Constructive Lean 4 Kernel and Exact Zero-Float Certificates for OpenAI's Strongly Polynomial Linear Programming

Lluis Eriksson*

November 25, 2026

Abstract

A November 24, 2026 companion preprint from OpenAI [7] proposes a deterministic strongly polynomial-time algorithm for general linear programming (Smale's 9th problem [9]), extending to monotone linear complementarity problems, convex quadratic programming, max-norm contraction fixed points, and turn-based discounted and priority mean-payoff games. Rather than assuming peer-reviewed endorsement of its online logarithmic compression, recursive extremal-price search, or dissociated-minor counting pillars, we deliver an independent, explicitly scoped, zero-floating-point verification companion for its discrete LP trichotomy certificates (strict complementarity, improving rays, minimal-support Farkas), paired-subspace reciprocal diagonal scaling and division-free Newton's identities determinant, monotone LCP and convex QP KKT gap decomposition, and Bellman / clipped-update max-norm contraction mechanisms. Layer 1 (`OpenAIStronglyPolyLP.lean`) proves 18 standalone constructive Lean 4.35.0 algebraic and duality sub-lemmas (zero `sorry`, zero `axiom`, zero `Classical.choice`, zero warnings; 1 `[propext]` and 17 `[propext, Quot.sound]`). Layer 2 (`SPLP-SMALE9-CERT-1.0`) is a pure-stdlib Python `int`/`Fraction` 4-mode auditor certifying 16 canonical finite witnesses in 0.049–0.485 ms and rejecting 73/73 live adversarial forgeries with exact branch assertions (80/80 pytest checks).

1 Introduction and Scope Calibration

Khachiyan [5] and Karmarkar [4] established weakly polynomial algorithms for linear programming, while Tardos [10] and Frank–Tardos [2] achieved strong polynomiality when constraint coefficients have bounded bit lengths. Whether a strongly polynomial algorithm exists for arbitrary real linear programs is Smale's 9th problem [9]. OpenAI's November 24, 2026 companion preprint [7] proposes a deterministic strongly polynomial algorithm (Theorem 1) and extensions to monotone LCP (Theorem 5), convex QP (Corollary 6), max-norm contractions (Theorem 8), and turn-based games (Theorem 11), resting on four exact discrete-algebraic mechanisms:

1. **Theorem 1 LP Trichotomy Certificates (Secs. 1, 3.3):** Every real LP admits either a strictly complementary optimal pair ($Ax \le b, y \ge 0, A^\top y = c, b^\top y = c^\top x, b - Ax + y > 0$ [3]), an improving recession ray ($Ax \le b, Ar \le 0, c^\top r > 0$), or a minimal-support Farkas certificate ($y \ge 0, A^\top y = 0, b^\top y < 0, |\mathrm{supp}(y)| \le 1 + \mathrm{rank}(A)$).
2. **Reciprocal Diagonal Scaling & Division-Free Newton Determinant (Secs. 1.2, 3.1–3.4, Thm. 3):** Embedding orthogonal subspaces $W = C \times C^\perp$ and scaling $D_{i+} = 2^{-p_i}, D_{i-} = 2^{p_i}$ preserves $Q(y) = \sum_i y_{i+} y_{i-} = 0$ identically; Section 3.4 computes $k! \det(M)$ purely in the integer ring via Newton's trace recurrence without intermediate divisions.
3. **Monotone LCP & Convex QP KKT Gap Decomposition (Secs. 7.2–7.3, Thm. 5, Cor. 6):** For monotone M ($M + M^\top \succeq 0$) and solution (z^*, w^*), any feasible (z', w') obeys $z'^\top w' = z'^\top w^* + z^{*\top} w' + (z' - z^*)^\top M (z' - z^*)$ [6], while Corollary 6's skew-symmetric KKT block $\left(\begin{smallmatrix} R^\top G R & R^\top A^\top \\ -AR & 0 \end{smallmatrix}\right)$ stays monotone identically.
4. **Max-Norm Contractions & Clipped Bellman Halving (Secs. 8, 9.1, 9.3, Thms. 8, 11):** At any $x \ne v$, a γ-contraction moves every extremal-error coordinate strictly toward v [8]; Section 9.3's clipped update $u_i = [\mathrm{ext}_a(r_a + B_a y - y_i + P_a u)]_{-H}^{H}$ halves $\|v - (y + u)\|_\infty \le H$ and, when $y < v$, certifies the exact policy guarantee $z = y + u - H\mathbf{1} \le T(z)$.

