Certified spectral approximation of transfer operators and the Gauss map

Isaia Nisoli

UFRJ

The Gauss–Kuzmin–Wirsing problem

The **Gauss map** $T(x) = \{1/x\}$ on $(0, 1]$ generates continued fractions.

Let $x_0 \sim \mathrm{Uniform}[0, 1]$ and $x_n = T^n(x_0)$. **Gauss (1812):** what is the distribution of x_n?

$$F_n(x) = \mathbb{P}(x_n \le x) \xrightarrow{n\to\infty} \log_2(1 + x).$$

The rate is governed by the **transfer operator**

$$(\mathcal{L}f)(x) = \sum_{k\ge1} \frac{1}{(x+k)^2} f\left(\frac{1}{x+k}\right), \qquad F_n(x) = \int_0^x (\mathcal{L}^n\mathbf{1})(t)\,dt.$$

Gauss–Babenko–Knuth problem *(Knuth TAOCP, 3rd ed., 1997)*

Compute the eigenvalues λ_j and eigenfunctions $\Psi_j(x)$ of the Gauss–Kuzmin–Wirsing (GKW) operator to high precision.

History: rate of convergence studied by Kuz'min, Lévy, Wirsing (1974). High-precision numerics by Babenko (1978), but rigorous certification historically limited to a few digits of the leading eigenvalue.

Main result: 50 eigenvalues to 90 digits

Theorem (Main result)

The first 50 *nonzero eigenvalues of the GKW operator on* $H^2(D_1)$ *are real and simple, with certified enclosures each of at least* **90 rigorous decimal digits**. *In particular, the subdominant eigenvalue (GKW constant) satisfies* $|\lambda_2 - \tilde{\lambda}_2| < 10^{-175}$, *where*

$$\tilde{\lambda}_2 = -0.30366\,30028\,98732\,65859\,74481\,\ldots \quad (\textit{160 digits shown})$$

Also certified: eigenvectors, Riesz projectors, spectral gap.

Validated spectral expansion for the Gauss–Kuzmin distributions:

$$\mathcal{L}^n\mathbf{1} = \sum_{j=1}^{50} \lambda_j^n \, \ell_j(\mathbf{1})\, v_j + R_{50}(n), \qquad \|R_{50}(n)\| \leq 3.72 \times 10^{21} \cdot (1.01 \times 10^{-21})^{n+1}.$$

Rigorous error below 10^{-20} already at $n = 1$.

The central idea: spectral gap as *output*

Standard approach: assume spectral gap $\Rightarrow$ derive quantitative consequences.

Our approach:

compute finite-rank data $\Rightarrow$ **certify** spectral gap a posteriori

The key observation

If a contour Γ in $\mathbb{C}$ satisfies

$$\varepsilon_N \cdot \sup_{z\in\Gamma} \|(zI - L_N)^{-1}\| < 1,$$

then $\Gamma \subset \rho(L)$: *the same contour certifies isolation of $\sigma(L)$ inside Γ, with no prior spectral gap assumption*.

Resolvent bounds along enclosing contours simultaneously establish:

- spectral separation
- eigenvalue multiplicity
- projector stability

The framework: three inputs

Given: bounded operator $L : \mathcal{B} \to \mathcal{B}$.

(A) Finite-rank approximation with explicit error bound:

$$\|L - L_N\| \le \varepsilon_N.$$

(B) Certified resolvent bounds along a contour Γ:

$$M_\Gamma := \sup_{z \in \Gamma} \|(zI - L_N)^{-1}\| < \infty.$$

(C) Key condition:

$$\varepsilon_N \cdot M_\Gamma < 1.$$

When (A)–(C) hold:

$\Gamma \subset \rho(L)$, and L, L_N have the same number of eigenvalues inside Γ (counted with multiplicity).

Two settings: compact operator on a single Banach space;
operator satisfying a Doeblin–Fortet–Lasota–Yorke (DFLY) inequality.

The resolvent perturbation bound

Lemma (Neumann factorization)

If $\|L - L_N\| \le \varepsilon_N$ *and* $\varepsilon_N \cdot \|(zI - L_N)^{-1}\| < 1$*, then* $z \in \rho(L)$ *and*

$$\|(zI - L)^{-1}\| \le \frac{\|(zI - L_N)^{-1}\|}{1 - \varepsilon_N \|(zI - L_N)^{-1}\|}.$$

Proof: Factor $zI - L = (I - (L - L_N)(zI - L_N)^{-1})(zI - L_N)$, invert via Neumann series. $\square$

Projector bound (Riesz projector transfer)

Under Assumption (A)–(C), writing $P = P_L(U)$, $P_K = P_{L_K}(U)$:

$$\|P - P_K\| \le \frac{\ell(\Gamma)}{2\pi} \cdot \frac{\varepsilon_K \cdot M_\Gamma^2}{1 - \varepsilon_K M_\Gamma}.$$

If $\|P - P_K\| < 1$: $\dim \mathrm{Ran}\, P = \dim \mathrm{Ran}\, P_K$.

