Anomaly Detection: Mapping the IIB Lamppost with Reinforcement Learning

Jim Halverson
Northeastern University

Based on:
180x.xxxxx with **Long**, Ruehle, Tian

Using geometries and results from:
1706.02299 with **Long** and Sung
1710.09374 with **Long** and Sung

Using RL techniques developed in:
180x.xxxxx with Nelson, Ruehle

Using supervised machine learning and intelligible AI techniques as in:
1707.00655 with Carifio, Krioukov, Nelson

This is work in progress.

But hopefully you like some of the **preliminary results.**

Talk's Two Big Q's

Q: How do we find and understand rare phenomena in the landscape? ***anomaly detection***.

Q: How do we find and understand the specific anomaly of **lampposts** / **safe havens**?

Landscape and Lampposts

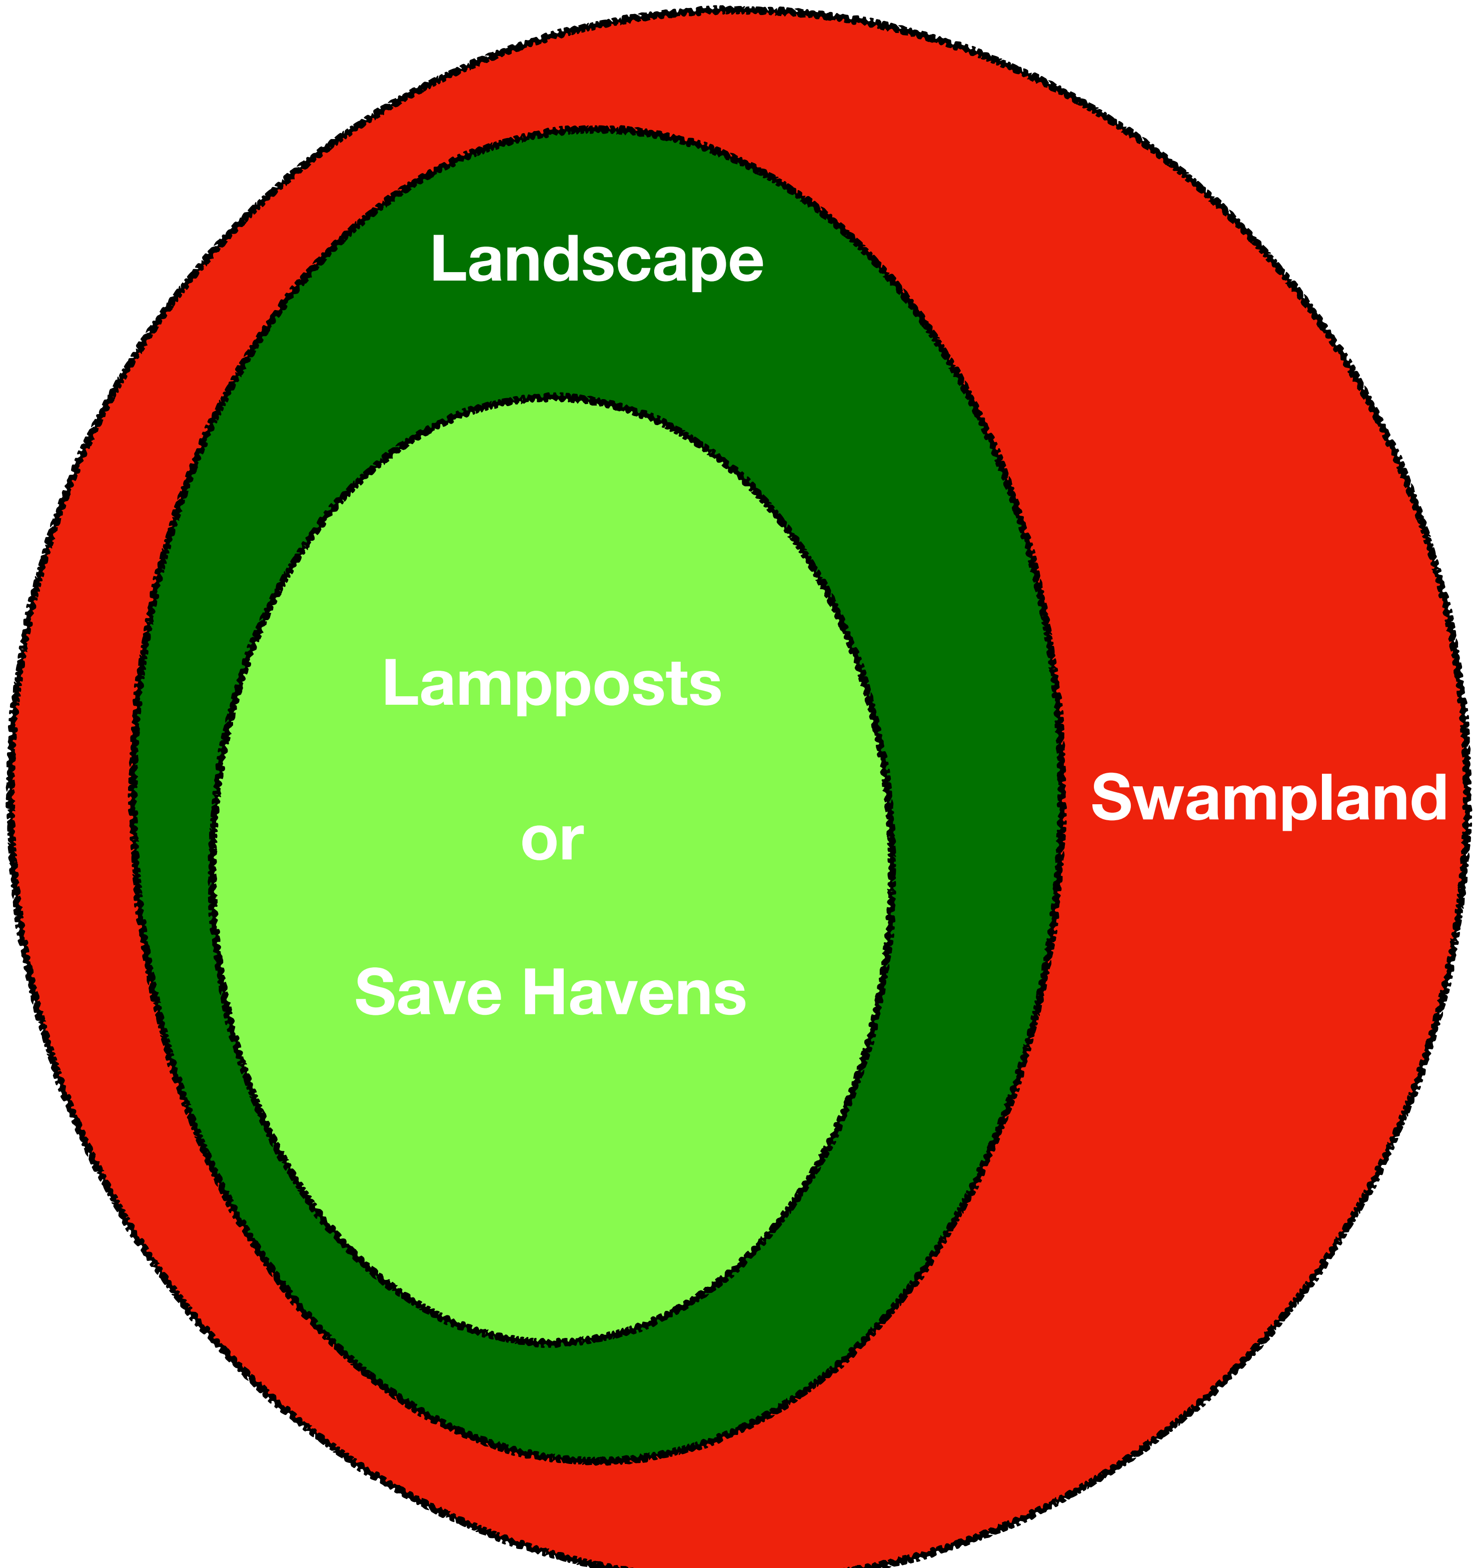

- Swampland:
  good for falsification.

- Landscape:
  good for verification?

- Well-controlled subregions:

  simultaneously

  lampposts, bad for breadth?

  safe havens,
  good for calculation?

  clearly: pros and cons.

Lamppost Questions

Q: size of a given lamppost?

Q: what is its boundary?

Q: how broadly do its **physical conclusions** apply in the landscape?

The IIB Lamppost

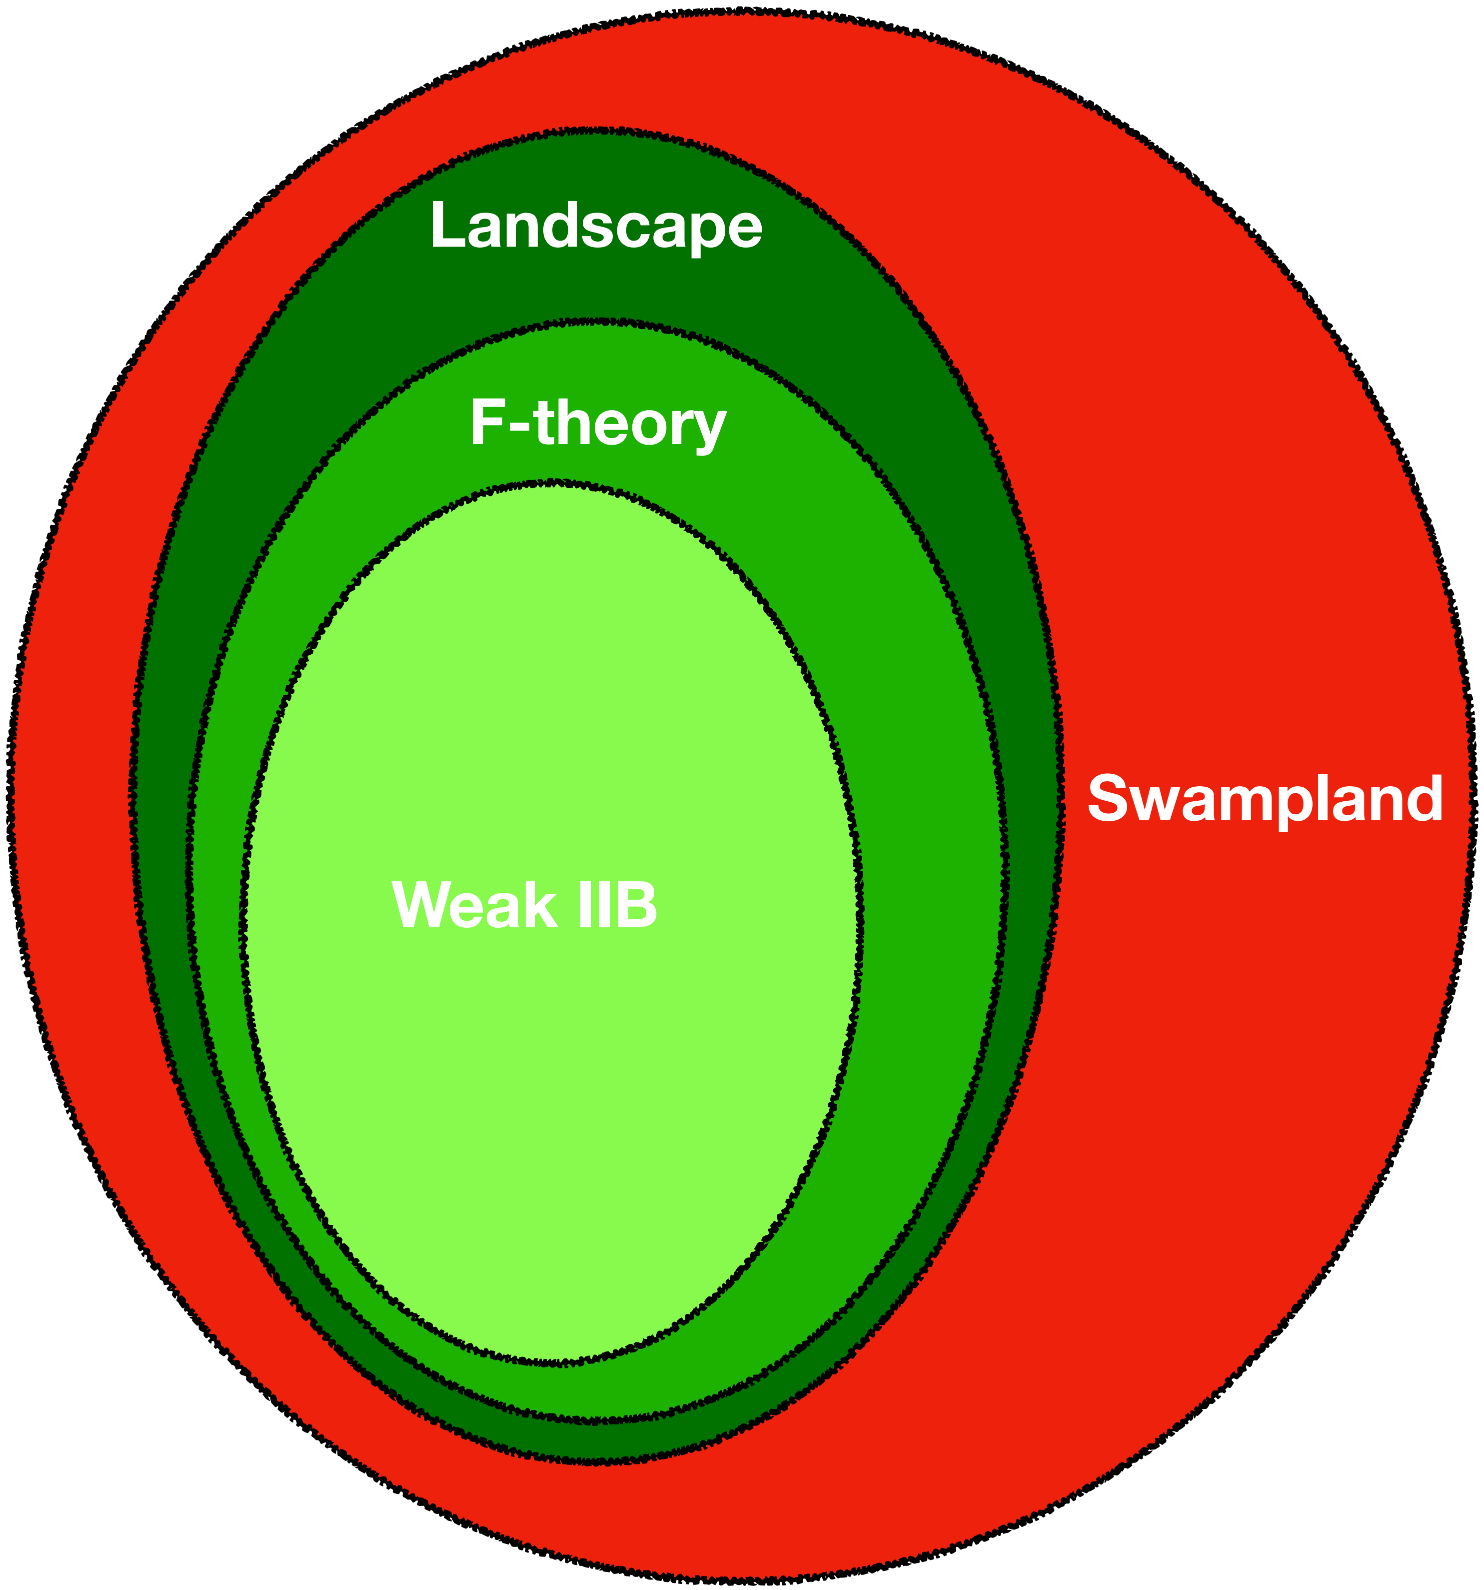

- Weakly Coupled IIB:

  most discussed as "lamppost"?

  exquisite control
  e.g. moduli stabilization.

- Broader context:

  sits inside intrinsically strongly coupled F-theory.

  miniscule subset.

