the perturbation.

Canonical polar coordinates

Away from equilibria and separatrices, use **polar coordinates adapted to the periodic energy curves** $C_j(E_j)$ of each pendulum:

$$I_j(E_j) = \frac{1}{2\pi}\oint_{C_j(E_j)} p_j \,\mathrm{d}q_j, \qquad \varphi_j = \frac{2\pi\, t_j}{T_j(E_j)} \pmod{2\pi}.$$

The action I_j labels the orbit; the angle φ_j measures time from a chosen reference point on it.

These are **canonical action–angle coordinates** $(q_j, p_j) = \Phi_j(\varphi_j, I_j)$: the product map $\Phi = \Phi_1 \times \Phi_2$ preserves Ω.

Writing each pendulum energy as $H_j(I_j)$, the full Hamiltonian becomes

$$\boxed{H_\varepsilon(\varphi, I) = H_1(I_1) + H_2(I_2) + \varepsilon h(\varphi_1, \varphi_2, I_1, I_2).}$$

Here $h = U \circ \Phi$ is the spring coupling expressed in the new coordinates.

Actions, angles, and integrability

In canonical coordinates (φ, I), suppose $H_0 = H_0(I)$ depends only on **half of the variables**. The I_j are called **actions**; the conjugate periodic variables $\varphi_j \in \mathbb{T}$ are **angles**.

For $\varepsilon = 0$ the Hamiltonian is **integrable**: we can solve its dynamics explicitly. Hamilton's equations give

$$\dot{I}_j = -\partial_{\varphi_j} H_0 = 0, \qquad \dot{\varphi}_j = \partial_{I_j} H_0 = \omega_j(I).$$

Thus every trajectory is known:

$$\boxed{I(t) = I(0), \qquad \varphi(t) = \varphi(0) + t\,\omega(I(0)) \pmod{2\pi}.}$$

For our uncoupled pendula,

$$H_0(I) = H_1(I_1) + H_2(I_2), \qquad \omega_j(I) = H_j'(I_j).$$

Fixed actions, uniformly rotating angles: motion on invariant tori.

Homework: preserve the symplectic form

Prove that $\Phi : (\varphi, I) \mapsto (q, p)$ preserves the symplectic two-form:

$$\boxed{\Phi^*\left(\sum_{j=1}^{2} \mathrm{d}q_j \wedge \mathrm{d}p_j\right) = \sum_{j=1}^{2} \mathrm{d}\varphi_j \wedge \mathrm{d}I_j.}$$

1. **Warm-up: ordinary polar coordinates.** For a normalized oscillator, use
$$q = \sqrt{2I}\cos\varphi, \qquad p = -\sqrt{2I}\sin\varphi, \qquad I > 0.$$
Compute $\mathrm{d}q \wedge \mathrm{d}p$.

2. **Pendulum: energy-adapted coordinates.** Prove the same identity for the coordinates $I(E)$ and $\varphi = 2\pi t/T(E)$ defined above.

Hint: first show that $I'(E) = T(E)/(2\pi)$ and express Ω in the coordinates (t, E).

WARNING

We are now entering a theorem–proof environment.

Beware of the dangers:
hidden hypotheses and
hungry counterexamples!

Please keep your assumptions inside the vehicle at all times.

BEWARE OF THE COUNTEREXAMPLES

Small arms. Strong assumptions.

Theorem: generic breakup at resonance

Fix a regular energy of an integrable Hamiltonian with two degrees of freedom. Let $\mathcal{T}_0$ be a regular $n : m$ torus, with $nT_1 = mT_2$ and $\gcd(n, m) = 1$.

Generic breakup of the periodic family

For a **generic sufficiently small Hamiltonian perturbation**, the torus filled with $n : m$ periodic orbits does not persist as such. In a compact neighbourhood, only **finitely many isolated $n : m$ periodic orbits** can remain.

Most periodic orbits in the resonant family do not persist.

Counted at fixed energy, modulo time translation, and for this resonance only.

Proof 1/4: periodic orbits are fixed points

Fix $H_\varepsilon = E$ and observe the flow when the phase φ_2 returns to 0 modulo 2π. The **Poincaré section**

$$\Sigma_\varepsilon = \{H_\varepsilon = E,\ \varphi_2 = 0 \pmod{2\pi}\}$$

has two coordinates $z = (\varphi, I)$: a phase and a transverse action. Let P_ε be its first-return map.

Since $nT_1 = mT_2$, use m **returns** to close the resonant orbit. Locally in cylinder coordinates, write

$$\boxed{F_\varepsilon = P_\varepsilon^m : \mathbb{T} \times \mathbb{R} \longrightarrow \mathbb{T} \times \mathbb{R}.}$$

$$\text{Nearby } n{:}m \text{ periodic orbit} \quad \Longleftrightarrow \quad F_\varepsilon(z) = z.$$

One such orbit meets the section at m points, all fixed by F_ε.

Fixing energy and taking a section removes the continuous freedoms.

Proof 2/4: a whole curve of fixed points

For the integrable system,

$$P_0(\varphi, I) = (\varphi + 2\pi\rho(I), I),$$

where $\rho(I)$ is the ratio of the two frequencies. The resonant action I_* satisfies

$$\rho(I_*) = \frac{n}{m}.$$

After m returns, the phase advances by $2\pi n$:

$$\boxed{F_0(\varphi, I_*) = (\varphi, I_*) \quad \text{for every } \varphi \in \mathbb{T}.}$$

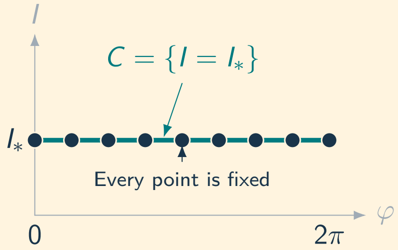

Identify $\varphi = 0$ and $\varphi = 2\pi$: the line is a circle.

The action-return equation is automatic: P_0 preserves I everywhere.

The fixed points are not isolated: they fill a whole circle.

Jordi-Lluís Figueras

Proof 3/4: two equations, two curves

In local coordinates $z = (x, y)$, choose the angular lift for this resonance and write $f(z) = F_\varepsilon(z) - z = (f_1(z), f_2(z))$.

$$F_\varepsilon(z) = z \quad \Longleftrightarrow \quad f_1(x, y) = 0 \quad \text{and} \quad f_2(x, y) = 0.$$

Each scalar zero set is locally a **curve** when its gradient is nonzero.

Generic perturbations make the curves meet **transversely** (transversality theorem):

$$\boxed{\det Df(z_*) = \det(DF_\varepsilon(z_*) - \mathrm{Id}) \neq 0.}$$

The inverse function theorem then makes each fixed point **isolated**.

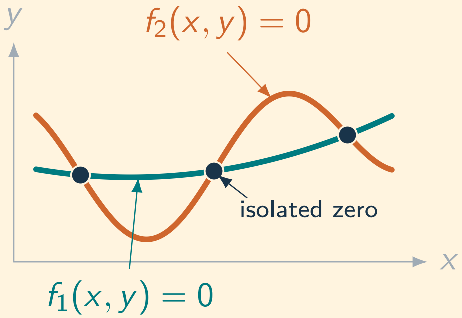

The dots solve both equations.

Proof 4/4: the periodic family breaks up

Why finitely many?

The fixed-point set is **closed**. In a compact annular neighbourhood, infinitely many fixed points would accumulate at another fixed point, contradicting isolation.

The original torus contributes a **whole circle of fixed points**. A finite set cannot carry this continuous family.

Each periodic orbit contributes m fixed points, so there are **at most finitely many** nearby $n : m$ periodic orbits.

Most periodic orbits of the resonant torus are destroyed. □

Resonance reorganizes the motion

Near a regular resonant torus, the generic weak-coupling picture is:

Before coupling: $\varepsilon = 0$

- Commensurate frequencies.
- A torus filled with periodic orbits.
- Any relative initial phase is allowed.

After weak coupling: $0 < \varepsilon \ll 1$

- The periodic family breaks up.
- Any surviving periodic orbits are isolated.
- The surrounding motion changes.

What happens to the tori whose frequencies are incommensurate?

Computing quasiperiodic tori

The Hamiltonians we consider

We consider an **integrable Hamiltonian plus a small perturbation**:

$$\boxed{H_\varepsilon(\varphi, I) = h(I) + \varepsilon f(\varphi, I), \qquad (\varphi, I) \in \mathbb{T}^n \times D, \quad D \subset \mathbb{R}^n.}$$

Here φ are angles, I are actions, and $|\varepsilon| \ll 1$.

For $\varepsilon = 0$, Hamilton's equations are

$$\dot{\varphi} = \nabla h(I), \qquad \dot{I} = 0.$$

Each torus $I = I_*$ is invariant and carries linear motion with frequency $\omega = \nabla h(I_*)$.

