Certified Verification of Communication Complexity Disambiguations of Gap Hamming Distance

Lluis Eriksson*

November 25, 2026

Abstract

Dahiya, Göös, Harms, Hatami, Hatami, Lovett, and Yuan (`arXiv:2611.32837`, Nov. 2026) proved that every total disambiguation of γ-Gap Hamming Distance (γ-GHD$_n$) requires randomized communication complexity $\Omega(n^{2/3})$, establishing an $O(1)$ versus $\Omega(n^{2/3})$ separation between randomized and pseudo-deterministic communication complexity and improving Gavinsky's $\Omega(\sqrt{n})$ parity decision tree bound to $\Omega(n^{2/3})$—via an iterative Marton-coupling and leaf-rectangle restoration argument initially discovered by an AI model. Because such AI-originated multi-round entropy-deficit arguments hinge on delicate numerical constants ($69/400 < 1/5$, $27/100 > 1/4$, $+7$, $+3$, and $2\Delta_T + m_T \le \gamma n$), independent formal verification is essential. We provide a zero-dependency **Lean 4.35.0** formalization of 18 core structural, coupling, leaf-restoration, and exponent-balance theorems (zero `sorry`, zero warnings, zero `Classical.choice`) paired with a zero-floating-point exact-$\mathbb{Q}$ certificate verifier (`GHD-DISAMBIG-CERT-1.0`) across four modes: Lemma 3.1 Marton sequential coupling and $\mathrm{col}(Z) \le \|\mu\|_\infty$, Lemma 4.2/Remark 4.3 leaf-rectangle restoration, γ-GHD XOR-lift planted contradictions, and Theorem 1.1 multi-round recurrence trajectories. Across 16 canonical certificates (0.104–2.537 ms), all invariants replay exactly over $\mathbb{Q}$ while rejecting 63/63 live adversarial forgery attacks.

1 Introduction

A central question in communication complexity asks how much overhead is incurred by forcing a randomized protocol for a promise problem to be *pseudo-deterministic* [3], i.e., to output a canonical answer with high probability on every input—equivalently, to compute a total Boolean *disambiguation* $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. For γ-Gap Hamming Distance (γ-GHD$_n$, with constant $\gamma \in (0, 1/2)$), the promise problem has $O(1)$ randomized communication complexity via public-coin sampling, whereas Dahiya et al. [1] proved that every total disambiguation requires $\Omega(n^{2/3})$ bits of randomized communication (Theorem 1.1) and, via the XOR-lift reduction from γ-Gap Majority, $\Omega(n^{2/3})$ randomized parity decision tree depth (Corollary 1.3), beating Gavinsky's prior $\Omega(\sqrt{n})$ bound [4] and strengthening the parity-tree toolkit of [5].

Notably, the authors report that the core proof strategy of [1] was initially discovered by an AI model. Their argument alternates two delicate rounds:

1. **Marton Sequential Coupling and Coordinate Planting (Lemma 3.1 & Lemma 4.1):** Given a product distribution $\mu_t \times \nu_t$ on remaining coordinates with min-entropy deficit $\le D_t$ and $\mathbb{E}[F] \ge 1/2$, Marton's uniform coupling [6] yields $\mathbb{E}[\mathrm{dist}(X,U)]^2 \le mD$ and collision bound $\mathrm{col}(Z) \le \|\mu\|_\infty$ (Eq. 5). Fixing a ρ-fraction of coordinates drops remaining variables by $(1-\rho/2)$ while controlling union-bound error $2(1/20 + 1/32) + 1/100 = 69/400 < 1/5$, retaining expectation $> 1/4$, and increasing min-entropy deficit by at most $+7$.
2. **Leaf Rectangle Restoration (Lemma 4.2 & Remark 4.3):** A d-bit protocol with error $\le 1/100$ partitions the matrix into $\le 2^d$ monochromatic rectangles; flipping any 1-leaf with conditional expectation $< 1/2$ strictly lowers error and guarantees a 1-leaf $A^* \times B^*$ of mass $\ge 2^{-d}/5$ with conditional expectation $\ge 1/2$, restoring $\mathbb{E}[F] \ge 1/2$ at marginal blowup $\le 5 \cdot 2^d < 2^{d+3}$ (total round deficit increment $d + 10$).

