Machine Learning the Refined de Sitter Conjecture

A No-Scale Supergravity Case Study

Hyunjoo Jung

Department of Physics, Yonsei University

SI2026 · Fuji Calm, August 2026

What a swampland classifier learns depends on the field-space metric you put in its labels.

Why?

what this project is, and why I did it

the error is **silent**: training, validation and test share the same labels

swampland criteria · machine-learned labels · field-space metric

Moduli — the size and shape of X_6

String theory is quantum-consistent only in **ten** dimensions (anomaly cancellation) $\Rightarrow$ six are compactified on an internal space X_6:

$$ds^2_{10} = g_{\mu\nu}\, dx^\mu dx^\nu + g_{mn}(y;\Phi)\, dy^m dy^n$$

Everything about X_6 lives in the internal metric g_{mn} — whatever g_{mn} is free to change becomes a 4D field.

shape — complex-structure modulus τ
$(T^2 = \mathbb{C}/(\mathbb{Z}+\tau\mathbb{Z}))$

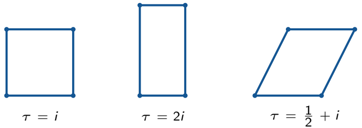

Fig. 1. deformations of Ω, counted by $h^{2,1}$;
both parts of τ are geometric (tilt + aspect ratio) $\to U$

size — volume-type modulus $x = \mathrm{Re}\,\Phi$

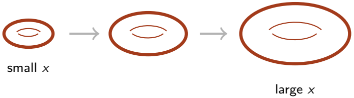

Fig. 2. the Kähler 2-form $J = \sum_{a=1}^{h^{1,1}} t^a \omega_a$: one volume per 2-cycle $\to x = \mathrm{Re}\,T$;
the gauge form-field on the same cycle $\to$ the axion $y = \mathrm{Im}\,T$

Every choice of (size, shape) is an allowed background — to a 4D observer each one is a **field**:

$$\Phi = x + iy \iff \underbrace{x}_{\text{size of the internal geometry (e.g. Re } T)} + \underbrace{i\,y}_{\text{axionic direction (e.g. Im } T)}$$

T — Kähler: **size/volume** · U — complex structure: **shape** · S — axio-dilaton: **string coupling**

With no potential they would be massless — **moduli stabilization**: fluxes must pick the vacuum via $V(\Phi)$.

The 4D data: K, W, V

Dimensional reduction leaves a 4D $\mathcal{N}=1$ supergravity fixed by **two functions** of the moduli:

$$S_4 = \int d^4x\sqrt{-g}\left[\frac{M_P^2}{2}R - K_{I\bar{J}}\,\partial_\mu\Phi^I\partial^\mu\bar{\Phi}^{\bar{J}} - V\right]$$

K **(Kähler potential)** — fixes the kinetic terms, i.e. **the metric on field space**, $G_{I\bar{J}} = \partial_I\partial_{\bar{J}}K$; *logarithmic* at tree level:

$$K = -\textstyle\sum_i n_i \ln(\Phi_i + \bar{\Phi}_i)$$

e.g. $-\ln(S+\bar{S}) - 3\ln(T+\bar{T}) - 3\ln(U+\bar{U})$

$\boxed{\textbf{our no-scale family}:\ K = -3\alpha\ln(\Phi+\bar{\Phi})}$

↑ the model I train on

W **(superpotential)** — the **discrete flux choice**; *linear in each quantum*:

$$W = \int_{X_6} G_3 \wedge \Omega \xrightarrow{T^6,\ +Q\text{-flux}} P_1(U) - i\,S\,P_2(U) + i\,T\,P_3(U)$$

F-term = each modulus's SUSY auxiliary field, $F^I \propto K^{I\bar{J}}\,\overline{D_J W}$ — V is the energy stored in the F's:

$$(K,\ W) \xrightarrow{\text{F-term formula}} V:$$

$$\boxed{V = e^K\left(K^{I\bar{J}}D_I W\,\overline{D_J W} - 3|W|^2\right), \quad D_I W = \partial_I W + K_I\,W}$$

One theory, many vacua — a landscape of V's

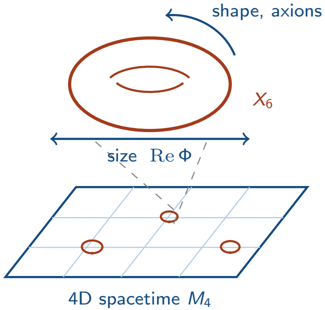

Fig. 3. compactification: a copy of X_6 over every point of M_4.