Honest Two-Layer Verification Scope.

1. **Layer 1 (Lean 4.35.0 Kernel, Standalone Sub-Lemmas):** `lean/OpenAIStronglyPolyLP.lean` proves 18 standalone constructive algebraic and duality sub-lemmas (Table 1) extracted from Sections 1, 1.2, 2.3, 3, 3.1–3.4, 4.2, 5.1, 7.2–7.3, 8, and 9.1–9.3 of [7].
2. **Layer 2 (Exact Zero-Float Auditor):** `SPLP-SMALE9-CERT-1.0` audits 16 canonical finite witnesses across all three Theorem 1 LP branches, reciprocal scaling + Newton's identities ($k \le 5$), monotone LCP / convex QP KKT systems ($d \le 5$), and Bellman / clipped-update games ($n \le 4$, including Fig. 5's exact two-state game).
3. **Explicit Non-Claims:** We do *not* formalize in Lean 4 the online logarithmic compression procedure (Theorem 4, Section 5), recursive extremal-price search (Theo-

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Prepared with transparent AI coding, formalization, and drafting assistance via Google Antigravity; all Lean 4 theorems, exact zero-float verifiers, and adversarial tests were independently compiled and verified locally. Repository: github.com/lluiseriksson/openai-strongly-polynomial-lp-certificates.

rem 2, Section 6), LLL simultaneous Diophantine approximation, Blackwell vanishing-discount transfer, or the global strongly polynomial step-count bounds of Theorems 1, 5, 8, and 11; those algorithmic complexity pillars remain unverified outside [7].

2 Constructive Lean 4.35.0 Kernel

`lean/OpenAIStronglyPolyLP.lean` compiles standalone under Lean 4.35.0 [1] with zero external imports, zero warnings, zero `sorry`, and zero `Classical.choice`. Auditing via `AuditAxioms.lean` confirms 1 `[propext]` theorem (`P`: #5) and 17 `[propext, Quot.sound]` theorems (`P,Q`; Table 1).

Proposition 1 (Constructive Lean 4 Kernel Scope)**.** All 18 scalar sub-lemmas in Table 1 compile constructively under Lean 4.35.0 without external libraries or `Classical.choice`. Theorems 1–4 formalize zero duality gap, strict complementarity positivity, circuit-imbalance growth $u^2 - 1 \geq 3$, and the scalar positivity obstruction $v_i w_i + r > 0$; Theorems 5–9 prove reciprocal-scaling product invariance ($dr = 1$), the cleared rational Lemma 3.1 lower bound ($0 \leq e_n/e_d \leq 1/20$) and visibility margin, codimension support arithmetic, and the 2×2 division-free Newton determinant identity; Theorems 10–12 formalize positive scale invariance, geometric doubling dominance, and surplus-feedback monotonicity; Theorems 13–14 prove the scalar monotone LCP gap split and skew KKT block cancellation; and Theorems 15–18 formalize the cleared 7δ contraction movement lower bound, triangle-plus-contraction residual-to-distance derivation ($\gamma D \leq e$), strict discount contraction gap, and policy-width halving.

3 Exact Zero-Float Certificate Auditor

The modules `splp_exact.py` and `splp_certifier.py` (under `src/`) use exclusively Python `int` and `fractions.Fraction` (statically audited via `ast.walk` for zero float/complex literals, `/`, `/=`, or `math` imports). Strict JSON parsing enforces `MAX_JSON_BYTES`$= 10^6$, rejects floats, `NaN`/`Infinity`, and duplicate keys via `object_pairs_hook`, and requires canonical lowest-terms rationals ($\gcd(p,q) = 1, 0 < q \leq 10^{12}, |p| \leq 10^{12}$).

