

Parameter Retuning in the Direct Reduction for the Euclidean Closest Vector Problem

$n^{1/30}$ -Hardness and a $1/28$ Retained-Parameter Frontier

Joshua Gay*

14 August 2026

Abstract

We revisit the parameter choices in OpenAI’s direct deterministic reduction from 3SAT to the Euclidean closest vector problem. The construction and its algebraic lemmas are retained. The retuning was developed in FPRD Lab through a resource audit that can be organized in four stages. First, the desired lattice gap is pulled backward through the parity lift, turning M^c into the binary threshold M^{2c} . Second, the source’s fixed powers are replaced by independent coordinates $q \asymp N^Q$, $K = N^\kappa$, and $T = N^\tau$. Reading the quantitative premises of the proof then gives

$$\kappa > 1 + 2c(1 + 2Q), \quad \tau > 1 + 5\kappa, \quad Q > 1 + 2\kappa + \tau.$$

Third, eliminating the internal coordinates exposes the feasible range $0 \leq c < 1/28$. Finally, an integral point is chosen with slack, the finite constants are restored, and the result is tested against the executable reduction. For $c = 1/30$, this procedure gives

$$q \asymp N^{242}, \quad K = N^{34}, \quad T = N^{172}$$

and output dimension $M \leq 40N^{485}$. It proves $n^{1/30}$ -hardness; the corresponding binary factor is $n^{1/15}$. Separate quarter-field loss budgets and the identity $N = 100 + s + m + \ell$ supply the exact finite- N margin. The audit also identifies an interface discrepancy: the source paper states coefficient-4 binary soundness, while one published Lean wrapper exposes coefficient 2. This does not affect OpenAI’s final theorem, which needs only coefficient 1. The retuned Lean development carries coefficient 4 through the physical interfaces and checks the complete concrete reduction and the retained parameter frontier. The value $1/28$ is a limit of this accounting, not a barrier for all direct reductions.

Contents

1	Introduction	2
1.1	Main result	2
1.2	Proof overview and relation to the source paper	3
1.3	Scope of the formal verification	3
2	Retuned parameters and retained construction	3
3	Encoding Boolean assignments by Reed–Solomon codes	4
3.1	Completeness	4
3.2	Dimension and construction size	5

*ChatGPT, Claude, and Codex were used for research exploration, source checking, exact computational auditing, formal verification, and editorial assistance. The author is responsible for the mathematical claims and the final manuscript.

4	Algebraic reconstruction of bounded polynomial-moment families	5
5	Soundness of the nearest-codeword reduction	5
5.1	Pointwise support sets of a low-weight solution	5
5.2	Applying the reconstruction lemma	6
5.3	Valuative recovery of local assignments	6
5.4	Clause matching	7
5.5	Completing soundness and exact constants	7
6	Transfer to CVP and complexity	7
7	Further consequences and verification	8
7.1	The retained parameter frontier	8
7.2	FPRD method and provenance	8
7.3	Lean verification	10
7.4	Limits of the conclusion	10

1 Introduction

OpenAI’s paper $n^{1/400}$ -*Hardness of the Euclidean Closest Vector Problem* gives a deterministic polynomial-time many-one reduction from 3SAT to Euclidean CVP with approximation factor $n^{1/400}$ [1]. Its main technical step is a binary affine system whose low-weight solutions encode satisfying assignments. The system uses one-hot evaluation tables over a characteristic-two field. Its soundness proof reconstructs algebraic root sets from polynomial moments, recovers Boolean values by valuations, and matches a common root across the clause tables. A dimension-preserving parity lift converts binary weight into squared Euclidean distance.

This note changes the numerical specialization of that construction. It is written as a companion to the OpenAI paper: the notation and the order of the encoding, reconstruction, soundness, and transfer arguments are kept so that the two proofs can be read in parallel.

The parameterization and elimination were developed in FPRD Lab—Finitely Presented Recognition Dynamics—an Insight Forge Boxlab project. Its mathematical contribution was to recast the source’s fixed hierarchy as a system of coupled resource coordinates and then eliminate the internal coordinates to expose the feasible exponent range. Section 7.2 states this procedure explicitly and distinguishes it from the source lemmas retained unchanged.

