Exploring Line Bundle Standard Models with Transformers

Alessandro **Mininno**

based on arXiv:2607.00078 [hep-th] with J. H. Yip and G. Shiu

GitHub: LB-Explorer

Department of Physics
UNIVERSITY OF WISCONSIN–MADISON

The Unreasonable Effectiveness of Toric Geometry:
Bridging Mathematics, Computation, and String Theory
ESI – Vienna

July 7, 2026

Motivations

Tackle the String Landscape

1. Exploring the String Landscape has always been a **challenging problem**.
2. We focus on $E_8 \times E_8$ heterotic compactifications on Complete Intersection Calabi–Yau (CICY) threefolds to realize SU(5) GUT dubbed **Line Bundle Standard Models** (LBSMs).[1]
3. Existing scans for relatively small values of $h^{1,1}(\mathbf{X}_3)$ or line bundle values.[2]
4. Machine Learning (ML) approaches, so far, involved Reinforcement Learning[3] (RL) and Genetic Algorithm (GA).[4]

LB-Explorer

Transformer-based RL architecture able to find LBSMs. Available on `LB-Explorer`.

[1] L. B. Anderson, J. Gray, A. Lukas, E. Palti, *PoS* **CORFU2011**, ed. by K. Anagnostopoulos, I. Antoniadis, D. Bahns, N. Irges, A. Kehagias, G. Lazarides, D. Luest, H. Steinacker, G. Zoupanos, 096; L. B. Anderson, J. Gray, A. Lukas, E. Palti, *Phys. Rev. D* **84**, 106005, arXiv: 1106.4804 (hep-th); L. B. Anderson, J. Gray, A. Lukas, E. Palti, *JHEP* **06**, 113, arXiv: 1202.1757 (hep-th); L. B. Anderson, A. Constantin, J. Gray, A. Lukas, E. Palti, *JHEP* **01**, 047, arXiv: 1307.4787 (hep-th)

[2] L. B. Anderson, A. Constantin, J. Gray, A. Lukas, E. Palti, *JHEP* **01**, 047, arXiv: 1307.4787 (hep-th); Y.-H. He, S.-J. Lee, A. Lukas, C. Sun, *JHEP* **06**, 077, arXiv: 1309.0223 (hep-th); A. Constantin, Y.-H. He, A. Lukas, *Phys. Lett. B* **792**, 258–262, arXiv: 1810.00444 (hep-th)

[3] M. Larfors, R. Schneider, *Fortsch. Phys.* **68**, 2000034, arXiv: 2003.04817 (hep-th)

[4] S. A. Abel, A. Constantin, T. R. Harvey, A. Lukas, L. A. Nutricati, *Fortsch. Phys.* **72**, 2300260, arXiv: 2306.03147 (hep-th); P. Berglund, G. Butbaia, Y.-H. He, E. Heyes, E. Hirst, V. Jejjala, *Phys. Lett. B* **860**, 139158, arXiv: 2405.21017 (hep-th)

LB-Explorer

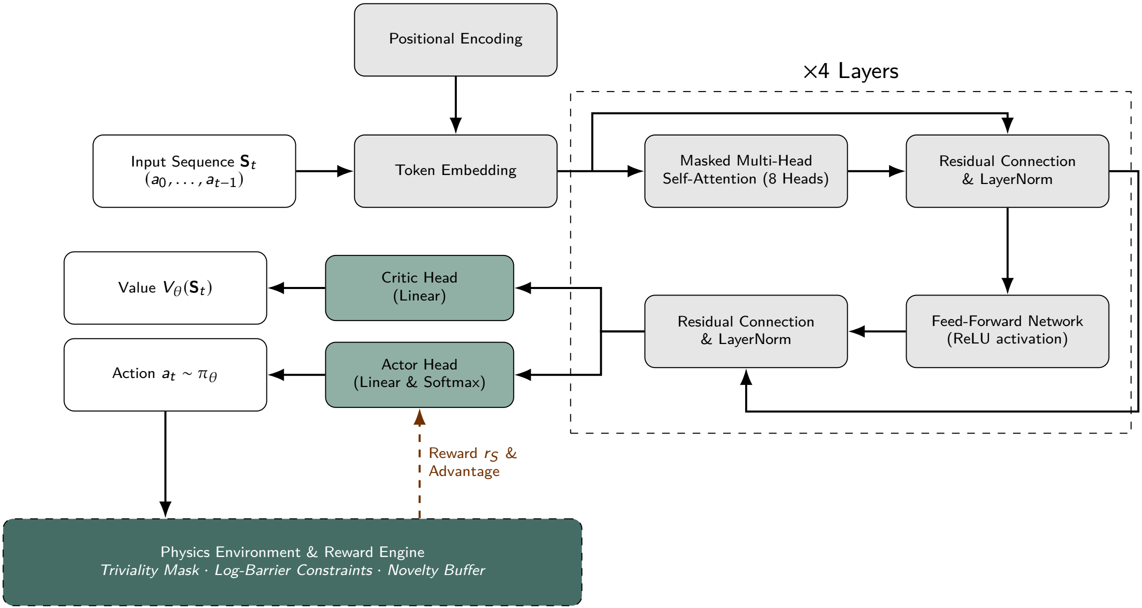
Positional Encoding
×4 Layers
Input Sequence $\mathbf{S}_t$ $(a_0, \dots, a_{t-1})$
Token Embedding
Masked Multi-Head Self-Attention (8 Heads)
Residual Connection & LayerNorm
Value $V_\theta(\mathbf{S}_t)$
Critic Head (Linear)
Residual Connection & LayerNorm
Feed-Forward Network (ReLU activation)
Action $a_t \sim \pi_\theta$
Actor Head (Linear & Softmax)
Reward r_S & Advantage
Physics Environment & Reward Engine
Triviality Mask · Log-Barrier Constraints · Novelty Buffer

Summary of the Results

	Conditions	LB-Explorer
E_8 embedding	$c_1(V) = 0\,,\ \sum_{a\in S} c_1(L_a) \neq 0$	✓ (Learned)
Anomaly cancellation	$c_2(T\mathbf{X}_3) - c_2(V) \in$ Mori cone	✓ (Learned)
Poly-stability	$\mu(L_a) = 0$	✓ (Learned)
Three chiral families	$\mathrm{ind}(V) = -3\lvert\Gamma\rvert$	✓ (Learned)
No exotic representations	$-3\lvert\Gamma\rvert \leq \mathrm{ind}(L_a) \leq 0\,,\ -3\lvert\Gamma\rvert \leq \mathrm{ind}(L_a \otimes L_b) \leq 0$	✓ (Learned)
Equivariance*	$\sum_{\text{distinct } L\subset V} m(L)\chi(\mathbf{X}_3, L) = 0 \mod \lvert\Gamma\rvert$	✗ (Filter-out later)
Spectrum	$h^1(\mathbf{X}_3, V) = 3\lvert\Gamma\rvert\,,\quad h^2(\mathbf{X}_3, V) = 0\,,$ $h^1(\mathbf{X}_3, \wedge^2 V) = 3\lvert\Gamma\rvert + n_h\,,\quad h^2(\mathbf{X}_3, \wedge^2 V) = n_h$	✗ (Filter-out later)

GitHub Repository

LB-Explorer Codebase

The complete codebase and datasets is available on GitHub:

`LB-Explorer`

Core Algorithms

- `LB-Explorer.py`: The main script.
- `CPSAT-closure.py`: CPSAT extension for extra constraints.
- `scripts/`: Auxiliary scripts for data processing and analysis.

