Descending into the Modular Bootstrap

Jesse Thaler

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
MIT CENTER FOR THEORETICAL PHYSICS – A LEINWEBER INSTITUTE

École Normale Supérieure Summer Institute, Paris — June 26, 2026

The NSF Institute for Artificial Intelligence and Fundamental Interactions (IAIFI /aɪ-faɪ/ iaifi.org)

New for 2026!

Artificial intelligence as a pathway to scientific insight

Physics intelligence as a pathway to AI innovation

Launched August 2020 & Renewed June 2026 (!!)

Progress in AI+Physics driven by early career talent with interdisciplinary training

Question I've been wrestling with on sabbatical:
What does it mean to do robust discovery science with AI?

Exactness guarantees?
e.g. lattice sampling

Explicit verification?
e.g. symbolic engines

Statistical verification?
e.g. calibrated confidence intervals

Sean Benevedes, PhD defense April 14, 2026!
[e.g. wifi ensembles in Benevedes, JDT, PRD 2025;
TAMMs in Alvarez, Benevedes, Szewc, JDT, arXiv 2026]

Today: Case study to show how my thinking has evolved
Searching for "new theoretical physics" through optimization

Punchline

By combining "direct AI/ML" and "indirect AI/ML", we find an apparent* obstruction to constructing two-dimensional conformal field theories with small central charges and large spectral gaps

If you are surprised that someone from the hep-ph community could make progress on a hep-th problem, you are not alone…

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026;
+ Langhoff, work in progress]







Descending into the Modular Bootstrap

Machine Learning through the Lens of Physics

AI is transforming physics research, but many foundational aspects of AI can be translated into the language of physics

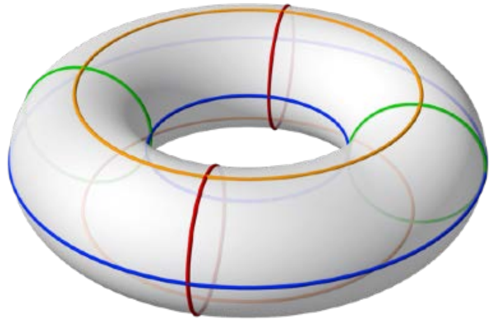

Revisiting the Modular Bootstrap

To confront thorny challenges in theoretical physics, we can reframe them as well-defined optimization problems

What Have We Actually Learned?

As we capitalize on emerging AI technologies, we still need to insist on robustness and reproducibility

Machine Learning through the Lens of Physics

AI is transforming physics research, but many foundational aspects of AI can be translated into the language of physics

Revisiting the Modular Bootstrap

To confront thorny challenges in theoretical physics, we can reframe them as well-defined optimization problems

What Have We Actually Learned?

As we capitalize on emerging AI technologies, we still need to insist on robustness and reproducibility

Confronting AI/ML Hype vs. Reality

"AI is a transformative technology" vs. "ML is just numerical optimization"

Both are true! And both are changing the scientific enterprise

Direct AI/ML: Enables scientific investigations of high-dimensional spaces
e.g. simulation-based inference

Indirect AI/ML: Capitalizes on emergent behaviors in computational systems
e.g. foundation models, agentic AI

ML is "Just" Optimization in Response to Data

Surprisingly powerful,
both directly and indirectly!

	ML Framing		Physics Translation	
Training Data	$\frac{1}{N}\sum_{i=1}^{N}(\cdots)$	$=$	$\int d\Phi\,(\cdots) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$	Monte Carlo Integration
Loss Function(al)	$\mathcal{L}_i[f(\Phi)]$	$\Leftrightarrow$	$\frac{\delta \mathcal{L}}{\delta f} = 0$	Lagrangian Mechanics *e.g. Euler-Lagrange equation*
Learnable Function *e.g. neural networks*	$f_{\mathrm{NN}}(\Phi; b)$	$\Rightarrow$	$f_{\mathrm{physics}}(\Phi; b)$	Physics Knowledge *e.g. symmetries*
Optimizer *e.g. gradient descent*	$\Delta b = -\eta\,\nabla_b \mathcal{L}$	$\Leftrightarrow$	$m\,\ddot{b} + \gamma\,\dot{b} + \frac{dV}{db} = 0$	Newtonian Dynamics *e.g. overdamped oscillator*

