Dimension-Free Differential Games via Latent Representation

Breaking the 6-Dimensional Barrier for Hamilton-Jacobi-Isaacs Equations

Spectral Latent methods solve HJI equations in high dimensions for smooth value functions; the Latent Number detects singular surfaces.

Tamás Nagy, Ph.D.

tamas.nagy@thelatent.space

Draft • April 2026

Executive Summary (Non-Technical)

A pursuit-evasion game: a drone chases an intruder across a 3D environment with obstacles. The drone chooses thrust; the intruder chooses direction. Each reacts to the other in real time. To compute the optimal strategy, you need to solve the **Hamilton-Jacobi-Isaacs (HJI) equation** — a partial differential equation that encodes the minimax value of the game at every possible configuration.

For 2D problems, this is routine. For 3D, expensive but doable. **Beyond 5-6 dimensions, it is currently impossible.** Grid-based methods need M^d points for d dimensions. A 10-dimensional problem (two 5D vehicles) needs M^{10} grid points — trillions for modest resolution. This barrier blocks applications in multi-vehicle autonomy, high-dimensional robust control, and multi-country climate negotiation.

This paper applies the **Latent framework** to differential games. The core result: **if the value function is smooth** (which it is away from the game's "barrier" surfaces), it has a Latent representation with $N^* = O(\log(1/\varepsilon)/\log\rho)$ modes, independent of dimension. The HJI equation becomes a system of ODEs for the Latent coefficients. Near singular surfaces — where optimal strategies switch discontinuously — the Latent Number ρ drops toward 1, providing an automatic singularity detector.

This unlocks three classes of problems:

1. **High-dimensional pursuit-evasion** — multi-vehicle, multi-obstacle scenarios in realistic dimensions.
2. **Robust control beyond H_∞** — worst-case controller design in systems with many state variables (power grids, chemical processes, financial portfolios).
3. **Climate negotiation games** — where the players are countries, the state is the climate system, and the time horizon is centuries. This is a differential game that is smooth, high-dimensional, and among the most consequential optimization problems of our era.

Abstract

We develop a spectral Latent method for Hamilton-Jacobi-Isaacs (HJI) equations arising from two-player and multi-player differential games in high dimensions. The value function $V(t, x)$ is represented in the Latent basis with N^* modes per dimension, reducing the PDE to an ODE system. For ρ-analytic value functions, we prove convergence: $\|V - V^{N^*}\| \leq C\rho^{-N^*}$, independent of the state dimension d. We address the key technical challenge — **singular surfaces** where the value function loses smoothness — via Latent Number monitoring: $\rho(t, x)$ is computed locally and drops toward 1 near barriers, dispersal surfaces, and focal lines, enabling adaptive resolution. For multi-player differential games ($N \geq 3$), we combine the HJI Latent solver with the grade decomposition (Nagy, 2026b): the N-player value function is approximated by its grade-R^* truncation, reducing the Nd-dimensional problem to a collection of coupled R^*d-dimensional problems. We demonstrate the method on three benchmark problems: a 3D pursuit-evasion game (exact solution available), a 12-dimensional multi-vehicle reach-avoid game, and a 50-dimensional stylized climate negotiation game. The Latent method solves the 12D problem in minutes on a single GPU; grid-based methods cannot solve it at all.

---

1. Introduction

1.1 The Hamilton-Jacobi-Isaacs Equation

A zero-sum two-player differential game has dynamics:

$$\dot{x} = f(x, u, v), \qquad x(0) = x_0 \in \mathbb{R}^d$$

where player 1 (the minimizer) controls $u \in U$ and player 2 (the maximizer) controls $v \in V$. The value function is:

$$V(t, x) = \min_u \max_v \mathbb{E}\left[\int_t^T L(x, u, v)\, ds + g(x(T))\right]$$

Under Isaacs' condition, V satisfies the HJI equation:

$$\partial_t V + \min_u \max_v \{f(x, u, v) \cdot \nabla V + L(x, u, v)\} = 0$$

with terminal condition $V(T, x) = g(x)$.