Weak IIB Lamppost, Concretely

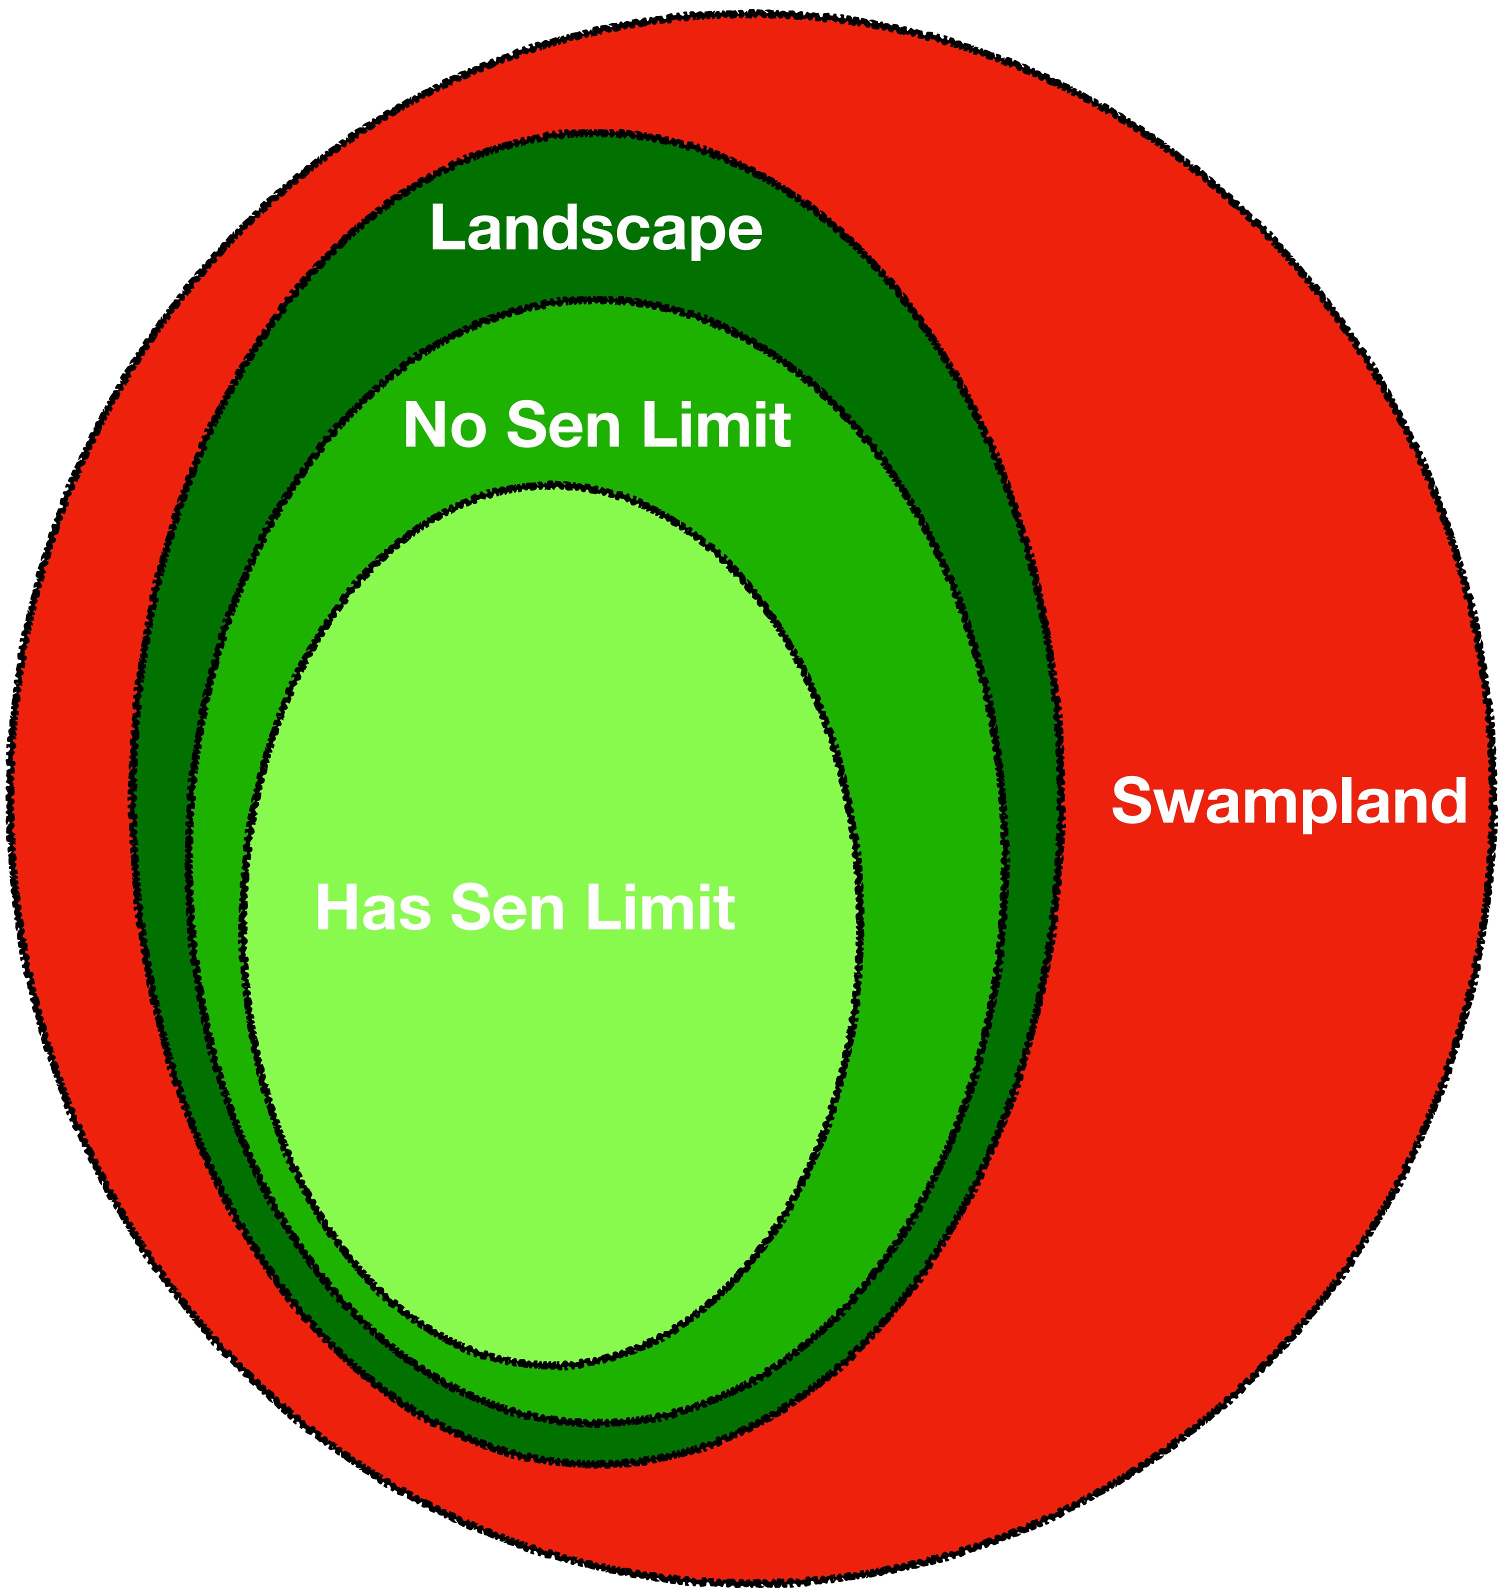

- Last 5 years:

  For fixed compact dimensions B, generic points in moduli space of 7-brane backgrounds are **strongly coupled**.

  e.g. complementary works

  [Taylor, Wang] [J.H., Long, Sung]

- Our Lamppost:

  Such geometries may admit **weak coupling limits** (a Sen limit).

  Q: is the existence of such limits rare?

Q: Can data science lead to stronger results about this concrete lamppost than humans have obtained?

Some humans and results

First theorem: due to

pretty good, but didn't make paper.

Second theorem: due to Cody

very good, this is in the paper.

[JH, Long, Sung] $\frac{N_{\text{Sen}}}{N_{\text{Total}}} \leq 3.0 \times 10^{-391}$

Third (strongest?) theorem due to

Goal: not just probability, but ***explore*** WC lamppost and probe its boundary.

[JH, Long, Ruehle, Tian]

Why not straight to boundary detection?

- Just a binary classification problem, in vs. out.

- **Caveat 1, dumb accuracy.**
  n in, m out, m >> n,
  then “always out” —> accuracy 1-n/m ~ 1.

- **Caveat 1, sample generation.**
  p(in) = 1/k, k>>1, then N in’s requires N*k samples.

  previous result: p(in) < 3×10^{-391}

  need efficient “in” generation! how?

- Previous paper: **[Carifio, Halverson, Krioukov, Nelson]**
  p(E6 on special divisor) ~ 1/1000, reasonable to just gen.

Outline

- Necessary F-theory / IIB / tree dataset review.
- 1) **Explore** lamppost, probe boundary.
  use: ***deep reinforcement learning***.
- 2) **Predict** the boundary.
  use: ***supervised machine learning.***
- 3) **Understand** the boundary.
  use: ***intelligible AI***

Critical Review

- *F-theory*
- *Weak IIB: what do we mean?*
- *Tree dataset: what is our data?*

F-theory & Bases

- **What:** IIb with gen. 7-branes, varying axiodilaton, strong coupling.

- **Math Description:** a Calabi-Yau ellip. fib. over base B, where B is the internal space. Seven-brane on discrim. Use Weierstrass form:

$$y^2 = x^3 + fx + g \qquad\qquad \Delta = 4f^3 + 27g^2 = 0$$

- **Basic data determined by B:**

$$f \in H^0(\mathcal{O}(-4K_B)), \;\; g \in H^0(\mathcal{O}(-6K_B)),$$

- **Non-Higgsable Clusters:** generic f,g for generic B give rise to networks of G-carrying int. 7-branes, can't be CS-Higgsed.

Some selective progress: Anderson, Braun, del Zotto, Halverson, Heckman, Grassi, Morrison, Schafer-Nameki, Shaneson, Taylor, Vafa, Wang

Weak IIB: What?

- **Sen Limit:** [Sen]

$$f = -3h^2 + \epsilon\eta$$

$$g = -2h^3 + \epsilon h\eta - \frac{\epsilon^2\chi}{12}$$

$$J = -\frac{64\left(3h^2 - \eta\epsilon\right)^3}{\epsilon^2\left(144h^3\chi - 144\eta^2h^2 - 72\eta h\chi\epsilon + 3\chi^2\epsilon^2 + 64\eta^3\epsilon\right)}$$

- **Simple Necessary Check:** given B,

  Non-Higgsable F4, E6, E7, E8 —> No Sen Limit due to O(1) gs

  Entire CS moduli space of ell. fib over B is F-theoretic

- **Later:** necessary and sufficient condition, given B.

Tree Language Review

- We'll refer to a sequence of blowups as a **"tree"**
- exceptional divisor from the sequence is a **"leaf"**
- Trees over edges = **"edge trees"**
- Trees over faces = **"face trees"**
- Points on polytope = leaves on ground = **"roots"**
- **Classify all trees** with h ≤ 6 for all leaves.
- Do so by exhaustively constructing the toric blowups.

Tree Dataset

[J.H., Long, Sung]

- Each “tree” is data representing a local sequence of blowups.
- **Form “forest”** (threefold base B) from trees by systematically adding trees to FRST of a 3d reflexive polytope. Face trees first, then edge.
- **Count:** polytopes whose FRST’s have the largest number of faces and edges dominate the ensemble.
- **Two polytopes dominate:** have 108 edges and 72 faces, very large facet.

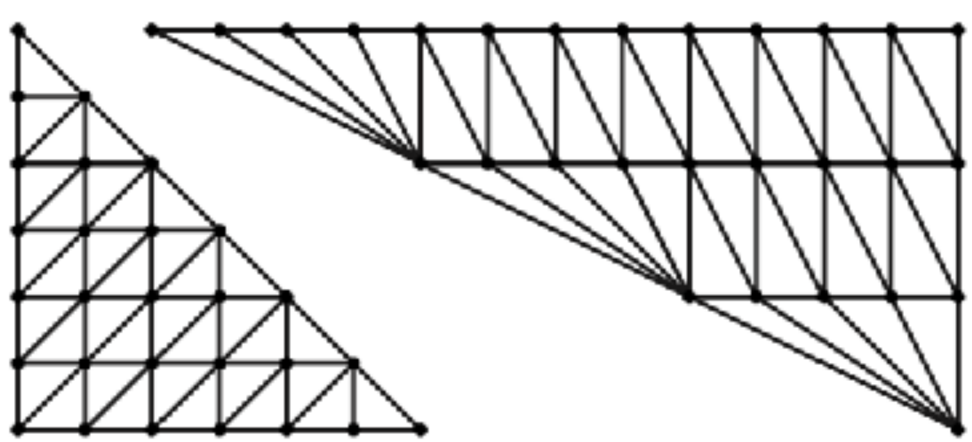

$$|S_{\Delta_1^\circ}| = \frac{2.96}{3} \times 10^{755} \qquad |S_{\Delta_2^\circ}| = 2.96 \times 10^{755}$$

large face of Δ_1° **large face of** Δ_2°

Explore Lampost, Probe Boundary

reinforcement learning,
have robots intelligently explore the space.

- *RL central ideas*
- *RL's most famous (?) result*
- *IIB Lamppost RL Setup*
- *IIB Lamppost RL Results*

Reinforcement Learning

supervised ML **predicts**, RL (AI) **explores / searches**
most famous examples: (?) AlphaGo & AlphaGo Zero

in strings: see [J.H., Nelson, Ruehle] to appear soon.

- an ***agent*** interacts in an ***environment***.
- it perceives a ***state*** from ***state space***.
- its ***policy*** picks and executes an action, given the state.
- agent arrives in new state, receives a ***reward.***
- successive rewards accumulate into ***return.***
- return may penalize future rewards via ***discount factor***.
- policy optimized to maximize reward, i.e. **agent learns how to act!**

Famous Example: AlphaGo Zero

"Mastering the game of Go without human knowledge."

Silver et al. (Google DeepMind), Nature Oct. 2017.

A long-standing goal of artificial intelligence is an algorithm that learns, tabula rasa, superhuman proficiency in challenging domains. Recently, AlphaGo became the first program to defeat a world champion in the game of Go. The tree search in AlphaGo evaluated positions and selected moves using deep neural networks. These neural networks were trained by supervised learning from human expert moves, and by reinforcement learning from self-play. **Here we introduce an algorithm based solely on reinforcement learning, without human data, guidance or domain knowledge beyond game rules.** AlphaGo becomes its own teacher: a neural network is trained to predict AlphaGo's own move selections and also the winner of AlphaGo's games. This neural network improves the strength of the tree search, resulting in higher quality move selection and stronger self-play in the next iteration. **Starting tabula rasa, our new program AlphaGo Zero achieved superhuman performance, winning 100–0 against the previously published, champion-defeating AlphaGo.**

Fact: Go has 10^{172} states, a "big" number, but for the task of playing excellently, superhuman progress achieved tabula rasa.

AlphaGo Zero: The Money Plot

Silver et al, Nature 2017.

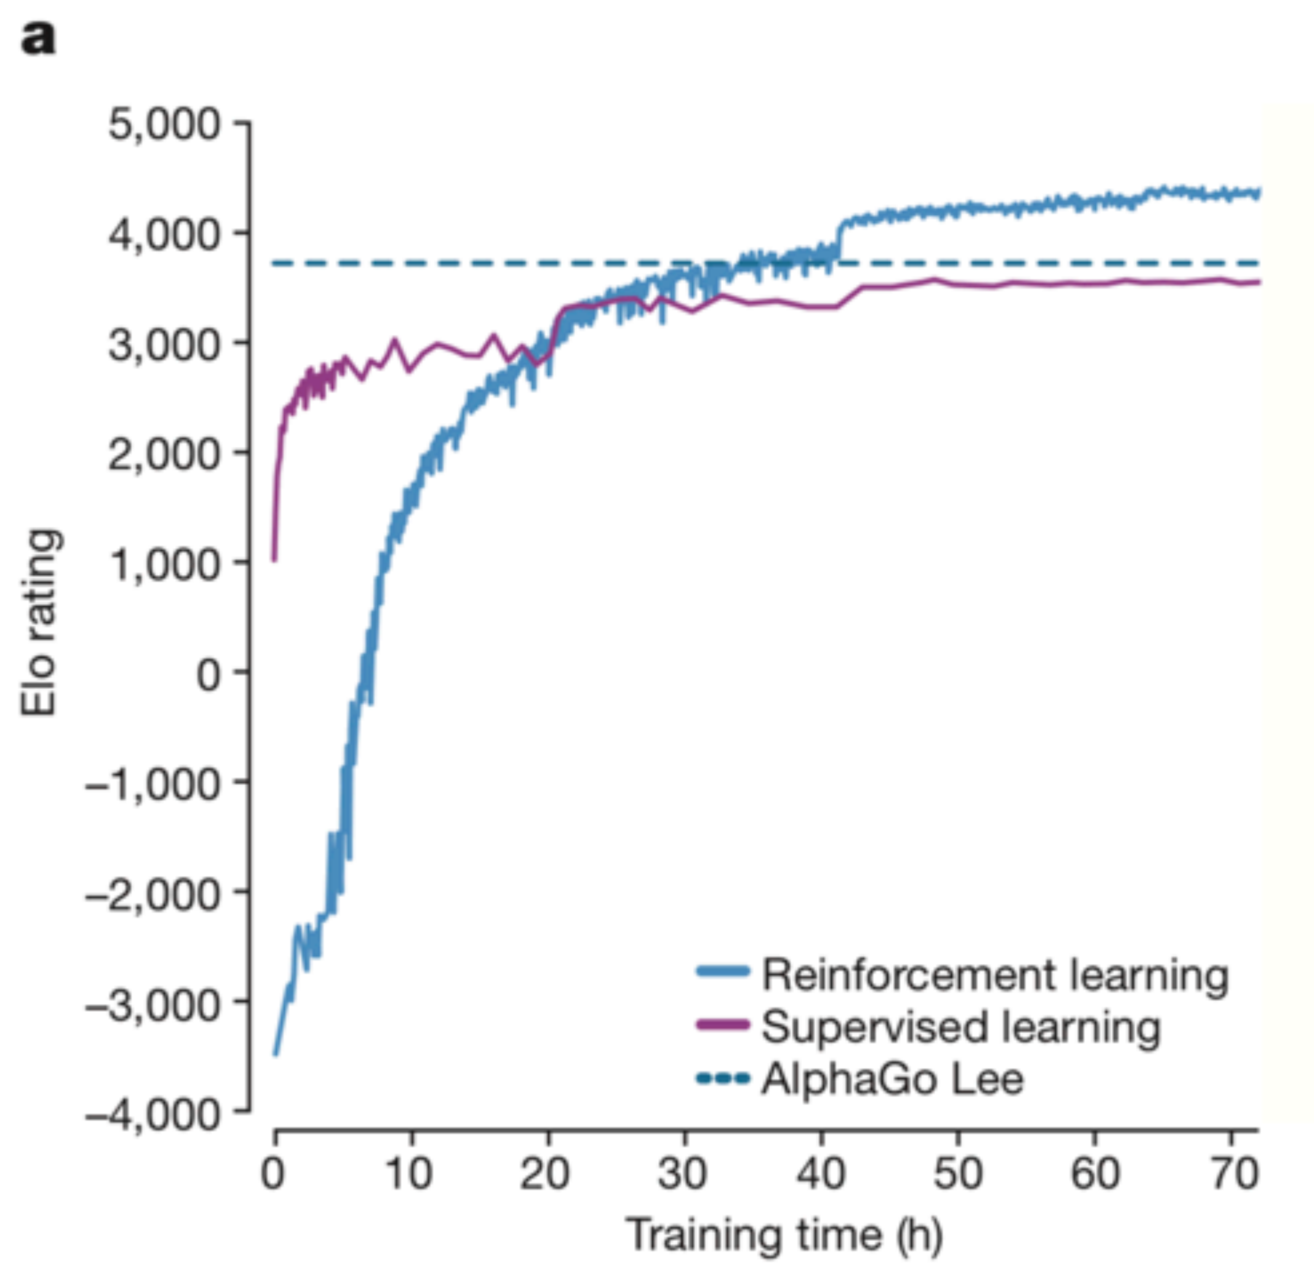

here:

supervised learning = training on human expert games.