Additional Aspects of Modern ML: Stochastic Optimization, Overparametrization, Generalization, …

[from my GGI lectures, January 2026]

Scenes from My Sabbatical: Direct AI/ML

Constraining 2D CFTs

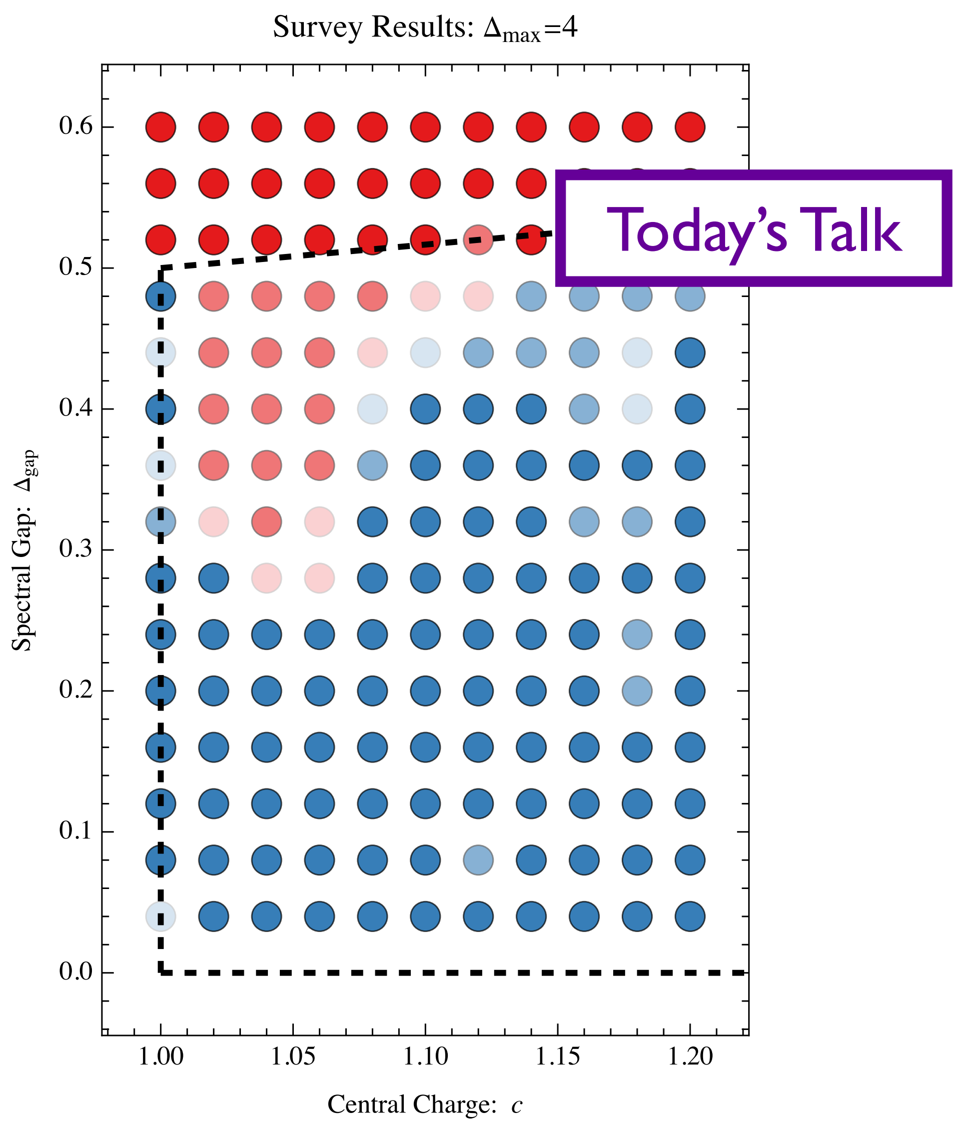

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
using Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Synthesizing QCD Predictions