This PDE is: - **Nonlinear** (the min-max operator is non-smooth at the saddle). - **High-dimensional** (d = state dimension of the dynamical system). - **Non-convex** in general (unlike optimal control, where the Hamiltonian is convex).

1.2 The Current State

Method	Max d achieved	Handles non-convex H?	Convergence guarantee?
Level set / grid (Osher-Sethian)	5–6	Yes	Yes ($O(h)$)
Hopf-Lax formula	Any d	Only convex/concave	Yes (exact for LQ)
Neural networks (DeepReach)	~10	Yes	No
Tensor decomposition	~15 (for HJB)	Partially	Rank-dependent
Latent spectral	**Any d (if smooth)**	**Yes**	**Yes ($O(\rho^{-N^*})$)**

The Hopf-Lax formula is the only existing method that avoids the curse of dimensionality, but it requires convex or concave Hamiltonians — a restriction that excludes most differential games. The Latent method has no convexity requirement; it only requires smoothness.

1.3 The Smoothness Structure of HJI Solutions

The value function of a differential game is NOT smooth everywhere. It develops **singular surfaces** where optimal strategies switch discontinuously (Isaacs, 1965). These are classified as:

- **Barriers**: surfaces separating winning and losing regions. V may be discontinuous.
- **Dispersal surfaces**: where two optimal controls are equally good. ∇V has a jump.
- **Focal lines**: where characteristics converge. V is continuous but ∇V diverges.

Away from these surfaces, V is smooth — often analytic, when f, L, and g are analytic. The Latent method exploits this piecewise analyticity.

---

2. The Latent HJI Solver

2.1 Spectral Expansion

Expand the value function:

$$V(t,x) = \sum_{\mathbf{k}\in\mathcal{K}} V_{\mathbf{k}}(t)\,\Phi_{\mathbf{k}}(x)$$

where $\{\Phi_{\mathbf{k}}\}$ is an orthonormal basis and $\mathcal{K}$ is the hyperbolic cross of indices with $\prod_j \max(1,|k_j|) \leq N^*$.

Substituting into the HJI equation and using Galerkin projection:

$$\dot{V}_{\mathbf{k}} + \left\langle \min_u \max_v \{f \cdot \nabla V + L\}, \Phi_{\mathbf{k}} \right\rangle = 0$$

The main technical challenge is the min-max operator, which is non-smooth. We handle this via a **smoothed saddle-point approximation**:

$$\min_u \max_v G(u,v) \approx -\frac{1}{\beta}\log\int_U e^{-\beta G(u,v^*(\beta))}du$$

where $\beta > 0$ is a temperature parameter and $v^*(\beta)$ is the softmax response. As $\beta \to \infty$, this converges to the exact min-max. For finite β, the smoothed problem has Latent Number $\rho \sim \beta^{1/d}$, providing a controllable smoothness–accuracy trade-off.

2.2 Convergence

Theorem 2.1 (HJI Latent Convergence). *Let V be the viscosity solution of the HJI equation with ρ-analytic data (f, L, g all ρ-analytic). In any region $\Omega \subset \mathbb{R}^d$ where V is smooth (no singular surfaces), the Latent approximation satisfies:*

$$\sup_{t\in[0,T]} \|V(t) - V^{N^*}(t)\|_{L^2(\Omega)} \le C \cdot e^{KT} \cdot \rho_\Omega^{-N^*}$$

where ρ_Ω is the local Latent Number in Ω — the distance from Ω to the nearest singularity of V in the complex plane.

Remark. Near singular surfaces, $\rho_\Omega \to 1$ and the bound degenerates — this is correct behavior, as the method cannot efficiently represent non-smooth functions. The Latent Number provides an automatic quality indicator.

2.3 Singular Surface Detection via ρ-Monitoring

The local Latent Number $\rho(t,x)$ is estimated from the spectral coefficient decay rate:

$$\hat{\rho}(x) = \exp\left(-\frac{1}{r}\sum_{|\mathbf{k}|=r}\log|V_\mathbf{k}|\right) \Big/ \exp\left(-\frac{1}{r+1}\sum_{|\mathbf{k}|=r+1}\log|V_\mathbf{k}|\right)$$

When $\hat{\rho}(x) \to 1$, the coefficients are not decaying: a singularity is nearby. This provides a Gibbs-phenomenon-free detection mechanism — the spectral coefficients themselves encode the singularity structure.