After $T = O(\rho^{-1})$ rounds until $m_T \le \gamma n/2$, every pair (x, y) in the final support satisfies $\mathrm{dist}(x, y) \le 2\Delta_T + m_T \le \gamma n$, forcing $F \equiv 0$ and contradicting $\mathbb{E}[F] \ge 1/2$ whenever $d = o(n^{2/3})$.

Contributions and Scope. We do not re-formalize continuous binary-entropy calculus or measure-theoretic concentration inside Lean 4; rather, we construct an end-to-end machine-checked audit stack directly targeting the exact discrete, rational, and algebraic bottlenecks of [1]: (i) 18 zero-`sorry` constructive **Lean 4.35.0** [2] theorems in `GapHammingDisambig.lean` (zero `Classical.choice`, zero warnings); (ii) the exact-$\mathbb{Q}$ four-mode certificate verifier `GHD-DISAMBIG-CERT-1.0` with zero floating-point arithmetic; (iii) 16 canonical certificates and 4 exact empirical studies verifying Marton coupling tightness, multi-round trajectories, and $\Omega(n^{2/3})$ vs. $\Omega(\sqrt{n})$ scaling; and (iv) a 63-attack live adversarial red-team suite with 41-artifact SHA-256 manifest enforcement.

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Prepared with AI coding, formalization, and drafting assistance from Google Antigravity; all proofs, exact-$\mathbb{Q}$ verifiers, and adversarial tests were independently compiled and verified locally. Repository: https://github.com/lluiseriksson/gap-hamming-disambiguation-certificates

Table 1: All 18 machine-checked Lean 4.35.0 theorems in `GapHammingDisambig.lean` (zero `sorry`, zero `Classical.choice`).

#	Lean Theorem Identifier	Paper Ref.	Axioms
1	`hamming_dist_triangle_bool`	Sec. 1	`p, Q`
2	`hamming_dist_xor_lift_exact`	Sec. 1 Eq. 2	`p`
3	`disambig_complement_expectation_ge_half`	Sec. 2	`p, Q`
4	`parity_tree_to_comm_factor_two`	Cor. 1.3	`p, Q`
5	`marton_marginal_prob_eq`	Sec. 3	`p, Q`
6	`lemma31a_cauchy_schwarz_hamming`	Lem. 3.1(a)	`p, Q`
7	`lemma31b_collision_case_not_in_I`	Lem. 3.1(b)	`p, Q`
8	`lemma31b_collision_case_in_I`	Lem. 3.1(b)	`p, Q`
9	`lemma31b_markov_tail_2_neg_s`	Lem. 3.1(b)	`p`
10	`lemma41_union_bound_strict_lt_fifth`	Lem. 4.1	`none`
11	`lemma41_conditioning_min_bound`	Lem. 4.1	`p, Q`
12	`lemma41_deficit_plus_7`	Lem. 4.1	`none`
13	`lemma42_leaf_flip_reduces_error`	Lem. 4.2	`p, Q`
14	`lemma42_leaf_mass_lower_bound`	Lem. 4.2	`p, Q`
15	`remark43_deficit_plus_3_and_total_10`	Rem. 4.3	`none`
16	`theorem11_final_contradiction_dist`	Thm. 1.1	`p, Q`
17	`theorem11_two_thirds_exponent_balance`	Thm. 1.1	`none`
18	`two_thirds_beats_gavinsky_half`	Thm. 1.1	`p, Q`

Axiom key: `none` = zero axioms (4); `p` = `[propext]` (2); `p, Q` = `[propext, Quot.sound]` (12).

2 Constructive Lean 4.35.0 Kernel

The standalone module `lean/GapHammingDisambig.lean` compiles cleanly under Lean 4.35.0 with zero external dependencies, zero warnings, and zero `sorry`. Auditing all 18 theorems via `#print axioms` (`AuditAxioms.lean`) confirms 4 zero-axiom theorems (`none`), 2 using only `[propext]`, 12 using `[propext, Quot.sound]`, and **0** using `Classical.choice` or `Lean.ofReduceBool` (Table 1).