AlphaGo Lee = previous version from 2016 that beat world Champ Lee Sedol.

stronger than AlphaGo Lee in **under 48 hrs**, beat **100-0**.

AlphaZero for Chess

similar architecture, arXiv preprint. **Silver, Hubert, Schrittwieser et al**

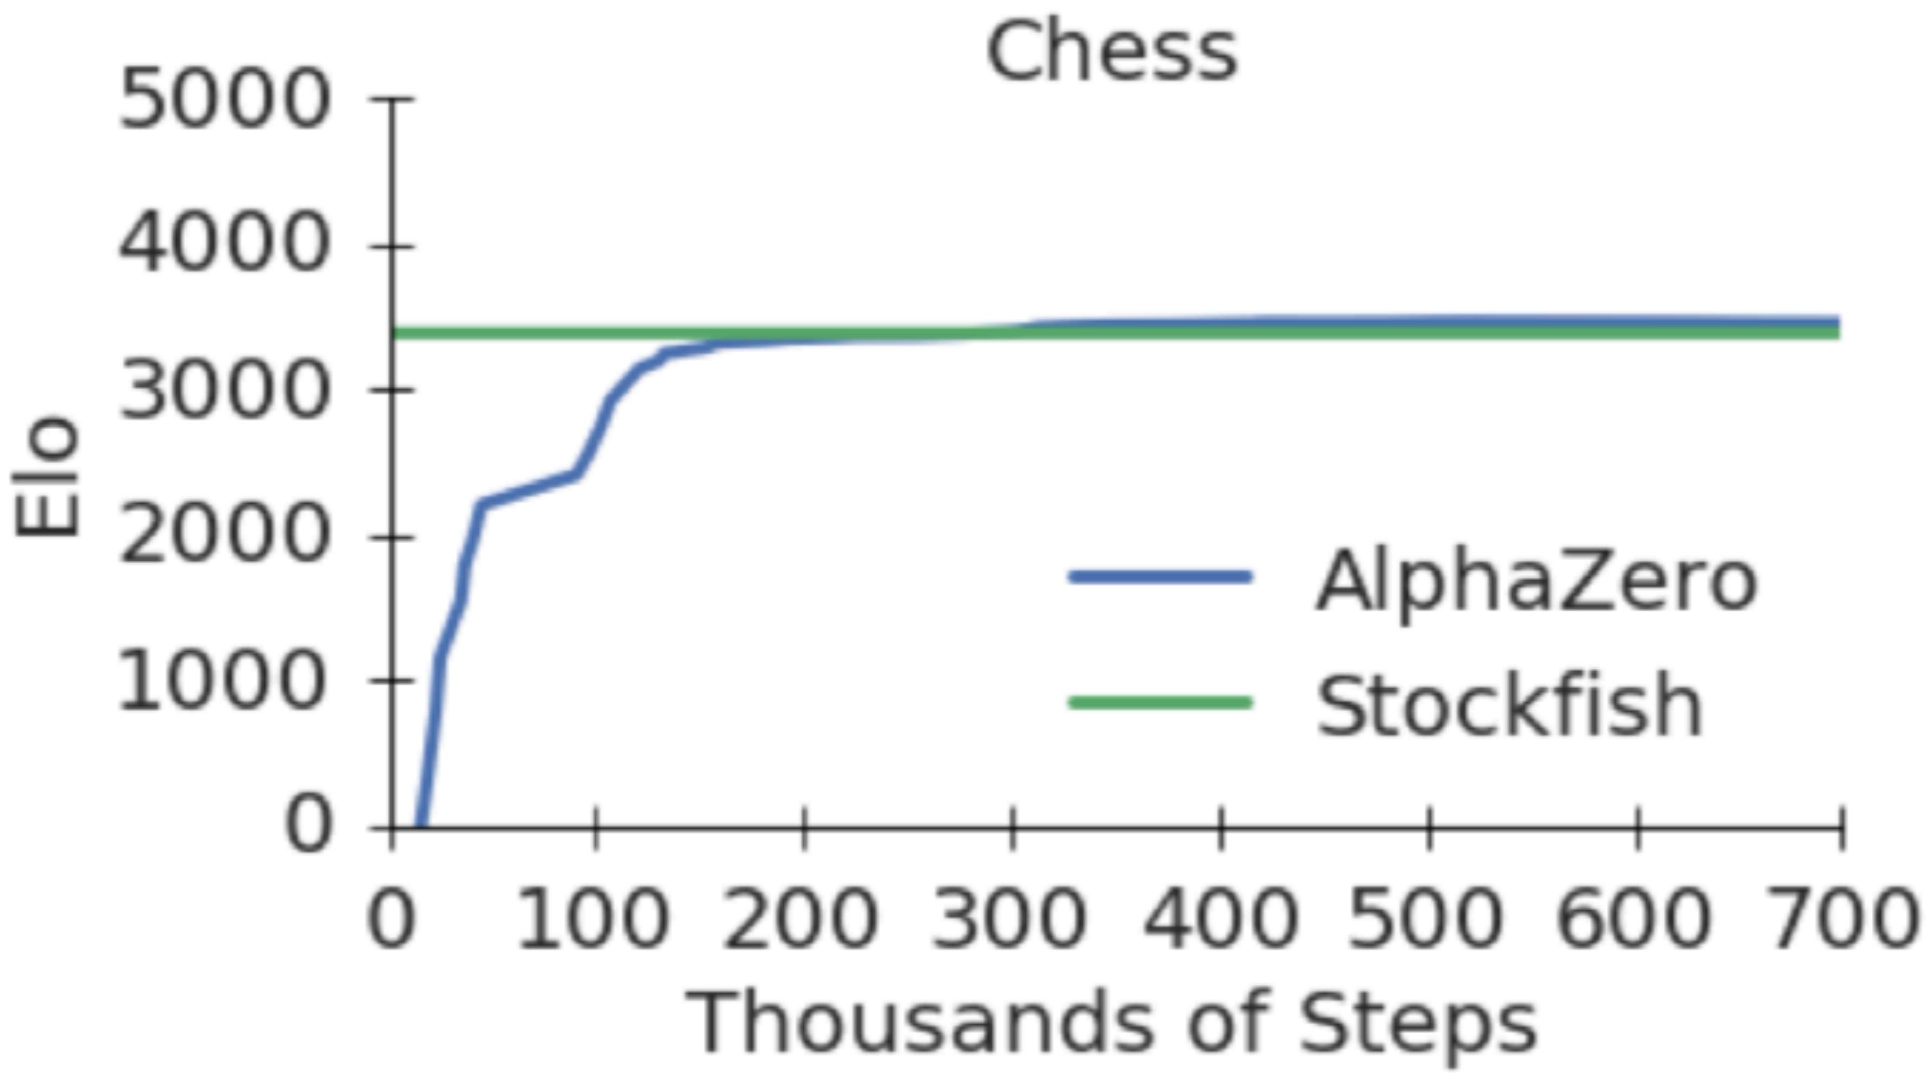

stronger than Stockfish in **under 4 hrs**, beat thoroughly.

RL to Explore the Lamppost

why might this work? Sen-possible geometries **connected subset** of our ensemble, can start at FRST of 3d refl. poly.

No Sen limit!

The Game:

move: place tree
goal: stay in bounds
reward: 100 if in bounds
game over: out of bounds

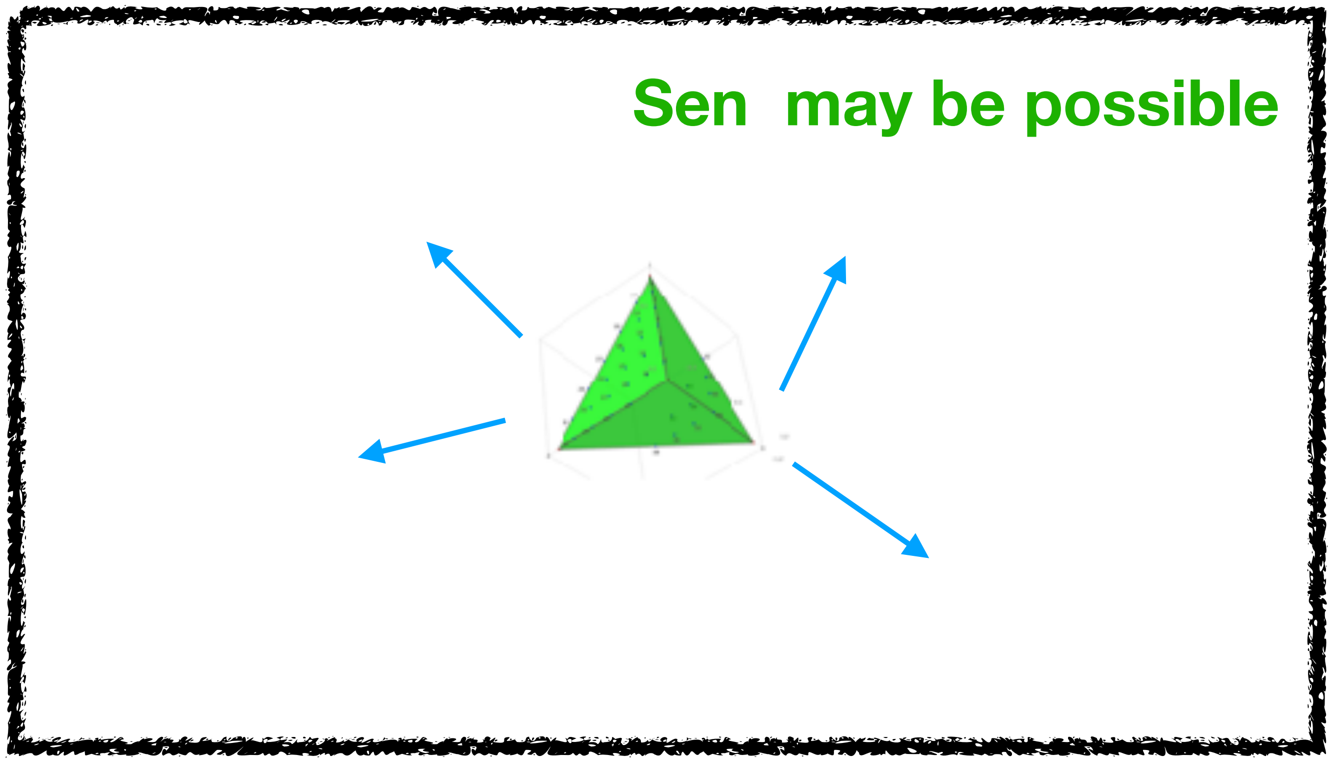

(presented in time series of results, to emphasize fun)

Implementation

model-free RL: want algorithms to work well regardless of environ.

means we can use CS-implemented algs!

three modules:

- Open AI (Musk) defines what an environ is and how to interface.
- ChainerRL provides RL algorithms and NN architecture.
- **Physicists provide:** the environment. two envs so far. **~50 new lines?**

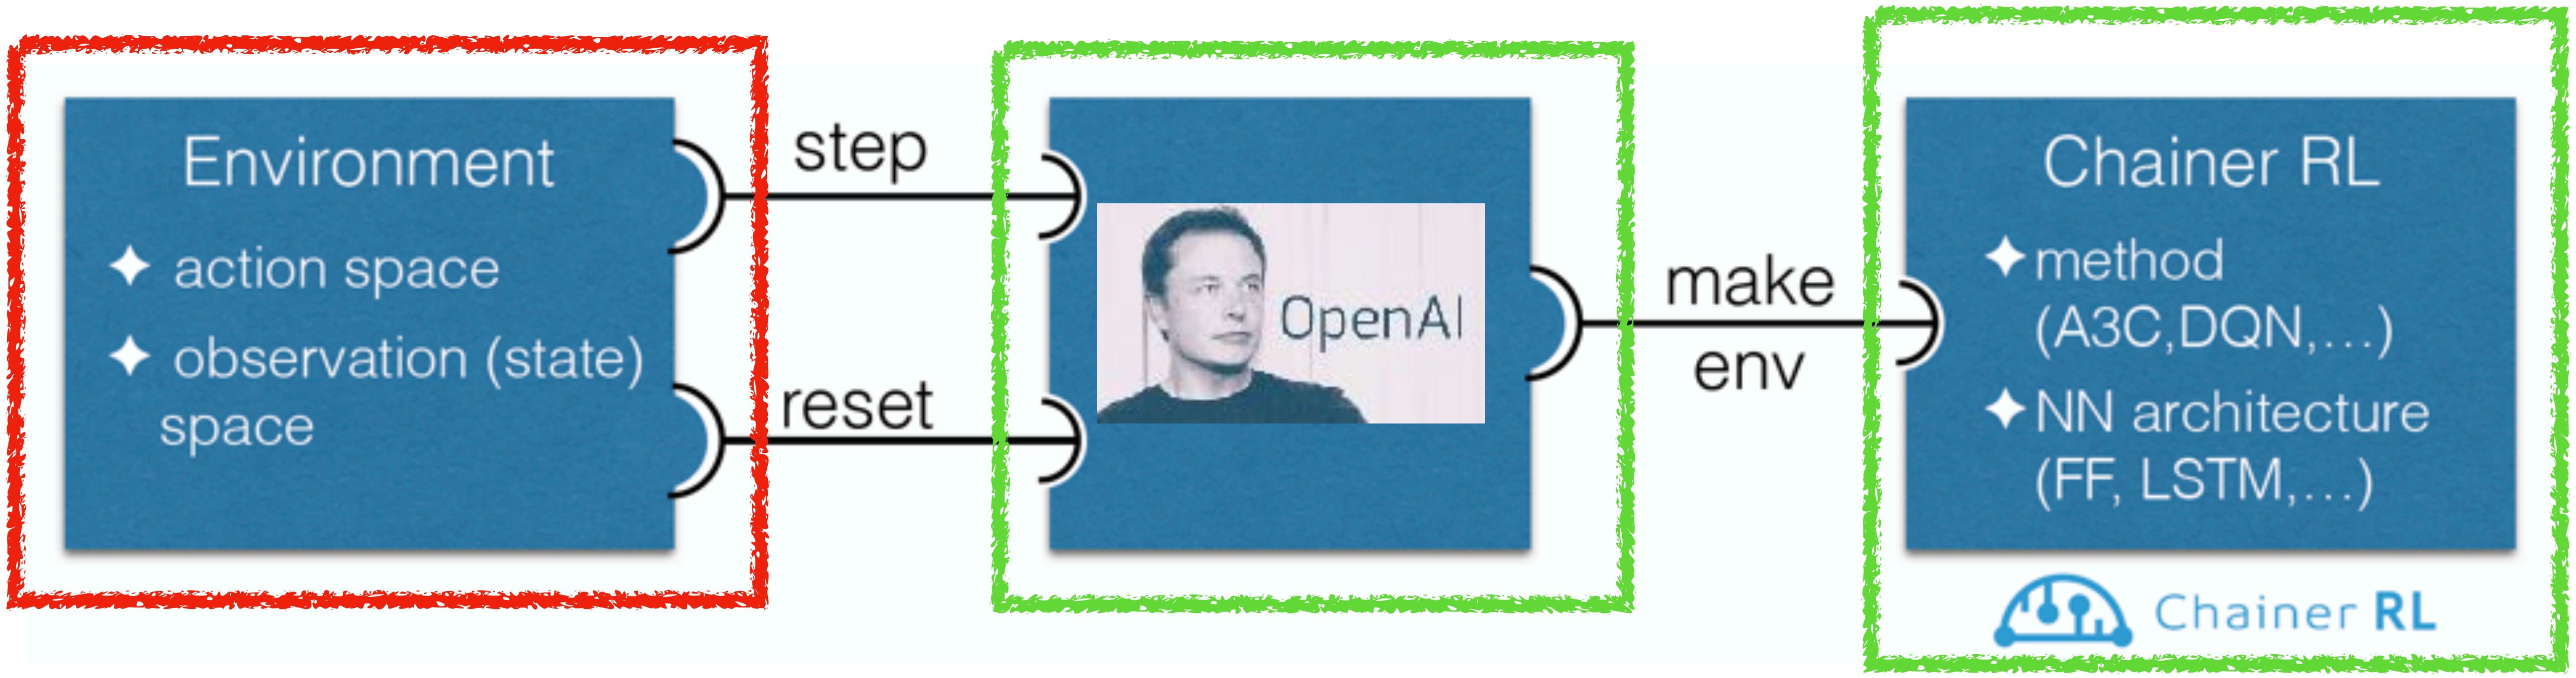

algorithm: asynchronous advantage actor-critic (A3C) [Minh et al 2016] (parallel CPU, not GPU)

Asynchronous Advantage Actor-Critic (A3C)

[Mnih et al, DeepMind 2016]

*"Our parallel reinforcement learning paradigm also offers practical benefits. Whereas previous approaches to deep reinforcement learning rely heavily on specialized hardware such as GPUs (Mnih et al., 2015; Van Hasselt et al., 2015; Schaul et al., 2015) or massively distributed architectures (Nair et al., 2015), **our experiments run on a single machine with a standard multi-core CPU**. When applied to a variety of Atari 2600 domains, on many games asynchronous reinforcement learning achieves better results, in **far less time than previous GPU-based algorithms**, using far less resource than massively distributed approached" - Mnih et al,* Asynchronous Methods for Deep RL

- **Actor-Critic Methods:** NN for determining both policy (actor) and value (critic).
- **Asynchronous:** many worker bees explore, report back to king (critic) and queen (actor) bee.

  i.e. use **communal knowledge.**
- > some 2016 GPU algs. Simple to run. Learns strat.