1.1 Main result

For a nonsingular integer matrix $B \in \mathbb{Z}^{n \times n}$, write $L(B) = B\mathbb{Z}^n$. An instance of $\text{GapCVP}_\gamma^{(2)}$ consists of B , a target $t \in \mathbb{Q}^n$, and a positive rational radius r . The promise is to distinguish

$$\text{dist}_2(t, L(B)) \leq r \quad \text{from} \quad \text{dist}_2(t, L(B)) > \gamma(n)r.$$

Theorem 1 (Retuned direct reduction). *For every fixed real number c with $0 \leq c < 1/28$, there is a deterministic polynomial-time mapping that assigns to each 3SAT formula φ a nonsingular matrix $B \in \mathbb{Z}^{n \times n}$, a target $t \in \mathbb{Z}^n$, and a radius $r \in \mathbb{Q}_{>0}$ such that*

$$\begin{aligned} \varphi \text{ satisfiable} &\implies \text{dist}_2(t, L(B)) \leq r, \\ \varphi \text{ unsatisfiable} &\implies \text{dist}_2(t, L(B)) > n^c r. \end{aligned}$$

Consequently, $\text{GapCVP}_{n^c}^{(2)}$ is NP-hard under deterministic polynomial-time many-one reductions. The same construction gives approximation factor n^{2c} for binary nearest codeword and binary syndrome decoding.

For $c = 1/30$, the proof has the explicit specialization

$$Q = 242, \quad \kappa = 34, \quad \tau = 172, \quad (1)$$

where $q \asymp N^Q$, $K = N^\kappa$, and $T = N^\tau$. The resulting binary block length and lattice rank satisfy

$$M = n \leq 40N^{485}. \quad (2)$$

The strict retained resource inequalities have no solution at $c \geq 1/28$. This establishes a frontier for the present one-hot representation and the present reconstruction and valuation bounds. It is not a lower bound on the exponent attainable by a different reduction.

1.2 Proof overview and relation to the source paper

The reduction continues to factor through binary nearest codeword:

$$3\text{SAT} \longrightarrow \text{binary nearest codeword} \longrightarrow \text{Euclidean CVP}.$$

The comparison with the OpenAI proof is as follows.

- (i) *Encoding and completeness (Section 3)*. The table types, one-hot coordinates, ordinary and shifted power-sum constraints, and satisfying-assignment solution are unchanged. The retuning changes only the fixed field, support, and moment exponents in the source's Section 3.5.
- (ii) *Algebraic reconstruction (Section 4)*. OpenAI's Lemma 10 is invoked without alteration. The retuning verifies its three numerical hypotheses with a larger support cutoff and moment horizon.
- (iii) *Nearest-codeword soundness (Section 5)*. As in the source, large pointwise support fibers are discarded separately for each table type. The available field is divided explicitly between those fibers and the anchor and Hankel-discriminant losses. The shifted-moment valuation and clause-root matching arguments are then unchanged.
- (iv) *Transfer to CVP (Section 6)*. The source parity lift identifies minimum binary weight with squared Euclidean distance. Thus binary exponent $2c$ becomes Euclidean exponent c , with lattice rank equal to binary block length.

1.3 Scope of the formal verification

The complete executable Lean reduction is checked for the concrete exponent $c = 1/30$, including the physical binary system, integer target, polynomial-time machines, and final NP-hardness statements. A separate Lean module proves that the retained resource inequalities admit natural-number witnesses for every real $0 \leq c < 1/28$ and admit none at or above $1/28$. The complete reduction has not been uniformly parameterized in Lean for arbitrary c ; the general statement in theorem 1 is the mathematical retuning proved below.

2 Retuned parameters and retained construction

After the preprocessing used in the source, let the formula have s symbols, m variables, and ℓ clauses, and put

$$N = 100 + s + m + \ell, \quad d = m. \quad (3)$$

Choose the least integer e for which $2^e \geq N^Q$, and let

$$q = 2^e, \quad N^Q \leq q < 2N^Q, \quad \mathbb{K} = \mathbb{F}_q. \quad (4)$$

The support cutoff and moment horizon are

$$K = N^\kappa, \quad T = N^\tau. \quad (5)$$

Here $\mathbb{K}$ denotes the field and K the integer support cutoff.