Spectral convergence guarantee

Theorem (Spectral convergence, compact case)

Let L be compact, (L_K) finite-rank with $\|L - L_K\| \leq \varepsilon_K \to 0$. Let $U \subset \mathbb{C}$ be a spectral window isolating finitely many eigenvalues of L.
Then for K large enough:

1. **No spectral pollution:** *L_K and L have the same number of eigenvalues in U.*
2. **Projector convergence:** $\|P_{L_K}(U) - P_L(U)\| \to 0$ *as* $K \to \infty$.
3. **A posteriori certification:** *the condition triggering (1)–(2) is checkable from computed data of L_K alone, with no spectral gap assumption.*

DFLY (Lasota–Yorke) version: same conclusions for eigenvalues outside the essential spectral radius, under a mild rate condition on $\delta_k \to 0$.

Li (1976): Ulam approximations converge to the invariant density.
This work: extends to the full discrete spectrum with explicit bounds.

The DFLY extension

Beyond compact operators: **Doeblin–Fortet–Lasota–Yorke** setting.

$$\|Lf\|_s \le a\|f\|_s + b\|f\|_w, \quad 0 < a < 1.$$

Examples: piecewise expanding maps (BV/Sobolev spaces), hyperbolic maps on anisotropic Banach spaces.

Challenge: the strong resolvent $\|(zI - L_k)^{-1}\|_s$ is hard to compute when the strong norm is, e.g., BV.

Solution: weak-to-strong lifting

1. Certify the **weak** resolvent $\|(zI - L_k)^{-1}\|_w$ (computable from a finite matrix).
2. Lift to the strong norm via the Lasota–Yorke inequality.
3. Certify spectral data outside the essential spectral radius.

Result: same convergence guarantees (no spectral pollution, projector convergence, a posteriori certification) for all eigenvalues outside any fixed $\mu > a$.

Hardy-space framework for the Gauss map

Work in $H^2(D_1)$, the Hardy space on the disc $D_1 = \{|w-1| < 1\}$.

The inverse branches of the Gauss map are $\tau_n(w) = \frac{1}{w+n}$.

Key geometric fact

For every $n \geq 1$: $\tau_n(D_{3/2}) \subset D_1$.

This **domain gain** $D_{3/2} \to D_1$ is the source of:

- compactness of $L = S\mathcal{L} : H^2(D_1) \to H^2(D_1)$
- **exponentially decaying** truncation error

Theorem (Truncation bound)

Let L_K be the degree-K Taylor truncation. Then

$$\|L - L_K\|_{H^2(D_1)} \leq C_2 \left(\frac{2}{3}\right)^{K+1}, \qquad C_2 = \|\mathcal{L}\|_{H^2(D_1) \to H^2(D_{3/2})} \approx 10.058.$$

Consequence: at $K = 512$, the truncation error is $\varepsilon_{512} \approx 10^{-90}$ — matching the digit target.

Certified resolvent bounds: the linear algebra

Input: matrix $A_K \in \mathbb{C}^{(K+1)\times(K+1)}$ representing L_K.

Step 1. Compute (numerical) Schur decomposition $A_K Q = QT + R$, certify defects:

$$\delta = \|I - Q^* Q\|, \quad r_{\mathrm{sch}} = \|R\|, \quad \|E\| \le r_{\mathrm{sch}}\sqrt{1+\delta} + \|A_K\|\,\delta.$$

Step 2. Sample contour Γ at m points; at each sample point z_ℓ certify lower bound $s_\ell \le \sigma_{\min}(z_\ell I - T)$ via validated SVD.
Lipschitz propagation (chord gap $\Delta_m = 2\rho \sin\frac{\pi}{2m}$):

$$s_* := \min_\ell s_\ell - \Delta_m > 0 \;\Rightarrow\; \sup_{z\in\Gamma} \|(zI - T)^{-1}\| \le \tfrac{1}{s_*}.$$

Step 3. Transfer resolvent bound from T to A_K (via Neumann):

$$\|(zI - A_K)^{-1}\| \le \frac{\kappa(Q)^2 \|(zI - T)^{-1}\|}{1 - \kappa(Q)^2 \|(zI - T)^{-1}\| \|E\|}.$$

All computations in **ball arithmetic** (`Arb`/`Julia`).