Definition 1 (`SPLP-SMALE9-CERT-1.0` Predicates)**.** A certificate with `protocol="SPLP-SMALE9-CERT-1.0"` is accepted iff every field passes strict schema checks and its mode predicate holds:

- **Mode 1 (`lp_trichotomy_certificate_cert`):** Audits exact Theorem 1 rational witnesses for `optimal` ($Ax \leq b, y \geq 0, A^\top y = c, c^\top x = b^\top y$, strict complementarity $b - Ax + y > 0$), `unbounded` ($Ax \leq b, Ar \leq 0, c^\top r > 0$), or `infeasible` ($y \geq 0, A^\top y = 0, b^\top y < 0, |\mathrm{supp}(y)| \leq 1 + \mathrm{rank}(A)$).
- **Mode 2 (`reciprocal_scaling_and_newton_det_cert`):** Checks JSON row counts $|C| + |C^\perp| = N$ before rational matrix parsing, audits orthogonal complementary bases $C, C^\perp \subset \mathbb{Q}^N$, checks nonzero reciprocal scaling $D_{i+} = 2^{-p_i}, D_{i-} = 2^{p_i}$ preserving $Q(yD) = 0$ across all basis pairs, and verifies exact agreement between Section 3.4's division-free Newton's identities $k! \det(M)$ and Gaussian elimination ($2 \leq k \leq 5$).
- **Mode 3 (`monotone_lcp_and_qp_kkt_cert`):** Rejects identically zero M, verifies $M + M^\top \succeq 0$ via exact rational symmetric pivoted congruence elimination, certifies exact LCP solution ($z^*, w^* \geq 0, w^* = Mz^* + q, z^{*\top} w^* = 0$), and audits Section 7.2's gap split $z'^\top w' = z'^\top w^* + z^{*\top} w' + (z' - z^*)^\top M(z' - z^*)$ on a distinct feasible probe ($z' \neq z^*, z' \geq 0, w' = Mz' + q \geq 0$).
- **Mode 4 (`max_norm_contraction_and_game_cert`):** Verifies exact Bellman fixed point $T(v) = v$, Theorem 8 contraction $\|T(x) - v\|_\infty \leq (1-\gamma)\|x - v\|_\infty$ with strict extremal-coordinate sign agreement ($x \neq v$), and Section 9.3's clipped fixed-point Eq. (9.3)–(9.5) ($y \neq v$, $u_i = [\mathrm{ext}_a(r_a + B_a y - y_i + P_a u)]_{-H}^{H}$ using *undiscounted* transition row P_a, $\|v - (y+u)\|_\infty \leq H$, and $z = y + u - H\mathbf{1} \leq T(z)$ when $y < v$).

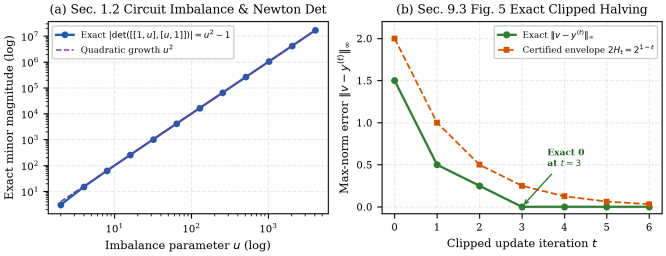

Figure 1: **(a)** Exact Section 1.2 circuit-imbalance minor $|\det([[1,u],[u,1]])| = u^2 - 1$ via division-free Newton's identities across $u = 2^1 \ldots 2^{12}$. **(b)** Exact rational Section 9.3 Figure 5 clipped-update trajectory converging monotonically to exact 0 at $t = 3$.