We seek quasiperiodic invariant tori when the perturbation is switched on.

The classical approach: Kolmogorov normal form

There are **several ways to find quasiperiodic invariant tori**. The classical **normal-form approach** constructs an **infinite sequence of canonical changes of variables**.

Under suitable nondegeneracy and smallness assumptions and a Diophantine condition on ω, they converge to **Kolmogorov normal form** in new variables (ψ, J):

$$\boxed{\widetilde{H}(\psi, J) = e + \omega \cdot J + R(\psi, J), \qquad R(\psi, J) = O(|J|^2).}$$

Thus $J = 0$ is an invariant torus with the prescribed frequency:

$$\dot{\psi} = \omega, \qquad \dot{J} = 0.$$

This approach began with **Kolmogorov (1954)** and was developed by **Arnold and Moser**, and many others: **KAM theory**.

Simplify the Hamiltonian by canonical transformations to reveal the torus.

Kolmogorov, Arnold and Moser

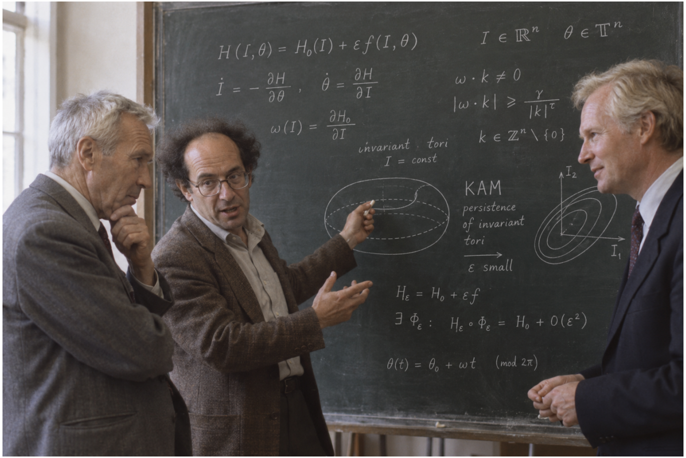

Illustration generated with ChatGPT

Our approach: the parameterization method

We will not follow the canonical-transformation approach. Instead, we use the **parameterization method**: solve directly for an embedding of the torus and its prescribed quasiperiodic dynamics.

- Introduced for invariant manifolds by **Cabré, Fontich and de la Llave** [1–3].
- Developed for **KAM tori** by **de la Llave, González, Jorba and Villanueva** [4].
- Many other researchers have since contributed to the field.

[1–3] X. Cabré, E. Fontich and R. de la Llave, *The parameterization method for invariant manifolds*:
[1] I. Manifolds associated to non-resonant subspaces. *Indiana Univ. Math. J.* **52** (2003), 283–328.
[2] II. Regularity with respect to parameters. *Indiana Univ. Math. J.* **52** (2003), 329–360.
[3] III. Overview and applications. *J. Differential Equations* **218** (2005), 444–515.

[4] R. de la Llave, A. González, À. Jorba and J. Villanueva, *KAM theory without action-angle variables*. *Nonlinearity* **18** (2005), 855–895.

Parameterizing an invariant torus

We seek a smooth **embedding**, called the **parameterization**,

$$K : \mathbb{T}^n \longrightarrow \mathbb{T}^n \times \mathbb{R}^n, \qquad \theta \longmapsto K(\theta) = \big(K_\varphi(\theta), K_I(\theta)\big),$$

whose image $\mathcal{T} = K(\mathbb{T}^n)$ is an invariant torus.

Invariance means that the Hamiltonian vector field is tangent to the torus at every point:

$$\boxed{X_H(K(\theta)) \in T_{K(\theta)}\mathcal{T} = \operatorname{Range} DK(\theta) \qquad \text{for every } \theta \in \mathbb{T}^n.}$$

Equivalently, there is a vector field v on $\mathbb{T}^n$ such that

$$X_H(K(\theta)) = DK(\theta)\, v(\theta).$$

The ambient vector field follows the torus instead of leaving it.

The geometry of the parameterization

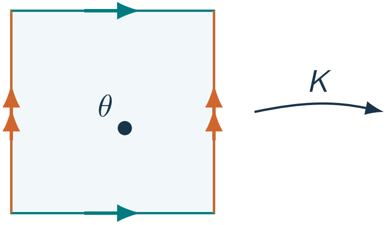

$\mathbb{T}^2$: opposite sides identified

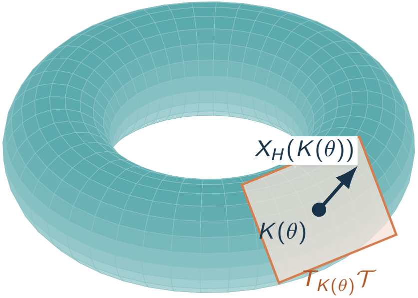

$\mathcal{T} = K(\mathbb{T}^2)$

$$X_H(K(\theta)) \in T_{K(\theta)}\mathcal{T} = \mathrm{Range}\, DK(\theta).$$

Schematic for $n = 2$: the embedded torus is drawn in three dimensions for visualization.

Pulling back the dynamics to the torus

The vector field $v(\theta)$ is the **pullback of the dynamics** of X_H restricted to $\mathcal{T} = K(\mathbb{T}^n)$:

$$X_H(K(\theta)) = DK(\theta)\, v(\theta).$$

Thus $\dot{\theta} = v(\theta)$ describes the motion in parameter coordinates.

To obtain **quasiperiodic internal dynamics**, we seek a parameterization for which $v(\theta) = \omega$ is constant:

$$\boxed{X_H(K(\theta)) = DK(\theta)\, \omega,} \qquad \theta(t) = \theta_0 + \omega t \pmod{2\pi}.$$

The components of ω must be **rationally independent**:

$$k \cdot \omega \neq 0 \qquad \text{for every } k \in \mathbb{Z}^n \setminus \{0\}.$$

For example, $(1, \sqrt{2}, \sqrt{3})$ is rationally independent, whereas $(1, \sqrt{2}, 2)$ is not: $2\omega_1 - \omega_3 = 0$.

The invariance equation

With our sign convention, define the **Lie derivative along** ω by

$$\mathcal{L}_\omega K(\theta) := -DK(\theta)\,\omega = -\sum_{j=1}^{n} \omega_j\, \partial_{\theta_j} K(\theta).$$

The **invariance equation** then reads

$$\boxed{\mathcal{F}_\omega(K) := \mathcal{L}_\omega K + X_H \circ K = 0.}$$

For a prescribed rationally independent frequency ω, our goal is to **find an embedding** K **satisfying this equation**.

In this course, we will learn how to

- **compute** an approximate solution numerically;
- **validate** the existence of a true solution nearby.

A Newton step for the invariance equation

Suppose $K_0(\theta)$ is an **approximate solution**, with residual

$$E(\theta) := \mathcal{F}_\omega(K_0)(\theta) = \mathcal{L}_\omega K_0(\theta) + X_H(K_0(\theta)).$$

We seek a correction $\Delta(\theta)$ and set $K_1 = K_0 + \Delta$. Expanding and retaining terms up to **first order** gives

$$\mathcal{F}_\omega(K_0 + \Delta) = E + \mathcal{L}_\omega \Delta + DX_H(K_0)\,\Delta + O(\Delta^2).$$

Cancel the constant and linear terms by solving the **Newton equation**:

$$\boxed{\mathcal{L}_\omega \Delta(\theta) + DX_H(K_0(\theta))\,\Delta(\theta) = -E(\theta).}$$

Then the new residual is **quadratic in the correction**:

$$E_1 := \mathcal{F}_\omega(K_1) = O(\Delta^2).$$

If the linear solve gives $\Delta = O(E)$ in suitable norms, then $E_1 = O(E^2)$: **quadratic error reduction**.

Computing the Newton correction

Our goal is to find $\Delta : \mathbb{T}^n \to \mathbb{R}^{2n}$ solving

$$\boxed{\mathcal{L}_\omega \Delta(\theta) + A(\theta)\,\Delta(\theta) = -E(\theta), \qquad A(\theta) = DX_H(K_0(\theta)).}$$

The difficulty: the matrix $A(\theta)$ **varies along the torus**.

The strategy: exploit the Hamiltonian geometry.

1. Express the correction in a frame adapted to the torus: $\Delta(\theta) = P(\theta)\,\xi(\theta)$.
2. Use the tangent directions and symplectic structure to simplify the linearized equation (**automatic reducibility**).
3. Solve the resulting **cohomological equations** using Fourier series, treating their averages separately.

Choose a geometric frame first; then solve for the correction.

Which change of frame should we choose?