Data & Setup

- `databases/`: datasets with properties of CICYs.[5]
- `setup_cy-explorer.sh`: Automated script to prepare the environment and install dependencies.

[5] L. B. Anderson, X. Gao, J. Gray, S.-J. Lee, *JHEP* **10**, 077, arXiv: 1708.07907 (hep-th); V. Braun, *JHEP* **04**, 005, arXiv: 1003.3235 (hep-th); A. Lukas, C. Mishra, *Commun. Math. Phys.* **379**, 847–865, arXiv: 1708.08943 (hep-th); R. Álvarez-García, F. Ruehle, arXiv: 2606.05280 (hep-th); L. B. Anderson, A. Constantin, J. Gray, Y.-H. He, S.-J. Lee, A. Lukas, arXiv: 2606.27588 (hep-th)

Summary of the Results

Scalability, Flexibility, & Transfer Learning

1. Evaluated on ~ 50 CICY manifolds over 10M episodes (across 5 seeds), outlining **strategies** to improve the discovery of solutions.
2. Significant **Transfer Learning** across CICYs, suitable for future generalization.

Hybrid LB-Explorer

LB-Explorer with **CP-SAT**[6] can be used to enforce extra constraints.

[6] L. Perron, F. Didier, *CP-SAT v9.12*, Google OR-Tools, https://developers.google.com/optimization/cp/cp_solver/, 2025

Systematic Scans: L. B. Anderson, A. Constantin, J. Gray, A. Lukas, E. Palti, *JHEP* **01**, 047, arXiv: 1307.4787 (hep-th); Y.-H. He, S.-J. Lee, A. Lukas, C. Sun, *JHEP* **06**, 077, arXiv: 1309.0223 (hep-th); A. Constantin, Y.-H. He, A. Lukas, *Phys. Lett. B* **792**, 258–262, arXiv: 1810.00444 (hep-th); L. B. Anderson, A. Constantin, J. Gray, Y.-H. He, S.-J. Lee, A. Lukas, arXiv: 2606.27588 (hep-th); **ML Searches**: M. Larfors, R. Schneider, *Fortsch. Phys.* **68**, 2000034, arXiv: 2003.04817 (hep-th); A. Constantin, T. R. Harvey, A. Lukas, *Fortsch. Phys.* **70**, 2100186, arXiv: 2108.07316 (hep-th); S. A. Abel, A. Constantin, T. R. Harvey, A. Lukas, L. A. Nutricati, *Fortsch. Phys.* **72**, 2300260, arXiv: 2306.03147 (hep-th); P. Berglund, G. Butbaia, Y.-H. He, E. Heyes, E. Hirst, V. Jejjala, *Phys. Lett. B* **860**, 139158, arXiv: 2405.21017 (hep-th)

Contents

Contents

Line Bundle Standard Models

SU(5) GUT from Heterotic Compactifications

We consider $E_8 \times E_8$ heterotic string theory general embeddings on Calabi–Yau $\mathbf{X}_3$,[7] with

1. V of E_8, **holomorphic vector bundle** over $\mathbf{X}_3$ with structure group H, i.e. $E_8 \supset G \times H$.
2. Existence of Wilson lines to break G to the MSSM gauge group.

Conditions

1. **Anomaly cancellation**: $\begin{cases} c_1(V) = 0 \\ c_2(T\mathbf{X}_3) - c_2(V) \in \text{Mori cone} \end{cases}$
2. **Poly-stability**: $\mu(L_a) = 0$.
3. **Equivariant structure**: $\exists \Gamma$ freely-acting symmetry on $\mathbf{X}_3$ and V is a vector bundle of $\widetilde{\mathbf{X}}_3 = \mathbf{X}_3/\Gamma$.
4. Extra conditions on particle spectrum.

[7] L. B. Anderson, J. Gray, A. Lukas, B. Ovrut, *JHEP* **02**, 088, arXiv: 1010.0255 (hep-th); L. B. Anderson, J. Gray, A. Lukas, B. Ovrut, *Phys. Rev. D* **83**, 106011, arXiv: 1102.0011 (hep-th); L. B. Anderson, J. Gray, A. Lukas, B. Ovrut, *JHEP* **10**, 032, arXiv: 1107.5076 (hep-th)

Ingredients for GUT Heterotic Compactifications

Assumptions

1. $V = \bigoplus_{a=1}^{5} L_a = \bigoplus_{a=1}^{5} \mathcal{O}_{\mathbf{X}_3}\left(\sum_{i=1}^{h^{1,1}} k_a^i J_i\right) = \bigoplus_{a=1}^{5} \mathcal{O}_{\mathbf{X}_3}(\mathbf{k}_a).$
2. Favorable CY threefolds X_3 with a **smooth** Kähler cone.
3. Known freely-acting symmetry Γ.[8]

Conditions on Particle Spectrum

1. Chirality: $\mathrm{ind}(V) = \sum_{a=1}^{5} \mathrm{ind}(L_a) = -3|\Gamma|$
2. No Exotic Representations: $\begin{cases} -3|\Gamma| \leq \mathrm{ind}(L_a) \leq 0 \\ -3|\Gamma| \leq \mathrm{ind}(L_a \otimes L_b) \leq 0 \end{cases}$

[8] V. Braun, *JHEP* **04**, 005, arXiv: 1003.3235 (hep-th); A. Lukas, C. Mishra, *Commun. Math. Phys.* **379**, 847–865, arXiv: 1708.08943 (hep-th)

Summary of Conditions

	Conditions	LB-Explorer
E_8 embedding	$c_1(V) = 0\,,\ \sum_{a\in S} c_1(L_a) \neq 0$	✓ (Learned)
Anomaly cancellation	$c_2(T\mathbf{X}_3) - c_2(V) \in$ Mori cone	✓ (Learned)
Poly-stability	$\mu(L_a) = 0$	✓ (Learned)
Three chiral families	$\mathrm{ind}(V) = -3\lvert\Gamma\rvert$	✓ (Learned)
No exotic representations	$-3\lvert\Gamma\rvert \leq \mathrm{ind}(L_a) \leq 0\,,\ -3\lvert\Gamma\rvert \leq \mathrm{ind}(L_a \otimes L_b) \leq 0$	✓ (Learned)
Equivariance*	$\sum_{\text{distinct } L\subset V} m(L)\chi(\mathbf{X}_3, L) = 0 \mod \lvert\Gamma\rvert$	✗ (Filter-out later)
Spectrum	$h^1(\mathbf{X}_3, V) = 3\lvert\Gamma\rvert\,,\quad h^2(\mathbf{X}_3, V) = 0\,,$ $h^1(\mathbf{X}_3, \wedge^2 V) = 3\lvert\Gamma\rvert + n_h\,,\quad h^2(\mathbf{X}_3, \wedge^2 V) = n_h$	✗ (Filter-out later)