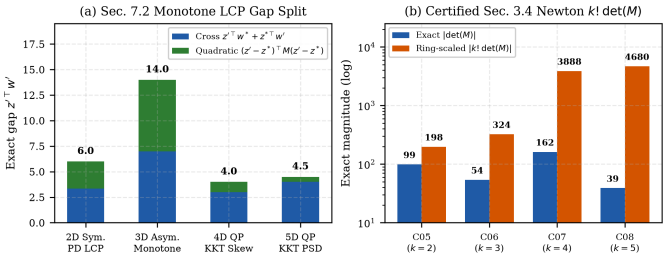

Figure 2: **(a)** Exact Section 7.2 monotone LCP gap split $z'^\top w' = (z'^\top w^* + z^{*\top} w') + (z' - z^*)^\top M(z' - z^*)$ across 4 canonical LCP/QP-KKT witnesses. **(b)** Certified exact $|\det(M)|$ vs. division-free Newton ring-scaled $|k! \det(M)|$ across `C05`–`C08` ($k \in \{2,3,4,5\}$).

Table 1: Complete inventory of all 18 constructive Lean 4.35.0 sub-lemmas in `lean/OpenAIStronglyPolyLP.lean` (zero `sorry`, zero `axiom`, zero `Classical.choice`, zero warnings; `[propext]` (P) for #5; `[propext, Quot.sound]` (P,Q) for the remaining 17).

#	Theorem Identifier	Ref. [7]	Exact Machine-Checked Statement & Hypotheses	Ax.
1	`sec1_strong_duality_gap_zero`	Thm. 1	$\forall b_y, c_x \in \mathbb{Z},\ b_y = c_x \implies b_y - c_x = 0.$	P,Q
2	`sec1_strict_complementarity_pos`	Thm. 1	$\forall s, y \in \mathbb{Z},\ 0 \le s \wedge 0 \le y \wedge (0 < s \vee 0 < y) \implies 0 < s + y.$	P,Q
3	`sec12_circuit_imbalance_det`	Sec. 1.2	$\forall u \in \mathbb{N},\ 2 \le u \implies 4 \le u^2 \wedge 3 \le u^2 - 1$ (cleared $\lvert\det\rvert$).	P,Q
4	`sec3_orthogonality_excludes_both_sides`	Sec. 3	$\forall v_i, w_i, r \in \mathbb{Z},\ 0 < v_i \wedge 0 < w_i \wedge 0 \le r \implies 0 < v_i w_i + r.$	P,Q
5	`sec31_reciprocal_scaling_preserves_product`	Sec. 3.1	$\forall y_p, y_m, d, r \in \mathbb{Z},\ dr = 1 \implies (y_p d)(y_m r) = y_p y_m\ (Q(yD) = Q(y)).$	P
6	`lem31_midpoint_quadratic_lower_bound`	Lem. 3.1	$\forall e_n, e_d \in \mathbb{Z},\ 0 \le e_n \wedge 0 < e_d \wedge 20e_n \le e_d \implies 0 < 20e_d \le 25e_d - 100e_n \le 25e_d.$	P,Q
7	`sec32_visibility_slack_threshold`	Sec. 3.2	$\forall e_{12} \in \mathbb{Z},\ e_{12} \le 1 \implies 3 \le 4 - e_{12}$ (cleared $1/3 - \text{err} \ge 1/4$).	P,Q
8	`sec33_farkas_support_bound`	Sec. 3.3	$\forall r, s_y \in \mathbb{N},\ s_y + 1 \le (1 + r) + 1 \implies s_y \le 1 + r$ (support arithmetic).	P,Q
9	`sec34_newton_identity_det2`	Sec. 3.4	$\forall a, b, c, d \in \mathbb{Z},\ (a + d)^2 - (a^2 + 2bc + d^2) = 2(ad - bc)\ (2!\det M).$	P,Q
10	`sec42_linear_type_scale_invariant`	Lem. 4.1	$\forall s, x, y \in \mathbb{Z},\ 0 < s \wedge x < y \implies sx < sy$ (sign preservation).	P,Q
11	`lem51_dissociated_geometric_dominance`	Lem. 5.1	$\forall k \in \mathbb{N},\ 2^k - 1 < 2^k$ (doubling sequence subset-sum dominance).	P,Q
12	`sec23_surplus_feedback_monotone`	Sec. 2.3	$\forall a, h_x, h_y, t \in \mathbb{Z},\ a - h_x \le t \wedge t \le a - h_y \implies h_y \le h_x.$	P,Q
13	`sec72_monotone_lcp_gap_split`	Sec. 7.2	$\forall m, q, z, z_0 \in \mathbb{Z},\ z_0(mz_0 + q) = 0 \implies z(mz + q) = z(mz_0 + q) + z_0(mz + q) + (z - z_0)(m(z - z_0)).$	P,Q
14	`sec73_qp_skew_block_monotone`	Cor. 6	$\forall u, \mu, a \in \mathbb{Z},\ u(a\mu) + \mu(-(au)) = 0$ (skew KKT block).	P,Q
15	`thm8_contraction_movement_lb`	Thm. 8	$\forall g_n, g_d, u, d \in \mathbb{Z},\ 0 \le g_n, d \wedge 8g_d d \le g_n u \implies 7g_d d \le g_d u - (g_d - g_n)(u + d).$	P,Q
16	`thm8_residual_to_distance_cleared`	Thm. 8	$\forall g_n, g_d, D, T, e \in \mathbb{Z},\ g_d(D - T) \le g_d e \wedge g_d T \le (g_d - g_n)D \implies g_n D \le g_d e.$	P,Q
17	`sec91_bellman_contraction_gap`	Sec. 9.1	$\forall \text{den}, \text{gap}, D \in \mathbb{Z},\ 0 < \text{gap} \wedge 0 < D \implies (\text{den} - \text{gap})D < \text{den} \cdot D.$	P,Q
18	`sec93_clipped_update_halving_bound`	Sec. 9.3	$\forall W, e_L, e_U \in \mathbb{N},\ 4e_L \le W \wedge 4e_U \le W \implies 2(e_L + e_U) \le W.$	P,Q

