AI-Guided Discovery for #EO:

From the f_{56} Cubic Potential to an FP Dichotomy

Jincheng Guan* **Shuai Shao*** **Ke Shi***

guanjincheng@mail.ustc.edu.cn shao10@ustc.edu.cn self.ke.shi@gmail.com

**School of Computer Science and Technology & Hefei National Laboratory, University of Science and Technology of China*

1. The Bottleneck in #EO

Complex-weighted $\#\mathrm{EO}(\mathcal{F})$ assigns local EO signatures from $\mathcal{F}$ to vertices and sums their weights over all Eulerian orientations. Via the standard NEQ_2 edge encoding, it is polynomial-time equivalent to $\mathrm{Holant}(\mathrm{NEQ}_2 \mid \mathcal{F})$, making it a key bottleneck toward complex-valued Holant dichotomies.

Prior work [MWX25a] gave an $\mathrm{FP}^{\mathrm{NP}}$ vs #P-hard dichotomy. The tractable side required:

- an XOR_3 up quasi-polymorphism or an XOR_3 down quasi-polymorphism;
- membership in the pairing class $\mathsf{EO}^{\mathscr{A}}$ or $\mathsf{EO}^{\mathscr{P}}$.

But the algorithm used an **NP oracle** for support identification.

Main Result: We remove the oracle.

Theorem. For every finite set $\mathcal{F}$ of complex-weighted EO signatures, if (i) all signatures in $\mathcal{F}$ satisfy the XOR_3 up quasi-polymorphism condition or all satisfy the XOR_3 down quasi-polymorphism condition, and (ii) $\mathcal{F} \subseteq \mathsf{EO}^{\mathscr{A}}$ or $\mathcal{F} \subseteq \mathsf{EO}^{\mathscr{P}}$, then $\#\mathrm{EO}(\mathcal{F}) \in \mathrm{FP}$. Otherwise $\#\mathrm{EO}(\mathcal{F})$ is #P-hard.

The oracle on the tractable side is replaced by a deterministic LP-based algorithm.

Where the oracle disappears

Old oracle task: decide whether a local support string can participate in some nonzero global EO configuration.

LP replacement: compute an LP-visible superset U_v by maximizing each $\lambda_{v,i}$ over $P(I)$. Every truly realizable local label lies in U_v; the product-pushforward lemma proves U_v is affine, which is enough to choose the pairings used in the final #CSP reduction. Since $\mathcal{F}$ is fixed, this uses only polynomially many LP solves.

2. Human–AI Discovery Path

AI transcript focuses on the special signature f_{56}
↓
First route: try to prove hardness
↓
Human spots a one-way reduction bug
↓
Pivot: search for a polynomial-time algorithm
↓
LP support-recovery viewpoint emerges
↓
Pinned-LP variants fail by counterexamples
↓
Recover **all** positive LP candidates U_v
↓
Cubic potential Ψ_{56} explains the 2-support phenomenon
↓
Generalize the degree-3 idea by product-pushforward
↓
Full tractable side becomes FP

The failed hardness route was useful: a human reviewer exposed the semantic gap, which redirected the search toward LP-based tractability. This is the poster's central human–AI pattern: propose, refute, repair, generalize.

Beyond #EO: odd-arity Holant

The removed oracle was also the remaining oracle source in the recent odd-arity Holant classification. Hence every finite complex-valued signature set $\mathcal{F}$ containing a non-trivial odd-arity signature has an ordinary dichotomy: $\mathrm{Holant}(\mathcal{F}) \in \mathrm{FP}$ or $\mathrm{Holant}(\mathcal{F})$ is #P-hard.

3. The f_{56} Cubic Potential

A Natural LP Relaxation

For each vertex v, enumerate $\mathrm{supp}(f_v) = \{\alpha_{v,1}, \ldots, \alpha_{v,m_v}\}$. The LP keeps a local distribution on this support:

$$\lambda_{v,i} \geq 0, \quad \sum_i \lambda_{v,i} = 1.$$

For a port h of v, set

$$S_{v,h} = \{i : \alpha_{v,i}(h) = 1\}$$

and define

$$p_{v,h} = \sum_{i \in S_{v,h}} \lambda_{v,i}.$$

If an edge pairs h with h', impose

$$p_{v,h} + p_{w,h'} = 1.$$

Let $P(I)$ be the resulting global LP polytope. Every nonzero global EO configuration gives an integral point of $P(I)$; the LP allows convex local mixtures and is therefore a relaxation.