Let $a_1, \dots, a_m \in \mathbb{K}$ be distinct anchors and $P = \mathbb{K} \setminus \{a_1, \dots, a_m\}$. Let Θ be the set consisting of the global table type and the satisfying local assignment types for each clause. As in the source,

$$|\Theta| \leq 1 + 8\ell \leq 9N, \quad M = |\Theta||P|q.$$

Quantity	Meaning	OpenAI specialization	Concrete retuning
q	field size	$N^{200} \leq q < 2N^{200}$	$N^{242} \leq q < 2N^{242}$
K	support cutoff	N^4	N^{34}
T	moment horizon	N^{30}	N^{172}
M	binary block length	$M \leq 40N^{401}$	$M \leq 40N^{485}$
binary factor	nearest codeword	$M^{1/200}$	$M^{1/15}$
Euclidean factor	closest vector	$M^{1/400}$	$M^{1/30}$

Table 1: The source parameters and the concrete specialization. The construction itself is unchanged.

For a desired Euclidean exponent c , four asymptotic costs control the retuning:

$$M \leq 40N^{1+2Q}, \quad (6)$$

$$\kappa > 1 + 2c(1 + 2Q), \quad (7)$$

$$\tau > 1 + 5\kappa, \quad (8)$$

$$Q > 1 + 2\kappa + \tau. \quad (9)$$

The first line is the dense one-hot cost. The second makes K dominate the support discarded from a binary solution of weight $O(M^{2c}R)$. The third provides all shifted moments used by the valuation argument. The fourth keeps the cleared reconstruction degree below the field size. The exact constants needed in addition to these exponent inequalities are recorded in section 5.

3 Encoding Boolean assignments by Reed–Solomon codes

For every table type $\theta \in \Theta$, evaluation point $p \in P$, and value $w \in \mathbb{K}$, the source construction has a binary coordinate $x_{\theta,p,w}$. Its ordinary and shifted power-sum constraints, expanded in an $\mathbb{F}_2$ -basis of $\mathbb{K}$, form the same explicit affine system

$$Hx = b \quad \text{over } \mathbb{F}_2. \quad (10)$$

No coefficient or equation in this system is changed by the retuning.

3.1 Completeness

OpenAI’s Lemma 8 applies verbatim. If φ is satisfiable, the one-hot graph of its interpolation polynomial, together with one active local table for each clause, gives a solution of weight

$$R = (\ell + 1)|P|. \quad (11)$$

Because equation (3) gives $\ell + 1 \leq N$ and $|P| < q$,

$$R \leq Nq. \quad (12)$$

For the Euclidean reduction, set $r = \lceil \sqrt{R} \rceil$. Since $R \geq 1$,

$$r^2 \leq 4R \leq 4Nq. \quad (13)$$

This exact estimate is used in the finite- N soundness calculation.

3.2 Dimension and construction size

From $|P| < q < 2N^Q$ and $|\Theta| \leq 9N$,

$$M = |\Theta||P|q \leq 40N^{1+2Q}. \quad (14)$$

The source's row count becomes

$$|P| + \ell|P|q + 4|\Theta|(T+1)|P|e = O\left(N^{1+2Q} + N^{1+Q+\tau} \log N\right). \quad (15)$$

For fixed Q and τ , both the system and the field construction remain deterministic polynomial time. At equation (1), the two exponents in equation (15) are 485 and 415, respectively.

4 Algebraic reconstruction of bounded polynomial-moment families

OpenAI's Lemma 10 is independent of the Boolean encoding. For a family of sets $S(p) \subseteq \mathbb{K}$ with $|S(p)| \leq K$, it reconstructs a separable polynomial over $\mathbb{K}(X)$ from power sums through degree T . We use that lemma exactly as stated.

For a retained evaluation set P_θ , its numerical hypotheses are

$$T \geq 2K - 1, \quad dT < |P_\theta|, \quad |P_\theta| - dK(K-1) > 2dK^2T. \quad (16)$$

The term $dK(K-1)$ bounds the zeros of the Hankel determinant. The term $2dK^2T$ is the degree range needed after clearing denominators. These are the only places where the retuned values enter Lemma 10's application. The lemma's construction of a common splitting field and its root-moment identities are unchanged.