At every point of 4D spacetime sits a copy of X_6, carrying the moduli Φ^I (T, U, S).

Each discrete flux choice gives a different W — hence a different potential $V(\Phi)$.

Many discrete flux choices $\Rightarrow$ a landscape of V's to grade.

The de Sitter conjecture

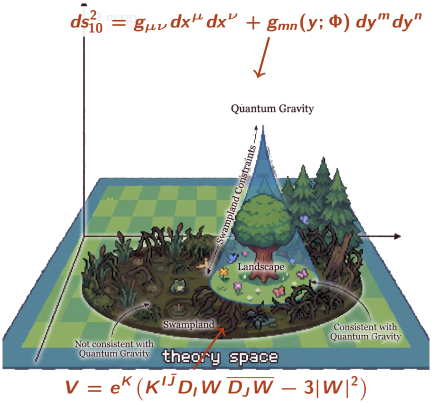

Fig. 4. theory space: quantum-gravity constraints carve the landscape out of the swampland.

- **Swampland**: EFTs with **no quantum-gravity UV completion** (landscape = the complement).
- **de Sitter conjecture** (Obied–Ooguri–Spodyneiko–Vafa 1806.08362): $|\nabla V| \geq \frac{c}{M_P} V, \quad c \sim O(1)$ — forbids a **(meta)stable dS vacuum**.
- controlled string constructions essentially never produce stable dS;
- if it holds, dark energy must **roll** — which is where DESI's $w(z)$ enters (appendix).

The criterion comes with its own ruler

For moduli in 4D $\mathcal{N}=1$ supergravity,

$$\mathcal{L}_{\rm kin} = -K_{I\bar{J}}\,\partial_\mu\Phi^I\partial^\mu\bar{\Phi}^{\bar{J}}, \qquad K = -3\ln(\Phi+\bar{\Phi})$$

$$\Rightarrow\ K_{\Phi\bar{\Phi}} = \partial_\Phi\partial_{\bar{\Phi}}K = \tfrac{3}{(\Phi+\bar{\Phi})^2} = \tfrac{3}{4x^2} \qquad (\Phi = x+iy)$$

$$\Rightarrow\ G_{ab} = 2K_{\Phi\bar{\Phi}}\,\delta_{ab} = \tfrac{3}{2x^2}\,\delta_{ab}$$

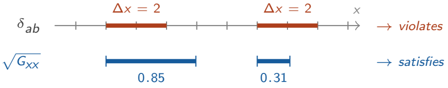

Fig. 5. the same V, graded with the flat ruler δ_{ab} (top) vs. the model's own ruler G_{ab} (bottom).

The coefficient depends on where you stand: **field space is curved** (a hyperbolic plane, $R=-4/3$), and the inverse metric $G^{ab} = \frac{2x^2}{3}\,\delta^{ab}$ *amplifies* derivatives by x.
The refined de Sitter criterion (Ooguri–Palti–Shiu–Vafa 1810.05506) asks, at each point,

$$\frac{|\nabla V|}{V} \ \geq\ c \qquad or \qquad \frac{-\lambda_{\rm min}(\nabla\nabla V)}{V}\ \geq\ c'.$$

The refined version: a $V>0$ point must be steep enough to *roll* **or** tachyonic enough to *decay*; acceleration ($\ddot{a}>0 \iff \dot{\phi}^2 < V$) is a property of a *region*, so the label lives on a window.
Collapse both branches into one pointwise number, and grade the window by its worst point:

$$D(\phi) = \max\!\left(\frac{|\nabla V|_G}{V},\ -\frac{\lambda_{\rm min}(G^{-1}\mathrm{Hess}\,V)}{V}\right), \qquad D_{\rm min} < c \Rightarrow violate.$$