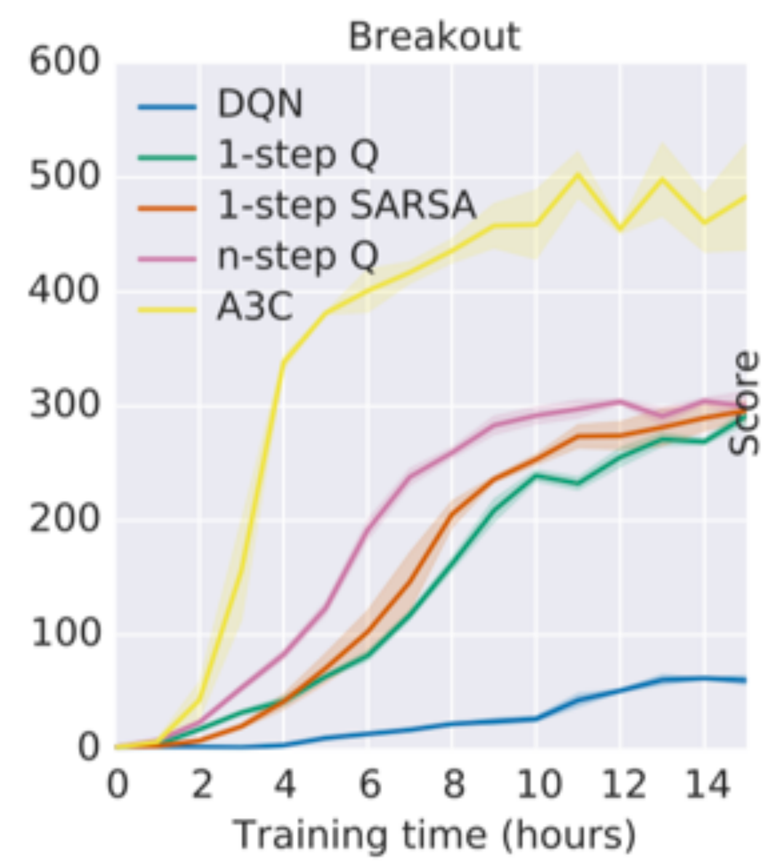

For Comparison: Random Walk

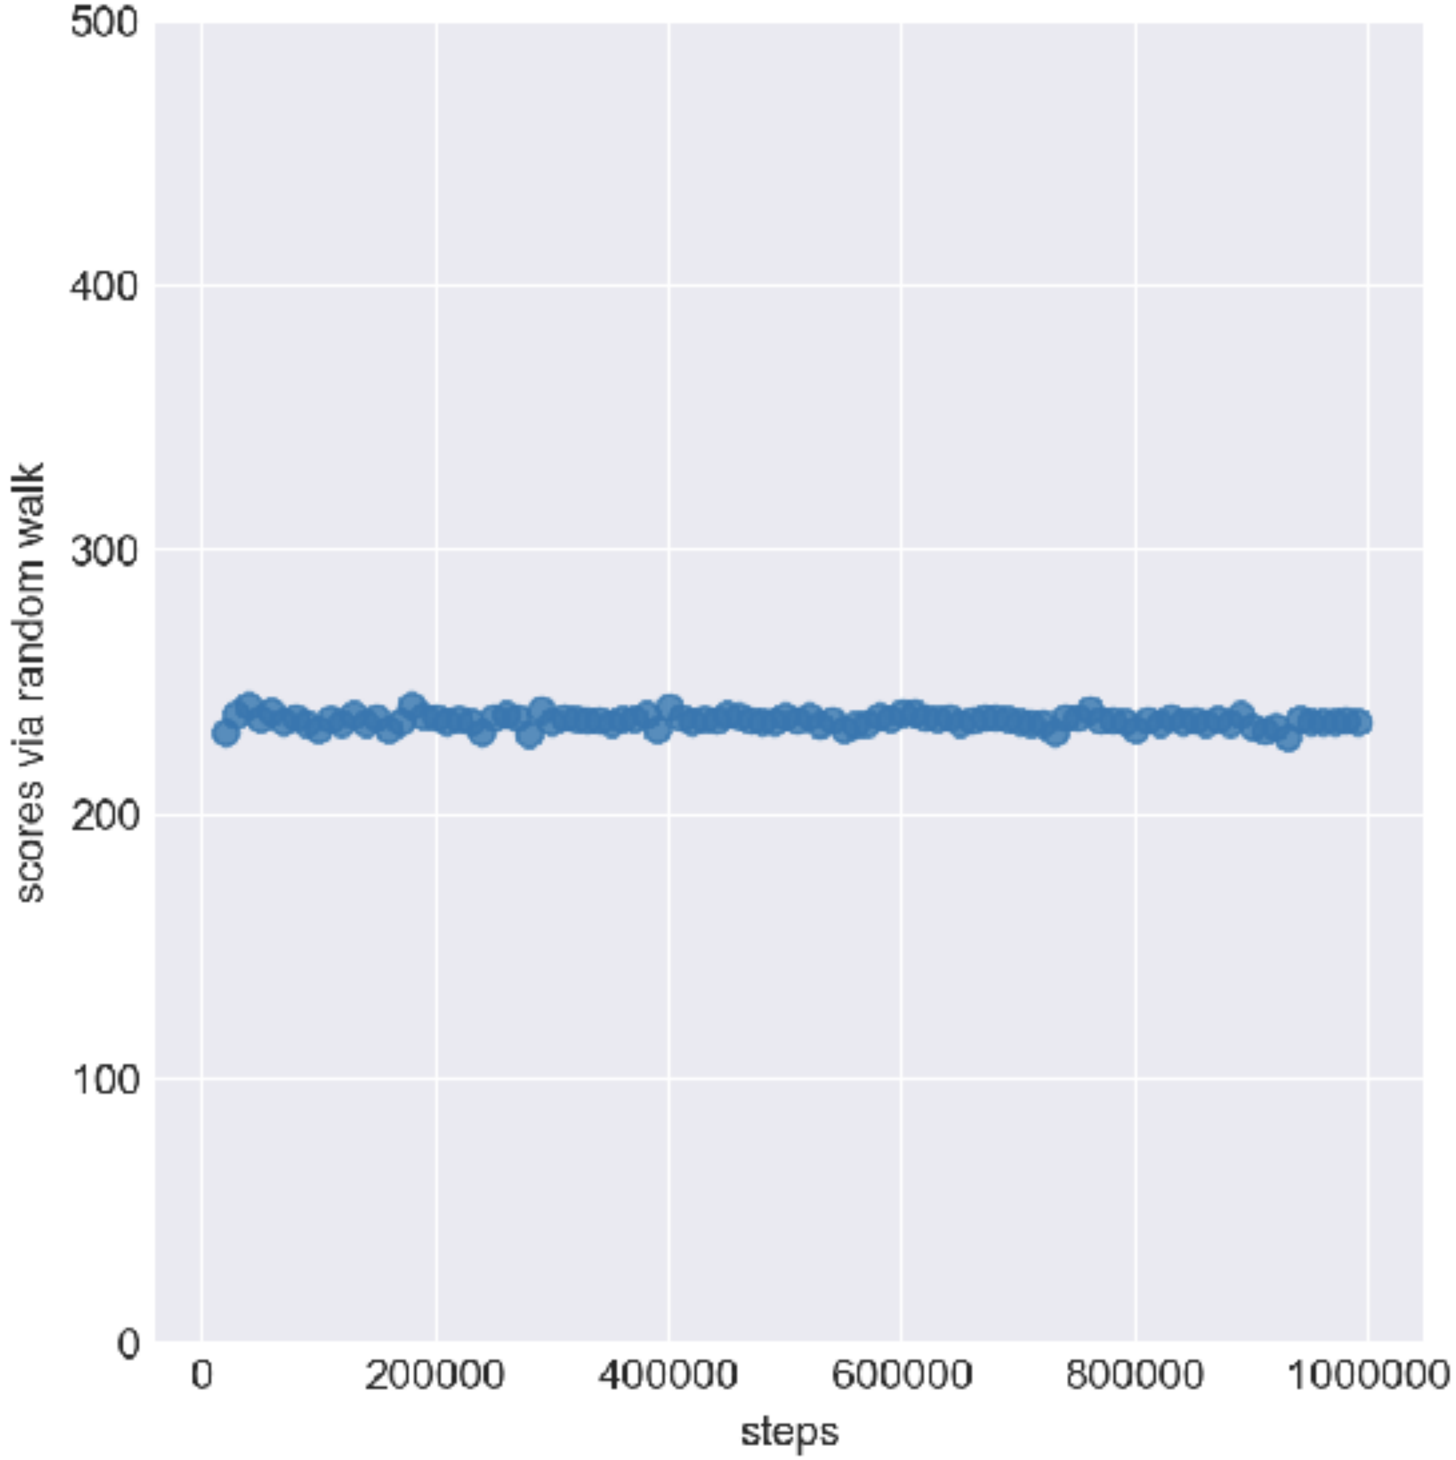

note scale: random walk takes 2-3 steps before NoSen

First Try: It Learns Quickly!