See: L. B. Anderson, Y.-H. He, A. Lukas, *JHEP* **07**, 104, arXiv: 0805.2875 (hep-th); L. B. Anderson, Y.-H. He, A. Lukas, *JHEP* **07**, 049, arXiv: hep-th/0702210; L. B. Anderson, Other thesis, University of Oxford, Aug. 2008, arXiv: 0808.3621 (hep-th); V. Braun, *JHEP* **04**, 005, arXiv: 1003.3235 (hep-th); L. B. Anderson, J. Gray, A. Lukas, E. Palti, *Phys. Rev. D* **84**, 106005, arXiv: 1106.4804 (hep-th); L. B. Anderson, J. Gray, A. Lukas, E. Palti, *JHEP* **06**, 113, arXiv: 1202.1757 (hep-th); L. B. Anderson, A. Constantin, J. Gray, A. Lukas, E. Palti, *JHEP* **01**, 047, arXiv: 1307.4787 (hep-th); A. Lukas, C. Mishra, *Commun. Math. Phys.* **379**, 847–865, arXiv: 1708.08943 (hep-th)

Contents

LB-Explorer Scheme

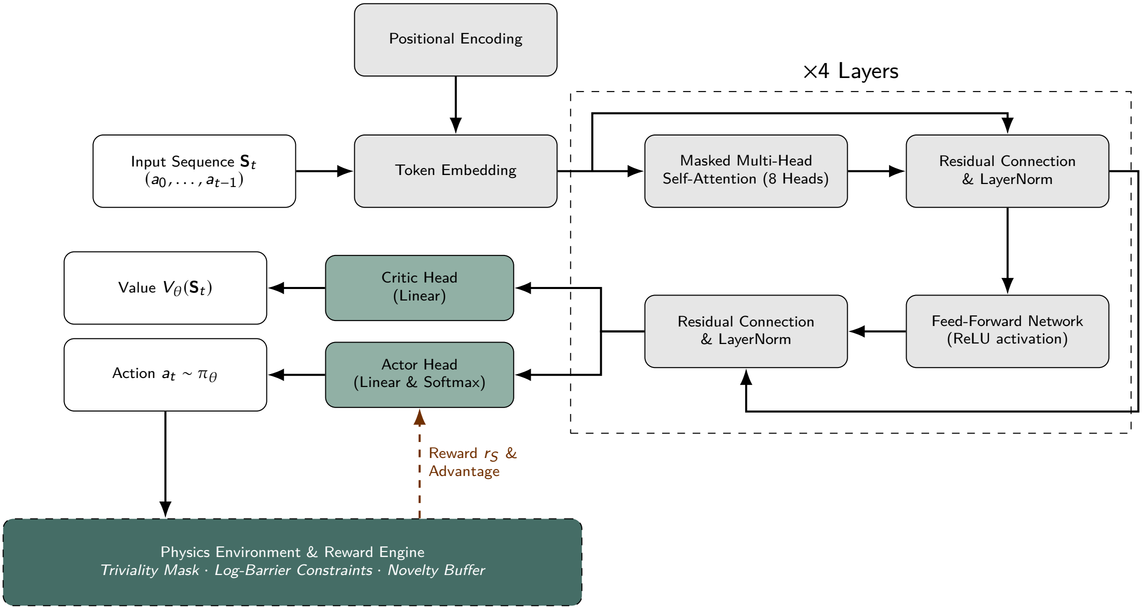

LB-Explorer Architecture: Input Processing

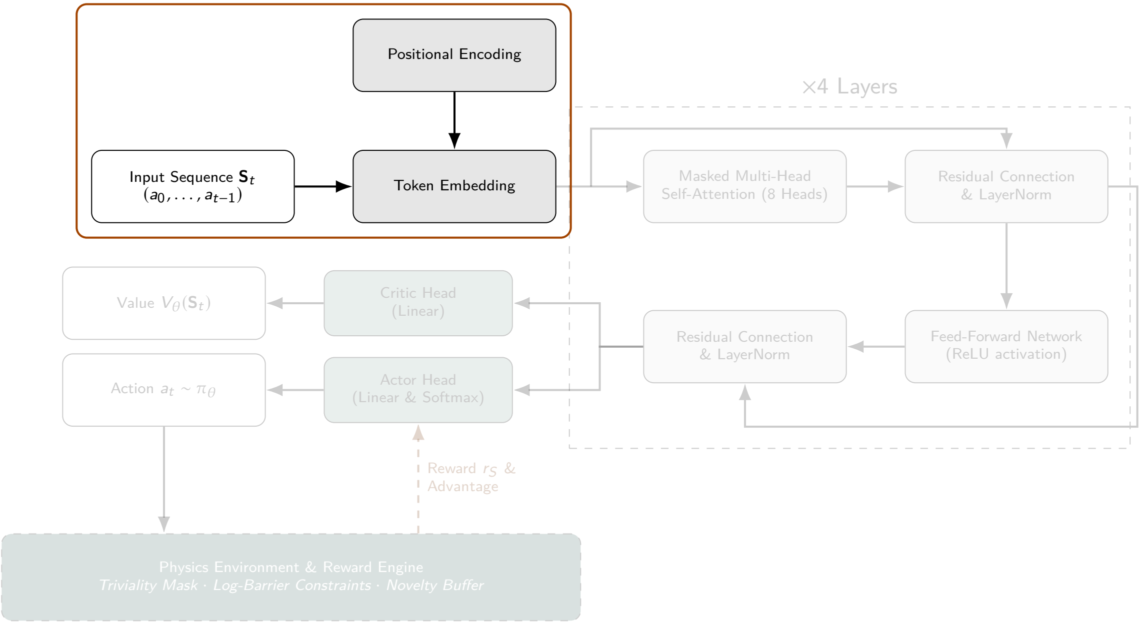

Input Processing

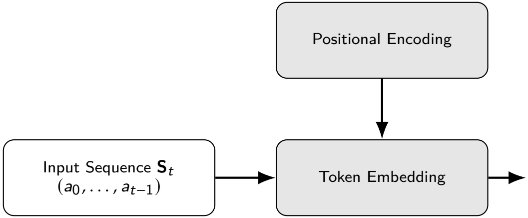

Input Initialization

1. **Sequence and Tokenization**: $\mathbf{K} = \{k_a^i\}_{i=1,\dots,h^{1,1}}^{a=1,\dots,5}$ is a flattened sequence $\mathbf{S} = \{s_t\}_{t=0,\dots,S-1}$, tokenized so that $s_t = a_t - k_{\max}$ and $a_t \in [0, 2k_{\max}] \equiv \mathbf{A}$, entries of a **vocabulary**.
2. **Blank Slate**: LB-Explorer starts with no prior knowledge of the CY. At step $t = 0$, the input sequence consists only of a special `[START]` token.