We seek an invertible matrix $P(\theta)$ such that the substitution $\Delta(\theta) = P(\theta)\,\xi(\theta)$ transforms

$$\mathcal{L}_\omega \Delta + A(\theta)\Delta = -E$$

into a linear equation with a **simpler matrix structure** that makes the equations easier to solve. Using the product rule,

$$\mathcal{L}_\omega(P\xi) = (\mathcal{L}_\omega P)\xi + P\mathcal{L}_\omega \xi,$$

we obtain

$$\mathcal{L}_\omega \xi + B(\theta)\xi = -P(\theta)^{-1}E(\theta), \qquad B(\theta) = P^{-1}(\mathcal{L}_\omega P + AP).$$

Our target is therefore a suitable structured matrix $B(\theta)$ satisfying

$$\boxed{\mathcal{L}_\omega P(\theta) + A(\theta)P(\theta) = P(\theta)B(\theta).}$$

Which $P(\theta)$ can achieve this?

Reducibility

Building the frame $P(\theta)$

We now construct an invertible matrix

$$P(\theta) \in \mathbb{R}^{2n\times 2n}$$

adapted to the geometry of the torus, and write

$$\Delta(\theta) = P(\theta)\,\xi(\theta).$$

Our aim is to give the transformed linear equation a **structure that we can solve easily**.

We start with the tangent directions: the columns of $DK(\theta)$.

Tangent directions satisfy the linearized equation

First suppose K parameterizes an **exact invariant torus**:

$$\mathcal{L}_\omega K(\theta) + X_H(K(\theta)) = 0.$$

Differentiate with respect to θ. Since ω is constant, D commutes with $\mathcal{L}_\omega$, and the chain rule gives

$$\boxed{\mathcal{L}_\omega DK(\theta) + DX_H(K(\theta))\, DK(\theta) = 0.}$$

We will denote the tangent frame by

$$\boxed{L(\theta) := DK(\theta).}$$

Completing the frame: symplectic conjugates

Since K is an **embedding**, $L = DK$ supplies n independent tangent vectors: a basis of the tangent space to the torus, but only **half a basis of the ambient tangent space**.

We complete them with n transverse vectors, the columns of $N(\theta)$, chosen as **symplectic conjugates** of the columns of $L(\theta)$. Writing $\Omega(u, v) = u^T J v$, where

$$J = \begin{pmatrix} 0 & \mathrm{Id}_n \\ -\mathrm{Id}_n & 0 \end{pmatrix}, \qquad \boxed{\Omega(L_i, N_j) = \delta_{ij}, \quad \Omega(L_i, L_j) = \Omega(N_i, N_j) = 0,}$$

these relations mean that $P = (L\ N)$ is a **symplectic frame**: $P^T J P = J$.

For a Lagrangian torus ($L^T J L = 0$), a concrete choice is

$$\boxed{N(\theta) = -J\, L(\theta)\big(L(\theta)^T L(\theta)\big)^{-1}.}$$

The columns of N span the complementary bundle.

The magic: normal becomes tangent

For an **exact invariant Lagrangian torus**, we already know

$$(\mathcal{L}_\omega + DX_H \circ K)L = 0.$$

The symplectic structure gives the key identity:

$$\boxed{(\mathcal{L}_\omega + DX_H \circ K)N(\theta) = L(\theta)\, T(\theta), \qquad T(\theta)^T = T(\theta).}$$

The image of the normal frame lies **entirely in the tangent directions**! The matrix $T(\theta)$, called the **torsion matrix**, is

$$T = -N^T J(\mathcal{L}_\omega N + (DX_H \circ K)N).$$

Thus, for $P = (L\ N)$,

$$\boxed{\mathcal{L}_\omega P + (DX_H \circ K)P = P \begin{pmatrix} 0 & T(\theta) \\ 0 & 0 \end{pmatrix}.}$$

A variable matrix remains, but only in one off-diagonal block.

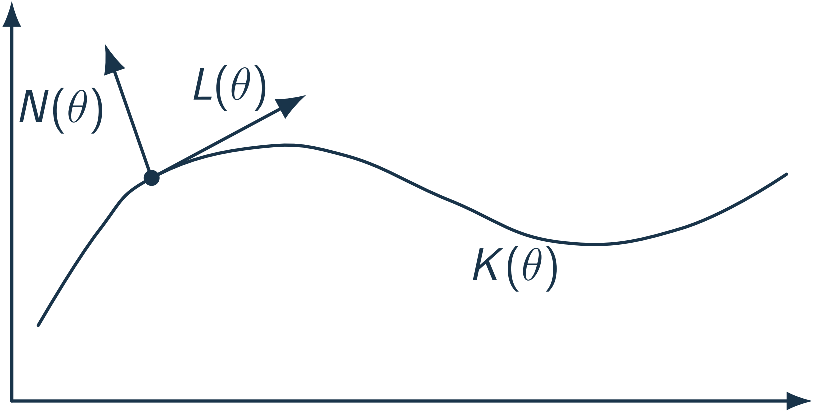

Geometric illustration of the tangent frame $L(\theta)$ and the complementary frame $N(\theta)$ along a torus.

Homework: the torsion identity

Let K parameterize an **exact invariant Lagrangian torus**, and let $P = (L\ N)$ be the symplectic frame constructed above, with $L = DK$.

Prove that applying the linearized operator to N gives only **tangent directions**:

$$\boxed{\mathcal{L}_\omega N + (DX_H \circ K)N = L\,T(\theta).}$$

The matrix $T(\theta)$ is called the **torsion matrix**.

1. Differentiate $L^T JN = \mathrm{Id}_n$ along ω and use $(\mathcal{L}_\omega + DX_H \circ K)L = 0$ to prove that the left-hand side belongs to Range L.
2. Deduce the formula $T = -N^T J(\mathcal{L}_\omega N + (DX_H \circ K)N)$.
3. Use $N^T JN = 0$ to prove that $T(\theta)^T = T(\theta)$.

Almost reducibility

Almost invariant means almost reducible

Suppose the torus is **almost invariant**, with a small residual

$$\mathcal{L}_\omega K + X_H \circ K = E(\theta).$$

Differentiating now gives

$$\mathcal{L}_\omega L + (DX_H \circ K)L = DE, \qquad L = DK.$$

The same geometric construction works **up to a small error**: for the adapted frame $P = (L\ N)$,

$$\boxed{\mathcal{L}_\omega P + (DX_H \circ K)P = P\left[\begin{pmatrix} 0 & T(\theta) \\ 0 & 0 \end{pmatrix} + E_{\mathrm{red}}(\theta)\right].}$$

Under the stated assumptions, E_{red} is controlled by the invariance defect.

The triangular structure survives, with a small remainder.

A useful inverse from the symplectic structure

The adapted frame $P(\theta) = (L(\theta)\ N(\theta))$ is **symplectic up to a small error**:

$$P^T J P = J + E_{\mathrm{symp}}(\theta).$$

If $E_{\mathrm{symp}} = 0$, its inverse is given directly by its transpose:

$$\boxed{P^{-1} = -JP^T J = \begin{pmatrix} -N^T J \\ L^T J \end{pmatrix}.}$$

For an almost invariant torus, define the **approximate inverse**

$$Q := -JP^T J, \qquad QP = \mathrm{Id}_{2n} - JE_{\mathrm{symp}}, \qquad P^{-1} = (\mathrm{Id}_{2n} - JE_{\mathrm{symp}})^{-1} Q.$$

Thus $P^{-1} \approx Q$ when the symplectic defect is small.

Very useful: the approximate inverse needs only transposes and multiplication by J.

A structured linear system: applying the adapted frame

The triangular correction system

Disregard E_{red} and use the approximate inverse $Q = -JP^T J$. With the change of unknown

$$\Delta = P\xi = L\xi^L + N\xi^N, \qquad \xi = \begin{pmatrix} \xi^L \\ \xi^N \end{pmatrix},$$

the Newton equation becomes the **triangular quasi-Newton system**

$$\mathcal{L}_\omega \xi + \begin{pmatrix} 0 & T(\theta) \\ 0 & 0 \end{pmatrix} \xi = \eta, \qquad \boxed{\eta := -QE = \begin{pmatrix} N^T JE \\ -L^T JE \end{pmatrix} = \begin{pmatrix} \eta^L \\ \eta^N \end{pmatrix}.}$$

Equivalently,

$$\boxed{\begin{aligned} \mathcal{L}_\omega \xi^L + T(\theta)\xi^N &= \eta^L, \\ \mathcal{L}_\omega \xi^N &= \eta^N. \end{aligned}}$$

We are almost there!

The correction now comes from the **triangular system**

$$\boxed{\begin{aligned} \mathcal{L}_\omega \xi^L + T(\theta)\xi^N &= \eta^L, \\ \mathcal{L}_\omega \xi^N &= \eta^N. \end{aligned}}$$

We have reached the end of the main reduction: **these equations are now easy to solve** — although it may not look like it yet!