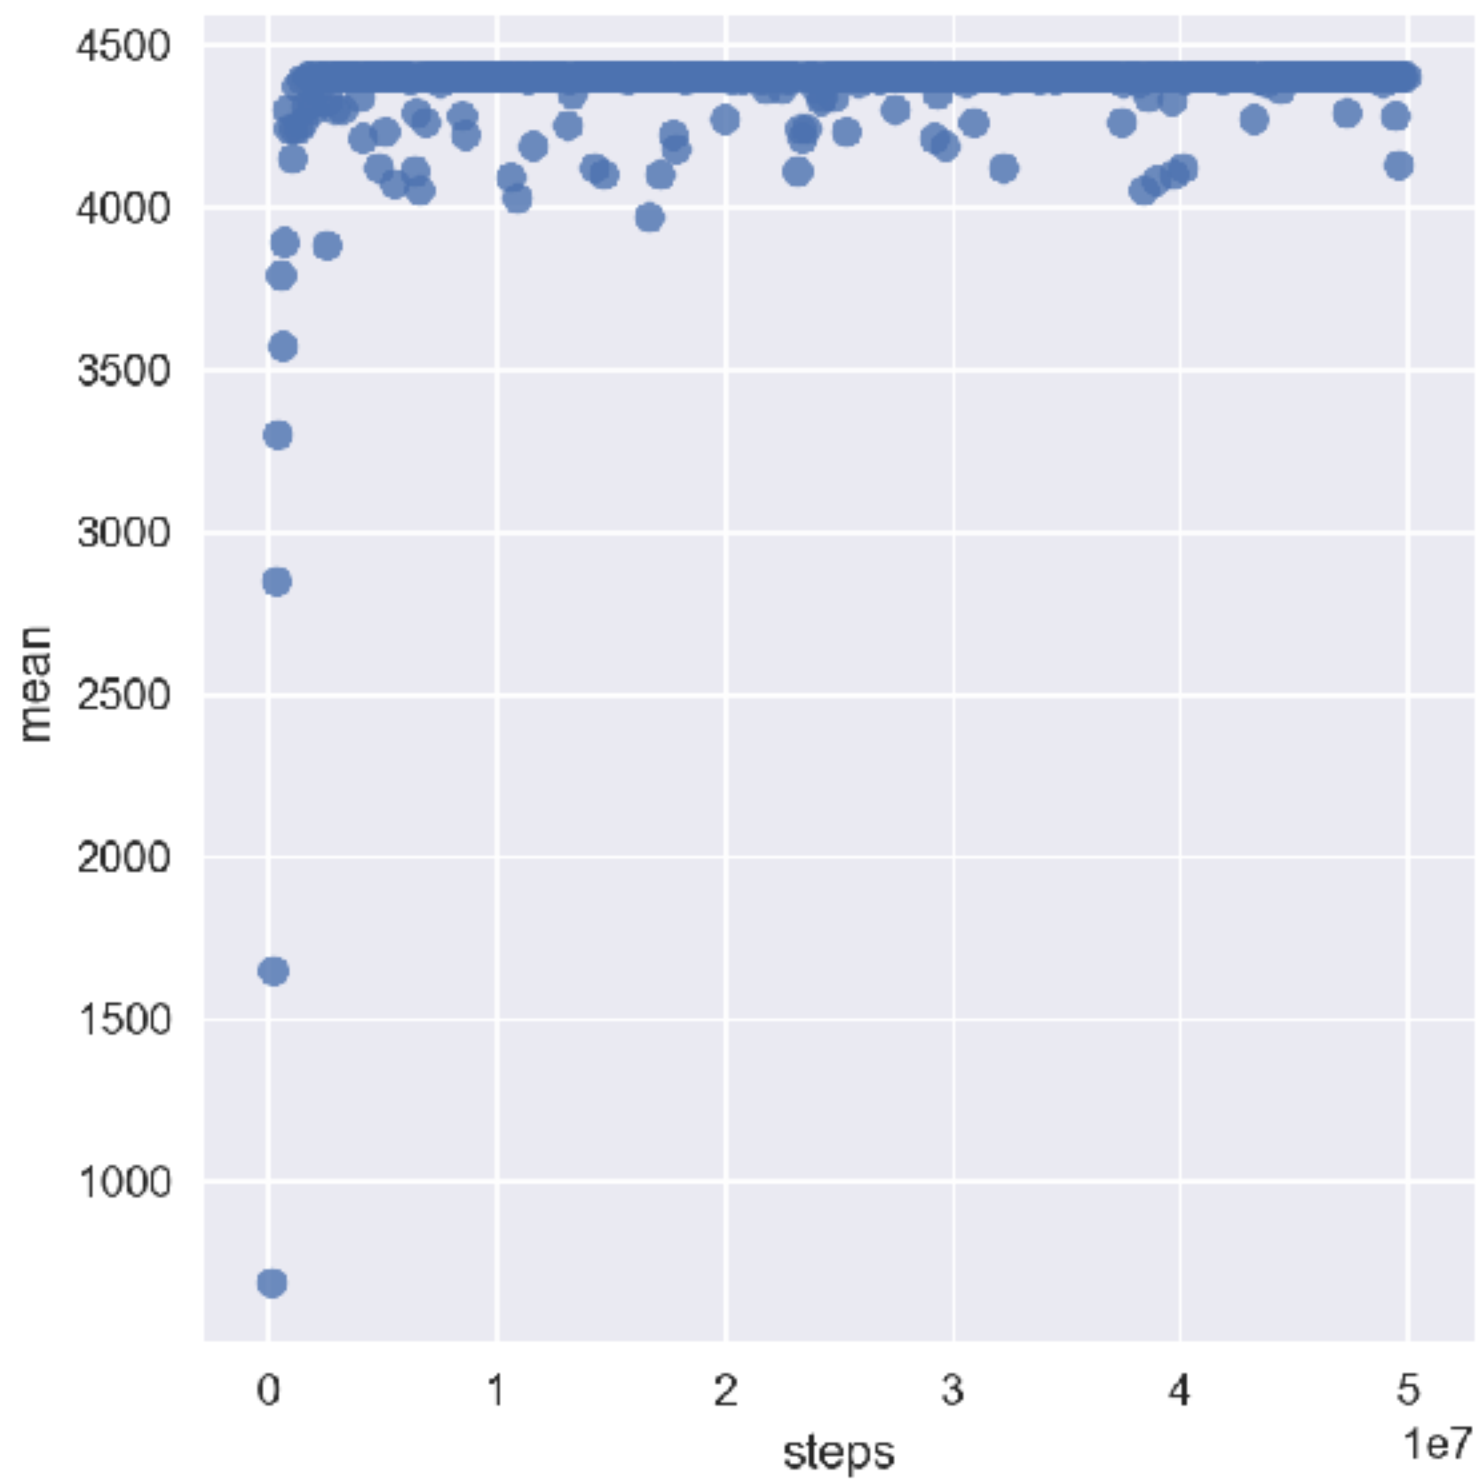

zoom in: decrease training time, increase eval interval

Second Try: See More Asymptote

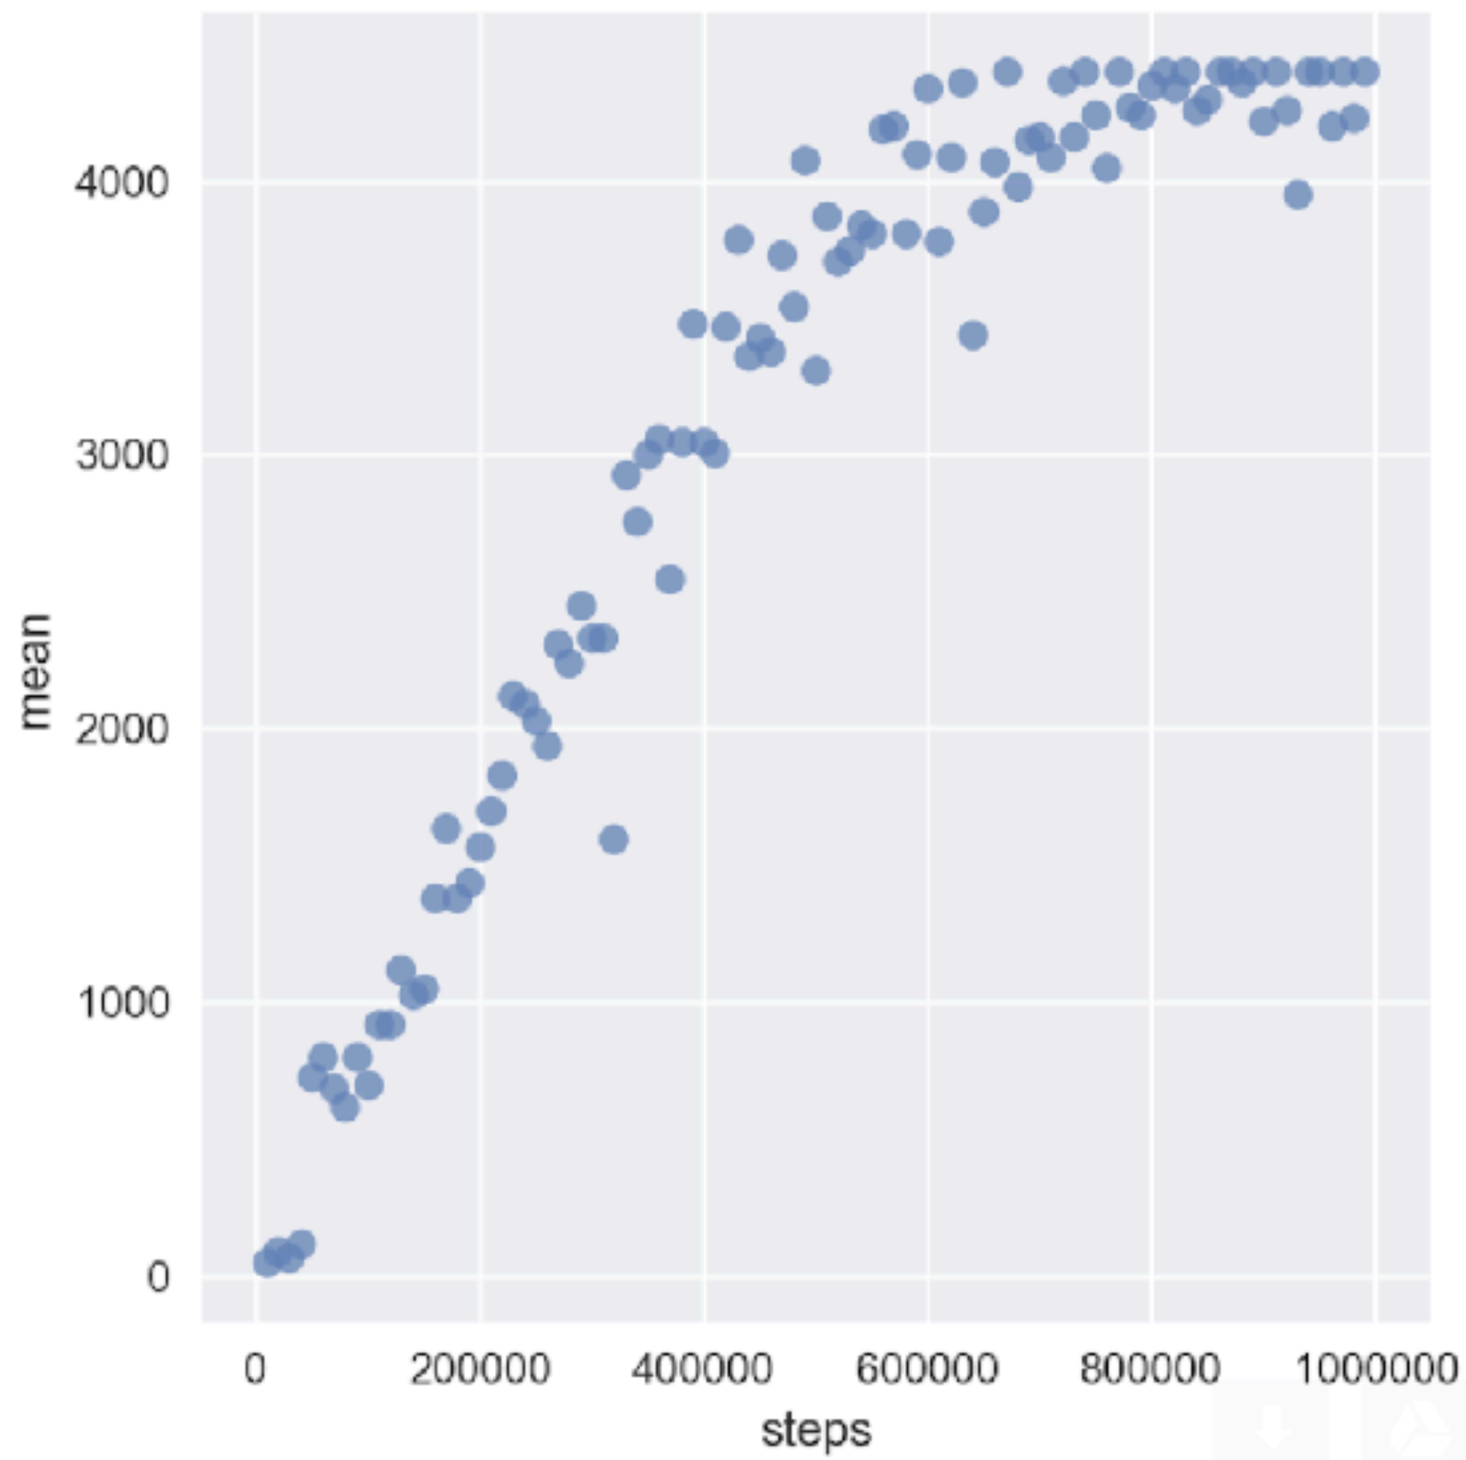

much better, but can we tweak so it does better?

Third Try: Different NN

use long short-term memory (LSTM) neural net

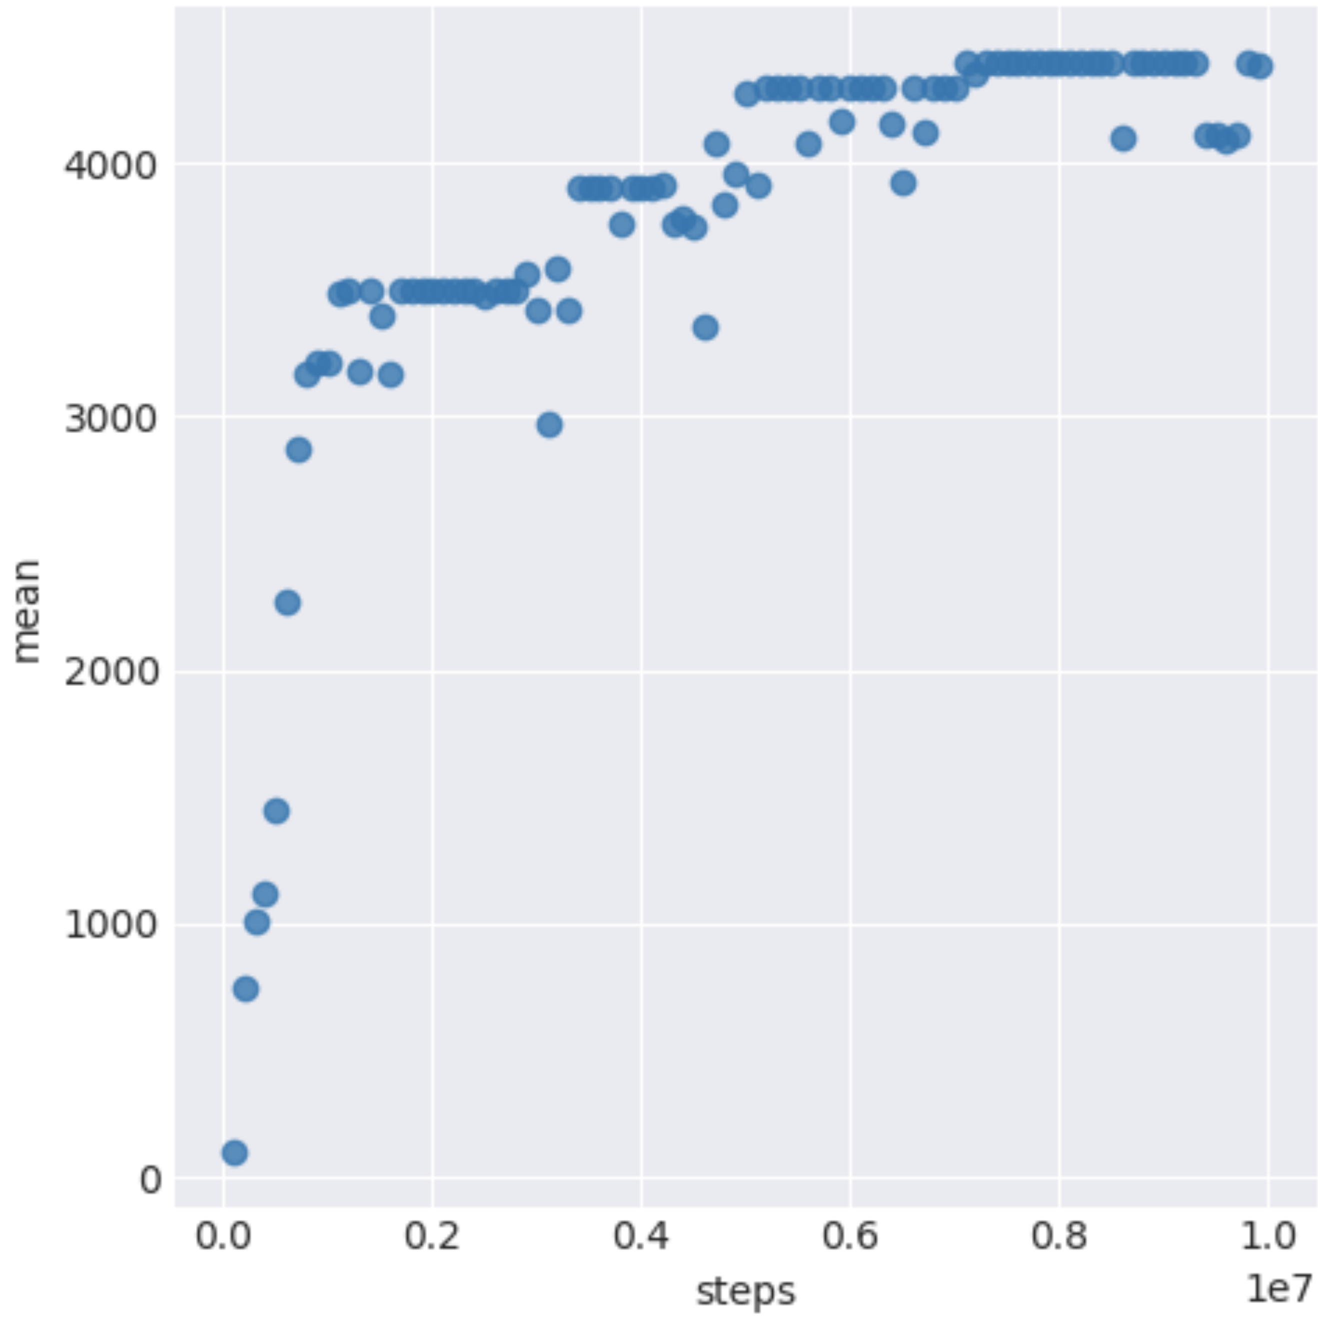

new feature: **four sharp plateaus**.

this is **punctuated equilibrium**, from evolution!

Fourth Try: Don’t Give Up

train a little longer, maybe it’s got more juice in it.

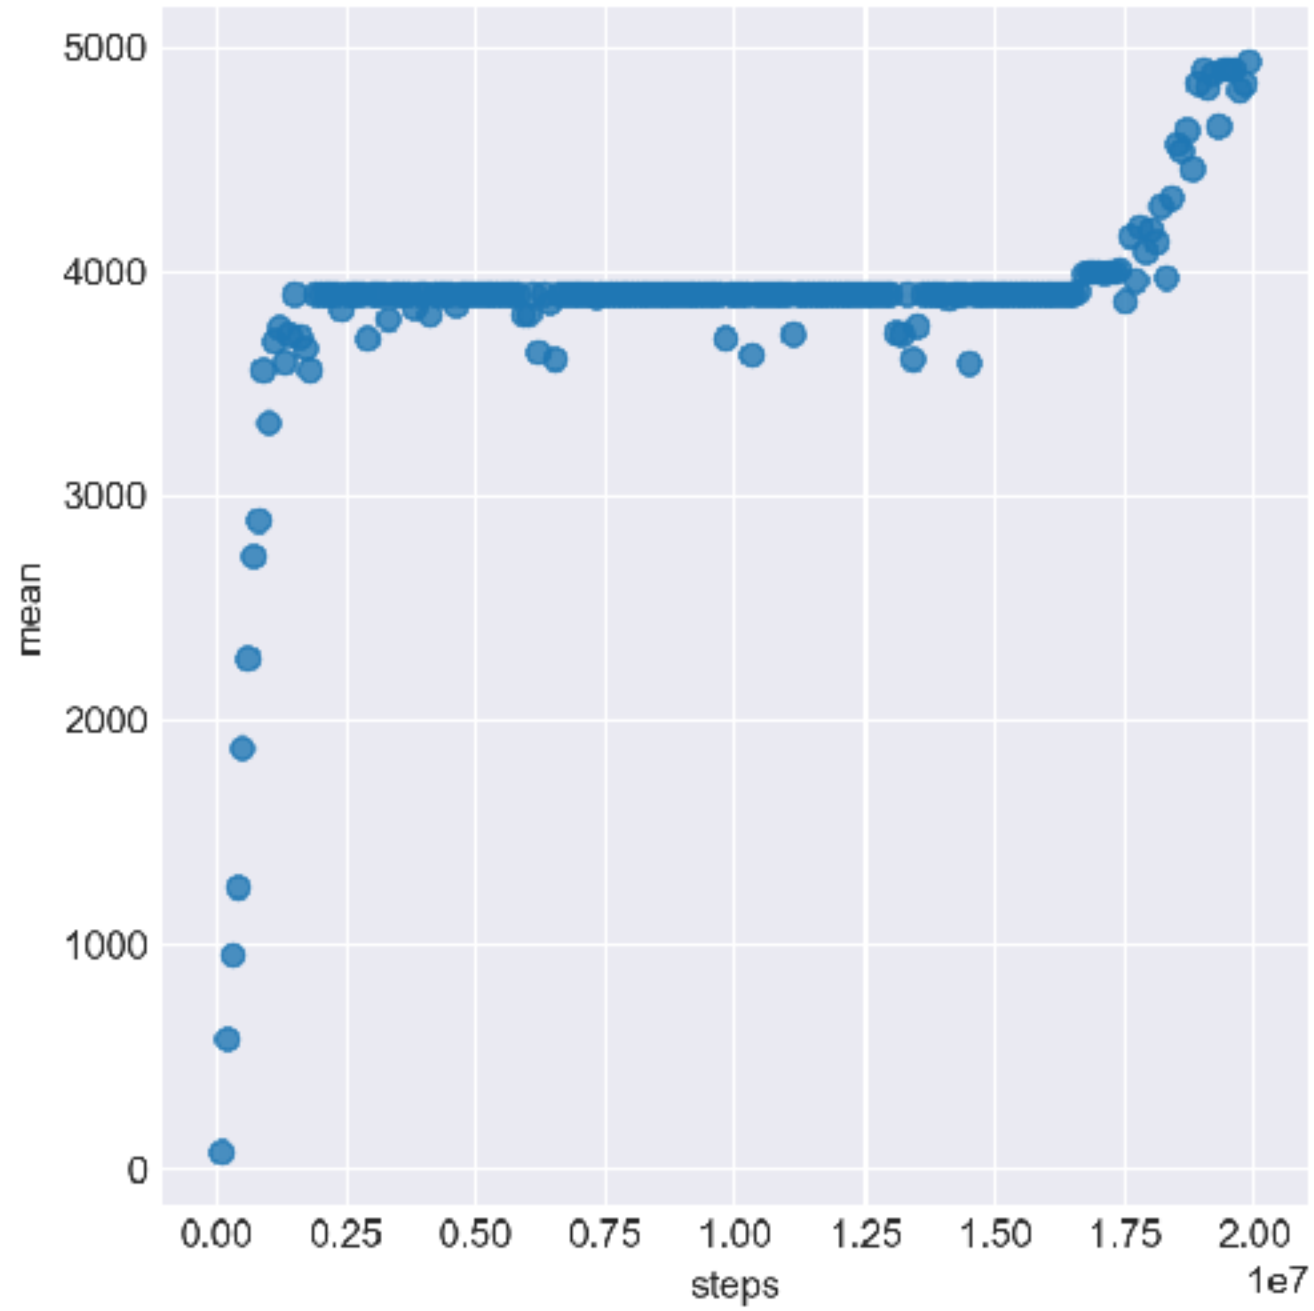

work work work, keep on training.

Fifth Try: The Best Yet

and there's clearly still room to grow.