LB-Explorer Architecture: Transformer Decoder

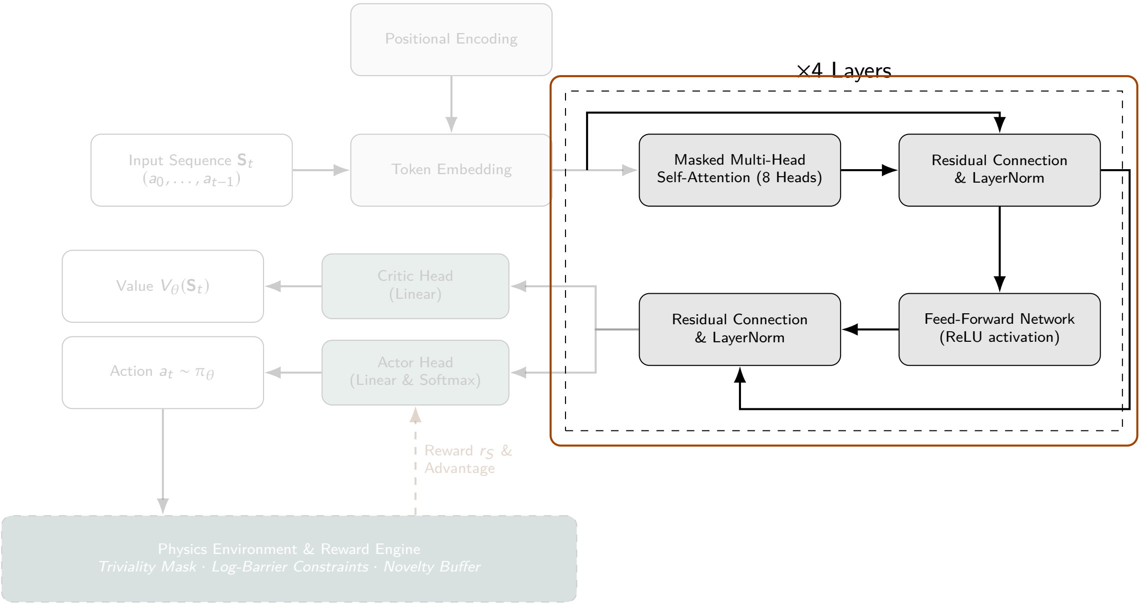

Transformer Decoder

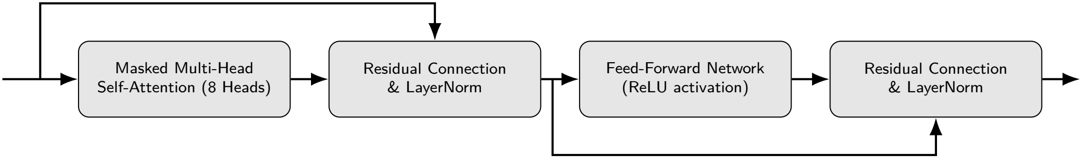

Contextualizing Past Tokens

1. **Self-Attention**:[9] LB-Explorer looks at the sequence generated **so far**. It analyzes patterns in the previous tokens using 8 attention heads to track different constraints.
2. **Masking**: The mask forces LB-Explorer to deduce the next step without looking into the ungenerated tokens.

[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, *arXiv e-prints*, arXiv:1706.03762, arXiv: `1706.03762 (cs.CL)`

LB-Explorer Architecture: Actor-Critic Generation

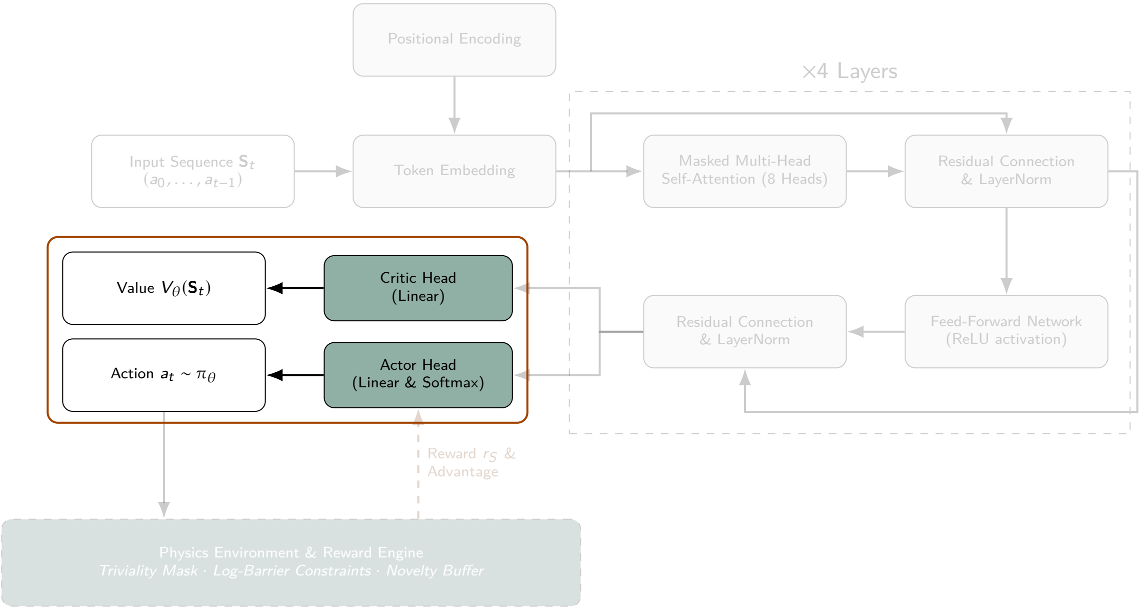

Actor-Critic: Autoregressive Matrix Generation

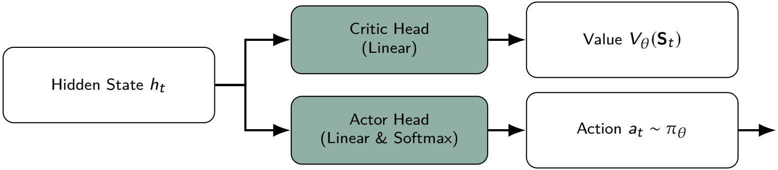

Predicting the Next Tokens

1. **Actor Head**:[10] The Actor outputs a probability distribution over the vocabulary.
2. **Critic Head**: The Critic guesses the final reward based on the current partial sequence.
3. **Loop**: All previous steps are repeated in loop until $\mathbf{K}$ is generated.

[10] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, *arXiv e-prints*, arXiv:1707.06347, arXiv: `1707.06347 (cs.LG)`