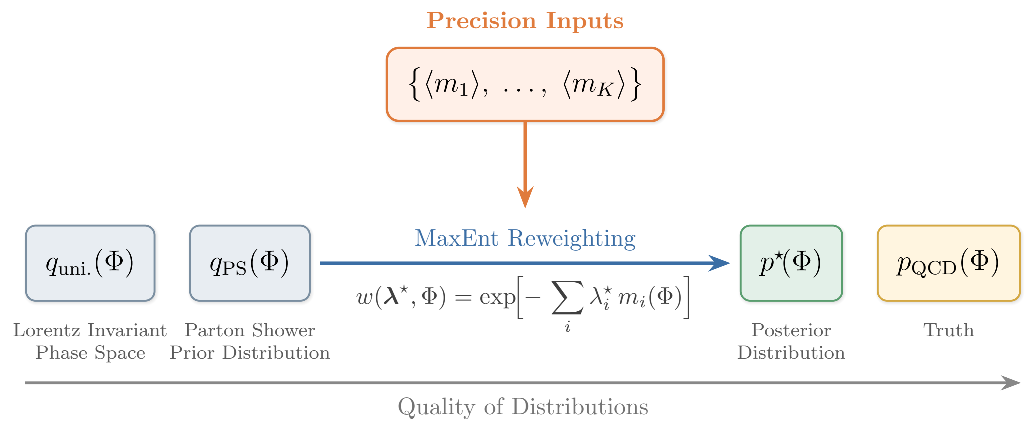

[Assi, Lee, JDT, arXiv 2026; based on Assi, Lee, Höche, JDT, PRL 2025]

Robust Inference Despite Domain Shift

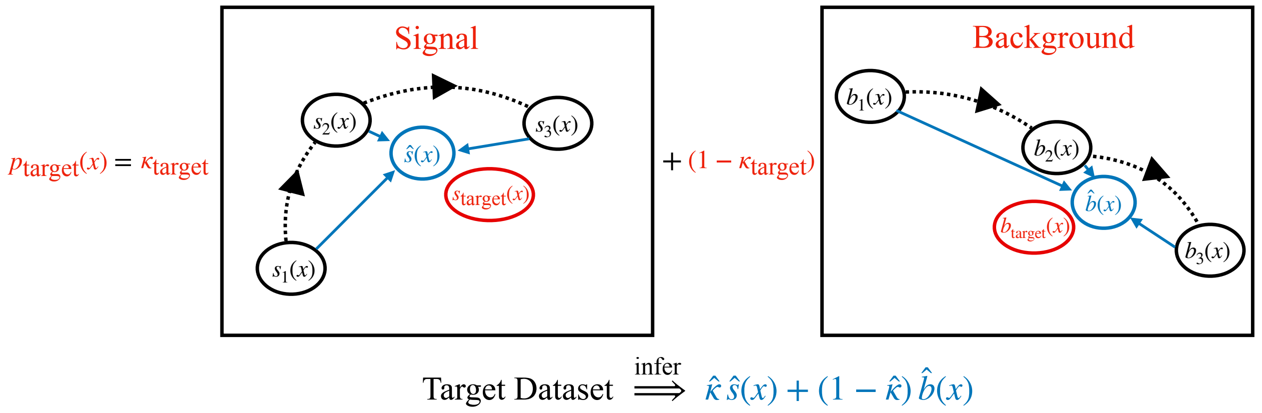

[Alvarez, Benevedes, Szewc, JDT, arXiv 2026; see also Benevedes, JDT, PRD 2025]

Scenes from My Sabbatical: Indirect AI/ML

AI Research Assistant

[Answer: use containers]

I am exploring the use of AI in fundamental physics. One of the challenges in this area is using simulation tools in a reliable way. A particularly difficult piece of software to even install properly is Herwig. Can you show how AI can get this tool up and running and run a demo event generation of 1000 events of e+e- to hadrons at the Z pole?