First solve for ξ^N, then use it to solve for ξ^L. Both steps reduce to the same basic equation:

$$\mathcal{L}_\omega u = g.$$

Fourier series will do the work; we must also handle the averages.

Solving the cohomological equation:

$$\mathcal{L}_\omega f = g$$

Solving one Fourier mode at a time

On $\mathbb{T}^n = (\mathbb{R}/2\pi\mathbb{Z})^n$, write

$$f(\theta) = \sum_{k\in\mathbb{Z}^n} \widehat{f}_k e^{ik\cdot\theta}, \qquad g(\theta) = \sum_{k\in\mathbb{Z}^n} \widehat{g}_k e^{ik\cdot\theta}.$$

Since $\mathcal{L}_\omega = -\omega \cdot \partial_\theta$,

$$\mathcal{L}_\omega e^{ik\cdot\theta} = -i(k\cdot\omega)e^{ik\cdot\theta}.$$

Thus $\mathcal{L}_\omega f = g$ becomes a scalar equation for each mode:

$$-i(k\cdot\omega)\widehat{f}_k = \widehat{g}_k, \qquad \boxed{\widehat{f}_k = \frac{\widehat{g}_k}{-i\,k\cdot\omega} \quad (k\neq 0).}$$

The zero mode requires $\widehat{g}_0 = \langle g\rangle = 0$. We choose $\widehat{f}_0 = \langle f\rangle = 0$ to fix the arbitrary constant.

The formula is simple. Does the resulting Fourier series converge?

Small denominators can be dangerous

$$|\widehat{f}_k| = \frac{|\widehat{g}_k|}{|k \cdot \omega|}.$$

Rational independence says $k \cdot \omega \neq 0$, but does not tell us **how close it can be to zero**.

Even for $\omega = (1, \sqrt{2})$, the mode $k = (-99, 70)$ gives

$$|k \cdot \omega| = |70\sqrt{2} - 99| = \frac{1}{99 + 70\sqrt{2}} \approx 0.00505.$$

Division by a small number can amplify a tiny Fourier coefficient. For some irrational frequencies, this amplification can even destroy the convergence of the formal solution.

We need quantitative control of the small divisors. Which frequencies provide it?

Diophantine frequencies

A frequency $\omega \in \mathbb{R}^n$ is **Diophantine** with constants $\gamma > 0$ and $\tau \geq n-1$ if

$$\boxed{|k \cdot \omega| \geq \frac{\gamma}{|k|_1^\tau} \quad \text{for every } k \in \mathbb{Z}^n \setminus \{0\},} \qquad |k|_1 = \sum_{j=1}^n |k_j|.$$

Small divisors are still allowed, but their reciprocals grow at most **polynomially**:

$$\frac{1}{|k \cdot \omega|} \leq \frac{|k|_1^\tau}{\gamma}, \qquad |\widehat{f}_k| \leq \frac{|k|_1^\tau}{\gamma} |\widehat{g}_k|.$$

- γ measures the lower bound; τ controls the allowed decay.
- For each fixed $\tau > n-1$, almost every frequency satisfies this condition for some $\gamma > 0$.

Exponential Fourier decay dominates polynomial amplification.

Analyticity on a complex strip

We measure analyticity by the width $\rho > 0$ of a complex strip:

$$\mathbb{T}^n_\rho = \{\theta \in \mathbb{C}^n/(2\pi\mathbb{Z})^n : |\operatorname{Im}\theta_j| < \rho \text{ for all } j\}, \qquad \|g\|_\rho = \sup_{\theta\in\mathbb{T}^n_\rho} |g(\theta)|.$$

If g is holomorphic and bounded there, its Fourier coefficients satisfy

$$|\widehat{g}_k| \le \|g\|_\rho \, e^{-\rho|k|_1}.$$

For a Diophantine frequency, on a smaller strip of width $\rho - \delta$,

$$|\widehat{f}_k| e^{(\rho-\delta)|k|_1} \le \frac{\|g\|_\rho}{\gamma} |k|_1^\tau e^{-\delta|k|_1}, \qquad k \neq 0.$$

The remaining exponential decay makes the series summable for every $0 < \delta < \rho$.

The solution is analytic too, at the cost of a smaller strip width.

Rüssmann's inequality (without proof)

Let ω be Diophantine with constants (γ, τ), and let g be bounded and analytic on $\mathbb{T}^n_\rho$, with $\langle g \rangle = 0$.

The equation

$$\mathcal{L}_\omega f = g, \qquad \langle f \rangle = 0,$$

has a unique zero-average analytic solution on every smaller strip. For $0 < \delta < \rho$,

$$\boxed{\|f\|_{\rho-\delta} \leq \frac{C_{n,\tau}}{\gamma\, \delta^\tau} \|g\|_\rho.}$$

Here $C_{n,\tau}$ depends only on the dimension and Diophantine exponent.

Analytic input gives analytic output, with a controlled loss of width.

Rüssmann's estimate; see also A. Haro and A. Luque, *J. Differential Equations* **266** (2019), 1605–1674, Lemma 4.1.

Back to the (upper-triangular) linear system:

$$\boxed{\begin{aligned} \mathcal{L}_\omega \xi^L + T(\theta)\xi^N &= \eta^L, \\ \mathcal{L}_\omega \xi^N &= \eta^N. \end{aligned}}$$

First solve the normal equation

Start with
$$\mathcal{L}_\omega \xi^N = \eta^N.$$

Solvability requires $\langle \eta^N \rangle = 0$. In our exact symplectic Hamiltonian setting, this follows from the geometry for $\eta^N = -L^T J E$.

Let $\mathcal{R}_\omega g$ denote the **zero-average solution** of $\mathcal{L}_\omega f = g$ for zero-average g. Then
$$\boxed{\xi^N(\theta) = \widetilde{\xi}^N(\theta) + c, \qquad \widetilde{\xi}^N = \mathcal{R}_\omega \eta^N, \qquad c = \langle \xi^N \rangle \in \mathbb{R}^n.}$$

All nonzero Fourier modes are determined:
$$\widehat{\widetilde{\xi}^N}_k = \frac{\widehat{\eta^N}_k}{-i\, k \cdot \omega}, \qquad k \neq 0.$$

The normal equation leaves its average c free. Do not set it to zero yet!

The tangent equation fixes the normal average

Insert $\xi^N = \widetilde{\xi}^N + c$ into the first equation:

$$\mathcal{L}_\omega \xi^L = \eta^L - T(\theta)\widetilde{\xi}^N(\theta) - T(\theta)c.$$

The left-hand side has zero average, so the right-hand side must too:

$$0 = \langle \eta^L \rangle - \langle T\widetilde{\xi}^N \rangle - \langle T \rangle c.$$

Thus we choose c by solving the **finite-dimensional system**

$$\boxed{\langle T \rangle c = \langle \eta^L - T\widetilde{\xi}^N \rangle.}$$

If the average torsion matrix is invertible,

$$c = \langle T \rangle^{-1} \langle \eta^L - T\widetilde{\xi}^N \rangle.$$

The free normal average removes the obstruction in the tangent equation.

Nondegeneracy: invertible average torsion

The required **torsion nondegeneracy condition** is

$$\boxed{\det\langle T\rangle \neq 0.}$$

It lets us determine the normal average c uniquely and hence solve the tangent compatibility condition.

For the integrable Hamiltonian $H(\varphi, I) = h(I)$ and the torus $K(\theta) = (\theta, I_*)$,

$$L = \begin{pmatrix} \mathrm{Id}_n \\ 0 \end{pmatrix}, \qquad N = \begin{pmatrix} 0 \\ \mathrm{Id}_n \end{pmatrix}, \qquad T(\theta) = D^2 h(I_*).$$

Thus the condition becomes the familiar **Kolmogorov nondegeneracy**:

$$\det D^2 h(I_*) \neq 0.$$

The average torsion measures how normal displacements change the frequencies.

Finish the tangent equation and fix the phase

With c fixed, we know $\xi^N = \widetilde{\xi}^N + c$, and

$$g^L := \eta^L - T\xi^N, \qquad \langle g^L \rangle = 0.$$

We can now solve the tangent equation completely:

$$\boxed{\xi^L = \mathcal{R}_\omega g^L + a, \qquad a = \langle \xi^L \rangle \in \mathbb{R}^n.}$$

The constant a represents an **infinitesimal phase shift**:

$$K(\theta + a) = K(\theta) + DK(\theta)a + O(|a|^2) = K(\theta) + L(\theta)a + O(|a|^2).$$

A phase shift changes the parameterization, not the torus. We therefore fix the irrelevant phase freedom by choosing $a = 0$.

$$\boxed{\Delta = L\xi^L + N\xi^N, \qquad K_{\text{new}} = K + \Delta.}$$

Both components are determined: we have our quasi-Newton correction.