From 10D to the label: one silent choice

$\text{10D on } X_6 \longrightarrow \Phi = x + iy \longrightarrow K,\ W \longrightarrow G_{ab},\ V \longrightarrow D_{\min} \gtrless c \longrightarrow \boxed{\text{label}} \longrightarrow \text{classifier}$

compactification | moduli | two functions | ruler + potential | refined criterion | flat / covariant | supervised ML

the one silent choice:
δ_{ab} or G_{ab}

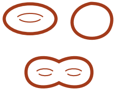

internal space X_6

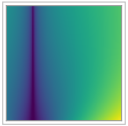

window of V (input image)

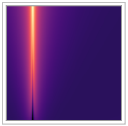

D map — dashed: $D = c$ boundary

violate

satisfies

label = $[\,D_{\min} < c\,]$

CNN $\rightarrow \hat{p}$

classifier

Flat labels: one-class collapse

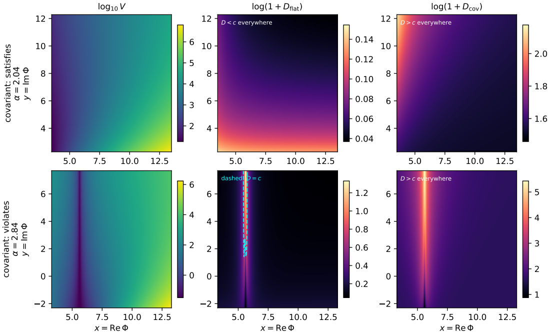

Fig. 6. experiment: render V on a 128^2 grid; compute D pointwise with flat and covariant derivatives; label the window by its worst point, $D_{\min} \lessgtr c$ (dashed contour).

Testbed: single-modulus no-scale model, $K = -3\alpha \ln(\Phi + \bar{\Phi})$, $W = a(\Phi^{n_-} - \Phi^{n_+})$, $n_\pm = \frac{3}{2}(\alpha \pm \sqrt{\alpha})$, and the standard $\mathcal{N} = 1$ F-term potential

$$V = e^K \left(K^{\Phi\bar{\Phi}} \, |D_\Phi W|^2 - 3|W|^2 \right)$$

(the scale a drops out of D). 3 000 windows, $\alpha \in [0.4, 3]$.

violate fraction	
flat labels	**3000 / 3000**
covariant labels	**0.550**

The flat diagnostic decays as $3\alpha/x$ and misses the amplification $\sqrt{G^{yy}} \propto x$ in the tachyon branch. Result: **one-class degeneracy**.

So a *constant* classifier scores 100% accuracy.

Wrong labels do not merely give wrong answers — **they hide the fact that they are wrong.**

The $\sqrt{2\mathcal{N}}$ plateau

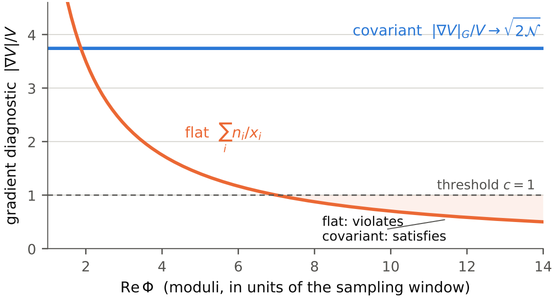

Fig. 7. experiment: the curves are analytic ($\mathcal{N} = 7$); the numerical check measures the median covariant slope on 200 tadpole+Bianchi-consistent integer flux arrays at random $V > 0$ points.

For *any* $K = -\sum_i n_i \ln(\Phi_i + \bar{\Phi}_i)$, write $\ln V = K + \ln(\text{bracket})$. The Kähler norm of the first piece is coordinate-independent:

$$|K_I|_G = \sqrt{2\mathcal{N}}, \qquad \mathcal{N} \equiv K^{I\bar{J}} K_I K_{\bar{J}} = \sum_i n_i.$$

covariant $\rightarrow$ plateau $\sqrt{2\mathcal{N}}$
flat $\rightarrow$ decays as $\sum_i n_i/x_i$

$\mathcal{N} = 3$ *is* the no-scale condition: one geometric invariant fixes the asymptotic verdict.