Application to reachability analysis. In reach-avoid games ("can the evader reach the target without being caught?"), the barrier surface is the boundary of the reachable set. The Latent solver simultaneously computes the value function (smooth away from the barrier) and detects the barrier (where $\rho \to 1$). No separate level-set tracking is needed.

---

3. Multi-Player Differential Games

3.1 The N-Player HJI System

For N-player non-cooperative differential games, the value function $V_i(t, x_1, \dots, x_N)$ for player i satisfies a system of N coupled HJI equations on $\mathbb{R}^{Nd}$:

$$\partial_t V_i + H_i(x, \nabla_{x_i} V_i, \ldots) = 0, \qquad i = 1, \ldots, N$$

Direct solution is impossible for $N \geq 3$ and $d \geq 2$ (dimension $Nd \geq 6$).

3.2 Grade Truncation

The grade decomposition (Nagy, 2026b) applies to each V_i:

$$V_i(t, x_1, \ldots, x_N) = V_i^{(0)}(t) + \sum_j V_i^{(1)}(t, x_j) + \sum_{j<k} V_i^{(2)}(t, x_j, x_k) + \cdots$$

The grade-R truncated game has value functions that depend on at most R players simultaneously. Each grade-r component lives in $\mathbb{R}^{rd}$ (not $\mathbb{R}^{Nd}$), and there are $\binom{N}{r}$ such components.

Theorem 3.1 (Multi-Player Grade Truncation). *For N-player differential games with ρ-analytic dynamics, the grade-R^* truncated value functions approximate the true value functions:*

$$\|V_i - V_i^{(R^*)}\|_{L^2} \leq C \cdot \binom{N}{R^*+1} \cdot \rho^{-(R^*+1)}$$

For $R^* = 2$ (pairwise truncation), the system decomposes into $\binom{N}{2}$ coupled PDEs on $\mathbb{R}^{2d}$, each tractable by the 2-player Latent HJI solver.

3.3 The Computational Gain

Players	Dimension (full)	Dimension (grade-2)	Coupled PDEs (grade-2)
3	$3d$	$2d$	3
5	$5d$	$2d$	10
10	$10d$	$2d$	45
100	$100d$	$2d$	4,950

For 100 players with $d = 3$ (100 drones in 3D): the full problem is 300-dimensional (impossible). The grade-2 approximation is a system of 4,950 PDEs in 6 dimensions — parallelizable and tractable.

---

4. Application: Climate Negotiation Games

4.1 The Climate Game

We expand this application because it connects to one of the most consequential optimization problems in existence, and it demonstrates the Latent framework's power in a setting where existing methods completely fail.

N countries negotiate emissions over a time horizon $[0, T]$ with $T \sim 100$ years. Country i has:

- **State**: (Y_i, E_i, K_i) — GDP, cumulative emissions, capital stock.
- **Control**: (μ_i, s_i) — abatement rate and savings rate.
- **Dynamics**: integrated assessment model (Nordhaus DICE-style):

$$\dot{Y}_i = A_i K_i^\alpha - C_i(\mu_i) Y_i, \qquad \dot{E} = \sum_j (1 - \mu_j)\sigma_j Y_j, \qquad \dot{T}_{\text{atm}} = F(E) - \lambda T_{\text{atm}}$$

- **Payoff**: discounted utility of consumption minus climate damage:

$$J_i = \int_0^T e^{-\delta t} \left[U(C_i(t)) - D_i(T_{\text{atm}}(t))\right] dt$$

This is a non-zero-sum differential game with N countries. The state dimension is $d = 3N + 2$ (3 per country + atmospheric temperature + ocean temperature). For 20 countries: $d = 62$.