Summing up: numerical algorithm

What have we achieved?

We started with a nonlinear equation for the embedding:

$$\mathcal{F}_\omega(K) = \mathcal{L}_\omega K + X_H \circ K = 0.$$

Our **geometric quasi-Newton step** reduces the correction to

1. **Pointwise matrix operations** to build the adapted frame and transform the residual.
2. **Two Fourier solves** of cohomological equations.
3. **One $n \times n$ linear solve** using the average torsion.

Then update $K \leftarrow K + \Delta$ and repeat.

Geometry turns a large variable-coefficient problem into structured computations.

Step 0: the Hamiltonian and the problem

We consider the nearly integrable Hamiltonian

$$\boxed{H_\varepsilon(\varphi, I) = h(I) + \varepsilon f(\varphi, I), \qquad (\varphi, I) \in \mathbb{T}^n \times D.}$$

Its Hamiltonian vector field is

$$X_{H_\varepsilon}(\varphi, I) = \begin{pmatrix} \partial_I H_\varepsilon \\ -\partial_\varphi H_\varepsilon \end{pmatrix}.$$

For a prescribed **Diophantine frequency** ω, find an embedding $K : \mathbb{T}^n \to \mathbb{T}^n \times D$ satisfying

$$\boxed{\mathcal{L}_\omega K + X_{H_\varepsilon} \circ K = 0, \qquad \mathcal{L}_\omega K = -DK\,\omega.}$$

At $\varepsilon = 0$, a starting torus is $K_0(\theta) = (\theta, I_*)$ with $\nabla h(I_*) = \omega$; for small coupling, use it as an initial guess.

Input: the Hamiltonian, the target frequency, and an approximate torus.

Step 1: represent K and compute E

Fix a Diophantine ω. For a torus homotopic to the standard one, write

$$K(\theta) = (\theta + u(\theta),\ I_* + v(\theta)),$$

where u, v are **periodic**, represented on a uniform grid and by truncated Fourier series. Use the FFT to switch representations and differentiate spectrally:

$$\widehat{\partial_{\theta_j} u}_k = i k_j \widehat{u}_k, \qquad L = DK = \begin{pmatrix} \mathrm{Id}_n + Du \\ Dv \end{pmatrix}.$$

Evaluate the vector field on the grid and form

$$\boxed{E = X_H \circ K - L\omega.}$$

Fourier differentiation acts on the periodic parts, not on the lift θ.

Multiplying periodic functions with the FFT

Given Fourier coefficients of two periodic functions $a(\theta)$ and $b(\theta)$, we want the coefficients of $c(\theta) = a(\theta)b(\theta)$.

1. **Inverse FFT:** evaluate both functions on the same uniform grid.
2. **Multiply pointwise:** $c(\theta_\ell) = a(\theta_\ell)b(\theta_\ell)$.
3. **FFT back:** transform the product values into Fourier coefficients.

$$\boxed{\widehat{c} = \mathrm{FFT}\left[\mathrm{IFFT}(\widehat{a}) \cdot \mathrm{IFFT}(\widehat{b})\right].}$$

Use a sufficiently fine, **zero-padded grid** to avoid aliasing, then truncate back to the retained Fourier modes.

Differentiate in Fourier space; multiply in physical space.

Step 2: build L, N, P and transform E

At each grid point, construct

$$G = L^T L, \qquad \boxed{N = -JLG^{-1}, \qquad P = (L\ N).}$$

The embedding condition ensures that G is invertible. Use the symplectic approximate inverse

$$Q = -JP^T J = \begin{pmatrix} -N^T J \\ L^T J \end{pmatrix}$$

to transform the residual into tangent and normal components:

$$\boxed{\eta = -QE = \begin{pmatrix} N^T JE \\ -L^T JE \end{pmatrix} = \begin{pmatrix} \eta^L \\ \eta^N \end{pmatrix}.}$$

All these operations are small matrix computations at each grid point.

Step 3: compute the torsion and its average

Differentiate N spectrally and evaluate $DX_H \circ K$ on the grid. Using $\mathcal{L}_\omega N = -DN\,\omega$, compute

$$\boxed{T = -N^T J(\mathcal{L}_\omega N + (DX_H \circ K)N).}$$

Compute the zero Fourier mode $\langle T \rangle$ and check that it is **invertible and numerically well conditioned**.

Neglecting the small geometric defects, solve

$$\boxed{\begin{aligned} \mathcal{L}_\omega \xi^L + T\xi^N &= \eta^L, \\ \mathcal{L}_\omega \xi^N &= \eta^N. \end{aligned}}$$

The Fourier divisors and the average torsion are the two key inversions.

Step 4: solve the normal modes and fix their average

Jordi-Lluís Figueras

First solve the normal equation at every nonzero Fourier mode:

$$\widehat{\widetilde{\xi}^N}_k = \frac{\widehat{\eta^N}_k}{-i\,k\cdot\omega} \quad (k \neq 0), \qquad \widehat{\widetilde{\xi}^N}_0 = 0.$$

Transform $\widetilde{\xi}^N$ back to grid values and compute the product $T\widetilde{\xi}^N$ pointwise.

Determine the missing average c from the torsion:

$$\boxed{\langle T\rangle c = \langle \eta^L - T\widetilde{\xi}^N\rangle, \qquad \xi^N = \widetilde{\xi}^N + c.}$$

In exact arithmetic, $\langle \eta^N\rangle = 0$ by the Hamiltonian geometry. Numerically, monitor this mode and project out its discretization error.

Nonzero modes come from division; the zero mode comes from the torsion.

Step 5: solve the tangent modes and update K

Form the compatible right-hand side on the grid:

$$g^L = \eta^L - T\xi^N, \qquad \langle g^L \rangle = 0.$$

Solve in Fourier space, fixing the phase by a zero tangent average:

$$\widehat{\xi^L}_k = \frac{\widehat{g^L}_k}{-i\,k\cdot\omega} \quad (k \neq 0), \qquad \widehat{\xi^L}_0 = 0.$$

Return to grid values and update:

$$\boxed{\Delta = L\xi^L + N\xi^N, \qquad K_{\mathrm{new}} = K + \Delta.}$$

Recompute the residual and repeat until the desired accuracy is reached. Increase the Fourier resolution if the tail is not resolved or the error stalls.

In the convergence regime, expect quadratic reduction until numerical errors dominate.

Numerics: computing an invariant torus

Our numerical example

We compute tori for two weakly coupled rotators:

$$\boxed{H_\varepsilon(\varphi, I) = \frac{1}{2}(I_1^2 + I_2^2) + \varepsilon\left[\cos\varphi_1 + \cos\varphi_2 + \cos(\varphi_1 - \varphi_2)\right].}$$

Prescribe the frequency and start from the integrable torus:

$$\omega = (1, (1+\sqrt{5})/2), \qquad K_0(\theta) = (\theta, \omega).$$

- Compare $\varepsilon = 0.01$ and 0.03, using double precision.
- $N = 64$ points per angle; nonlinear operations on a 128^2 grid.
- Check the final residual independently on a 256^2 grid.

From `codes/KAMExample/`, reproduce a run with

```
./exampleKAM.x -e 0.03 -n 64 -t 1e-12
```

Watch the quasi-Newton method converge

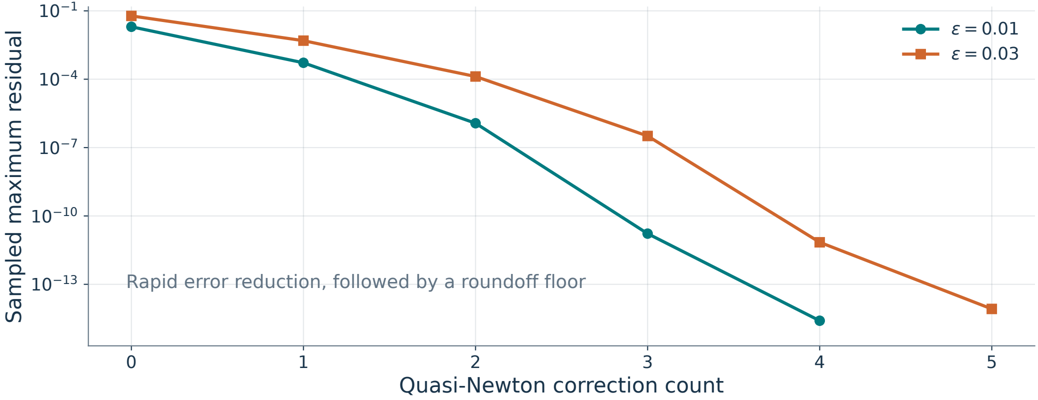

Four corrections for $\varepsilon = 0.01$; five for $\varepsilon = 0.03$.