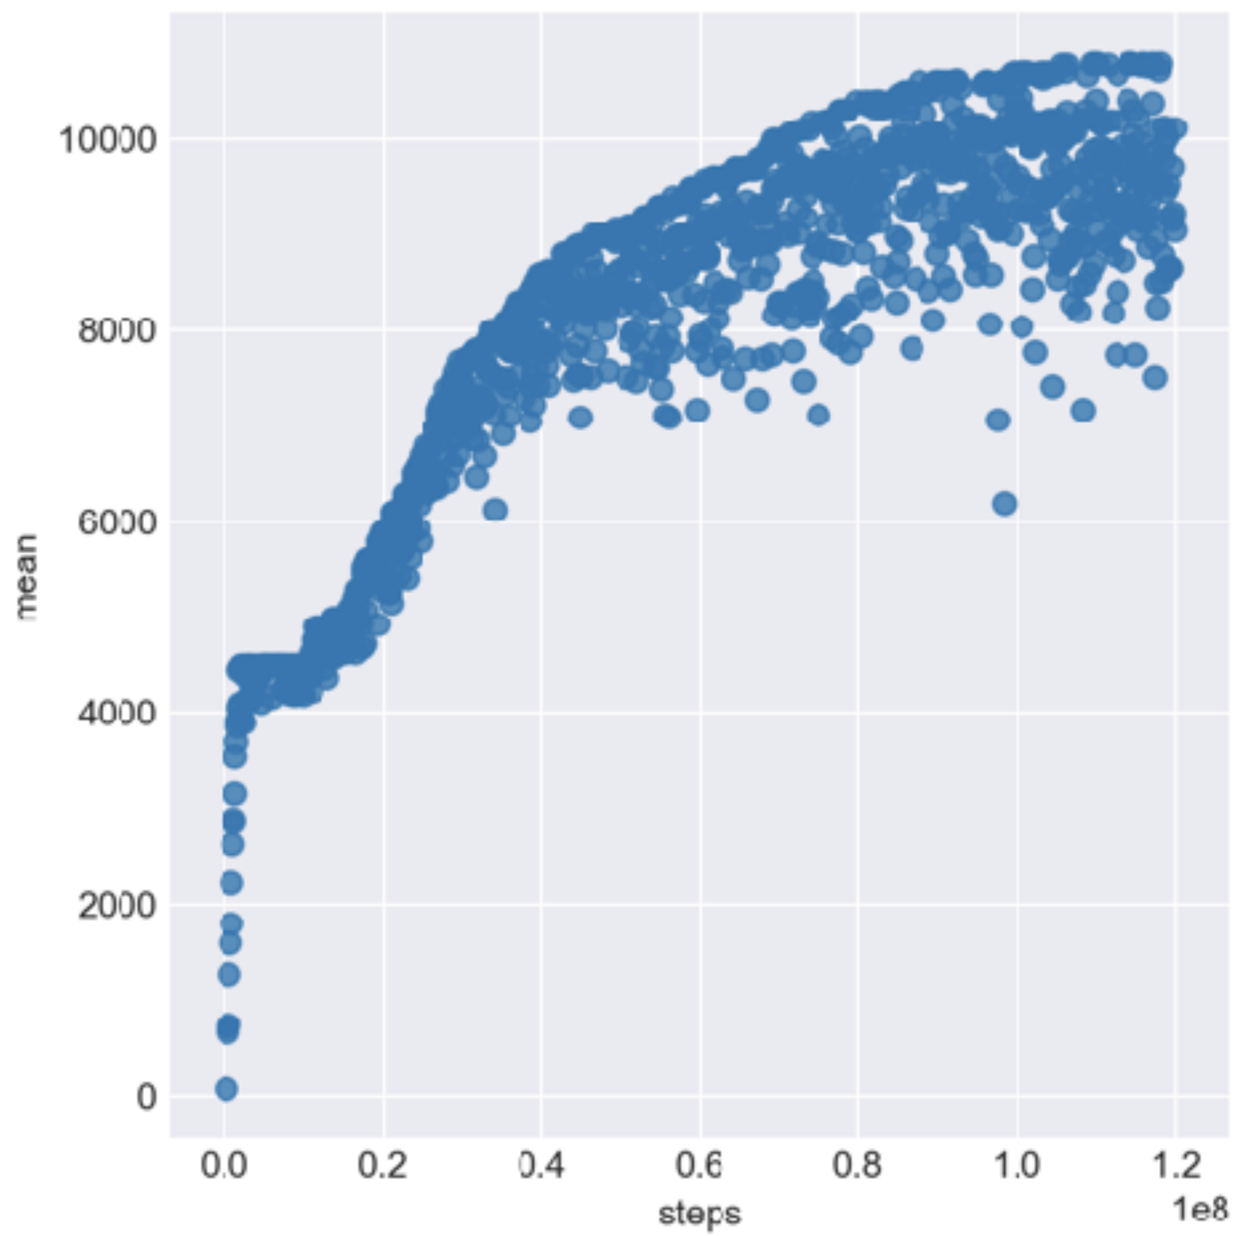

this is **training**, just phase 1. phase 2 and 3 to start soon.

Improving the Game

First game: out of bounds, no Sen limit. in bounds, maybe.

(rule: if F4, E6, E7, E8 in G, out of bounds. if not, in bounds.)

New game: out of bounds, no Sen limit. in bounds, has it.

(extra rule: if there are < I0 NH7's, must also be able to tune all to (2,3)-type I0* without forcing F4, E6, E7, E8 somewhere.* ***if not, out of bounds. otherwise, in bounds.****)*

Key points:

Necessary and sufficient for Sen limit.

Move reduction to 555.
<= 2.018 Googol states.

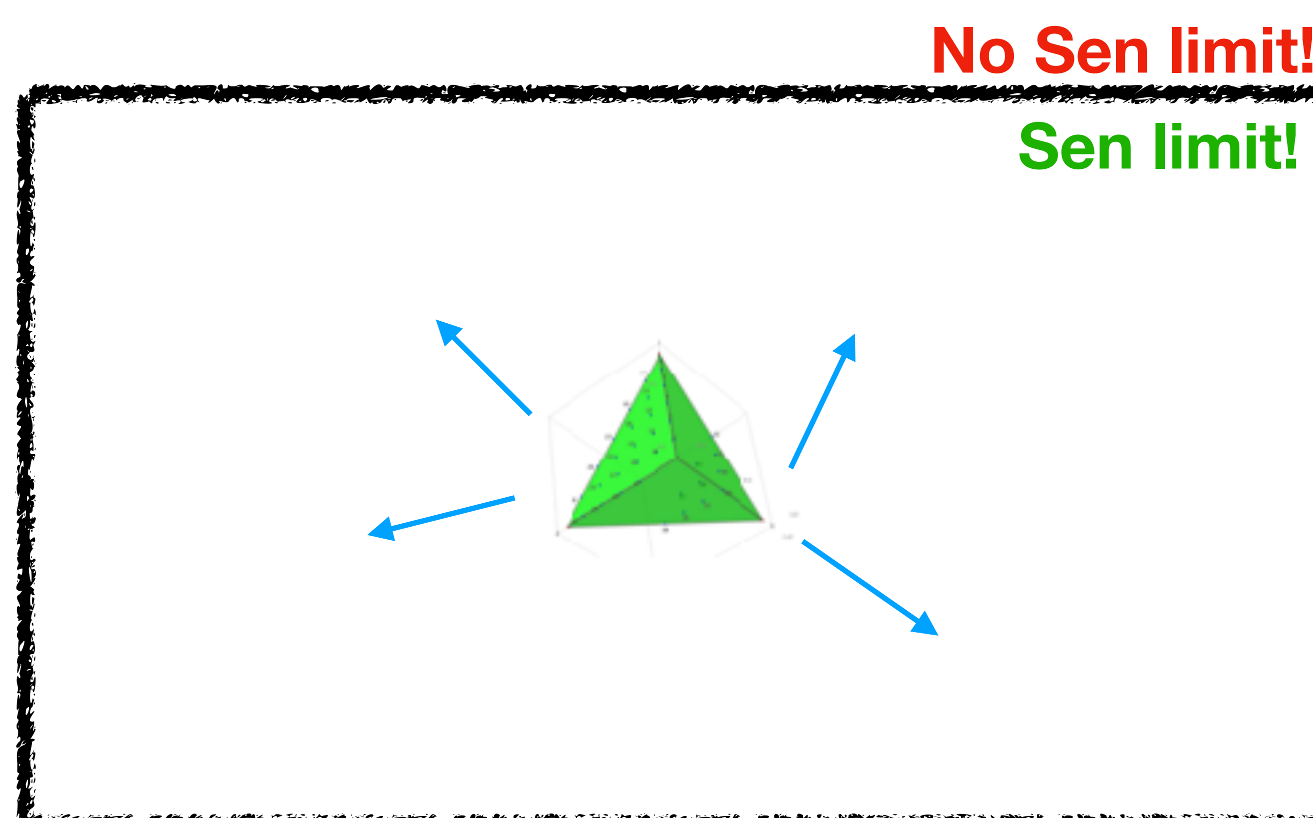

Random Walk, New Rules

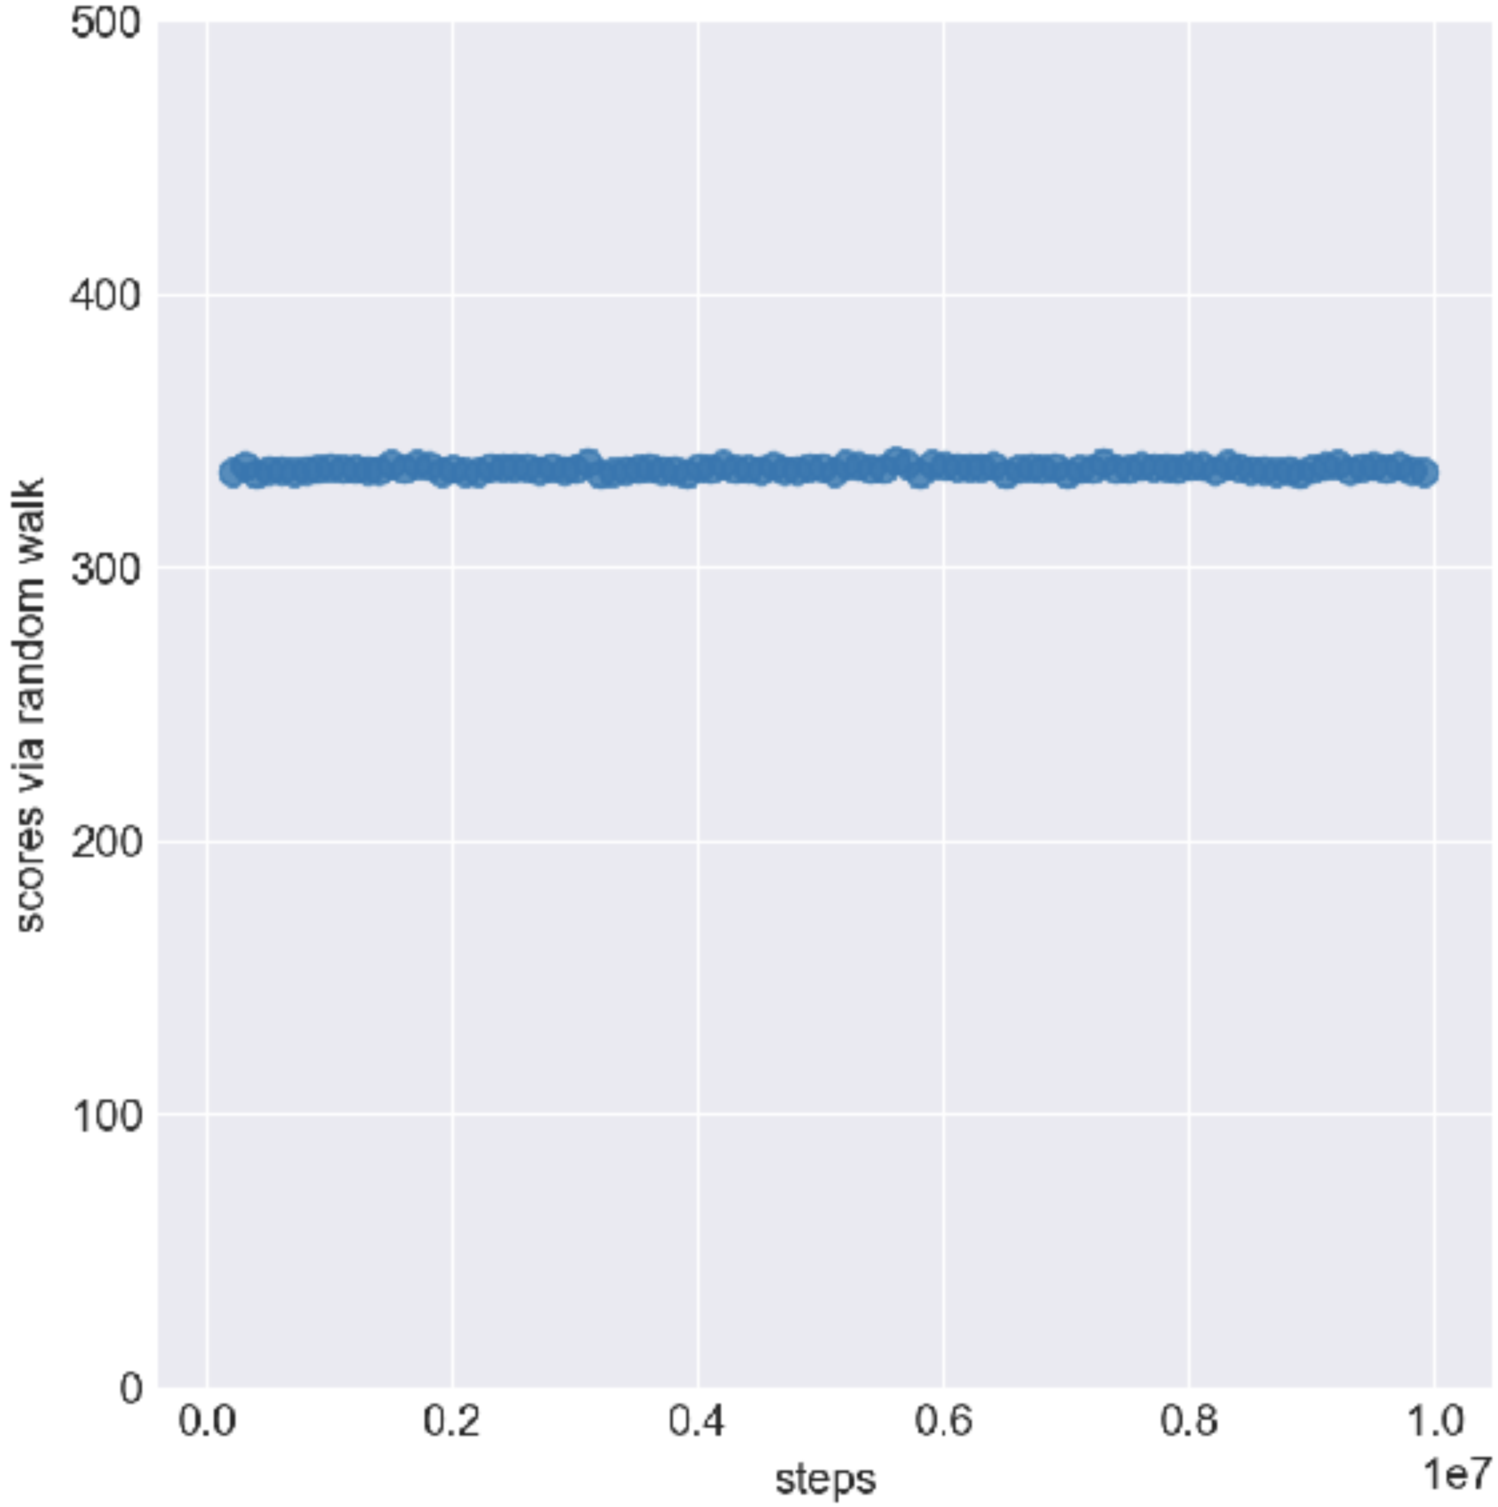

Note: now 3-4 steps before out of bounds instead of 2-3.