Check outside this family — 2006.07290 ($\mathcal{N} = 7$), physical integer flux: median slope = **99.7% of $\sqrt{14}$**.

Not a hand-picked-W artifact: **random** superpotentials disagree just as much — 68–71% (appendix).

ML method on this model

- **Input**: a window of V, evaluated analytically on a 128×128 grid; per-image normalization (kills the overall scale).
- **Method**: one plain CNN — 3 conv blocks $\to$ global average pooling $\to$ one logit z (an unbounded raw score), $\sim 5 \times 10^5$ parameters; binary cross-entropy, AdamW, **3 seeds**, frozen splits. Ground truth $y = [\,D_{\min} < c\,]$, computed **twice**: flat vs. covariant derivatives.
- **Output**: $\hat{p}(\mathit{violate}) = \sigma(z) = \frac{1}{1+e^{-z}} \in [0,1]$ — one number per window; the threshold $\frac{1}{2}$ turns it into the verdict, *violate* or *satisfies*.
- **Score**: rank-based AUC ($1 =$ perfect ranking, $0.5 =$ coin flip), always with a c-scan.

For the flux track: no images — the 8 integer quanta go into **lightweight gradient-boosted trees**, evaluated with 5-fold cross-validation (appendix).

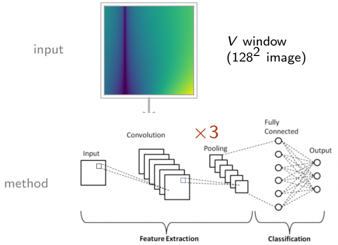

Fig. 8. CNN f_θ: conv blocks $\times 3 \to$ GAP $\to$ one logit z

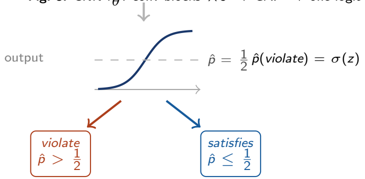

trained & scored against the ground truth $y = [\,D_{\min} < c\,]$ — **flat or covariant**

Same images, different labels

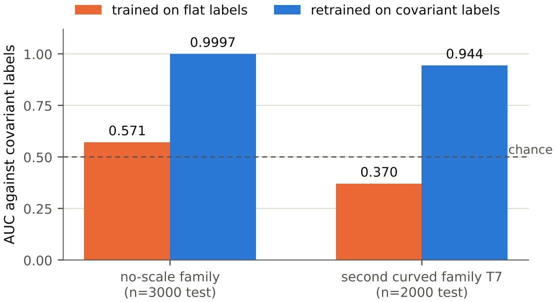

Fig. 9. experiment: two networks, same architecture — trained on flat vs. covariant labels — scored on the *same* test windows against covariant truth (bars = AUC).

Same CNN, same input images, same split. **Only the ground truth changes.**

AUC against covariant labels. T7 = a second engineered family: two complex moduli, $K = -n_1 \ln(\Phi_1+\bar{\Phi}_1) - n_2 \ln(\Phi_2+\bar{\Phi}_2)$, with a Rasulian-type W realizing the same K_I cancellation.

Label-only control (kills the distribution-shift objection): same T7 images, split, seeds — *only* the label set exchanged:

	test flat	test cov
train flat	0.850	0.675
train cov	0.518	0.941

Diagonal learnable, off-diagonal dead: the two criteria are *different functions*, and the failure is **in the labels**.

Not knife-edge in c: across the whole scan $c \in [0.25, 3]$, the flat-trained AUC stays within 0.54–0.60 — never usable.

What the metric cannot change

At a critical point the physical mass matrix is a *congruence* of the flat Hessian, $G^{-1/2}(\nabla\nabla V)G^{-1/2}$ with $G \succ 0$. By **Sylvester's law of inertia** the signature is preserved.