LB-Explorer Architecture: Environment & Reward

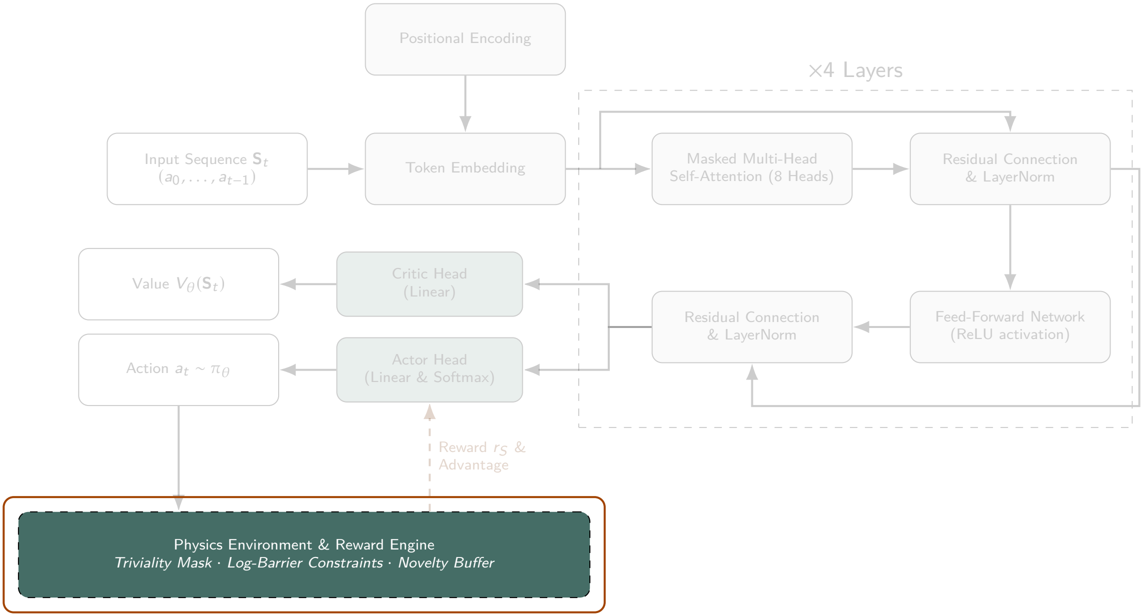

Environment & Reward

Generated Matrix

$\downarrow$

Physics Environment & Reward Engine $\dashrightarrow$ Reward r_S & Advantage

Terminal Reward r_S

$$r_S = \begin{cases} B+5 & \forall\, C_i \text{ \& novel} \\ p_{\text{pen}}(B+5) & \forall\, C_i \text{ but dup.} \\ B & \exists C_i \text{ \& novel} \\ p_{\text{pen}}B & \exists C_i \text{ but dup.} \\ 0 & \text{otherwise} \end{cases}$$

Base Score:

$$B = \sum_i w_i C_i$$

Optimization & GAE

Advantage Estimation:

$$\hat{A}_t = \sum_{l=0}^{S-t-1} (\gamma\lambda)^l \delta_{t+l}$$

$$\delta_t = r_t + \gamma V_{\text{old}}(\mathbf{S}_{t+1}) - V_{\text{old}}(\mathbf{S}_t)$$

Total PPO Loss:

$$\mathcal{L}^{\text{total}} = \mathcal{L}^{\text{CLIP}} + c_1 \mathcal{L}^{\text{VF}} - c_2 \mathcal{S}[\pi_\theta]$$

Contents

Complete Intersection Calabi-Yau Threefolds (CICYs)

CICYs

Obtained as zero loci in projective spaces $\mathcal{A} = \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_m}$,[11] encoded in a **Configuration Matrix**.

Property	CICY 6927
Conf.	$\left[\begin{array}{c\|cccccc} \mathbb{P}^1 & 1 & 1 & 0 & 0 & 0 & 0 \\ \mathbb{P}^1 & 0 & 0 & 1 & 1 & 0 & 0 \\ \mathbb{P}^1 & 0 & 0 & 0 & 0 & 0 & 2 \\ \mathbb{P}^1 & 0 & 0 & 0 & 0 & 2 & 0 \\ \mathbb{P}^5 & 1 & 1 & 1 & 1 & 1 & 1 \end{array}\right]^{5,37}_{-64}$
$c_2(T\mathbf{X}_3) \cdot J_i$	$\{24, 24, 24, 24, 64\}$
$\mathcal{I}$	$2J_1J_2J_3 + 2J_1J_2J_4 + 2J_1J_3J_4 + 2J_2J_3J_4 + 4J_1J_2J_5 + 4J_1J_3J_5 + 4J_2J_3J_5 + 4J_1J_4J_5 + 4J_2J_4J_5 + 4J_3J_4J_5 + 8J_1J_5^2 + 8J_2J_5^2 + 8J_3J_5^2 + 8J_4J_5^2 + 16J_5^3$
Sym.	$G_{\mathbf{X}_3} = S_4 = \langle (34), (23), (12) \rangle, \quad G_{\text{conf}} = \mathbb{Z}_2 \times \mathbb{Z}_2 = \langle (34), (12) \rangle, \quad \Gamma = \mathbb{Z}_2 \times \mathbb{Z}_2 \quad (\|\Gamma\| = 4)$

[11] P. Candelas, A. M. Dale, C. A. Lutken, R. Schimmrigk, *Nucl. Phys. B* **298**, 493; L. B. Anderson, Y.-H. He, A. Lukas, *JHEP* **07**, 049, arXiv: hep-th/0702210; L. B. Anderson, Y.-H. He, A. Lukas, *JHEP* **07**, 104, arXiv: 0805.2875 (hep-th); L. B. Anderson, X. Gao, J. Gray, S.-J. Lee, *JHEP* **10**, 077, arXiv: 1708.07907 (hep-th); V. Braun, *JHEP* **04**, 005, arXiv: 1003.3235 (hep-th)

Some Results

Set-up

1. We considered $\sim$ **50** favorable CICYs with $h^{1,1}(\mathbf{X}_3) \in [4, 15]$ and $|\Gamma| \in [2, 4]$.
2. We ran the LB-Explorer over **five seeds** for **10 million** episodes.
3. We collected the solutions found and filter them to impose equivariant structure.

#	$h^{1,1}$	$\lvert\Gamma\rvert$	Sol.	Total	Equivariance
7447	5	2	N	43903	12670
			N_{S_5}	12658	3919
			$N_{G_{\mathbf{X}_3}}$	2590	804
			N_{full}	554	204
7300	8	2	N	2403949	372913
			N_{S_5}	2365092	369556
			$N_{G_{\mathbf{X}_3}}$	2289114	357146
			N_{full}	2232797	348247
5257	10	2	N	2903468	482530
			N_{S_5}	2899680	482427
			$N_{G_{\mathbf{X}_3}}$	2898620	482420
			N_{full}	2878537	481878

Learning Pattern: Staged Discovery

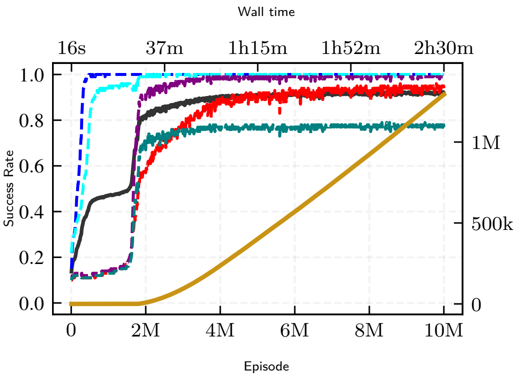

CICY 3413, $|\Gamma| = 3$, Seed 42

Staged Learning

The LB-Explorer learns the constraints in sequential stages:

- **Early Stage**: Anomaly cancellation and poly-stability are learned first.
- **Critical Threshold**: First solutions found when chirality condition is rewarded.
- **Fine-Tuning**: Extra conditions improve the search.

Exploration Status: Exhaustive or Not?

Observed Scenarios

The accumulation of solutions for different CICYs generally falls into three regimes.

Exploring (CICY 7300)

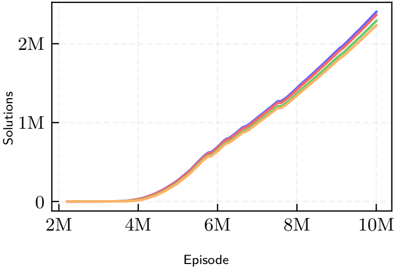

Consistently finding new, non-redundant solutions.

Plateaued (CICY 2639)

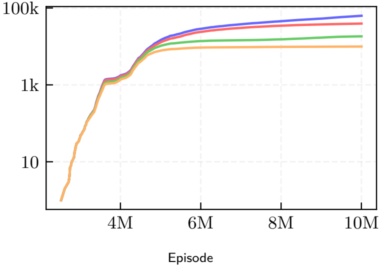

Slowdown in finding any valid solutions.