Interactive Advice Document

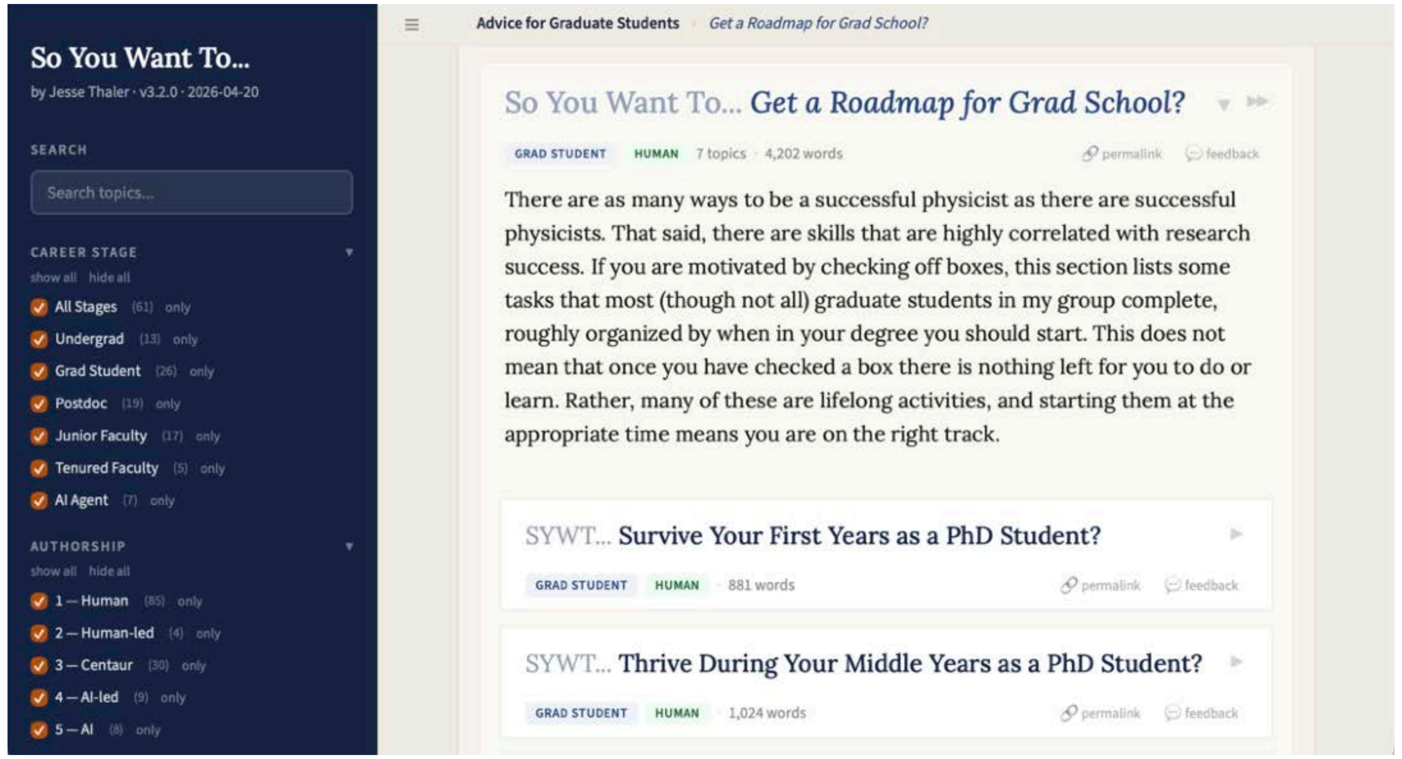

[JDT, not yet public]