Independent 256^2-grid residuals: 2.2×10^{-15} and 7.1×10^{-15}, respectively.

The computed torus is deformed

$$K(\theta) = (\theta + u(\theta), \omega + v(\theta)), \qquad \varepsilon = 0.03.$$

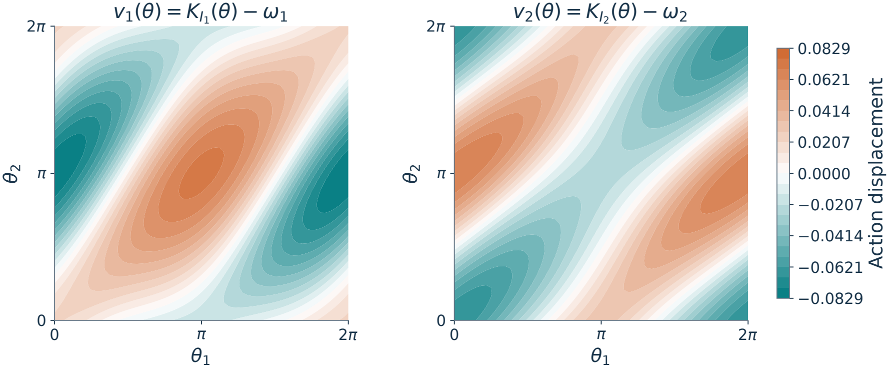

Opposite edges of each parameter square are identified.

Actions vary over the torus; the prescribed frequency stays fixed.

Have we resolved the torus?

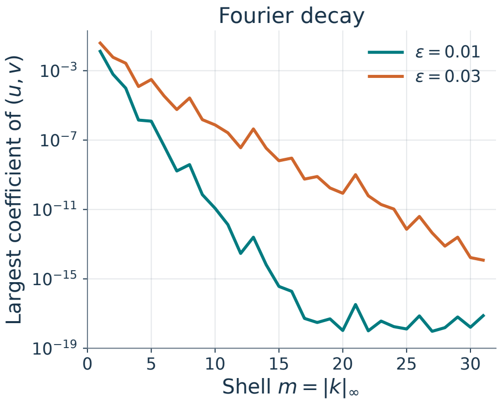

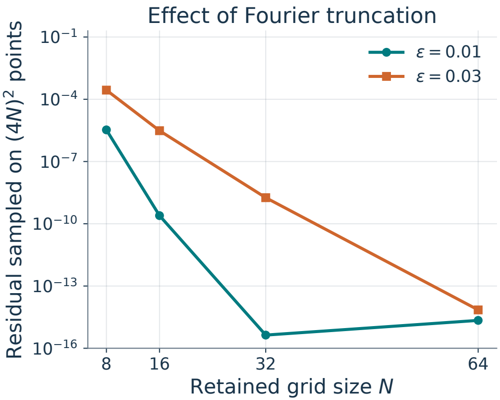

Left: largest coefficient of (u, v) in each shell $|k|_\infty = m$.
Right: truncate the same converged torus to size N; check its residual on $(4N)^2$ points.

Fourier decay tells us how much resolution we need.

Does the torus reproduce the dynamics?

Compare $K(\theta_0 + \omega t)$ with an independently integrated orbit starting at $K(\theta_0)$; $\varepsilon = 0.03$.

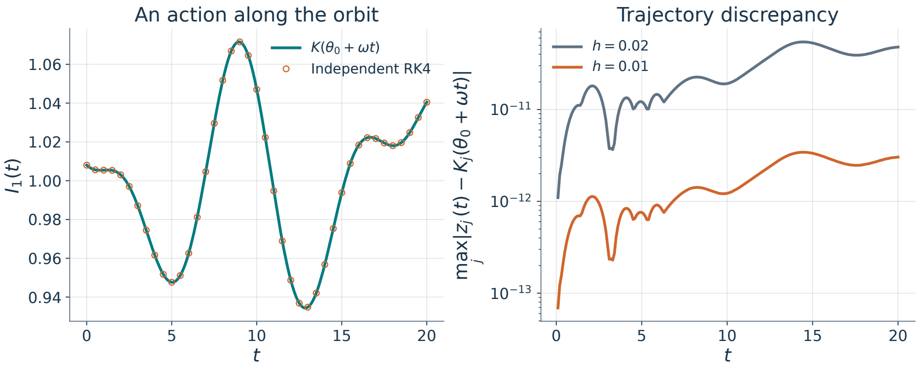

For RK4 with $h = 0.01$, the largest sampled discrepancy on $[0, 20]$ is 3.4×10^{-12}.

Strong numerical evidence. How do we turn it into a proof?

Validating the existence of these tori with the help of Computers

in two steps:

Write the Theorem
and Perform the Validation

Step 1: turn the algorithm into a theorem

Use the numerically computed torus as the initial approximation K_0 for the proof. Recall the iteration:

$$E_j = \mathcal{F}_\omega(K_j), \qquad K_{j+1} = K_j + \Delta_j.$$

We work with analytic objects and a fixed Diophantine frequency ω. A proof must control

1. the **linear correction** and the **nonlinear remainder**;
2. the **loss of analyticity** and the sum of all corrections;
3. the **domain and nondegeneracy bounds** throughout the iteration, including the frame inverses and $\langle T \rangle^{-1}$.

Replace observed convergence by explicit estimates that can be iterated.

The nonlinear part returns after the update

Starting from $E = \mathcal{F}_\omega(K)$, our correction satisfies

$$\mathcal{L}_\omega \Delta + (DX_H \circ K)\Delta = -E + r_{\mathrm{lin}}.$$

Here r_{lin} is the defect in solving the full linearized equation. The **new invariance error** is

$$E_{\mathrm{new}} = \mathcal{F}_\omega(K + \Delta) = r_{\mathrm{lin}} + \mathcal{N}_K(\Delta),$$

$$\mathcal{N}_K(\Delta) := X_H(K + \Delta) - X_H(K) - (DX_H \circ K)\Delta.$$

If $\|D^2 X_H\| \leq M_2$ along the correction, then on a common smaller strip

$$\boxed{\|E_{\mathrm{new}}\|_{\rho'} \leq \|r_{\mathrm{lin}}\|_{\rho'} + \frac{M_2}{2}\|\Delta\|_{\rho'}^2.}$$

An exact Newton solve has $r_{\mathrm{lin}} = 0$; our quasi-Newton proof must also bound it quadratically in the previous error.

The next residual includes both the linear-solve defect and the Taylor remainder.

Quadratic improvement has an analytic cost

Recall the losses in the linear solve and in differentiation:

$$\|\mathcal{R}_\omega g\|_{\rho-\delta} \le \frac{C_{n,\tau}}{\gamma\delta^\tau}\|g\|_\rho, \qquad \|\partial_{\theta_\ell} g\|_{\rho-\delta} \le \frac{1}{\delta}\|g\|_\rho.$$

Write $e_j = \|E_j\|_{\rho_j}$ and let $\delta_j = \rho_j - \rho_{j+1} > 0$ be the **total width lost in one step**. With controlled geometric bounds, a full step gives estimates of the form

$$\|\Delta_j\|_{\rho_{j+1}} \le C_\Delta \delta_j^{-\alpha} e_j, \qquad \boxed{e_{j+1} \le C_E \delta_j^{-\beta} e_j^2.}$$

Here $\alpha, \beta > 0$ summarize the losses; the constants include the Diophantine and nondegeneracy bounds.

Smaller cuts preserve more analyticity, but **make the error bound worse**.

The error is still quadratic, multiplied by an increasing loss factor.

Spend a finite amount of analyticity

Choose **geometrically decreasing cuts**:

$$\delta_j = \frac{\delta_0}{2^j}, \qquad \rho_{j+1} = \rho_j - \delta_j, \qquad 0 < 2\delta_0 < \rho_0.$$

Although there are infinitely many steps, the total loss is finite:

$$\boxed{\rho_\infty = \rho_0 - \sum_{j=0}^{\infty} \delta_j = \rho_0 - 2\delta_0 > 0.}$$

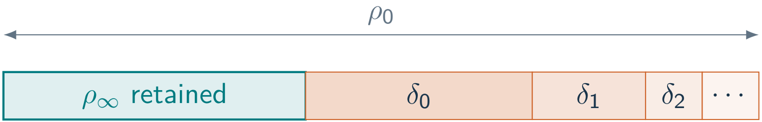

Schematic partition of the initial strip width

Infinitely many cuts can leave a positive analytic strip.