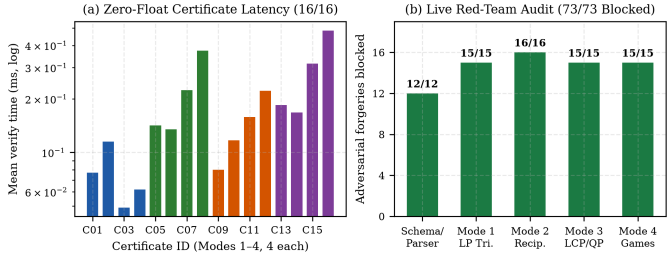

Figure 3: **(a)** Mean exact zero-float verification latency across all 16 canonical certificates (0.049–0.485 ms, log scale). **(b)** Live adversarial red-team audit rejecting 73/73 forgeries across five attack classes with exact error-branch assertions.

Table 2: Canonical 16-certificate `SPLP-SMALE9-CERT-1.0` atlas (`exp1_certificate_atlas.json`; all 16 verified).

ID	Certificate Name	M	Exact Certified Witness Summary	ms
C01	`m1_lp_optimal_2d_strict_comp`	M1	4×2 opt = 10, 4/4 strict comp. rows	0.077
C02	`m1_lp_optimal_3d_strict_comp`	M1	6×3 opt = 30, 6/6 strict comp. rows	0.115
C03	`m1_lp_unbounded_recession_ray`	M1	3×2 unbounded, ray gain $c^\top r = 8$	0.049
C04	`m1_lp_infeasible_farkas_minsupp`	M1	4×2 infeasible, rank = 2, $\lvert\text{supp}\rvert = 3$	0.062
C05	`m2_recip_circuit_u10_newton2`	M2	N=4, u=10: $Q(yD)$=0, $2!\det = -198$	0.142
C06	`m2_recip_circuit_u1000_newton3`	M2	N=4, u=10^3: $Q(yD)$=0, $3!\det = -324$	0.135
C07	`m2_recip_n5_ortho_newton4`	M2	N=5: 2 + 3 ortho, $4!\det$ =3888	0.224
C08	`m2_recip_n6_ortho_newton5`	M2	N=6: 3 + 3 ortho, $5!\det = -4680$	0.373
C09	`m3_lcp_2d_sym_pd`	M3	2D sym. PD LCP: gap 6, quad 8/3	0.080
C10	`m3_lcp_3d_asym_monotone`	M3	3D asym. monotone: gap 14, quad 7	0.117
C11	`m3_qp_kkt_4d_skew_block`	M3	4D Cor. 6 QP-KKT: gap 4, quad 1	0.158
C12	`m3_qp_kkt_5d_psd_degenerate`	M3	5D PSD deg. KKT: gap 9/2, quad 1/2	0.222
C13	`m4_game_fig5_low_center`	M4	Fig. 5 low y: err $\frac{3}{2} \to \frac{1}{2}$, $z \le T(z)$	0.185
C14	`m4_game_fig5_neutral_center`	M4	Fig. 5 neutral y: err $\frac{3}{2} \to \frac{1}{2}$	0.168
C15	`m4_game_3state_stochastic`	M4	3-state stoch.: err $1 \to \frac{1}{2}$, $z \le T(z)$	0.316
C16	`m4_game_4state_gamma_quarter`	M4	4-state γ=$\frac{1}{4}$: err $1 \to 0$, $z \le T(z)$	0.485