Best Training Yet

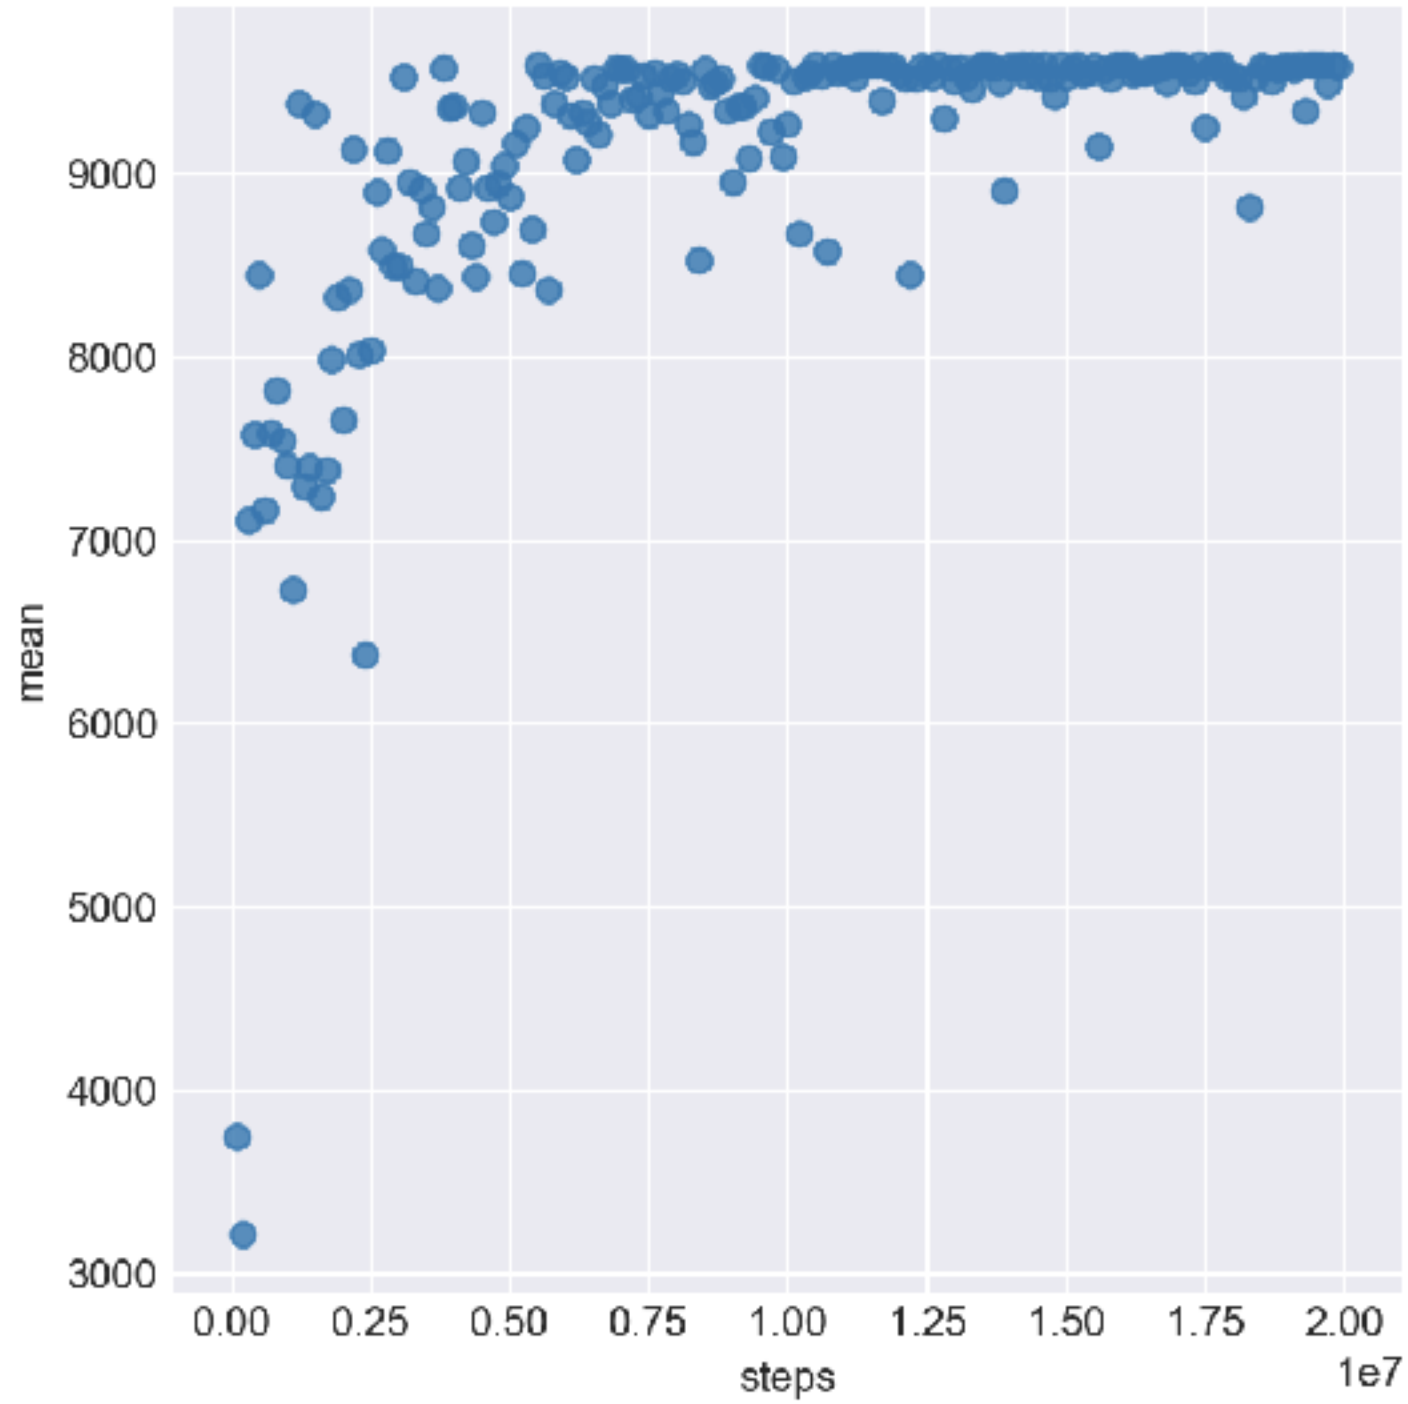

Note: now 9600 at 1m steps. old game ~ 4400. **Learns faster**.
max not yet old max, classic **exploration vs. exploitation problem.**

Predict the Boundary

supervised machine learning
for **simple** predictions

- *Generating the data*
- *10-fold cross validated:*
  - logistic regression
  - linear discriminant analyses
  - decision trees.

Supervised ML: Training

- **Generate boundary pairs:**
  let agent run for awhile. track last Sen, first no Sen

  407487 last Sen states
  413264 first no Sen states

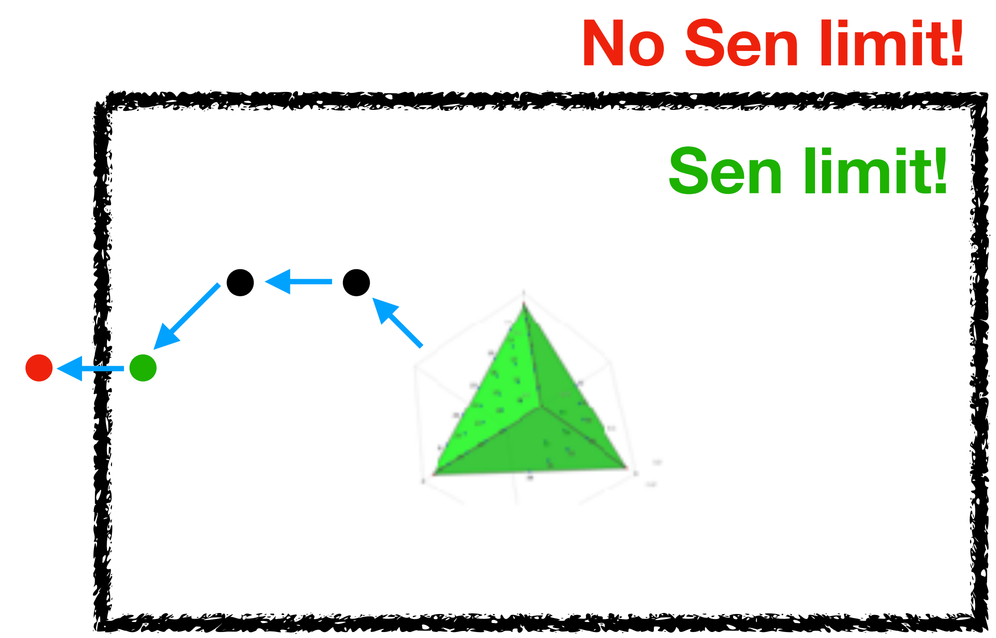

- **Goal:** given state (added trees),
  predict Sen vs. no Sen.

- **Training:** **simple algorithms** using scikit-learn
  can be more understandable / interpretable

Algorithms: logistic regression, linear discriminant analysis
decision trees.

Logistic Regression

an example of **"simple" machine learning.**

- For binary classification, drops optimal hyperplane,

$$\vec{c} \cdot \vec{s} + b = 0$$

- Decision: which side of plane?
- All info in coeffs and intercept.
- Probability determined by sigmoid of distance from plane.

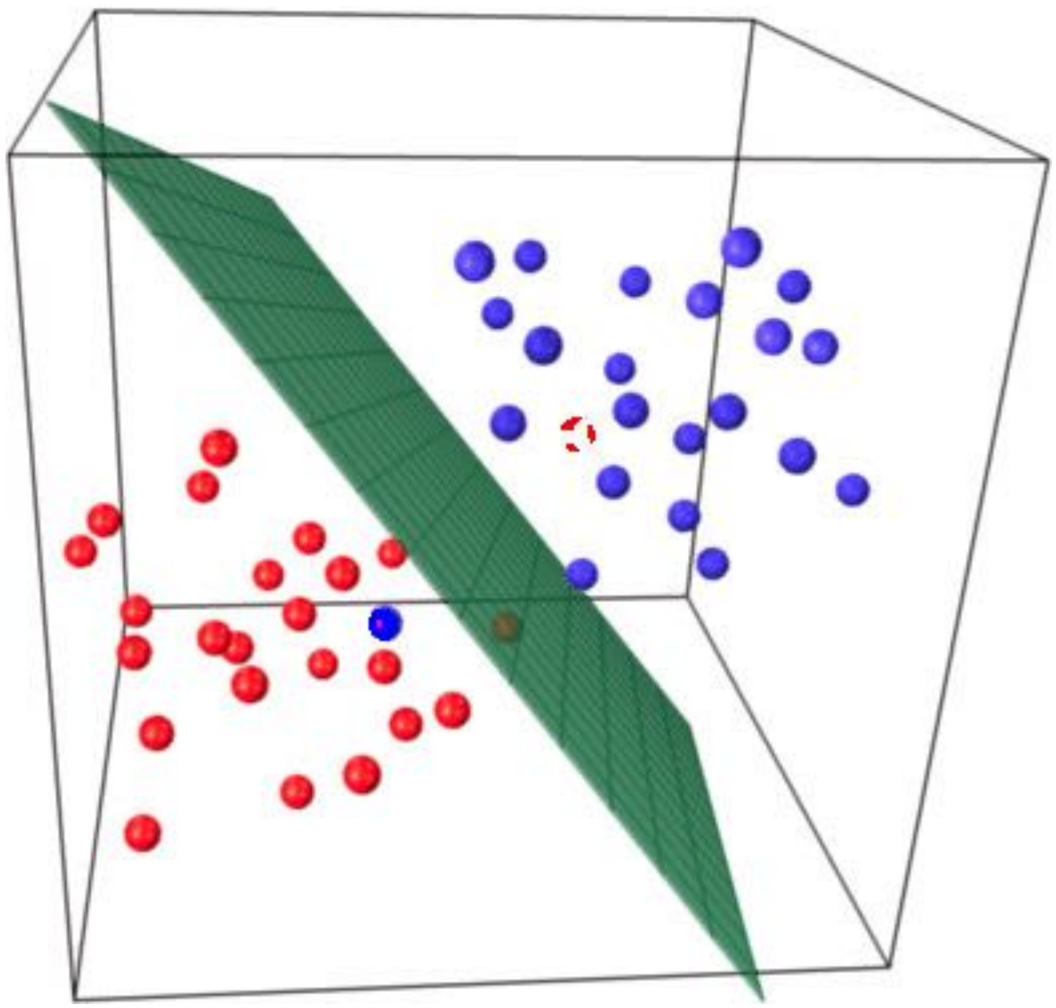

(seriously!? that's supposed to work!?)

Supervised ML: Results

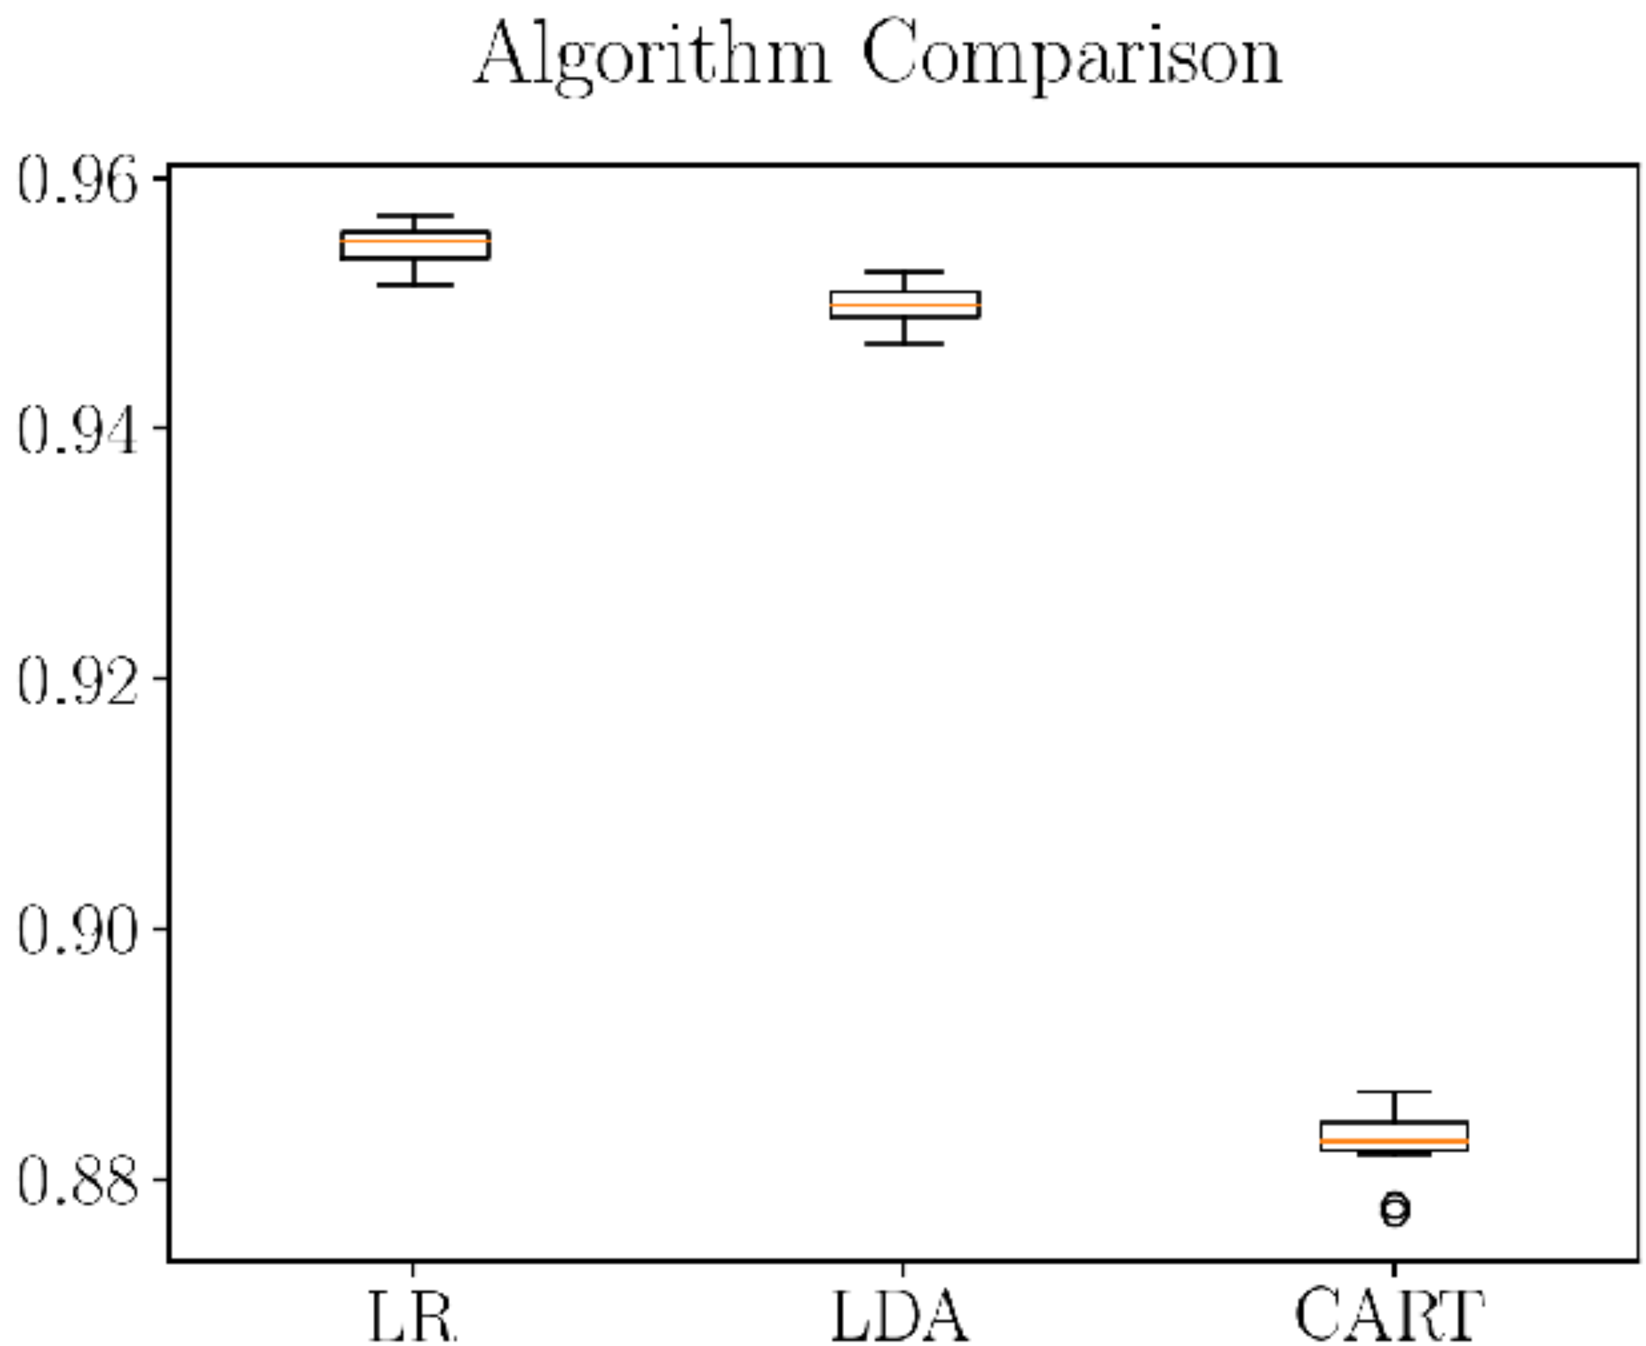

Note:
~8-fold data increase
in a moment

Data: 66201 first no Sen, and 65161 last Sen

10-fold cross-validated.

Understand the Boundary

intelligible AI and conjecture generation

- *characteristics of the boundary.*
- *understanding level of dangerousness of trees.*
- *visualizing dangerous trees.*

Brane Changes at Boundary

- **Boundary:** tune up to strongly coupled (SCFT?) point in moduli space with tensionless strings or actionless instantons. Blow ups in base.