Quadratic decay wins over geometric losses

With these cuts, the error estimate becomes

$$e_{j+1} \le C_E \delta_0^{-\beta} 2^{\beta j} e_j^2.$$

The loss factor grows geometrically. Rescale the error by setting

$$s_j := C_E \delta_0^{-\beta} 2^{\beta(j+1)} e_j.$$

Then

$$\boxed{s_{j+1} \le s_j^2, \qquad s_0 < 1 \implies s_j \le s_0^{2^j}.}$$

This **doubly exponential decay** makes the corrections summable:

$$\sum_{j=0}^{\infty} \|\Delta_j\|_{\rho_\infty} < \infty, \qquad K_j \longrightarrow K_\infty.$$

The proof also keeps the domain margins and nondegeneracy bounds under control.

Quadratic convergence lets us afford progressively smaller cuts.

An a posteriori KAM theorem

Assume a real-analytic Hamiltonian and an exact symplectic structure.

- K_0 is an analytic embedding on $\mathbb{T}^n_{\rho_0}$; ω is Diophantine.
- Quantitative domain, frame and torsion bounds hold.
- $e_0 = \|\mathcal{F}_\omega(K_0)\|_{\rho_0}$ satisfies the smallness inequalities.

Conclusion (overview)

There is a nearby invariant Lagrangian torus with the same frequency, analytic on a positive strip $\rho_\infty < \rho_0$:

$$\mathcal{F}_\omega(K_\infty) = 0, \qquad \|K_\infty - K_0\|_{\rho_\infty} \leq C_{\mathrm{dist}}\, e_0.$$

Figueras–Haro's modified update gives sharper bounds.

J.-Ll. Figueras and A. Haro, *A modified parameterization method for invariant Lagrangian tori for partially integrable Hamiltonian systems*, *Physica D* **462** (2024), 134127. Theorem 2.18 ($d = n$).

Step 2: certify the theorem's input

Using the theorem's frame and torsion, certify **explicit bounds**:

- Bound E, its projections η^L, η^N, the frame inverses, the average torsion inverse and derivatives of X_H.
- Certify Diophantine constants and positive complex-domain margins.

For Figueras–Haro, the smallness test has the form

$$\boxed{\frac{\mathcal{C}}{\gamma \delta_{\mathrm{th}}^{\tau+1}} \max \left\{ \|\eta^L\|_{\rho_0}, \frac{\|\eta^N\|_{\rho_0}}{\gamma \delta_{\mathrm{th}}^{\tau}} \right\} < 1.}$$

Here $0 < \delta_{\mathrm{th}} < (\rho_0 - \rho_\infty)/3$ is a substep strip loss; $\mathcal{C}$ is computed from the theorem's bounds and margins.

Certify the upper bounds, then check the inequalities.

Figueras–Haro (2024), Theorem 2.18 and Appendix B; constants rescaled to our conventions.

Interval arithmetic

Freeze the numerical Fourier coefficients: the periodic parts of K_0 are **real trigonometric polynomials**.

For each real quantity, store a guaranteed enclosure

$$x \in [x] = [\underline{x}, \overline{x}].$$

Operations use **outward rounding**, so that, for example,

$$x \in [x],\ y \in [y] \quad \Longrightarrow \quad x + y \in [x] + [y], \qquad xy \in [x][y].$$

- Enclose the residual, elementary functions and matrix operations.
- Include the exact parameters and target frequency.
- Complex quantities use enclosures for real and imaginary parts.

Numerics supplies a candidate; intervals certify the bounds.

Rigorous FFTs for the residual

The quantity to bound is still

$$E(\theta) = X_H(K_0(\theta)) - DK_0(\theta)\,\omega.$$

1. Evaluate K_0 and DK_0 on a grid using interval FFTs.
2. Enclose $E = X_H \circ K_0 - DK_0\,\omega$ at every grid point.
3. Use a rigorous FFT to enclose its **discrete Fourier coefficients**: $|\widetilde{E}_k| \leq b_k$.

For the trigonometric interpolant $\widetilde{E}$ with modes $\mathcal{K}$,

$$\boxed{\|\widetilde{E}\|_{\rho_0} \leq \sum_{k\in\mathcal{K}} b_k\, e^{\rho_0 |k|_1}.}$$

Rigorous FFTs enclose the finite Fourier computation.

From Fourier data to a strip bound

Obtain a bound $\|E\|_{\widehat{\rho}} \leq M_E$ on a **wider strip** $\widehat{\rho} > \rho_0$. Analyticity gives

$$|\widehat{E}_k| \leq M_E e^{-\widehat{\rho}|k|_1}.$$

Explicit interpolation estimates control both **omitted modes and aliasing**:

$$\|E - \widetilde{E}\|_{\rho_0} \leq C_{\mathrm{grid}}(\rho_0, \widehat{\rho})\, M_E.$$

Combining this with the interval FFT gives the rigorous bound

$$\boxed{\|E\|_{\rho_0} \leq \sum_{k \in \mathcal{K}} b_k e^{\rho_0 |k|_1} + C_{\mathrm{grid}}(\rho_0, \widehat{\rho}) M_E =: b_E.}$$

At fixed strip widths, C_{grid} decays exponentially with grid resolution.

Intervals and analytic tails control the whole strip.

J.-Ll. Figueras, A. Haro and A. Luque, *Rigorous Computer-Assisted Application of KAM Theory: A Modern Approach*, *Found. Comput. Math.* **17** (2017), 1123–1193. Theorem 3.3 and Section 5.

From bounds to an invariant torus

1. **Numerics:** supply an accurate finite Fourier approximation K_0.
2. **Rigorous computation:** enclose the residual, the geometric quantities and all other inputs to the theorem.
3. **Apply the theorem:** verify its explicit inequalities and obtain an existence and distance bound.

A successful validation yields a certified statement of the form

$$\boxed{\mathcal{F}_\omega(K_\infty) = 0, \qquad \|K_\infty - K_0\|_{\rho_\infty} \leq r_{\text{validated}}.}$$

The limiting torus has the prescribed frequency ω and a positive analyticity width ρ_∞.

Finite computations certify an infinite iteration.

Our example: from approximation to proof

We validate the numerical example for the **exact** parameters

$$H_\varepsilon(\varphi, I) = \frac{1}{2}(I_1^2 + I_2^2) + \varepsilon\left[\cos\varphi_1 + \cos\varphi_2 + \cos(\varphi_1 - \varphi_2)\right],$$

$$\varepsilon = \frac{1}{100}, \qquad \omega = \left(1, \frac{1+\sqrt{5}}{2}\right).$$

Validation setup	**Value**
Candidate Fourier array / evaluation grid	64×64 / 256×256
Interval arithmetic	MPFI, 256 bits
Initial / final analytic strip	$\rho_0 = 0.10$ / $\rho_\infty = 0.05$ radians
Wider strip for interpolation bounds	0.50 radians

K_0 is the **frozen Fourier candidate after two multiprecision refinement steps**, with exact action offset ω.

`codes/validation/validateKAM.cpp`; certificate: `codes/validation/results/eps0p01.json`.

The certified inequalities pass

For the frozen candidate K_0, let $E_0 = X_H \circ K_0 - DK_0\,\omega$.

Quantity	**Certified upper bound**
Invariance defect $\|E_0\|_{\rho_0}$	1.149×10^{-19}
Approximate-symplecticity quantity ν	$1.907 \times 10^{-12} < 1$
Theorem smallness quantity	$0.08691 < 1$
Convergence quantity κ	$0.08666 < 1$

All displayed upper bounds are rounded upward.

VALIDATED

All specialized analytic, arithmetic, geometric, domain and smallness checks passed.

The residual is small enough relative to the certified theorem constants.

Figueras–Haro (2024), Theorem 2.18 and Appendix B; `codes/validation/results/eps0p01.json`.

A computer-assisted existence result

For $\varepsilon = 1/100$, there exists a **real-analytic invariant Lagrangian torus** with frequency

$$\omega = \left(1, \frac{1+\sqrt{5}}{2}\right).$$

Certified conclusion. Its embedding K_∞ satisfies

$$\boxed{X_{H_\varepsilon} \circ K_\infty = DK_\infty\,\omega, \qquad \|K_\infty - K_0\|_{0.05} \le 3.753 \times 10^{-10}.}$$

The distance is a **uniform maximum norm on the complex strip** $|\operatorname{Im}\theta_j| < 0.05$ radians, in the original coordinates $(\varphi_1, \varphi_2, I_1, I_2)$, with a consistent angular lift.

Frozen candidate independently rechecked at 320 bits.

An exact quasiperiodic torus lies within the certified distance.

Figueras–Haro (2024), Theorem 2.18, $d = n = 2$; `codes/validation/results/eps0p01.json`.

Bibliography and further exploration

Classical proofs, geometry, numerics and validation,
and infinite-dimensional KAM

The literature is vast: this is a personal, incomplete selection.

A personal reading list

The literature on KAM theory is **vast**, and I am sure that I have not managed to gather all of it here.

I am particularly fond of **Rafael de la Llave's Tutorial**. Rafa is one of my mentors, and his text really resonates with me. I warmly recommend it.