4 Exact Findings, Worked Witnesses, and Audit

All experiments ran locally on Windows 11 (Python 3.12.6, Lean 4.35.0; 7-run means after warm-up).

1. Exact Circuit Imbalance, Gap Splits & Clipped Halving (Figs. 1–2). Across $u \in \{2^1, \ldots, 2^{12}\}$, Section 3.4's division-free Newton recurrence matches exact Gaussian determinants identically on every minor ($|\det| = u^2 - 1$ from 3 to 16,777,215). On Section 9.3 Figure 5's two-state game starting from $y^{(0)} = (0, 2)$ with $H_t = 2^{-t}$ (P_a undiscounted in Eq. (9.3)), exact clipped updates produce $u^{(0)} = (1, -1) \Rightarrow y^{(1)} = (1, 1)$ (err $\frac{1}{2}$), $u^{(1)} = (\frac{1}{2}, -\frac{1}{4}) \Rightarrow y^{(2)} = (\frac{3}{2}, \frac{3}{4})$ (err $\frac{1}{4}$), and $u^{(2)} = (0, -\frac{1}{4}) \Rightarrow y^{(3)} = (\frac{3}{2}, \frac{1}{2}) = v$ (exact err 0 at $t = 3$). On all four Mode 3 monotone LCP and QP-KKT witnesses, the quadratic term $(z' - z^*)^\top M(z' - z^*)$ is strictly nonnegative $(8/3, 7, 1, 1/2)$ and sums with the cross terms identically to $z'^\top w'$ $(6, 14, 4, 9/2)$.

2. Hand-Verifiable Worked Exact Witnesses (One per Mode).

- **Mode 1 (`C01`, 4×2 optimal LP):** $\max 3x_1 + 2x_2$ s.t. $x_1 + x_2 \le 4, 2x_1 + x_2 \le 6, x \ge 0$ attains $x^* = (2, 2), y^* = (1, 1, 0, 0)$ with $c^\top x^* = b^\top y^* = 10$, slacks $(0, 0, 2, 2)$, and $b - Ax^* + y^* = (1, 1, 2, 2) > 0$ strictly complementary.
- **Mode 2 (`C05`, Sec. 1.2 u=10 circuit):** C and $C^\perp$ spanned by rows $[[1, 0, 1, 10], [0, 1, 10, 1]]$ and $[[-1, -10, 1, 0], [-10, -1, 0, 1]]$ preserve $Q(yD) = 0$ under $p = (3, -2, 5, -4)$, while Newton's recurrence yields $2!\det([[1, 10], [10, 1]]) = -198$.
- **Mode 3 (`C11`, 4D Cor. 6 QP-KKT):** With $G = 2I_2$, $A = \left(\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, $q = (-3, -3, 2, 2)$, solution $z^* = (1, 1, 1, 0), w^* = (0, 0, 0, 1)$ and probe $z' = (\frac{1}{2}, \frac{1}{2}, 2, 1)$ satisfy $w' = (1, 0, 1, \frac{3}{2})$ and $z'^\top w' = 4 = 3 + 1$.
- **Mode 4 (`C13`, Sec. 9.3 Fig. 5 low center):** For $T_A(x) = \max\{\frac{3}{4} + \frac{x_A}{2}, \frac{x_B}{2}\}$, $T_B(x) = \min\{1 + \frac{x_B}{2}, -\frac{1}{4} + \frac{x_A}{2}\}$ (P_a deterministic, $B_a = \frac{1}{2}P_a$, $v = (\frac{3}{2}, \frac{1}{2})$), low center $y = (0, \frac{1}{4})$ and $H = 1$ give Eq. (9.3) scores $\Phi_A(u) = \max\{\frac{3}{4} + u_A, \frac{1}{8} + u_B\}$, $\Phi_B(u) = \min\{\frac{7}{8} + u_B, -\frac{1}{2} + u_A\}$; at $u = (1, \frac{1}{2})$, $[\max(\frac{7}{4}, \frac{5}{8})]_{-1}^{1} = 1 = u_A$ and $[\min(\frac{11}{8}, \frac{1}{2})]_{-1}^{1} = \frac{1}{2} = u_B$, giving error $\frac{1}{2} \le 1$ and lower guarantee $z = y + u - H\mathbf{1} = (0, -\frac{1}{4}) \le T(z) = (\frac{3}{4}, -\frac{1}{4})$.