Exhausted (CICY 5425)

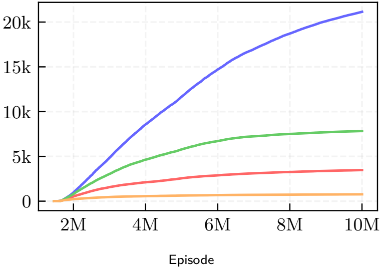

Finding "new" solutions that are redundant.

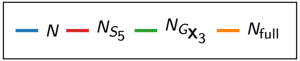

Improving Exploration: Increasing Episodes

Extending Exploration

Extending the search increases the solutions, but ultimately the learning plateaus.

CICY 7300 ($|\Gamma| = 2$)

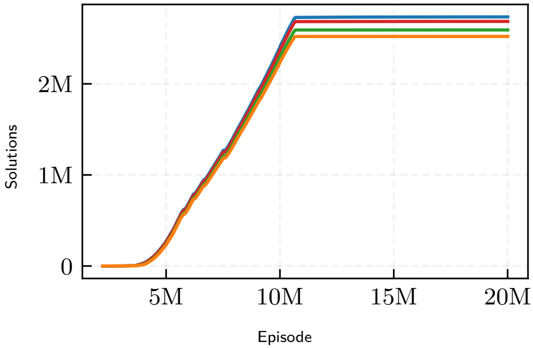

Sol.	10M Eps.	20M Eps.	Equivariant
N	2403949	2734631	427908
N_{S_5}	2365092	2685187	423916
$N_{G\mathbf{x}_3}$	2289114	2591548	407864
N_{full}	2232797	2520239	396819

- Finds 14% more solutions at early stages.
- No overlaps across seeds (duplication $< 0.01\%$).

Improving Exploration: Increasing Seeds

Increasing Number of Seeds

A more effective strategy is **increasing** the number of seeds.

CICY 6 ($h^{1,1} = 15$, 10 Seeds)

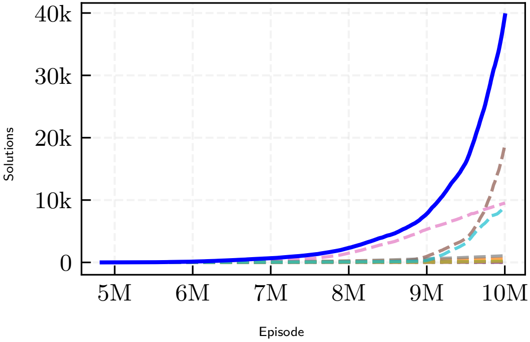

Sol.	Seeds 42-46	Seeds 42-51	Equivariant
N	846	39628	5166
N_{S_5}	828	37891	4941
$N_{G_{X_3}}$	845	39623	5164
N_{full}	827	37882	4940

- Significant increment.
- Negligible duplication of solutions.

Transfer Learning: Scenarios & Methods

Transfer Learning Scenarios

We analyze the NN's ability to generalize in three regimes:

1. **Within Bucket**: Same $h^{1,1}(\mathbf{X}_3)$ and $|\Gamma|$, different CICYs, i.e. different $G_{\mathbf{X}_3}$.
2. **Across $|\Gamma|$**: Same CICY but different $|\Gamma|$.
3. **Across $h^{1,1}(\mathbf{X}_3)$**: Same $|\Gamma|$, different CICY.

Evaluation Methods

1. **Zero-shot Reward** and **Episodes to First Solution**.
2. **Sample Efficiency**: Episodes required to discover 1000 unique solutions.
3. **Weight Distance & CKA Heatmap**: Restructuring of internal weights.
4. **Policy Divergence**: Token-by-token behavioral shift of the autoregressive policy.
5. **Mean CRS (Continuous Reward Score)**: Measure catastrophic forgetting.

Within Bucket: Average Best

CICY 5967 → CICY 3413

- **Zero-Shot Reward**: +260.6% improvement (0.155 to 0.558).
- **Initial Discovery**: Drops from 1.6M to 90K episodes (+94.4% speedup).
- **Sample Efficiency**: Finds 1000 solutions in 245K episodes (+86.4% speedup).

Sample Efficiency

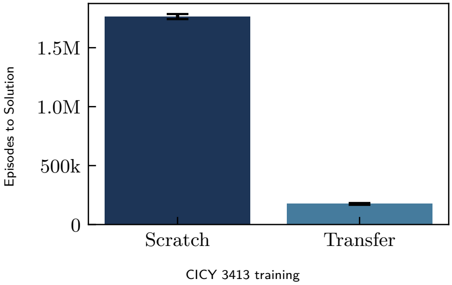

CKA Heatmap

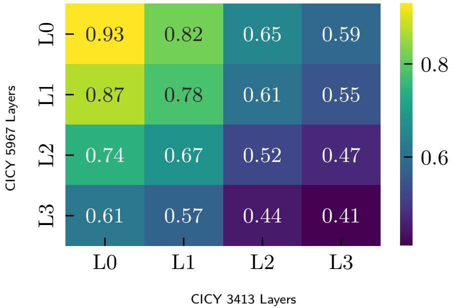

Within Bucket: Directional Asymmetry

Trade-off based on $|G_{\mathbf{X}_3}|$

1. Transfer large $|G_{\mathbf{X}_3}|$ to small $|G_{\mathbf{X}_3}|$: **Fast** but **local** exploration.
2. Transfer small $|G_{\mathbf{X}_3}|$ to large $|G_{\mathbf{X}_3}|$: **Fast** and **broad** exploration.

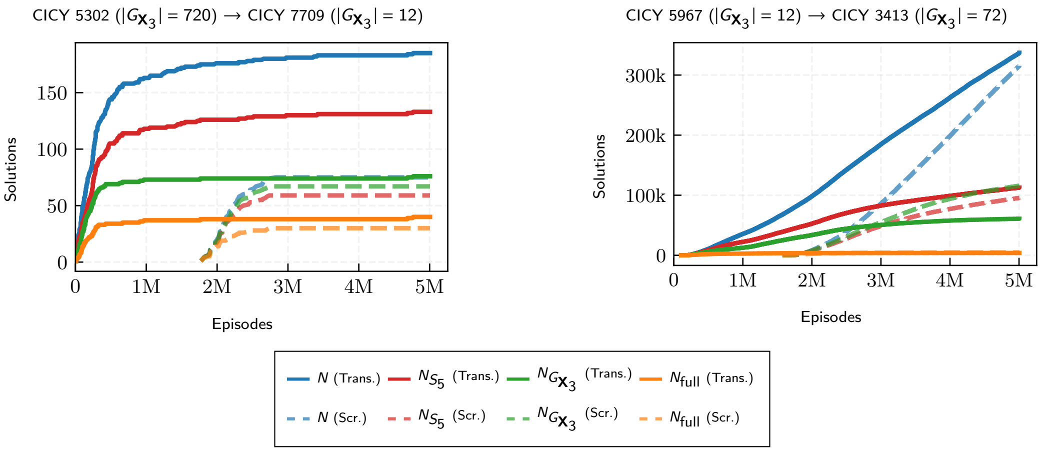

Across $|\Gamma|$: Average Best

CICY 5302 ($|\Gamma| = 4$) $\rightarrow$ CICY 5302 ($|\Gamma| = 2$)

- **Zero-Shot Reward**: +197.2% improvement (0.275 to 0.811).
- **Initial Discovery**: Drops from 1.4M to ~19K episodes (77.6× speedup).
- **Sample Efficiency**: Finds 1000 solutions in 196K episodes (8.7× speedup).