5 Soundness of the nearest-codeword reduction

Suppose that $x \in \mathbb{F}_2^M$ solves equation (10) and

$$\text{wt}(x) \leq 4M^{2c}R. \quad (17)$$

For each type θ and point p , let

$$S_\theta(p) = \{w \in \mathbb{K} : x_{\theta,p,w} = 1\}.$$

The proof follows the five parts of the source's Section 5.

5.1 Pointwise support sets of a low-weight solution

For each table type separately, retain

$$P_\theta = \{p \in P : |S_\theta(p)| \leq K\}.$$

Markov's inequality and equations (12) and (17) give

$$|P \setminus P_\theta| \leq \frac{4M^{2c}R}{K} \leq \frac{4M^{2c}Nq}{K}. \quad (18)$$

Thus the exact sufficient condition

$$16M^{2c}N < K \quad (19)$$

makes the discarded set smaller than $q/4$. The sets P_θ may differ; they are never intersected over all table types.

At the concrete parameters, equation (19) follows from

$$16(40N^{485})^{1/15}N < N^{34}. \quad (20)$$

Raising both sides to the fifteenth power reduces this to

$$16^{15} \cdot 40 < N^{10}, \quad (21)$$

which holds already at $N = 100$ and remains true for larger N .

5.2 Applying the reconstruction lemma

Reserve a second quarter of the field for the anchors and the zeros of the Hankel determinant. It is enough that

$$4(m + dK(K - 1)) < q. \quad (22)$$

Together with equations (18) and (19), this gives

$$|P_\theta| - dK(K - 1) > q/2. \quad (23)$$

The remaining reconstruction and interpolation requirements follow from

$$4dK^2T < q, \quad 2dT < q, \quad T \geq 2K - 1. \quad (24)$$

Indeed, equations (23) and (24) imply all three conditions in equation (16).

For equation (1), using $d, m \leq N$ and $q \geq N^{242}$, the exact checks are

$$4(N + N^{69}) < N^{242}, \quad (25)$$

$$4N^{241} < N^{242}, \quad (26)$$

$$2N^{173} < N^{242}, \quad (27)$$

$$N^{172} \geq 2N^{34} - 1. \quad (28)$$

All hold for $N \geq 100$. Hence Lemma 10 reconstructs the root sets by the same per-table argument as in the source proof and places them in one common splitting field.

5.3 Valuative recovery of local assignments

OpenAI's Lemma 12 studies, for a reconstructed root α and a clause bit β_i ,

$$y = \frac{\alpha - \beta_i}{X - a_i}.$$

If $h \leq K$ is the number of roots of the table polynomial, the proof defines

$$D_0 = dh^2 + h, \quad U_h = h^2D_0 + dh^2 + 1, \quad z = hU_h + 1.$$

The unchanged estimates give

$$D_0 \leq 2N^{1+2\kappa}, \quad U_h \leq 4N^{1+4\kappa}, \quad z + h - 1 \leq 5N^{1+5\kappa}. \quad (29)$$

Thus $5N^{1+5\kappa} < T$ supplies every shifted moment required by the valuation argument. At the concrete values,

$$5N^{171} < N^{172} = T \quad (30)$$

for $N \geq 100$. The conclusion of Lemma 12, that each local root recovers the prescribed bit at the corresponding anchor valuation, is therefore unchanged.

5.4 Clause matching

The source parity argument uses ordinary moments to match every global root to at least one satisfying local root set for each clause. Since a clause has at most eight satisfying local assignments, the Vandermonde system has size at most

$$(1 + |B_C|)K \leq 9K.$$

The sufficient condition is $9K < T + 1$. At the concrete parameters the stronger inequality

$$9N^{34} < N^{172} = T \tag{31}$$

holds. OpenAI’s Lemma 13 therefore applies without alteration.

5.5 Completing soundness and exact constants

The preceding bounds give the retuned form of OpenAI’s Proposition 14.