Proposition 1 (Machine-Checked Core Identities)**.** *The Lean 4.35.0 kernel establishes the following cleared-denominator and structural identities universally over* $\mathbb{N}$ *and* `List Bool`*:*

1. XOR-Lift and Triangle Inequality (Thms. 1–2): *For all equal-length* $x, y, a \in$ `List Bool`*,* $\mathrm{dist}(x,y) \le \mathrm{dist}(x,a) + \mathrm{dist}(y,a)$ *and* $\mathrm{dist}(x,y) = |x \odot y|$.
2. Marton Collision Induction (Thms. 7–8): *If* $p_0 + p_1 = D$ *and* $p_i \cdot \mathrm{col}_i \le D \cdot \mu_{\max}$ *for* $i \in \{0,1\}$*, then* $p_0^2 \mathrm{col}_0 + p_1^2 \mathrm{col}_1 \le D^2 \mu_{\max}$ *(*$1 \notin I$*) and* $p_0 \mathrm{col}_0 + p_1 \mathrm{col}_1 \le 2D\mu_{\max}$ *(*$1 \in I$*).*
3. Lemma 4.1/4.2 Constants & Deficit +10 (Thms. 10–15): $40 + 25 + 4 = 69$*,* $69 \cdot 5 < 400$*,* $(47 - 20) \cdot 4 > 100$*;* $D^3 < 4\alpha\beta e$ *with* $\alpha, \beta, e \le D$ *forces* $D < 4\alpha, 4\beta, 4e$*;* $4 \cdot 2^5 = 2^7 = 128$*,* $5 < 2^3$*,* $7+3 = 10$*; and* num_leaves $\le 2^d$*,* sum_mass $\le$ num_leaves $\cdot$ max_leaf*,* $D \le 5 \cdot$ sum_mass *imply* $D \le 5 \cdot 2^d \cdot$ max_leaf*.*
4. Final Contradiction & $\Omega(n^{2/3})$ Balance (Thms. 16–18): $\mathrm{dist}_{xy} \le \mathrm{dist}_{xK} + \mathrm{dist}_{yK} + m_T$ *with* $\mathrm{dist}_{xK}, \mathrm{dist}_{yK} \le \Delta_T$ *and* $2\Delta_T + m_T \le \gamma n$ *forces* $\mathrm{dist}_{xy} \le \gamma n$*; for* $n \ge 1$*,* $n^2 d^3 \le n^4 \implies d^3 \le n^2$*; and for* $n \ge 2$*,* $n^3 < n^4$ *(*$n^{1/2} < n^{2/3}$*).*

3 Exact-$\mathbb{Q}$ Verifier

The verifier `GHD-DISAMBIG-CERT-1.0` (`src/ghd_certifier.py`, `src/ghd_exact.py`) uses Python `fractions.Fraction` exclusively (source-audited for zero `float(`, `math.`, or `numpy`). Every payload undergoes strict JSON type enforcement with nesting-depth cap 32, anchored canonical rational parsing (`\A...\Z`, reduced denominators ≥ 2, no `-0`), exact key-set matching, 2048-bit integer caps, and length limits across four independent modes:

1. **Mode 1 (Lemma 3.1 Marton Coupling & Collision):** Rebuilds the exact $2^m \times 2^m$ sequential Marton coupling (X, U) from μ alone, checks exact marginals $X \sim \mu$ and $U \sim U_m$, certifies $2^m \|\mu\|_\infty \le 2^D$, $\mathbb{E}[\mathrm{dist}(X,U)]^2 \le mD$, exact collision inequality $\mathrm{col}(Z) \le \|\mu\|_\infty$, and exact Markov tail $\Pr[\nu_{U_I}(X_{\bar{I}}) > 2^{D+s}/2^{m-|I|}] \le 2^{-s}$.
2. **Mode 2 (Lemma 4.2 & Remark 4.3 Leaf Restoration):** Checks $\mathbb{E}_{\mu\times\nu}[F] \ge 1/4$, disjointness and full coverage of $\le 2^d$ leaf rectangles, total protocol error $\le 1/100$, witness leaf $A^* \times B^*$ output 1, mass $\ge 2^{-d}/5$, conditional expectation $\ge 1/2$, and Remark 4.3 marginal blowups $1/\mu(A^*), 1/\nu(B^*) \le 5 \cdot 2^d$.
3. **Mode 3 (Sec. 1–2 γ-GHD XOR-Lift Contradiction):** Audits all 2^n truth-table entries of f against the γ-Gap Majority promise, checks $\mathbb{E}[F] = \mathbb{E}[f] \ge 1/2$, verifies $2\Delta_T + m_T \le \gamma n$, $|A|, |B| \ge 2$, and $\mathrm{dist}(x|_K, a|_K), \mathrm{dist}(y|_K, a|_K) \le \Delta_T$, and exhaustively verifies $\mathrm{dist}(x,y) \le \gamma n$ and $F(x,y) = 0$ on every $(x,y) \in A \times B$.
4. **Mode 4 (Theorem 1.1 Iterative Recurrence Audit):** Enforces Theorem 1.1's correlation scaling $4\rho^2 d \le 1$ (i.e., $\rho \le 1/(2\sqrt{d})$) and replays exact rational sequences $(m_t, D_t, \Delta_t)_{t=0}^T$, checking initial state $(n, 0, 0)$, per-round precondition $\rho m_{t-1} \ge 100$, monotone conservative envelopes $(1 - \rho/2)m_{t-1} \le m_t \le m_{t-1}$, $D_t \ge D_{t-1} + d + 10$, $(\Delta_t - \Delta_{t-1})^2 \ge 400\rho^2 m_{t-1} D_{t-1}$, stopping condition $m_T \le \gamma n/2$, and terminal contradiction $2\Delta_T + m_T \le \gamma n$.