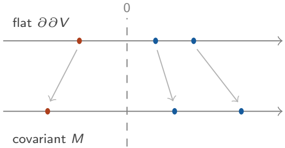

Fig. 10. magnitudes stretch ($\times \sim 1.7$) — **no eigenvalue crosses zero**.

	can the metric change it?	**evidence**
stability *sign* at a critical point	**no — theorem**	0 sign flips in 2 891 vacua
magnitudes ($m^2/\lvert\Lambda\rvert$, $\lvert\nabla V\rvert/V$)	**yes**	median ratio 1.69
verdicts that compare magnitudes (BF, scale sep.)	**yes**	the flux census
truncation to an isotropic sector	yes, but that is *not* the metric	—

This also protects the 2020 ML$\times$refined-dS work (Cabo Bizet *et al.*): verdicts from eigenvalue *signs* at critical points — **not exposed**. What has not been measured is the cost of the wrong metric as a *learning* problem — **that cost is the 0.571**.

Three layers get conflated in the literature: **Level 0** flat partials (coordinate-dependent, not physical) $\to$ **Level 1** covariant Hessian (*the definition*) $\to$ **Level 2** actual perturbation masses (Mizuno *et al.*). My claim lives entirely at $0 \to 1$.

From the toy to a real landscape

Both engineered families share the same K_I cancellation — needed: a landscape **nobody hand-picked** (where $\sqrt{2\mathcal{N}}$ is testable at $\mathcal{N}=7$).

toy family — engineered

$$K = -3\alpha \ln(\Phi + \bar{\Phi})$$

$$W = a\left(\Phi^{n_-} - \Phi^{n_+}\right)$$

one modulus; W chosen *by hand* to realize the cancellation

$$\Longrightarrow$$

real landscape — type IIB on the isotropic T^6, $\mathcal{N}=7$

$$K = -\ln(S+\bar{S}) - 3\ln(T+\bar{T}) - 3\ln(U+\bar{U})$$

$$W = \int F_3 \wedge \Omega \;-\; S\int H_3 \wedge \Omega \;=\; \underbrace{P_1(U)}_{\text{RR flux } f} - S\;\underbrace{P_2(U)}_{\text{NSNS flux } h}$$

$\downarrow$ + non-geometric Q-flux

$$\boxed{W = P_1(U) - i\,S\,P_2(U) + i\,T\,P_3(U)}$$

$P_{1,2,3}$ cubic in U; coefficients = the **14 flux parameters** $(f_0, f_*, f^*, f^0;\ h_0, h_*, h^*, h^0;\ b_1, \ldots, b_6)$; the 8 (f,h) are **integers** — fluxes are quantized on cycles ($\oint F_3, \oint H_3 \in \mathbb{Z}$); the 6 non-geometric b are solved with the SUSY conditions — **nobody picks W by hand**

Constraints (1904.11091), imposed exactly — here H = the NSNS 3-form flux H_3 (quanta h), F_3 = the RR flux (quanta f), and Q = the **non-geometric flux** (quanta b):

D3 tadpole (charge cancellation): $f_0 h^0 + 3 f_* h^* - 3 f^* h_* - f^0 h_0 = 16 = 2N_{O3}$
Jacobi $Q \cdot H = 0$ (4 eqs) · Bianchi $Q \cdot Q = 0$ (3 eqs) · one non-geometric tadpole
all self-tested against a known solution (10^{-15}).

How the landscape is built

Construction. SUSY vacua solve $D_I W = 0$ — *linear* in the flux — so vacua are built by a joint solve over (moduli, non-geometric flux), then the integer (f, h) box $[-3, 3]^8$ is enumerated **exhaustively**.

Why $[-3, 3]$: the largest integer box we could solve *exhaustively* — every consistent array inside is constructed, none sampled; the non-geometric b come out of the joint solve.

Yield. 64 064 consistent arrays → 2 619 vacua (all supersymmetric AdS); pooled, deduplicated: $n = 2\,891$.

Why all AdS? At a SUSY vacuum $D_I W = 0$, so the positive term of the F-term formula drops out:

$$V_* \;=\; e^K\big(\underbrace{K^{I\bar{J}} D_I W\, \overline{D_J W}}_{=0} - 3|W|^2\big) \;=\; -3\, e^K |W_*|^2 \;\leq\; 0$$

negative vacuum energy ⇒ **Anti-de Sitter** (Minkowski only if $W_* = 0$).