Proposition 2 (Retuned binary soundness). *Under the exact conditions above, if equation (10) has a solution x satisfying*

$$\text{wt}(x) \leq 4M^{2c}R,$$

then φ is satisfiable.

Proof. For each table type, equation (16) permits the source reconstruction lemma to be applied to the retained support family. The global root set is nonempty by the source’s zeroth-moment parity argument. Lemma 13 matches a chosen global root to a satisfying local root set for each clause. Lemma 12, applied with the same anchor valuation across all clauses containing a variable, shows that the local Boolean values agree. They therefore assemble into a satisfying assignment. This is the proof of the source proposition with only its numerical hypotheses replaced. $\square$

The exponent inequalities control growth rates; the finite coefficients enter through the quarter-field allocation in equations (19) and (22) and the exact source estimate equation (13). Together they give the coefficient-16 condition in equation (20), which equation (21) verifies already at $N = 100$. This exact bookkeeping does not change the affine system, reconstruction lemma, valuation argument, or parity lift.

The OpenAI paper states Proposition 14 with coefficient 4. The imported Lean snapshot exposed coefficient 2 at one binary wrapper, although its factor- $n^{1/200}$ output theorem only needed coefficient 1. This is a source/interface discrepancy, not a failure of OpenAI’s final theorem. The retuned formalization carries coefficient 4 through the explicit and physical binary interfaces so that proposition 2 is available for the Euclidean transfer.

6 Transfer to CVP and complexity

OpenAI’s parity lift maps a consistent affine system $Hx = b$ to an integer lattice $L(B_H)$ and a binary target u with

$$\text{dist}_2(u, L(B_H))^2 = \min\{\text{wt}(x) : x \in \mathbb{F}_2^M, Hx = b\}. \tag{32}$$

The lattice has rank $n = M$. The preprocessing and inconsistent-system branches use the same fixed one-dimensional instances as the source.

If φ is satisfiable, equations (11) and (32) give distance at most $\sqrt{R} \leq r$. If φ is unsatisfiable and the distance were at most $M^c r$, then

$$\text{dist}_2(u, L(B_H))^2 \leq M^{2c} r^2 \leq 4M^{2c} R.$$

By equation (32), this would give a binary solution excluded by proposition 2. Hence

$$\text{dist}_2(u, L(B_H)) > M^c r = n^c r,$$

with the strict inequality required by the promise problem.

At the concrete values, the field extension degree is $e = O(\log N)$, the affine system has $M \leq 40N^{485}$ variables and $O(N^{485})$ rows, and the square parity-lift basis has fixed-polynomial encoding length. Field construction, Reed–Solomon checks, Gaussian elimination, serialization, and basis construction therefore remain deterministic polynomial-time operations. The exponents are large, but none depends on the input.

This proves the concrete case of theorem 1. The binary gap in proposition 2 also gives the following source-parallel consequence.

Corollary 3. *Binary nearest codeword and binary syndrome decoding are NP-hard to approximate within $n^{1/15}$ under deterministic polynomial-time many-one reductions.*

7 Further consequences and verification

7.1 The retained parameter frontier

Eliminating τ and κ from equations (7) to (9) gives

$$Q > 2 + 7\kappa > 2 + 7(1 + 2c(1 + 2Q)),$$

and therefore

$$(1 - 28c)Q > 9 + 14c. \tag{33}$$

For $c \geq 1/28$ and $Q \geq 0$, the left side is nonpositive and the right side is positive. Thus the retained strict inequalities are infeasible at and above $1/28$.

Conversely, fix $0 \leq c < 1/28$. For sufficiently large integer Q , the interval

$$\left[2 + 2c(1 + 2Q), \frac{Q - 4}{7} \right] \tag{34}$$

contains an integer κ . Set $\tau = 2 + 5\kappa$. Then

$$\kappa > 1 + 2c(1 + 2Q), \quad \tau > 1 + 5\kappa, \quad 1 + 2\kappa + \tau < Q.$$

These choices leave at least one full power of N in the support, valuation, and reconstruction comparisons. Since $N \geq 100$ and $16 \cdot 40^{2c} < 100$ for $c < 1/28$, the exact support coefficient is also absorbed. The lower-order anchor, Hankel, interpolation, and parity conditions then follow. This proves theorem 1 for every fixed $c < 1/28$.