Sample Efficiency

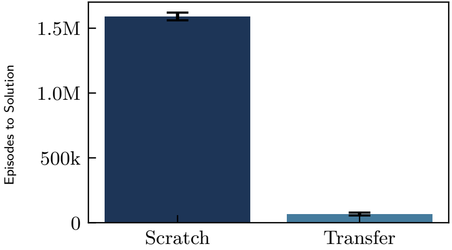

CICY 5302 ($|\Gamma| = 2$) training

Weight Distance

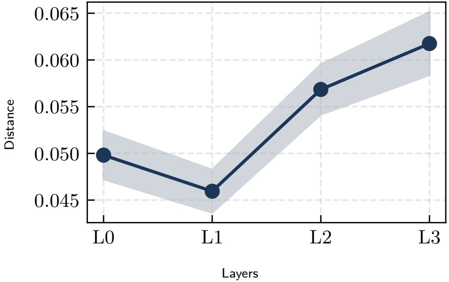

Across $|\Gamma|$: Directional Asymmetry

Trade-off based on $|\Gamma|$

1. Transfer large $|\Gamma|$ to small $|\Gamma|$: **Fast** and **broad** exploration.
2. Transfer small $|\Gamma|$ to large $|\Gamma|$: **Fast** but **local** exploration.

CICY 5302 ($|\Gamma| = 2$) $\rightarrow$ CICY 5302 ($|\Gamma| = 4$) | CICY 5452 ($|\Gamma| = 4$) $\rightarrow$ CICY 5452 ($|\Gamma| = 2$)

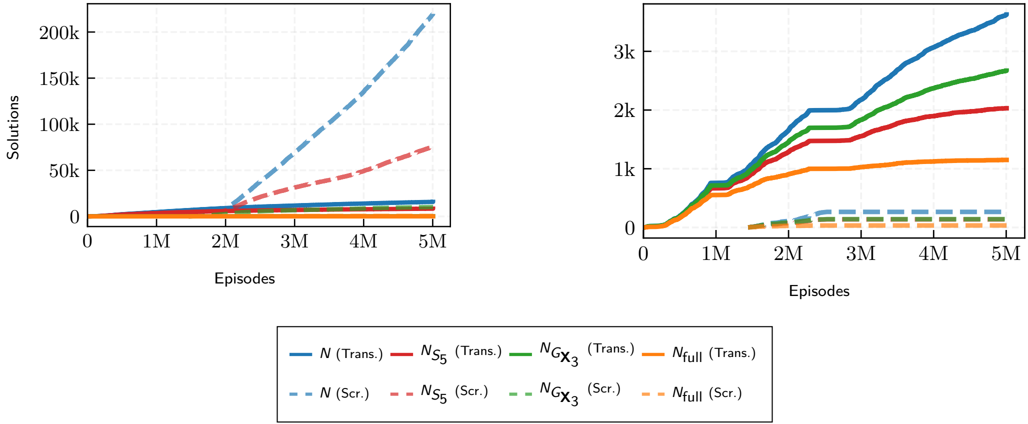

Across $h^{1,1}(\mathbf{X}_3)$: Average Best

CICY 3413 ($h^{1,1} = 6$) → CICY 6024 ($h^{1,1} = 5$)

- **Zero-Shot Reward**: +303.7% improvement (0.162 to 0.654).
- **Initial Discovery**: Drops from 1.25M to 8K episodes (+99.3% speedup).
- **Sample Efficiency**: Finds 1000 solutions in 54K episodes (+96.3% speedup).

Sample Efficiency

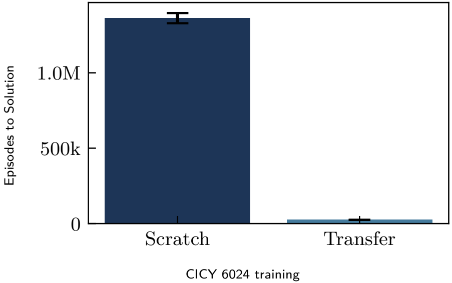

Policy Divergence

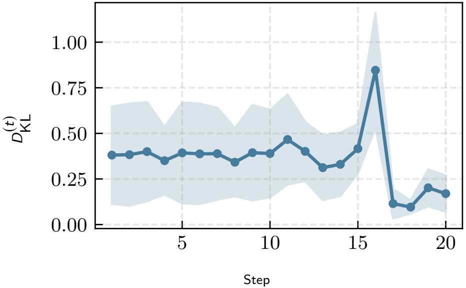

Across $h^{1,1}(\mathbf{X}_3)$: Directional Asymmetry

Trade-off based on Dimensional Topology

1. Transfer large $h^{1,1}(\mathbf{X}_3)$ to small $h^{1,1}(\mathbf{X}_3)$: **Fast** and **broad** exploration.
2. Transfer small $h^{1,1}(\mathbf{X}_3)$ to large $h^{1,1}(\mathbf{X}_3)$: **Fast** but **local** exploration.

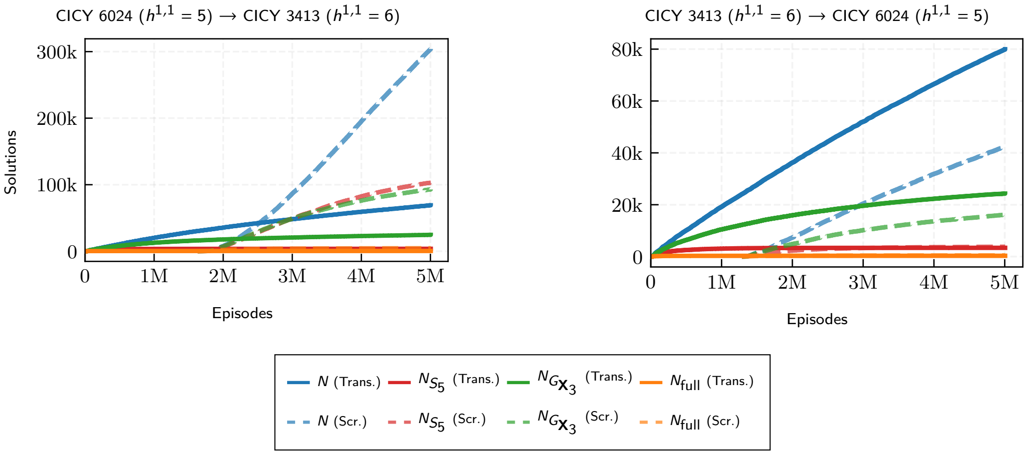

Summary of Transfer Learning

Transfer Learning

All Transfer Learnings show **speed up** in discovering solutions, with some **restructuring** of internal NN parameters.

Optimal Transfer

Small $\lvert G_{\mathbf{X}_3}\rvert$	$\Longrightarrow$	**Large** $\lvert G_{\mathbf{X}_3}\rvert$
Large $\lvert\Gamma\rvert$	$\Longrightarrow$	**Small** $\lvert\Gamma\rvert$
Large $h^{1,1}(\mathbf{X}_3)$	$\Longrightarrow$	**Small** $h^{1,1}(\mathbf{X}_3)$

Sub-optimal Transfer

Larger $\lvert G_{\mathbf{X}_3}\rvert$	$\Longrightarrow$	**Small** $\lvert G_{\mathbf{X}_3}\rvert$
Small $\lvert\Gamma\rvert$	$\Longrightarrow$	**Large** $\lvert\Gamma\rvert$
Small $h^{1,1}(\mathbf{X}_3)$	$\Longrightarrow$	**Large** $h^{1,1}(\mathbf{X}_3)$

Contents

Stringent Filtering Constraints

Beyond LB-Explorer Conditions

- **Equivariant Structure**: Imposing equivariance truncates the solutions.
- **Spectrum Analysis**:[12] Compute line bundle cohomologies is computational expensive.