Diagnostics. Every vacuum graded twice (flat vs. covariant):

$$M^i{}_j = g^{ik}\big(\partial_k \partial_j V - \Gamma^l_{kj} \partial_l V\big),$$

spectrum via the congruence $g^{-1/2} H g^{-1/2}$; BF verdict $\lambda_0/|V| \geq -3/4$.

Beyond BF: scale sep. flips 7.4% (147/67, both ways), tracking volume ($\mathrm{Re}\, T$ +0.81).

The exhaustive census and the BF bound

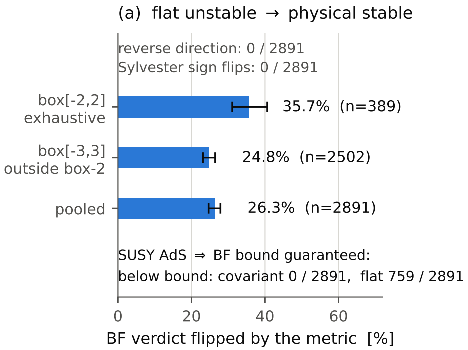

Fig. 12. experiment: solve *every* array in the box via the linear SUSY equations; compute each vacuum's spectrum twice (flat Hessian vs. covariant congruence); compare the BF verdicts.

Type IIB isotropic T^6, $\mathcal{N}=7$, tadpole+Bianchi-consistent integer flux; vacua *constructed* via SUSY $D_I W = 0$ (linear in flux; appendix). Box$[-3,3]$ **exhaustive**: 64 064 arrays $\to$ pooled $n = 2\,891$.

- BF verdict flips **26.3%**, **reverse** 0
- Sylvester sign flips $0/2\,891$
- scale separation, volume correlations: appendix

An independent check. SUSY AdS $\Rightarrow$ BF bound is a *theorem*: **covariant** $0/2\,891$ **below** (min exactly at the bound), **flat** $759/2\,891$ **below** (min -13.1).

So flat is not *a different convention* — it **contradicts a theorem**, on a metric-independent test. And exhaustive means there is no counterexample to find.

Summary

- Flat labels collapse a curved family to **one class** — and every metric still reports success.
- Relabeled covariantly, the same architecture goes 0.571 $\to$ 0.9997.
- The mechanism is a derivation: $|\nabla V|_G / V \to \sqrt{2\mathcal{N}}$.
- On an exhaustive flux landscape, a SUSY theorem rules the flat verdicts **wrong** (0 vs 759).

A wrong field-space metric does not degrade a swampland classifier — it dissolves the problem while every metric still reports success.

Q&A

A wrong field-space metric does not degrade a swampland classifier — it **dissolves the problem while every metric still reports success**.

Appendix — contents

The label error in observational units

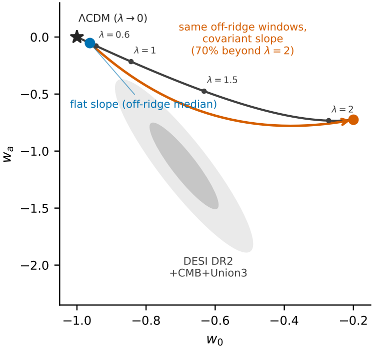

Fig. 13. the gradient-only slope λ mapped to (w_0, w_a); the DESI contour sets the scale.

For an exponential-attractor slope $\lambda = |\nabla V|_G / V$,

$$w \simeq \frac{\lambda^2}{3} - 1, \qquad \ddot{a} > 0 \iff \lambda < \sqrt{2}.$$

Gradient-only, both metrics — D has a Hessian branch and cannot be read as a slope.

- ridge windows (65%): gradient $= 0$, so $\lambda = 0$ in *both* metrics — **no discrepancy there**
- the other **35%**: flat median $\lambda\,0.49 \to$ covariant 2.8; **70% land beyond $\lambda = 2$** — too steep to accelerate in the constant-λ proxy
- same-window $|\Delta w_0| = 0.74 = 7.4\times$ the DESI w_0 resolution; second family reproduces it (0.72)

Not a claim about cosmology. An illustrative scale estimate — the true multi-field $w(z)$ needs the full trajectory, not one local slope.