7.2 FPRD method and provenance

The resource decomposition and frontier calculation are the specific mathematical contribution of FPRD Lab—Finitely Presented Recognition Dynamics. They form a parameter audit of the OpenAI construction, not a new construction. OpenAI supplies the one-hot encoding, moment reconstruction, valuation argument, clause matching, and parity lift. FPRD Lab supplied the resource decomposition, the elimination argument, and the concrete specialization. Exact and formal audits verify these outputs. In the lab’s terminology, this is an algebraic-sewing and resource-compression audit: local proof obligations are retained, while their shared resource coordinates are composed and then eliminated according to the final readout.

Resource-audit procedure. The analysis is organized as the following reusable procedure.

1. *Fix the external readout.* State the desired approximation factor in the final dimension, then propagate it backward through every transfer. Here the squared-distance parity lift turns the Euclidean factor M^c into a binary weight scale M^{2c} .
2. *Expose independent coordinates.* Replace the source’s chosen powers by one coordinate for each proof budget. Here q measures the available evaluation field, K bounds each retained support, T is the moment horizon, and M is the external representation size. No relation among their growth rates is assumed at this stage.
3. *Extract the source obligations.* Read the proof in dependency order and record every quantitative premise that connects two coordinates. Keep leading exponents separate from fixed constants and from losses at the smallest input size.
4. *Eliminate and instantiate.* Write $q \asymp N^Q$, $K = N^\kappa$, and $T = N^\tau$. Eliminate the internal coordinates to obtain the feasible range for c , then choose an integral tuple with explicit slack.
5. *Certify the endpoint.* Restore the finite constants, check the smallest allowed N , and push the tuple through the actual executable construction. Report the verified concrete endpoint separately from the symbolic frontier.

Application to this reduction. The first step gives the low-weight threshold $4M^{2c}R$ in equation (17). The next two steps extract the leading dependency chain

$$M \asymp q^2, \quad K \gtrsim M^{2c}, \quad T \gtrsim K^5, \quad q \gtrsim K^2 T.$$

The suppressed powers of N become the constant terms in equations (7) to (9). Eliminating κ and τ makes $(1 - 28c)Q > 9 + 14c$ visible. It gives the retained range $c < 1/28$ and guides the choice $(Q, \kappa, \tau) = (242, 34, 172)$ at $c = 1/30$.

The certification step uses the quarter-field allocation in equations (19) and (22) and the exact radius estimate in equation (13) to obtain the coefficient-16 margin. The Lean development checks the complete concrete endpoint. It separately checks the symbolic retained frontier, without claiming one uniformly parameterized executable reduction for every $c < 1/28$.

The same elimination also identifies which retained costs control the answer. If a related construction has

$$M \asymp q^a, \quad T \gtrsim K^v, \quad q \gtrsim K^r T,$$

the same leading-exponent elimination gives

$$c < \frac{1}{2a(r+v)}. \tag{35}$$

The present values are $a = 2$, $r = 2$, and $v = 5$. The factor 28 therefore records three separate costs: squaring in the Euclidean transfer, the dense one-hot q^2 representation, and the combined seventh power of K in valuation and reconstruction. This also separates three attack seams. A more economical representation reduces a ; a sharper valuation lemma reduces v ; and a lower-degree cleared-reconstruction identity reduces r .

FPRD Lab is an Insight Forge Boxlab project (<https://gitlab.com/insight-forge/boxlab>) that begins with recognition-theoretic constructions and asks how their structural and resource tradeoffs transfer to other settings. Here that lens was applied to the parameter dependencies in the direct GapCVP reduction; no source reconstruction or valuation lemma was changed.

The lab also served as a testing and training bed for Sidecat (<https://gitlab.com/sidecat-dev/sidecat>), the agent-work system used to coordinate the AI-assisted work and preserve its evidence and publication record. That operational provenance affected the mathematical outcome. The FPRD parameter search produced the 1/30 specialization and 1/28 frontier. The retained dependency record, exact arithmetic audit, and independent Lean development established the finite- N margins and verified the complete concrete endpoint.