#	$h^{1,1}$	$\lvert\Gamma\rvert$	**Sol.**	**Total**	**Equivariance**	**Spectrum** ($\lvert\mathbf{k}_a\rvert \leq 2$)
5302	6	2	N	187426	26594	6095 (0)
			N_{S_5}	30547	4644	852 (0)
			$N_{G_{\mathbf{X}_3}}$	2882	379	743 (0)
			N_{full}	194	32	76 (0)

[12] L. B. Anderson, J. Gray, A. Lukas, E. Palti, *Phys. Rev. D* **84**, 106005, arXiv: 1106.4804 (hep-th); L. B. Anderson, A. Constantin, J. Gray, A. Lukas, E. Palti, *JHEP* **01**, 047, arXiv: 1307.4787 (hep-th); M. Larfors, R. Schneider, *Fortsch. Phys.* **67**, 1900083, arXiv: 1906.00392 (hep-th)

Augmenting the LB-Explorer

Where does the LB-Explorer struggle?

1. LB-Explorer is efficient at navigating the space of configurations.
2. It struggles with equalities (chirality conditions) for sparse solutions.

Imposing Extra Constraints

1. We take CICYs where LB-Explorer did **not** find solutions: CICYs with $h^{1,1}(\mathbf{X}_3) = 4$.
2. **Phase 0 Trigger**: We halt training when all scores are learned **except for** chirality, i.e. $\mathbf{K}_0$.
3. **CP-SAT** Module:[13] We perturb $\mathbf{K}_0$ to impose chirality using exact numerical methods.

[13] L. Perron, F. Didier, *CP-SAT v9.12*, Google OR-Tools, `https://developers.google.com/optimization/cp/cp_solver/`, 2025

Solving Constraints via CP-SAT

CP-SAT Optimization

CP-SAT is an **optimization solver** that explores the solution tree using logic and constraint propagation.

Objective

Find a matrix $\mathbf{K}'$ that satisfies chirality while minimizing:

$$\min_{\mathbf{K}'} ||\mathbf{K}' - \mathbf{K}_0|| = \min_{\mathbf{K}'} \sum_{i=1}^{h^{1,1}(\mathbf{X}_3)} \sum_{a=1}^{5} |\mathbf{K}'^{\,i}_a - \mathbf{K}^i_{0|a}|$$

Results with CP-SAT

Hybrid LB-Explorer

Pure LB-Explorer failed to find solutions? Hybrid LB-Explorer is the **answer**.

#	$h^{1,1}$	$\lvert\Gamma\rvert$	Sol.	Total	Equivariance
7862	4	2	N	7222	1984
			N_{S_5}	2868	884
			$N_{G_{\mathbf{X}_3}}$	1237	590
			N_{full}	440	150
7862	4	4	N	23403	2715
			N_{S_5}	1206	88
			$N_{G_{\mathbf{X}_3}}$	3954	290
			N_{full}	72	12

Upshot

Hybrid LB-Explorer can be used to learn **more conditions**.

Contents

Summary of the Results

The LB-Explorer Framework

Transformer-based RL architecture designed to find heterotic **line bundle standard models** on CICYs.

Scalability, Flexibility, & Transfer Learning

1. Evaluated on ~ 50 CICY manifolds over 10M episodes (across 5 seeds), outlining **strategies** to improve the discovery of solutions.
2. Significant **Transfer Learning** across CICYs, suitable for future generalization.

Hybrid LB-Explorer

LB-Explorer with **CP-SAT** can be used to enforce extra constraints.

GitHub Repository

LB-Explorer Codebase

The complete codebase and datasets is available on GitHub:

`LB-Explorer`

Core Algorithms

- `LB-Explorer.py`: The main script.
- `CPSAT-closure.py`: CPSAT extension for extra constraints.
- `scripts/`: Auxiliary scripts for data processing and analysis.

Data & Setup

- `databases/`: datasets with properties of CICYs.[14]
- `setup_cy-explorer.sh`: Automated script to prepare the environment and install dependencies.

[14] L. B. Anderson, X. Gao, J. Gray, S.-J. Lee, *JHEP* **10**, 077, arXiv: 1708.07907 (hep-th); V. Braun, *JHEP* **04**, 005, arXiv: 1003.3235 (hep-th); A. Lukas, C. Mishra, *Commun. Math. Phys.* **379**, 847–865, arXiv: 1708.08943 (hep-th); R. Álvarez-García, F. Ruehle, arXiv: 2606.05280 (hep-th); L. B. Anderson, A. Constantin, J. Gray, Y.-H. He, S.-J. Lee, A. Lukas, arXiv: 2606.27588 (hep-th)

Future Directions

Possible Generalizations

- Integrate **extra constraints** into the LB-Explorer or its Hybrid version.
- Integrate ML techniques to compute line bundle cohomologies.[15]
- Going beyond, introducing **non-Abelian** line bundles,[16] computing Yukawa couplings...[17]
- Test the architecture on different set-ups, e.g. **PFFVs** in Type IIB compactifications[18]

[15] D. Klaewer, L. Schlechter, *Phys. Lett. B* **789**, 438–443, arXiv: 1809.02547 (hep-th); C. R. Brodie, A. Constantin, R. Deen, A. Lukas, *Fortsch. Phys.* **68**, 1900087, arXiv: 1906.08730 (hep-th); C. R. Brodie, A. Constantin, A. Lukas, F. Ruehle, *JHEP* **03**, 024, arXiv: 2108.10323 (hep-th); A. Constantin, T. R. Harvey, A. Lukas, *Fortsch. Phys.* **70**, 2100186, arXiv: 2108.07316 (hep-th); A. Constantin, arXiv: 2401.14463 (math.AG)

[16] E. I. Buchbinder, A. Constantin, A. Lukas, *JHEP* **03**, 025, arXiv: 1311.1941 (hep-th)

[17] L. B. Anderson, J. Gray, D. Grayson, Y.-H. He, A. Lukas, *Commun. Math. Phys.* **297**, 95–127, arXiv: 0904.2186 (hep-th)

[18] M. Demirtas, M. Kim, L. Mcallister, J. Moritz, *Phys. Rev. Lett.* **124**, 211603, arXiv: 1912.10047 (hep-th); R. Álvarez-García, R. Blumenhagen, M. Brinkmann, L. Schlechter, *Fortsch. Phys.* **68**, 2000088, arXiv: 2009.03325 (hep-th); M. Demirtas, M. Kim, L. McAllister, J. Moritz, *Fortsch. Phys.* **68**, 2000085, arXiv: 2009.03312 (hep-th); F. Carta, A. Mininno, P. Shukla, *JHEP* **02**, 205, arXiv: 2112.13863 (hep-th); F. Carta, A. Mininno, P. Shukla, *JHEP* **08**, 297, arXiv: 2201.10581 (hep-th)

Thank you!