The other texts are excellent too: each offers a different perspective, and **there is something to learn from all of them**.

The links are clickable. Free texts, code and videos are distinguished from books that generally require library or publisher access.

Rafa's tutorial: a broad perspective

Rafael de la Llave
A Tutorial on KAM Theory.

Proc. Sympos. Pure Math. **69** (2001), 175–292. Expanded notes: 182 pages.

- Small divisors and the geometric background.
- Different proof strategies and hard implicit-function theorems.
- Connections with Aubry–Mather theory and computer-assisted proofs.

A broad tutorial and a source of ideas, with different levels of detail across the topics.

Free expanded notes (DANCE):
https://www.dance-net.org/images/files/events/rtns2008/materiales/notes_de_la_llave.pdf

A first complete classical proof

Jürgen Pöschel
A Lecture on the Classical KAM Theorem.
Proc. Sympos. Pure Math. **69** (2001), 707–732; updated arXiv version, 2009 (33 pages).

A focused, complete proof following a traditional KAM iteration: **main ideas before optimal estimates**. A manageable reading-seminar text.

Free text: `https://arxiv.org/abs/0908.2234`.

Read alongside: Abed Bounemoura
Some remarks on the Classical KAM Theorem, following Pöschel (2020).

Discusses a technical correction to the proof and an improvement of the quantitative statement.

Free companion note: `https://arxiv.org/abs/2006.04453`.

Other introductory lectures

C. Eugene Wayne
An Introduction to KAM Theory.
Lecture chapter, 1996, pp. 3–29; author-hosted notes dated 2008.

A short alternative exposition to read alongside Pöschel.
Free author PDF: `http://math.bu.edu/people/cew/preprints/introkam.pdf`.

Alexey A. Glutsyuk
Introduction to KAM Theory.
Free HSE lecture notes, May 2025 (57 pages); recorded HSE lectures, 2026.

From **Siegel linearization and circle diffeomorphisms** to symplectic geometry, integrable systems and analytic KAM, with Newton estimates and the homological equation.

Course: `https://www.hse.ru/en/edu/courses/920973104`
Lecture 1: `https://www.youtube.com/watch?v=_NFsFk0fS7s`
Videos: `https://www.youtube.com/@mathematicsathse1021`.

Herman normal forms and counterterms

Jacques Féjoz
Introduction to KAM Theory, with a View to Celestial Mechanics.
In *Variational Methods in Imaging and Geometric Control*, Radon Series **18**, de Gruyter, 2016.

Organizes the analytic work around a **twisted-conjugacy theorem**, then derives KAM statements. Useful for understanding finite-dimensional obstructions and the organization of a proof.

Free author version:
https://www.ceremade.dauphine.fr/~fejoz/Articles/Fejoz_2016_introduction-KAM.pdf.

Jacques Féjoz
A Simple Proof of the Invariant Torus Theorem (2010 preprint).

A concentrated companion: Herman's normal form, an inverse-function theorem, Newton iteration and interpolation. Best read after a fuller introduction.

Free text: https://arxiv.org/abs/1005.5604.

History and dynamical perspective

H. Scott Dumas
The KAM Story: A Friendly Introduction to the Content, History, and Significance of Classical Kolmogorov–Arnold–Moser Theory.
World Scientific, 2014.

Historical context and motivation: why KAM changed our understanding of mechanics. A companion to a proof-oriented text.

https://doi.org/10.1142/8955 — library/publisher access.

Jürgen Moser
Stable and Random Motions in Dynamical Systems: With Special Emphasis on Celestial Mechanics.
Princeton University Press, 1973; later reissued.

A classic perspective on regular and complicated motion, to read alongside modern lecture notes.

https://www.jstor.org/stable/j.ctt1bd6kg5 — library/publisher access.

Hamiltonian methods and classical proofs

Antonio Giorgilli
Notes on Hamiltonian Dynamical Systems.
Cambridge University Press, 2022.

Hamiltonian mechanics, perturbation theory, normal forms, Lie series and Lie transforms, invariant tori and long-time stability. **Chapter 8: persistence of invariant tori.**

https://doi.org/10.1017/9781009151122 — library/publisher access.

Achim Feldmeier
Introduction to Arnold's Proof of the Kolmogorov–Arnold–Moser Theorem.
CRC Press, first published 2022.

A **step-by-step account of Arnold's classical proof**, for readers who want to unpack the original strategy in greater detail.

https://doi.org/10.1201/9781003287803 — library/publisher access.

Celestial mechanics and applications

Alessandra Celletti
Stability and Chaos in Celestial Mechanics.
Springer, 2010.

Invariant tori alongside numerical methods, the Kepler and three-body problems, rotational dynamics, perturbation theory and long-time stability. **Invariant-tori chapter: pp. 127–176.**

https://link.springer.com/book/10.1007/978-3-540-85146-2 — library/publisher access.

Ugo Locatelli
Introduction to KAM Theory and Applications.
I-CELMECH Training School, Milan, February 2020.

A recorded introductory lecture to accompany Giorgilli or Celletti in an applications-oriented reading programme.

Free lecture: https://www.youtube.com/watch?v=oTloSYaNuSc.

From rigorous results to computations

Àlex Haro, Marta Canadell, Jordi-Lluís Figueras, Alejandro Luque and Josep-Maria Mondelo
The Parameterization Method for Invariant Manifolds: From Rigorous Results to Effective Computations.
Springer, 2016. Applied Mathematical Sciences **195**.

Chapter 4: The Parameterization Method in KAM Theory. Chapter 3, on quasi-periodic systems and validated numerics, is complementary. The bridge between invariance equations, rigorous results and effective algorithms.

https://link.springer.com/book/10.1007/978-3-319-29662-3 — library/publisher access.

Companion C/C++ code
Implementations of algorithms and examples from the book, including **Chapter 4: Lagrangian invariant tori**. Read the construction, then inspect its implementation.

Free companion site: https://www.maia.ub.es/dsg/param/
Chapter 4 code: https://www.maia.ub.es/dsg/param/chapter4.html.

From approximate tori to validated tori

Jordi-Lluís Figueras, Àlex Haro and Alejandro Luque
Rigorous Computer-Assisted Application of KAM Theory: A Modern Approach.
Found. Comput. Math. **17** (2017), 1123–1193.

A posteriori verification for Lagrangian tori of **symplectic maps**: computable estimates, rigorous Fourier bounds and numerical examples. Read after an introduction to the parameterization method.

Free preprint: https://arxiv.org/abs/1601.00084
Published article: https://doi.org/10.1007/s10208-016-9339-3.

Jordi-Lluís Figueras and Alex Haro
A modified parameterization method for invariant Lagrangian tori for partially integrable Hamiltonian systems.
Physica D **462** (2024), 134127.

The **Hamiltonian a posteriori theorem** used in our validation: modified updates, strip losses and explicit constants.

Open-access article: https://doi.org/10.1016/j.physd.2024.134127.

Further directions in KAM theory

Henk W. Broer, George B. Huitema and Mikhail B. Sevryuk
Quasi-Periodic Motions in Families of Dynamical Systems: Order amidst Chaos.
Springer, 1996. Lecture Notes in Mathematics **1645**.

Parameter-dependent KAM theory: conjugacy, continuation and Whitney-smooth families. A research-oriented next step after a first complete proof.

https://link.springer.com/book/10.1007/978-3-540-49613-7 — library/publisher access.

Thomas Kappeler and Jürgen Pöschel
KdV & KAM.
Springer, 2003.

An advanced route into **infinite-dimensional Hamiltonian systems**, through periodic KdV and its KAM analysis.

https://link.springer.com/book/10.1007/978-3-662-08054-2 — library/publisher access.

Thank you very much!

jordi.lluis.figueras@upc.edu

Interested in pursuing a **PhD or a postdoc in dynamical systems**?

Feel free to contact me. I would be happy to help you explore possible opportunities and next steps.

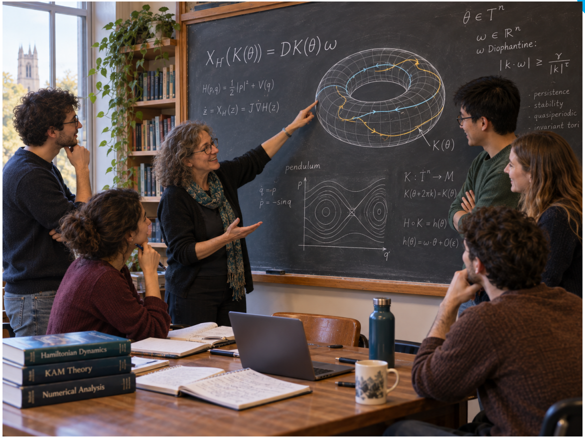

The next questions are waiting.
Illustration generated with ChatGPT