7.3 Lean verification

The independent formal audit uses OpenAI source commit `94bc0feb6a9f` and is published at fork commit `3ac608981eb4` [3]. The bibliography gives the complete immutable fork revision. The complete retuned build reaches the physical integer-target theorem `gapCVP30IsNPHard` and the two factor-1/15 binary endpoints. The frontier module proves natural-number witnesses below 1/28 and impossibility at and above it.

The principal axiom probes report only `propext`, `Classical.choice`, and `Quot.sound`. The modified proof files contain no `sorry`, `admit`, `sorryAx`, or new axiom declaration. The principal source-to-check correspondence is shown in table 2.

OpenAI paper	Retuned obligation	Checked endpoint
Section 3.5, equations (4)–(5)	field, dimension, and row bounds	$N^{242} \leq q < 2N^{242}$ and $M \leq 40N^{485}$
Section 5.1, equations (13)–(14)	per-table retained sets and finite constants	quarter-field heavy and anchor/Hankel bounds
Lemma 10	cleared reconstruction and moment range	$2dK^2T \leq 2N^{241} < q/2$
Lemma 12, (17)–(21)	shifted valuation window	$z + h - 1 \leq 5N^{171} < T$
Lemma 13	clause-root parity window	$(1 + B_C)K \leq 9K < T + 1$
Proposition 14	coefficient-four binary soundness	explicit and physical factor-1/15 interfaces
Lemma 7 and Section 6	parity lift and strict CVP promise	<code>gapCVP30IsNPHard</code>

Table 2: Crosswalk from the source proof to the retuned verification.

7.4 Limits of the conclusion

The audit establishes that the selected parameters satisfy every numerical hypothesis encountered in the source proof and its executable Lean specialization. It does not replace the source reconstruction and valuation lemmas with an independent proof. The complete formal reduction is presently instantiated only at $c = 1/30$; the formal result for arbitrary real $c < 1/28$ concerns the retained resource inequalities. A uniformly parameterized Lean version of the complete reduction would narrow this remaining distinction.

The value 1/28 is correspondingly an unattained supremum of the retained parameter system. It should not be interpreted as an optimal hardness exponent for Euclidean CVP or as an obstruction to a different direct construction.

Reproducibility

The exact integer-arithmetic certificate and its canonical output are published in the FPRD Lab reproduction package [2]. The independent Lean fork records the imported source revision, modified paths, complete build logs, source hashes, and axiom output [3]. These two checks are complementary: the small certificate evaluates the seven concrete finite- N side conditions reported in its output, while Lean checks that the concrete values reach the actual executable reduction and final promises. The certificate is `evidence/gapcvp-parameter-audit/check_gapcvp_parameters.py`; its recorded output is `evidence/gapcvp-parameter-audit/RUN-OUTPUT.txt`.

References

- [1] OpenAI, $n^{1/400}$ -Hardness of the Euclidean Closest Vector Problem, in *Ten Advances in Mathematics and Theoretical Computer Science*, pp. 183–218, 2026. <https://cdn.openai.com/pdf/ten-proofs-oai.pdf>.
- [2] Joshua Gay, *FPRD Lab: GapCVP Parameter-Retuning Reproduction Package*, source code and recorded output, FPRD Lab GitLab repository, commit 013f5eec, 8 August 2026. <https://gitlab.com/insight-forge/boxlab/fprd-lab/-/tree/013f5eec822c6f7b474991bfa6ac76e1cef2d0fe/evidence/gapcvp-parameter-audit>.
- [3] Insight Forge Boxlab, *ten-proofs-gapcvp-retuning*, independent audit fork of OpenAI’s *ten-proofs*, commit 3ac608981eb48c432e77b42b899689260351ca7d, 13 August 2026. <https://gitlab.com/insight-forge/boxlab/ten-proofs-gapcvp-retuning/-/tree/3ac608981eb48c432e77b42b899689260351ca7d>.

Afterword

Parameter retuning in the direct reduction for the Euclidean closest vector problem:
1/30-hardness and a 1/28 retained-parameter frontier

Methodological and contribution addendum. The paper's mathematical scope and review status remain as stated in the paper.