Why machine learning?

Every alternative stalls at landscape scale:

- **exhaustive enumeration**: the 64 064-array box = days of CPU; non-isotropic $\sim 10^{12}$ candidates
- **computer algebra**: 155/162 surviving dS candidates unchecked (2410.22444)
- **provable barrier**: flux vacua with prescribed Λ are **NP-complete** (Denef–Douglas)

And ML measurably works: quanta-only ranking recovers **61% of vacua from the top 10%** — and it holds outside the training box.

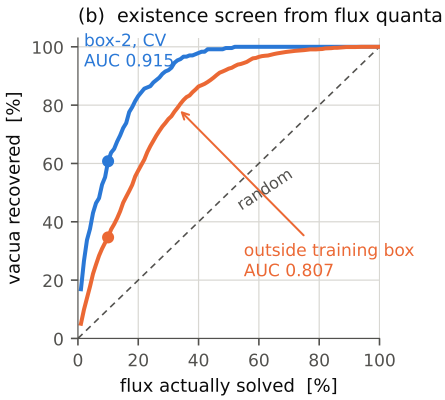

Fig. 11. the existence screen: predict from the integer quanta which flux arrays admit a vacuum, and solve only the top-ranked slice. Classifier: lightweight gradient-boosted trees, 5-fold cross-validation (appendix).

How the ML is implemented, experiment by experiment

2D no-scale track (images → CNN).

- **Data**: 3 000 windows; V evaluated analytically on a 128×128 grid; per-image normalization (kills the overall scale).
- **Labels**: the *same* windows labeled twice — $D_{\min} < c$ with flat vs. covariant derivatives; everything else frozen.
- **Model**: 3 conv blocks (7×7 filters, BatchNorm, ReLU, $2\times$ pool; channels $32 \to 64 \to 128$) $\to$ global average pooling $\to$ dropout $\to$ one logit; $\sim 5 \times 10^5$ parameters.
- **Training**: binary cross-entropy, AdamW (10^{-3}), early stopping on validation AUC; **3 seeds**; single GPU.
- **Evaluation**: rank-based AUC + c-scan; blocked-α split for extrapolation; 2×2 label-only control.

Flux track (integers → boosted trees).

- **Inputs**: the 8 RR/NSNS integer quanta (f, h) — tabular, no images. Gradient-boosted trees (sklearn), 5-fold CV.
- **Existence screen**: positive = the SUSY solve yields a vacuum; negatives are *well defined* because the box was solved **exhaustively**.
- **BF-flip pilot**: metric sensitivity from quanta alone (AUC 0.79) vs. with the solved vacuum (0.99).
- **Feature importance** independently recovers $\mathrm{Re}\, T$ — the $g^{-1} \sim x^2$ volume mechanism.

Everything else is controls: the *same* CNN class is re-used for every retraining (covariant, second family); seeds and splits frozen across label swaps.

Reproducibility: every run seeded (20260526 / 777 / 20260723); train on one RTX 3080, eval on CPU; figures recomputed from stored per-sample data. Script↔result map: theory note §5.

No architecture search, no tuning contest — one fixed pipeline, so that the only moving part is the physics in the labels.

Robustness in the threshold c

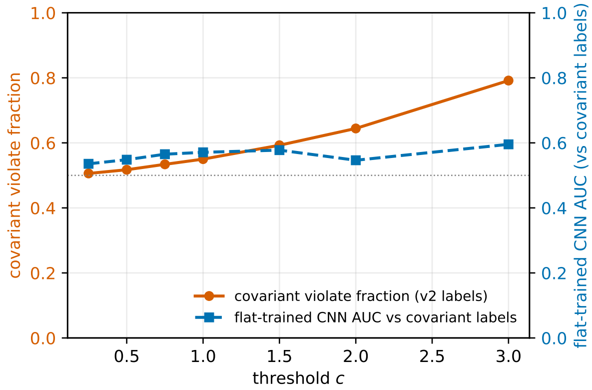

experiment: recompute labels at each c from the stored $D_{\min}$; re-rank the *same* flat-trained network's scores against each label set.

c is microscopically undetermined, so every headline carries a scan.