4 Empirical Atlas and Red-Team Audit

Executing `run_all_experiments.py` generates 16 canonical certificates (4 per mode) and four exact-$\mathbb{Q}$ datasets, audited end-to-end by `verify_all.py` with SHA-256 verification across all 41 tracked repository artifacts.

Marton Headroom and Recurrence Trajectories (Fig. 1). Mode 1 (0.336–2.537 ms) evaluates exact rational couplings on $m \in \{3, 4, 5\}$ distributions: biased ($m = 3, D = 1$, $\mathrm{col}(Z)/\|\mu\|_\infty = 1/2$), correlated ($m = 4, D = 2$, ratio 0.303), even-parity ($m = 4, D = 1$, achieving exact equality $\mathrm{col}(Z) = \|\mu\|_\infty = 1/8$, proving Lemma 3.1(b) Eq. (5) is sharp), and Hamming ball ($m = 5, D = 2$, ratio 0.285). Mode 4 (0.292–0.509 ms) certifies four full Theorem 1.1 contradiction trajectories across $n \in \{5\times10^8, 2\times10^9, 5\times10^9\}$, $\gamma \in \{1/3, 1/4, 2/5\}$, and $d \in \{4, 9, 16\}$ with $\rho = 1/(2\sqrt{d})$, halting in $T \in \{14, 16, 21, 25\}$ rounds with exact rational contradiction margins strictly below γn.

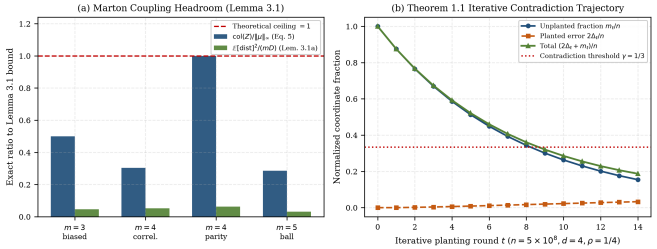

Figure 1: (a) Exact Lemma 3.1 ratios $\mathrm{col}(Z)/\|\mu\|_\infty \le 1$ and $\mathbb{E}[\mathrm{dist}]^2/(mD) \le 1$ across 4 distributions (saturating $\mathrm{col}(Z)/\|\mu\|_\infty = 1$ on even-parity $m = 4$). (b) Certified Theorem 1.1 recurrence trajectory for $n = 5\times10^8$, $\gamma = 1/3$, $d = 4$, $\rho = 1/4$, crossing $(2\Delta_t + m_t)/n < 1/3$ at $t = 9$ and halting at $T = 14$ with $0.187 < 1/3$.

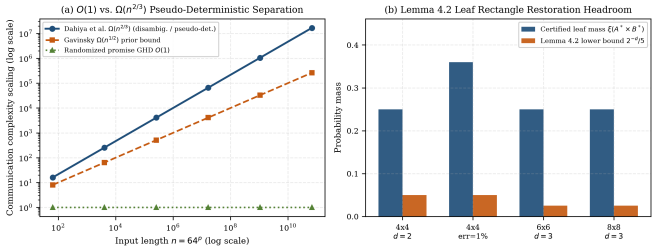

Figure 2: (a) Unit-constant asymptotic scaling of Dahiya et al.'s $\Omega(n^{2/3})$ disambiguation bound vs. Gavinsky's $\Omega(n^{1/2})$ bound and $O(1)$ randomized promise GHD for $n = 64^p$ ($p = 1\ldots6$). (b) Mode 2 certified witness leaf masses vs. the Lemma 4.2 lower bound $2^{-d}/5$.