The Lab's approach

FPRD means **Finitely Presented Recognition Dynamics**. Its starting point is to distinguish a finite description, the possibly infinite process it generates, and what an observer can learn from that process. Equal answers need not come from equal internal structures. Conversely, an infinite state space may have a short, exact description and useful coordinates; failure to compress it into finitely many states need not end the investigation.

A small example: balanced parentheses

The grammar $S \rightarrow (S)S \mid \epsilon$ presents the language of balanced parentheses. To recognize a prefix, keep the number of unmatched opening parentheses, or a failure state if a closing parenthesis arrived too early. An opening parenthesis increments the count. A closing parenthesis decrements a positive count; at zero it enters failure. Failure persists. At the end, accept exactly when the count is zero.

The prefixes $($ and $(($ leave the same future task: any suffix that completes one completes the other. The prefix $(($ leaves a different task: $)$ completes the first two prefixes but not this one. These future tasks are **residuals**. There are infinitely many counts, yet a finite set of update rules generates their dynamics. This is a classical automata example, not a discovery of FPRD Lab. It illustrates the research question: which coordinates preserve all relevant futures, and what does updating or reading those coordinates cost?

From viewpoint to proof

In practice, the Lab fixes the observable and the allowed operations, compares presentations of the same problem, and seeks an exact translation or invariant. A failed shortcut is useful when a counterexample identifies the missing coordinate or hypothesis. Small computations test particular uncertainties; a proof must explain why the result holds beyond those tests. Costs, relation sides, boundary conditions, and the direction of each translation remain explicit.

AI-assisted exploration, source comparison, executable checks, and adversarial reconstruction support this work. Their agreement is not automatically independent evidence: shared assumptions must be exposed. A formal development certifies its stated formal theorem, not every interpretation attached to it. Each paper retains its own evidence and review status.

The approach is most useful when a poor representation hides a reusable construction, a compatibility condition, or a sharp obstruction. Outside recognition theory, it can guide the search for a structural reformulation without producing a literal residual dynamical system. The account below distinguishes that methodological influence from the ingredients of the proof.

Contribution to this publication

This addendum describes the current HTML reading edition. The preceding PDF pages retain their earlier edition; the afterword does not silently revise them.

FPRD's contribution to this paper

Here the useful change of presentation was numerical. The source reduction fixed several large parameters. The Lab treated them instead as coupled resource coordinates: field size, support cutoff, and moment horizon. It then asked which inequalities the retained proof actually needed, rather than changing a displayed exponent in isolation.

Eliminating those internal coordinates exposed a frontier at $1/28$ for the retained strict inequalities. The paper gives a concrete $1/30$ specialization and checks the finite constants as well as asymptotic exponents. These numbers describe different claims: feasibility below the retained frontier, an explicit specialization, and a barrier for this particular system of obligations. The frontier is not an impossibility theorem for other reductions.

The representation audit mattered at several interfaces. Binary Hamming weight becomes squared Euclidean distance, so the approximation exponents differ by a factor of two. Support must be discarded separately for each table type, rather than by taking an unjustified common intersection. The coefficient needed in binary soundness must survive into the Euclidean wrapper. A discrepancy between one imported formal wrapper and the paper's coefficient requirement was addressed explicitly instead of assuming a literal substitution preserved the proof.

Inherited ingredients and the boundary of the contribution

OpenAI's direct reduction supplies the one-hot encoding, reconstruction from polynomial moments, valuation and clause-matching arguments, and parity lift. The Lab did not replace those constructions. Its contribution is the resource parameterization, elimination, finite-constant bookkeeping, and corresponding retuning of the verification interfaces.

This is an application of the Lab's habit of separating representation costs and proof obligations. It is not a construction of a new language-residual system or an independent derivation of every imported lemma. The general written retuning and the archived executable Lean specialization also have different scopes: the latter establishes the concrete $1/30$ reduction, while a separate formal result treats feasibility of the parameter inequalities. The full reduction is not uniformly formalized for every exponent below $1/28$.

Where to inspect the contribution: the paper, Sections 1.2–1.3, 2, 5.5, and 7. These distinguish retained source mathematics, the repaired interfaces, and the exact formalization boundary.