- **Breaking the Sen Limit:** branes enhance in a no Sen way. data: 62089 (last Sen, first no Sen pairs)

	II	A_1	A_{1m}	A_2	G_2	$SO(7)$	$SO(8)$
F_4	0	147	625	5937	18526	1	0
E_6	0	17	11	20	88	0	0
E_7	0	9	4	24	156	0	0
E_8	0	1	9	5	73	0	0

- Future work (?): physics of strongly coupled points, cosmological implications of passing boundary?

Intelligible AI from Decision Tree “Importances” Histogram

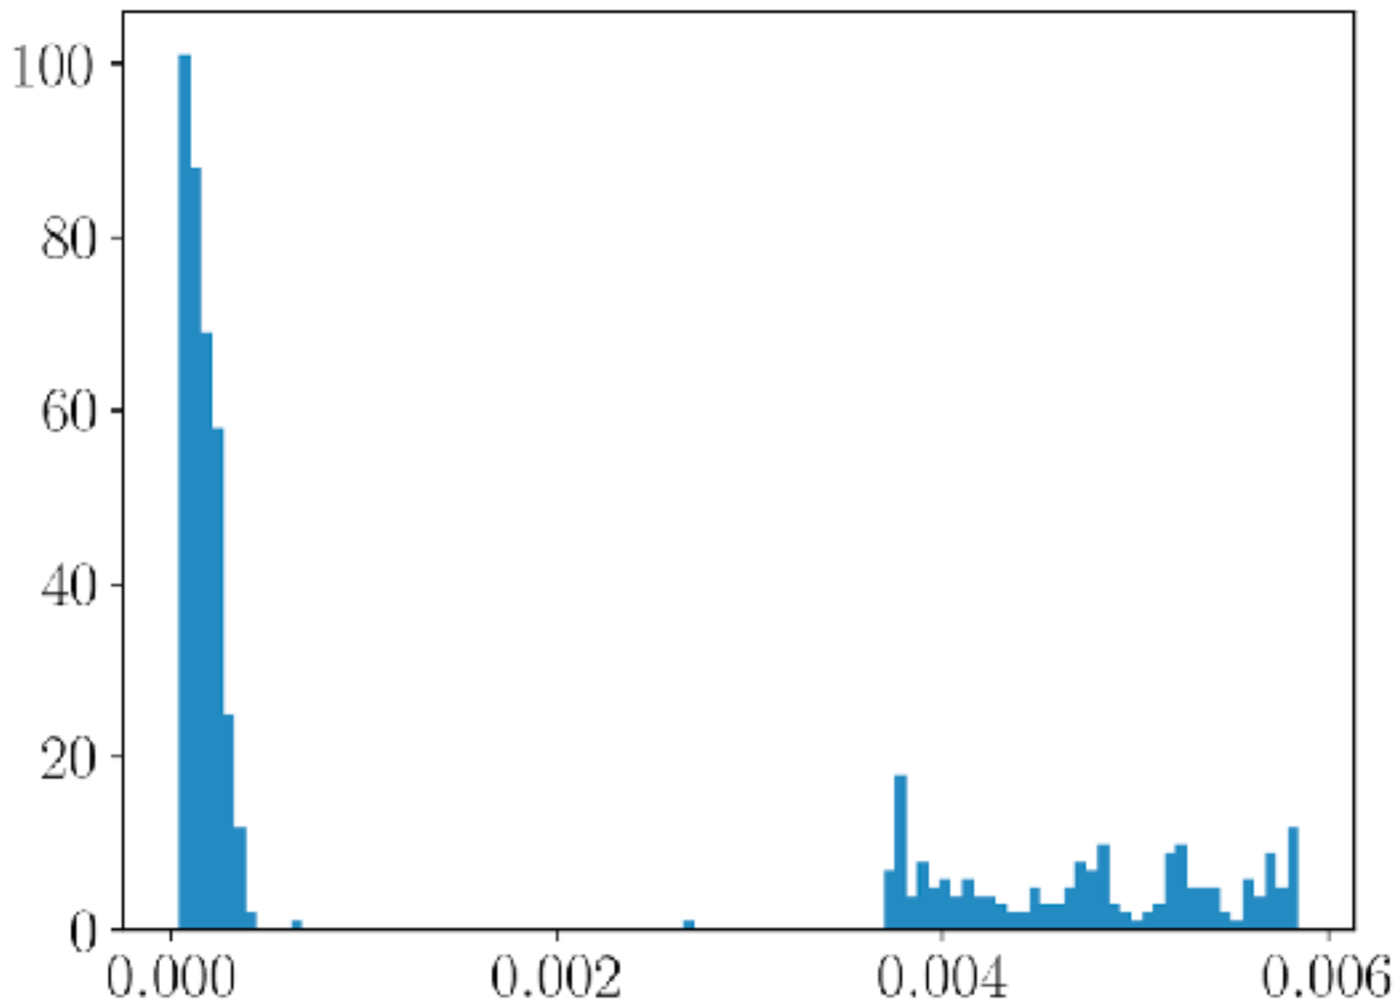

Data:
413264 First No Sen
407487 Last Sen

Accuracy: 96.6%

Higher importance means no Sen is more likely.

Ranks trees by most likely to force no Sen.

Intelligible AI from Logistic Regression Coefficients Histogram

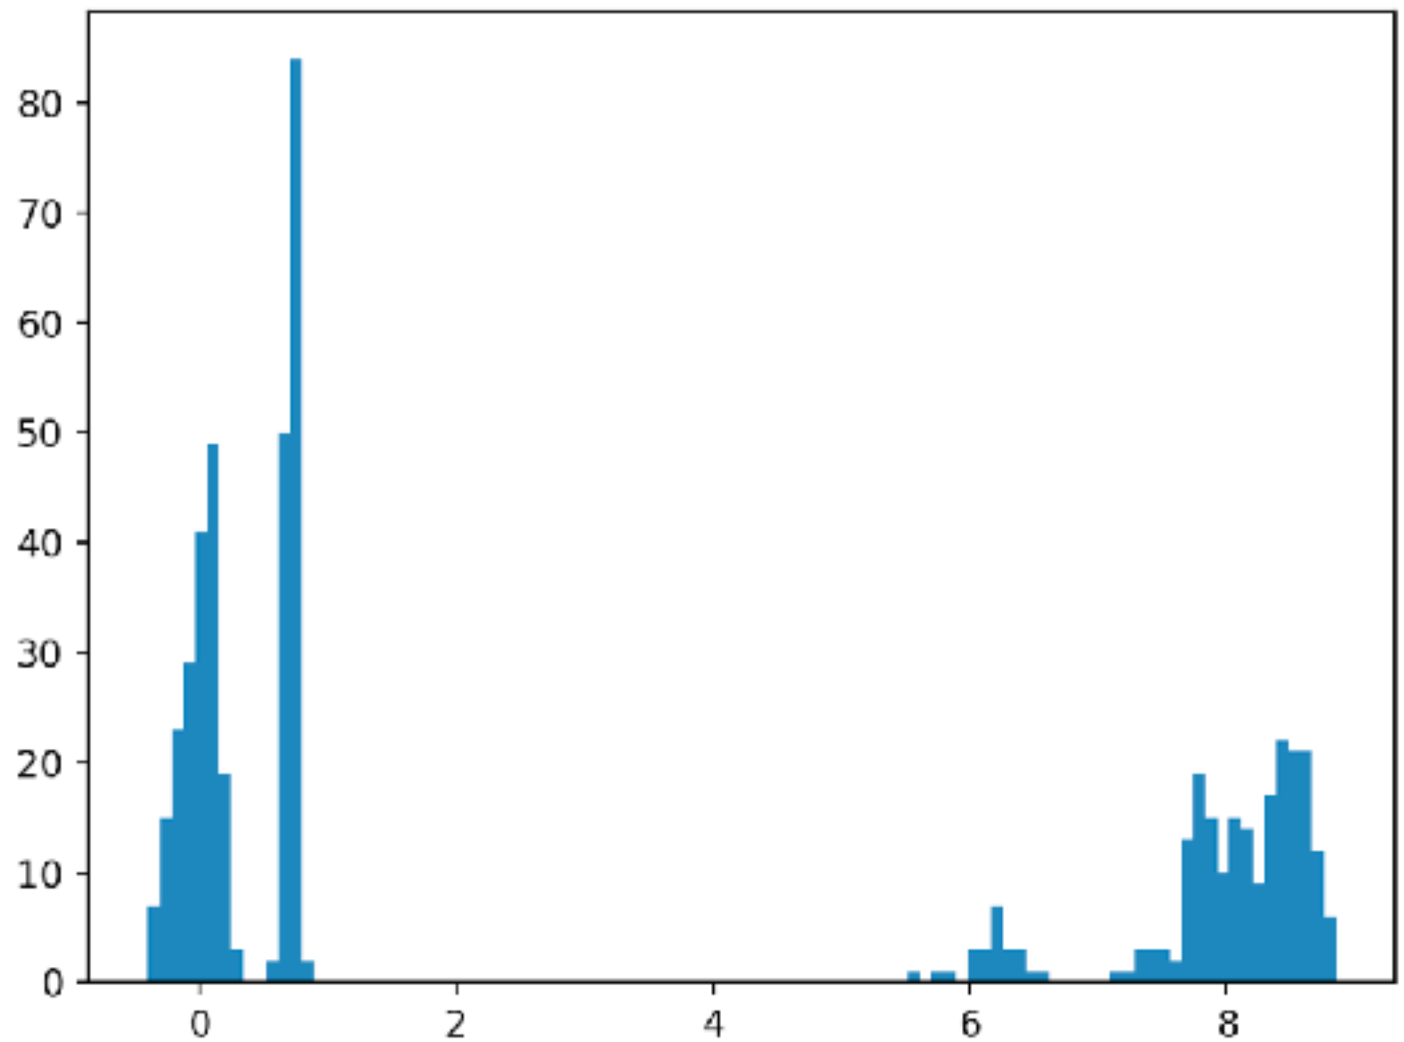

Data:
413264 First No Sen
407487 Last Sen

Accuracy: 98.5%

Similar, but **intuitive push-pull game**. Intercept -7, trees with coefs > 0 pull to right. Only need to add a few trees with coefficient > 5 to be in serious danger of no Sen Limit.

Visualizing Dangerous Trees, Working Towards Conjecture

- **Red edges:** any trees there or on triangle wraparounds are in the set of dangerous trees.

  i.e. 12 bad ones associated with red edges.

- **25 other bad:** 18 of which are connected to vertices.

- remaining dangerous trees are of understandable type.

  intuitive, but **ML —> linchpins.**

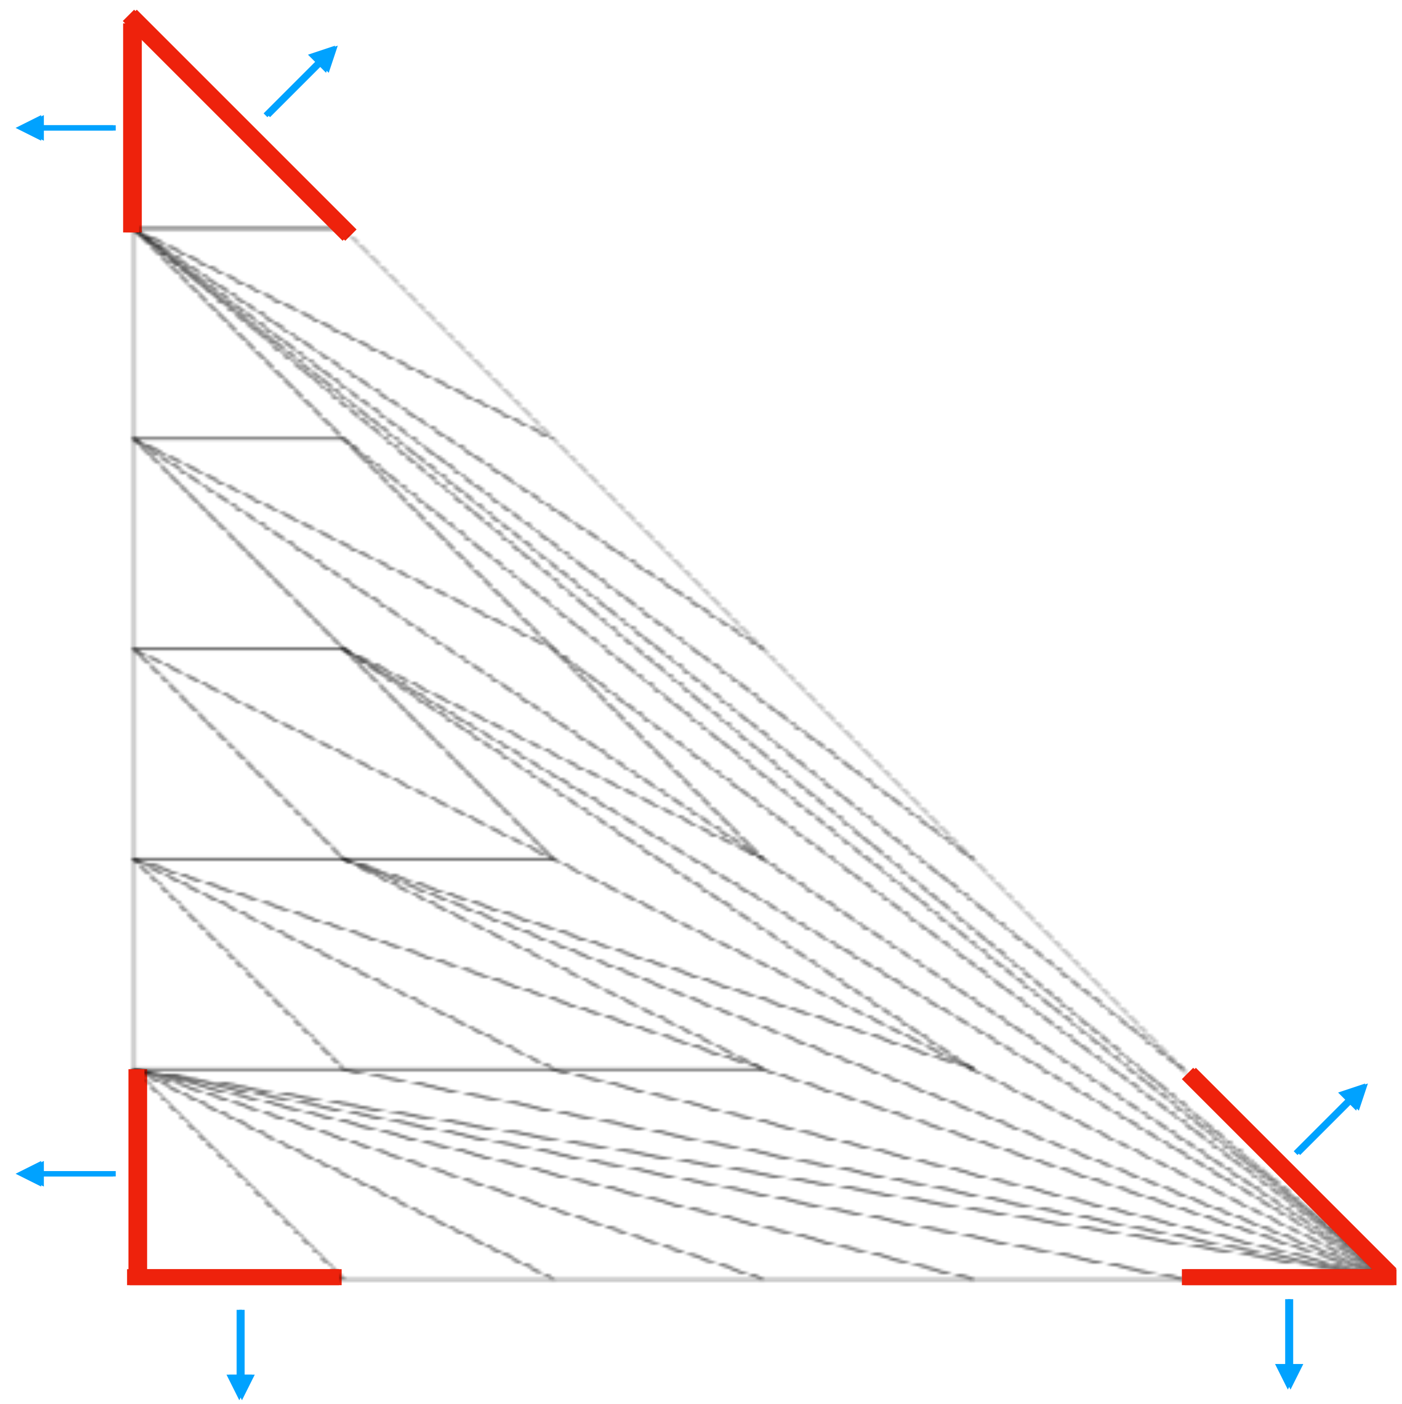

Linchpins —> Theorem?

we don't know,
but it's a promising idea,
and probabilities should be computable.
stay tuned!

Concluding Thoughts

- 1) **Explore** lamppost, probe IIB-F boundary.

  use: ***deep reinforcement learning***.

  New class of anomaly detection problem?

- 2) **Predict** the boundary.

  use: ***supervised machine learning.***

- 3) **Understand** the boundary.

  use: ***intelligible AI.***

  - branes enhance through strong coupling SCFT (?) to strongly coupled brane systems.

  - (implicit) nearly all F-theory backgrounds with Sen limits have strongly coupled at generic CS with intersecting I0* (O7).

  - Rigorous theorem on boundary is likely.

IIB is ε

This work provides even more motivation for doing the **hard formal and data science work**

to understand which of the many detailed lessons from the IIB **safe haven** genuinely carry over into F-theory,

and which are merely artifacts of being under this **lamppost**.