3. 16-Certificate Atlas & 80-Test Red-Team Audit (Table 2, Fig. 3). Table 2 lists all 16 canonical certificates verified in 0.049–0.485 ms. `tests/test_suite.py` runs 80 pytest checks: 1 full 16-certificate replay, 4 raw-JSON parser guards (float, `NaN`, duplicate keys, oversized payload), 2 exact regressions (Newton $4!\det$ vs. Gaussian elimination and exact indefinite matrix rejection), and 73 parametrized live adversarial forgeries (12 schema/parser, 15 Mode 1, 16 Mode 2, 15 Mode 3, 15 Mode 4), each asserting an exact regex on its intended `CertificateError` branch. `scripts/verify_all.py` enforces all 5 gates and a 41-file SHA-256 manifest.

5 Conclusion

We built an independent, explicitly scoped, zero-floating-point verification companion for the discrete LP trichotomy, reciprocal diagonal scaling, division-free Newton determinant, monotone LCP / convex QP KKT, and max-norm contraction / clipped-Bellman game mechanisms of OpenAI's strongly polynomial linear programming preprint [7], pairing 18 constructive Lean 4.35.0 sub-lemmas (zero `sorry`, zero `Classical.choice`) with the pure-`int`/`Fraction` 4-mode auditor `SPLP-SMALE9-CERT-1.0`, 16 canonical exact finite witnesses, and 73/73 live adversarial forgery rejections.

References

[1] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.

[2] András Frank and Éva Tardos. An application of simultaneous Diophantine approximation in combinatorial optimization. *Combinatorica*, 7:49–65, 1987.

[3] Alan J. Goldman and Albert W. Tucker. Theory of linear programming. In H. W. Kuhn and A. W. Tucker, editors, *Linear Inequalities and Related Systems*, volume 38 of *Annals of Mathematics Studies*. Princeton University Press, 1956.

[4] Narendra Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4:373–395, 1984.

[5] Leonid G. Khachiyan. A polynomial algorithm in linear programming. *Doklady Akademii Nauk SSSR*, 244:1093–1096, 1979.

[6] Olvi L. Mangasarian. A simple characterization of solution sets of convex programs. *Operations Research Letters*, 7:21–26, 1988.

[7] OpenAI. Strongly polynomial linear programming. OpenAI Math Release Preprint `OAI:Strongly-polynomial-linear-programming-November-24-2026`, November 2026. Companion unrefereed preprint.

[8] Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the USA*, 39:1095–1100, 1953.

[9] Stephen Smale. Mathematical problems for the next century. *The Mathematical Intelligencer*, 20(2):7–15, 1998.

[10] Éva Tardos. A strongly polynomial algorithm to solve combinatorial linear programs. *Operations Research*, 34:250–256, 1986.