Computational pipeline

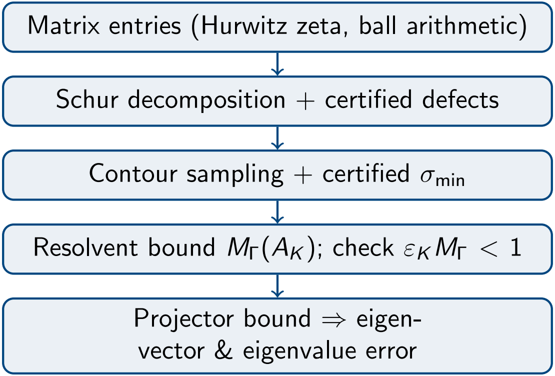

Single desktop workstation. $K = 1024$, 2048-bit arithmetic for 90-digit precision.

The coarse-fine strategy

Two levels, two roles:

Coarse level K:

– Small matrix, fast to certify.
– Establishes that contour $\partial U \subset \rho(L)$.
– Certifies spectral isolation.

Fine level K':

– Large matrix, high precision.
– *Reuse the same contour ∂U.*
– Projector bound improves by $\varepsilon_{K'}/\varepsilon_K$.

Key inequality (coarse-fine propagation)

For $K' \geq K$, the resolvent at the fine level satisfies

$$\mathcal{R}(z, L_{K'}) \leq \frac{\mathcal{R}(z, L_K)}{1 - \varepsilon_K \mathcal{R}(z, L_K)} \quad \forall z \in \partial U.$$

No recomputation of resolvents at the fine level.

Practical benefit: certify isolation cheaply, then boost precision for free.

Certified eigenfunctions

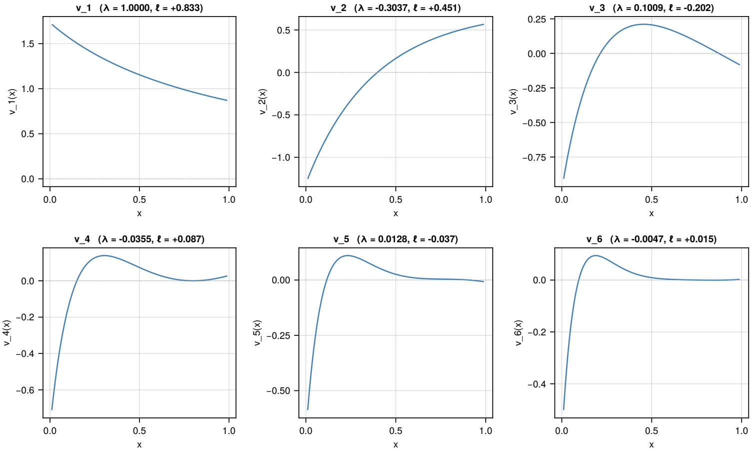

First six certified eigenfunctions ($K = 512$, 1024-bit precision). Invariant density $v_1 = \frac{1}{\ln 2} \cdot \frac{1}{1+x}$ (top left) is positive; higher eigenfunctions oscillate, reflecting the Markov partition.

Connection to the SCI hierarchy

Solvability Complexity Index (Colbrook–Hansen, 2023): for general bounded operators, computing the discrete spectrum and determining spectral gaps are Σ_2^A-complete.

No algorithm taking only matrix entries can certify spectral data at any finite step.

Our result: for transfer operators, the discrete spectrum is certifiable by a *single-limit* algorithm with explicit error bounds at every finite K where

$$\varepsilon_K \cdot \sup_{z \in \Gamma} \|(zI - L_K)^{-1}\| < 1.$$

What breaks the barrier?

The computable truncation bound ε_K is not available for abstract compact operators given only by matrix entries, but is furnished by the **dynamical structure** of the transfer-operator setting (domain gain, Lasota–Yorke inequality).

DFLY together with a compatible discretization hierarchy provides the oracle $(\mathcal{B}_s, \mathcal{B}_w, \delta_k)$ needed to break the SCI barrier — unavailable from quasi-compactness alone or from finite-dimensional matrix data without dynamical structure.

Related work I: high-precision spectral computations

All entries below assume a spectral gap a priori.