LP-visible local candidates:

$$\Gamma_v = \left\{ i \in [m_v] : \max_{\lambda \in P(I)} \lambda_{v,i} > 0 \right\},$$

$$U_v = \{\alpha_{v,i} : i \in \Gamma_v\} \subseteq \mathrm{supp}(f_v).$$

Equivalently, U_v is the set of local support strings that can appear with positive mass in some feasible LP solution.

Let $\mathrm{supp}(f_{56}) = \{a_1, \ldots, a_5\}$. In sign form, write

$$\sigma_i = 2a_i - 1 \in \{\pm 1\}^{56}$$

and for an LP distribution $\lambda = (\lambda_1, \ldots, \lambda_5)$ define the port marginal

$$m_h(\lambda) = \sum_{i=1}^{5} \lambda_i \sigma_i(h).$$

Sign-domain dictionary

$\alpha \oplus \beta \oplus \gamma \leftrightarrow \sigma \cdot \tau \cdot v$ coordinatewise.
EO strings satisfy $\Delta(\sigma) = 0$; up/down failures have $\Delta(\sigma) > 0$ or $\Delta(\sigma) < 0$.
Edge complementarity becomes $\sigma_{v(h)} + \sigma_{w(h')} = 0$.

Key discovered invariant

$$\Psi_{56}(\lambda) = \frac{1}{8} \sum_{h=1}^{56} m_h(\lambda)^3 = 3 \sum_{1 \leq i < j < k \leq 5} \lambda_i \lambda_j \lambda_k \geq 0.$$

Edge complementarity gives $m_{v,h} + m_{w,h'} = 0$, so cubic terms cancel edge-by-edge:

$$\sum_v \Psi_{56}(\lambda_v) = 0.$$

Therefore every feasible local LP distribution has support of size at most 2. By convexity of $P(I)$, the LP-visible candidate set U_v for f_{56} also has size at most 2: otherwise averaging three witnessing LP points would create a 3-supported local distribution.

The f_{56} cubic calculation was the seed.

The full proof keeps the same degree-three idea, but replaces this scalar potential by a **product-pushforward** argument for arbitrary tractable EO signatures.

4. LP-based Algorithm

Input #EO instance $I = \Omega(G, \pi)$
↓
Solve the edge-complementarity LP $P(I)$
↓
Compute LP-visible sets U_v
↓
Product-pushforward → ternary-XOR closure of U_v
↓
U_v is support-affine
↓
Choose pairings P_v with $U_v \subseteq \mathrm{EOP}(P_v)$
↓
Form $g_v = (f_v)_{(P)_v} \in \mathscr{A}$ or $\mathscr{P}$
↓
Original edges become EQ_2 or NEQ_2 constraints
↓
Boolean #CSP → polynomial-time evaluation

Core Lemma: Product-Pushforward Closure

If all signatures satisfy the XOR_3 up quasi-polymorphism condition, or all satisfy the XOR_3 down quasi-polymorphism condition, then for every vertex v:

$$\alpha, \beta, \gamma \in U_v \Rightarrow \alpha \oplus \beta \oplus \gamma \in U_v.$$

Thus U_v is an affine subspace of $\mathbb{F}_2^{2d_v}$ contained in $\mathrm{supp}(f_v)$.

Why this proves tractability: the affine set U_v gives a pairing P_v with $U_v \subseteq \mathrm{EOP}(P_v)$. Since the original signature f_v lies in $\mathsf{EO}^{\mathscr{A}}$ or $\mathsf{EO}^{\mathscr{P}}$, the induced signature $g_v = (f_v)_{(P)_v}$ lies in $\mathscr{A}$ or $\mathscr{P}$. Every truly realizable local string lies in U_v, so this Boolean #CSP reduction preserves the original partition function.

5. What Did the AI Actually Do?

Route generation

AI proposed LP-based routes, support-recovery viewpoints, counterexample searches, and cubic/degree-three invariants around f_{56}.

Adversarial debugging

Counterexamples ruled out pinned-LP variants; the human caught a one-way-reduction bug in the initial hardness attempt — exactly the kind of semantic gap that motivates semantic auditing.

6. QR Links

Paper **Other materials**

References

[MWX25a] Meng, Wang, Xia. The $\mathrm{FP}^{\mathrm{NP}}$ versus #P-hard dichotomy for #EO.

[MWXZ25] Meng, Wang, Xia, Zheng. From an odd arity signature to a Holant dichotomy.

[CFS20] Cai, Fu, Shao. Beyond #CSP: counting weighted Eulerian orientations with ARS.

[CLX14] Cai, Lu, Xia. The complexity of complex weighted Boolean #CSP.