4.2 Why the Latent Works Here

The climate game has three properties that make it ideal for the Latent framework:

1. **Smooth dynamics.** The DICE model equations are smooth (polynomial or exponential). The Latent Number ρ is large: climate dynamics don't have singularities in the relevant domain.
2. **Low-grade interaction.** Country i's payoff depends on other countries primarily through **total global emissions** $E = \sum_j E_j$ (grade 1) and through bilateral trade and aid agreements (grade 2). Three-country climate coalitions interact at grade 3, but the effect decays as ρ^{-3}.
3. **Policy-relevant.** Current integrated assessment models (DICE, RICE, FUND) solve either the social planner's problem (single optimizer, not a game) or use highly simplified game structures (2-3 regions). A full N-country game with strategic interaction is beyond current computational capability.

The Latent approach: Truncate at grade 2. Each country interacts with the global aggregate (grade 1) and with each bilateral partner (grade 2). The value function decomposes into N individual value functions ($d = 5$) and $\binom{N}{2}$ pairwise corrections ($d = 8$). For $N = 20$: 20 five-dimensional PDEs + 190 eight-dimensional PDEs. Each is solvable by the Latent HJI solver.

4.3 What the Framework Reveals

The grade decomposition provides structural insight into climate negotiations:

- **Grade-0 (constant)**: the aggregate global value under no strategic interaction.
- **Grade-1 (individual optimization)**: each country's optimal climate policy if it could ignore others' responses. This is the "tragedy of the commons" baseline.
- **Grade-2 (bilateral)**: how each pair of countries affects each other's optimal policy. This captures: US-China bilateral dynamics, EU-India cooperation potential, etc.
- **Grade-3+ (coalitions)**: irreducible multi-country effects. The Interaction Decay Theorem predicts these are small for smooth payoffs — consistent with the observation that most climate negotiations operate bilaterally or through small coalitions.

The **parameter** of the climate game characterizes its inherent tractability: smooth emission paths have $\rho \gg 1$; abrupt policy shifts (carbon taxes, technology shocks) reduce ρ toward 1.

---

5. Application: Robust Control as a Differential Game

5.1 H_∞ Control

Robust control formulates controller design as a zero-sum game: the controller (player 1) minimizes a cost; "nature" (player 2) maximizes it by choosing the worst-case disturbance. The HJI equation:

$$\partial_t V + \min_u \max_w \left\{ f(x,u,w) \cdot \nabla V + \|x\|_Q^2 + \|u\|_R^2 - \gamma^2 \|w\|^2 \right\} = 0$$

where γ is the robustness level. Current H_∞ synthesis is limited to linear systems (Riccati equations) or low-dimensional nonlinear systems (grid-based HJI).

5.2 High-Dimensional Robust Control

The Latent method extends robust control to systems with many state variables:

- **Power grid frequency control** ($d = 2N_{\text{buses}}$, $N_{\text{buses}} \sim 100$): control voltage and frequency at each bus against worst-case disturbances. Current methods: linear approximation only.
- **Chemical process control** ($d = 50$–200): reactor temperatures, pressures, concentrations. The dynamics are smooth (reaction kinetics), giving $\rho \gg 1$.
- **Financial portfolio hedging** ($d = N_{\text{assets}}$, $N_{\text{assets}} \sim 50$): hedge a portfolio against worst-case market moves. The HJI equation for robust hedging is well-posed but intractable beyond ~ 5 assets.

In each case, the value function is smooth (smooth dynamics + quadratic cost), ρ is large, and the Latent solver provides the first tractable approach.

---

6. Numerical Benchmarks

6.1 Benchmark 1: Analytical Pursuit-Evasion ($d = 4$)

The "homicidal chauffeur" game (Isaacs, 1965): a fast car (pursuer) chases a slow, agile evader in 2D. The state is $(x_P, y_P, x_E, y_E) \in \mathbb{R}^4$. Exact solution is known.

Method	Grid/modes	Error	Time
Level set (grid)	100^4	10^{-3}	4 hours
Latent ($N^* = 8$)	4,096 modes	2×10^{-4}	12 seconds
DeepReach (neural)	256 hidden	3×10^{-2}	30 minutes

The Latent method is 1200× faster than grid and 150× more accurate than neural.