Over $c \in [0.25, 3]$:

- covariant violate fraction 0.506–0.792 — **two-class throughout**
- flat-trained AUC 0.54–0.60 — **never usable**

What I do not claim: that 0.550 is a physical number. The kinetic normalization is anchored to Rasulian *et al.*'s convention (numerical identity to 5.5×10^{-16}), but an $O(1)$ convention dependence remains.

The physical statement stops at: *two-class, and robust in c*.

the two curves are different quantities — both lie in [0, 1], so they share one scale (color-matched axes).

No conclusion here sits on a knife edge in c.

Why $D_I W = 0$ is the SUSY vacuum condition

In 4D $\mathcal{N} = 1$ supergravity, the SUSY variation of the modulus fermion χ^I in a vacuum is

$$\delta\chi^I \;\propto\; F^I \epsilon, \qquad F^I \;\propto\; e^{K/2} K^{I\bar{J}} \overline{D_J W},$$

with ϵ the SUSY parameter, F^I the auxiliary F-term, and $D_I W = \partial_I W + K_I W$ the Kähler-covariant derivative.

Unbroken SUSY $\Rightarrow$ the vacuum is invariant under the variation:

$$\delta\chi^I = 0 \;\Rightarrow\; F^I = 0 \;\Rightarrow\; \boxed{D_I W = 0}$$

(the last step because $K^{I\bar{J}}$ is invertible).

Warning-light picture. F^I is the "SUSY is broken" light:

- $D_I W \neq 0 \;\Rightarrow\; F^I \neq 0$ — light ON, SUSY broken
- $D_I W = 0 \;\Rightarrow\; F^I = 0$ — SUSY preserved

With a gauge sector the D-terms must vanish too, $D^a = 0$; our flux setup uses the F-term condition $D_I W = 0$.

$D_I W$ determines the chiral F-terms; unbroken supersymmetry requires the fermion variations — hence all F-terms — to vanish.

Boosted trees and five-fold cross-validation

Gradient-boosted trees. Many small decision trees, added in sequence — each new tree focuses on the samples the previous ones got wrong:

$$\mathrm{Tree}_1 + \mathrm{Tree}_2 + \cdots \;\rightarrow\; \text{one strong classifier.}$$

One tree asks simple questions about the quanta $(f_1, \ldots, f_4,\ h_1, \ldots, h_4)$:

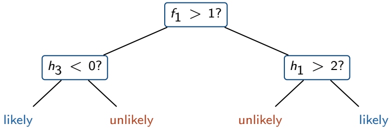

"vacuum likely / unlikely" — per flux array

Five-fold cross-validation. Split the data into five parts. Train on four, test on the held-out one, rotate — so *every* sample is tested exactly once:

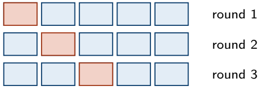

$\cdots$ $(\times 5)$; test / train

So "box-2, CV AUC 0.915" means: ranked by the quanta alone, vacuum existence is predicted at AUC 0.915, measured fold by fold.

The flux track therefore uses a lightweight boosted-tree classifier, scored by five-fold cross-validation.

Why ten dimensions?

- The string worldsheet is a 2D quantum theory, and its Weyl (scale) symmetry is a **gauge** symmetry — it must survive quantization.
- Each ingredient contributes a conformal anomaly (central charge); consistency requires the total to **cancel**:

$$\underbrace{D \cdot \left(1 + \tfrac{1}{2}\right)}_{X^\mu \;+\; \psi^\mu} + \underbrace{(-26 + 11)}_{\text{ghosts}} = \tfrac{3}{2} D - 15 \stackrel{!}{=} 0 \quad \Longrightarrow \quad \boxed{D = 10}.$$

- Bosonic string (no worldsheet fermions): $D - 26 = 0 \Rightarrow D = 26$ — but tachyonic, and no spacetime fermions.
- Same logic as gauge-anomaly cancellation: $D = 10$ is not an assumption but the **only** dimension where the quantum theory is consistent.

Ten is not a choice — it is forced by quantum consistency: the worldsheet conformal anomaly cancels only in $D = 10$.