Bandtlow–Jenkinson (2007/08) explicit *a priori* upper bounds on eigenvalue moduli for holomorphic transfer operators; no certified gaps.

Pollicott–Slipantschuk (2024), Vytnova–Wormell (2025) high-precision thermodynamic computations (zeta functions, Hausdorff dimensions) for one-dimensional maps.

Matheus–Moreira–Pollicott–Vytnova (2022) Hausdorff dimension of Gauss–Cantor sets in Lagrange–Markov spectra via transfer operators.

Crimmins–Froyland (2020) Fourier discretizations for transfer operators on **higher-dimensional anisotropic Banach spaces** (piecewise hyperbolic maps); spectral gap assumed, no finite-K error bounds.

Common limitation: gap is an input, not an output. No a posteriori certification at finite truncation level.

Related work II: theory and computability

Keller–Liverani (1999) qualitative spectral stability under DFLY: isolated eigenvalues of L persist for L_K at large K, so $\mathcal{R}_w(L_K, \Gamma)$ is eventually bounded along any isolating contour Γ.

Colbrook–Hansen (2023) SCI impossibility: for general operators, no algorithm can certify discrete spectrum or spectral gaps at finite K; only asymptotically correct results are achievable.

Blumenthal–Nisoli–Taylor-Crush (2025) validated linear algebra for eigenvalue enclosures; starting point of this work.

The key chain (DFLY setting)

KL $\Rightarrow \mathcal{R}_w(L_K, \Gamma)$ bounded for large K.

Approximation scheme $\Rightarrow \delta_K \to 0$ (rate depends on the setting).

Therefore $\delta_K \cdot \mathcal{R}_w(L_K, \Gamma) < 1$ is eventually satisfied — and **once verified at a specific finite** K, full certified spectral data (eigenvalues, eigenvectors, projectors) follows with explicit bounds.

What breaks the SCI barrier: DFLY supplies both δ_K and (via KL) bounded resolvents — neither is available for abstract operators.

Portability

The three-input interface applies beyond transfer operators:

- **Compact integral operators** (directly).
- **Markov operators** satisfying Doeblin–Fortet/Harris-type inequalities: certified spectral gaps $\Rightarrow$ rigorous mixing rates.
- **Renormalization operators:** domain gain makes DR_{f_*} compact on Hardy/Bergman spaces; could certify hyperbolicity of period-doubling renormalization from computed data (Lanford 1982, systematic version).
- **Koopman operators with noise-induced compactness**, hypocoercive PDEs.

Limitation: certifies isolated spectral data, not continuous spectrum. The condition $\varepsilon_N M_\Gamma < 1$ fails near the essential spectrum.

Summary

Two main contributions

1. **Abstract framework:** spectral certification from finite-rank data, a posteriori, no spectral gap input. Convergence guaranteed for all isolated eigenvalues (compact and DFLY settings).
2. **GKW benchmark:** first 50 nonzero eigenvalues of the Gauss–Kuzmin–Wirsing operator certified to ≥ 90 decimal digits, with eigenvectors, Riesz projectors, and spectral gap.

Key messages:

- Spectral gap is an *output*, not an input.
- The DFLY/domain-gain structure is the oracle that breaks the SCI computability barrier.
- Analysis work is confined to a few quantitative constants; the rest is certified numerical linear algebra.

Code and data:

`https://github.com/orkolorko/GKWExperiments.jl`

Data: `doi:10.7910/DVN/HKM3Y2`

Appendix: Schur-based resolvent bound (detail)

Given: $A_K Q \approx QT$ (numerical Schur).
Defects:

$$\delta = \|I - Q^* Q\|, \quad r_{\mathrm{sch}} = \|A_K Q - QT\|, \quad \|E\| \le r_{\mathrm{sch}}\sqrt{1+\delta} + \|A_K\|\delta.$$

Resolvent via proxy $S_0 = QTQ^*$:

$$\|(zI - S_0)^{-1}\| \le \kappa(Q)^2 \|(zI - T)^{-1}\|.$$

Transfer to A_K by Neumann: if $\kappa(Q)^2 \|(zI - T)^{-1}\| \|E\| < 1$, then

$$\|(zI - A_K)^{-1}\| \le \frac{\kappa(Q)^2 \|(zI - T)^{-1}\|}{1 - \kappa(Q)^2 \|(zI - T)^{-1}\| \|E\|}.$$

Certified $\sigma_{\min}$ of triangular $zI - T$ via Rump's SVD theorem and Lipschitz propagation between sample points.