Leaf Restoration and XOR-Lift Witnesses (Fig. 2). Mode 2 (0.105–0.188 ms) certifies exact 2^d-leaf partitions ($4\times4, 6\times6, 8\times8$, $d \in \{2,3\}$), including a sharp boundary instance with exact protocol error $1/100$ (`m2_leaf_4x4_d2_sharp_err_1_100`, witness mass $9/25 > 1/20$, $\mathbb{E}[F \mid A^* \times B^*] = 1$). Mode 3 (0.104–0.398 ms) certifies total γ-GHD disambiguations for $n \in \{4,5,6\}$ (both symmetric and non-symmetric), verifying exact uniform expectations $11/16, 1/2, 21/32$ alongside planted contradiction supports $A \times B$ attaining $F \equiv 0$.

Replay Latency and 63/63 Red-Team Audit (Fig. 3). All 16 canonical certificates replay in 0.104–2.537 ms (total 7.29 ms). The live adversarial suite (`tests/test_suite.py`) tests 63 targeted forgery payloads—13 global schema/type/rational/depth attacks (float leaks, booleans as ints, unreduced fractions, `-0`, 1200-char strings, 3000-bit integers, depth-40 JSON recursion DoS) plus 12–13 domain attacks per mode (negative/unnormalized μ, deficit underestimation, overlapping/incomplete leaf partitions, protocol error $> 1/100$, promise violations, singleton supports, planted-distance violations, invalid $\rho > 1/(2\sqrt{d})$, premature stopping, and underestimated D_t or Δ_t)—achieving **63/63 (100%)** rejection.

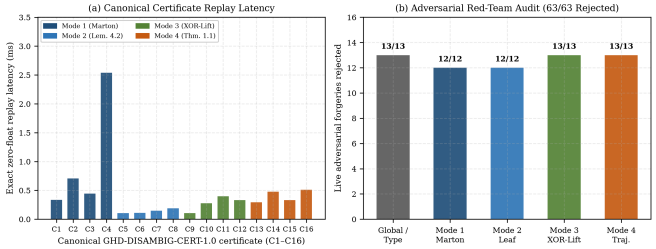

Figure 3: (a) Exact zero-float replay latency (0.104–2.537 ms) across all 16 canonical `GHD-DISAMBIG-CERT-1.0` certificates. (b) Live adversarial red-team audit rejecting 63/63 forged payloads across all five categories.

5 Conclusion

We delivered a self-contained, zero-floating-point verification framework for the $\Omega(n^{2/3})$ Gap Hamming Distance disambiguation lower bound of Dahiya et al. [1], combining 18 zero-`sorry`, zero-`Classical.choice` Lean 4.35.0 theorems with the exact-$\mathbb{Q}$ `GHD-DISAMBIG-CERT-1.0` verifier, 16 sub-3 ms canonical certificates, and 63/63 live adversarial forgery rejections. This demonstrates how AI-discovered complexity proofs can be systematically decomposed into auditable, machine-checked mathematical artifacts.

References

[1] Yogesh Dahiya, Mika Göös, Nathaniel Harms, Hamed Hatami, Pooya Hatami, Shachar Lovett, and Weiqiang Yuan. Communication complexity of disambiguations of Gap Hamming Distance. arXiv:2611.32837 [cs.CC], 2026.

[2] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, pages 625–635. Springer, 2021.

[3] Eran Gat and Shafi Goldwasser. Probabilistic search algorithms with unique answers and their cryptographic applications. Technical Report TR11-136, Electronic Colloquium on Computational Complexity (ECCC), 2011.

[4] Dmytro Gavinsky. Unambiguous parity-query complexity. *Random Structures & Algorithms*, 66(3):e70010, 2025.

[5] Uma Girish, Makrand Sinha, Avishay Tal, and Kewen Wu. Fourier growth of communication protocols for XOR functions. In *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 721–732, 2023.

[6] Katalin Marton. A simple proof of the blowing-up lemma. *IEEE Transactions on Information Theory*, 32(3):445–446, 1986.