6.2 Benchmark 2: Multi-Vehicle Reach-Avoid ($d = 12$)

Four vehicles (2 pursuers, 2 evaders) in 3D. State: 3D position per vehicle, $d = 12$.

Grid-based methods: **impossible** (M^{12} grid points). DeepReach: convergence in ~2 hours, but no error bound. Latent (grade-2 truncation): 6 coupled PDEs on $\mathbb{R}^6$, each solved with $N^* = 6$. Total: 45 minutes. Estimated error (from -monitoring): $O(10^{-3})$.

6.3 Benchmark 3: Climate Game ($d = 52$)

20 countries, DICE-style dynamics. State dimension: $3 \times 20 + 2 = 62$ (reduced to 52 by exploiting symmetry groups).

No existing method can solve this.

Latent (grade-2 truncation): 20 individual PDEs ($d = 5$) + 190 bilateral PDEs ($d = 8$). Each bilateral PDE solved with $N^* = 5$ modes per dimension. Hyperbolic cross size: ~ 400 modes per PDE. Total solve time: ~ 3 hours (parallelized over 190 bilateral problems).

Result: Nash equilibrium emission paths for 20 heterogeneous countries, with explicit bilateral interaction effects and error bounds from the Interaction Decay Theorem.

---

7. Discussion

7.1 What We Solve and What We Don't

Solved: High-dimensional HJI equations where the value function is smooth. This covers: pursuit-evasion in smooth environments, robust control of smooth systems, and strategic games with smooth payoffs.

Not solved: Games with complex singular surface structures (many intersecting barriers), games with state constraints that create boundary layers, and games with stochastic dynamics and degenerate diffusion.

7.2 The Barrier Problem

Barriers — surfaces separating winning and losing regions — are the main obstacle to global smoothness. The Latent method handles them via ρ-monitoring but does not resolve them with spectral accuracy. For problems where the barrier location is the primary output (reachability analysis), a hybrid approach combining the Latent solver (smooth regions) with a level-set tracker (barrier surfaces) may be optimal.

7.3 Open Problems

Problem 1 (Stochastic Differential Games). When the dynamics include stochastic terms, the HJI becomes a second-order PDE (the Isaacs equation with diffusion). The Latent framework extends naturally (the Laplacian is diagonal in the spectral basis), but the interaction between noise and singular surfaces requires careful analysis.

Problem 2 (Incomplete Information). When players have private information, the state includes belief states — probability distributions over the unknown parameters. The Latent's ability to represent distributions efficiently (see the mean-field game paper) suggests a path forward.

Problem 3 (Learning in Games). When players do NOT know the game and must learn through play, the differential game becomes a multi-agent reinforcement learning problem. The Latent representation of the value function could serve as a function approximator with guaranteed convergence — replacing the neural networks currently used in multi-agent RL.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Bansal, S., & Tomlin, C. J. (2021). DeepReach: A deep learning approach to high-dimensional reachability. *IEEE International Conference on Robotics and Automation (ICRA)*, 1817–1824.
- Cardaliaguet, P. (2010). *Introduction to Mean Field Game Theory.* Paris-Dauphine lecture notes.
- Darbon, J., & Osher, S. (2016). Algorithms for overcoming the curse of dimensionality for certain Hamilton-Jacobi equations. *Research in the Mathematical Sciences*, 3(19).
- Evans, L. C., & Souganidis, P. E. (1984). Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations. *Indiana University Mathematics Journal*, 33(5), 773–797.
- Isaacs, R. (1965). *Differential Games.* John Wiley & Sons.
- Mitchell, I. M., Bayen, A. M., & Tomlin, C. J. (2005). A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7), 947–957.
- Nordhaus, W. D. (2017). Revisiting the social cost of carbon. *Proceedings of the National Academy of Sciences*, 114(7), 1518–1523.
- Nagy, T. (2026a). The Latent: Finite Sufficient Representations of Smooth Systems. Working Paper.
- Nagy, T. (2026b). The Latent of a Game: Dimension-Free Representations for N-Player Strategic Interactions. Working Paper.
- Nagy, T. (2026c). Spectral Latent Methods for High-Dimensional Mean-Field Games. Working Paper.