AI and Academic Publishing Salon

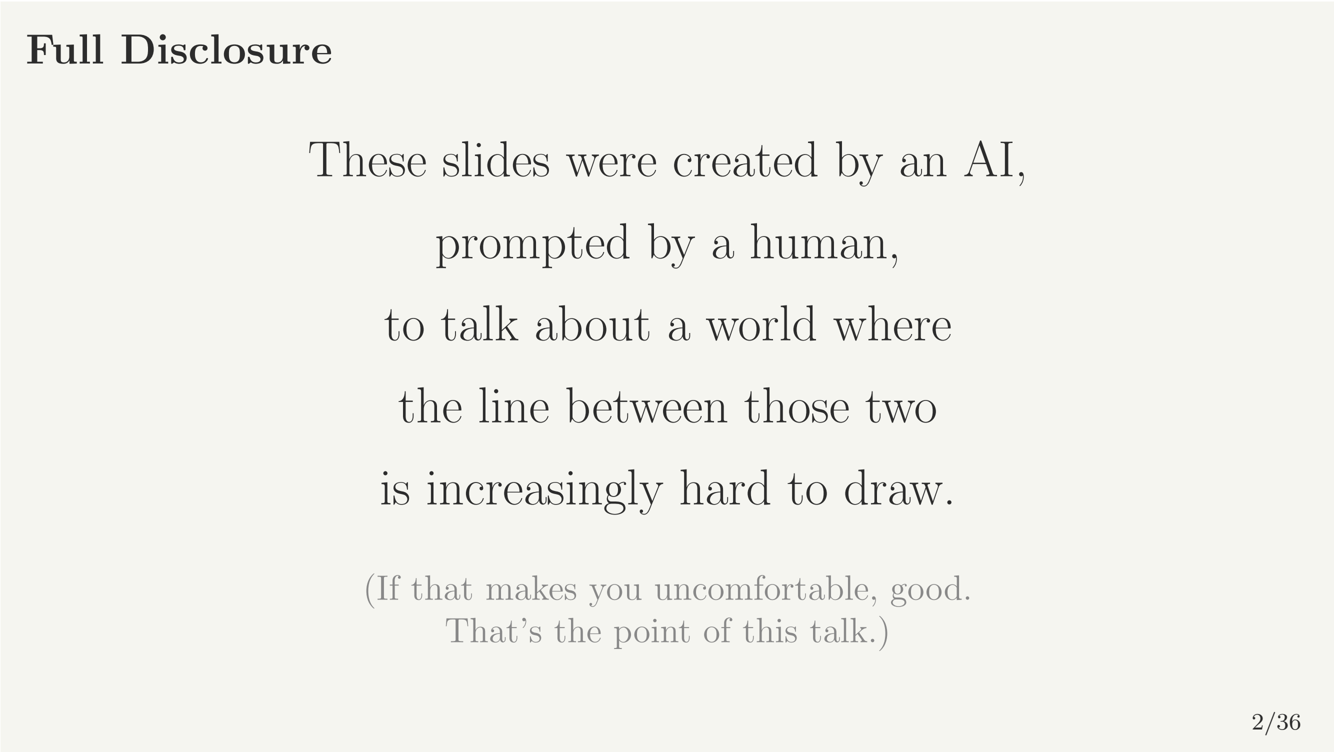

[JDT, "Predictably Uncertain", Annual Reviews Salon, March 2026]

Scenes from My Sabbatical: Indirect AI/ML

LHCP 2026: The most ambitious/controversial talk of my career (thus far…)

Machine Learning for the LHC Physics Program
A 2025–2026 stocktake

Concept & Guidance: Jesse Thaler *(MIT, IAIFI, IHES, IPhT)*
Implementation: Claude Opus 4.7 *(Anthropic, Max x20 plan, ∼72 hours of inference)*

LHCP 2026 | Paris | May 20, 2026

103 papers from the HEPML Living Review, read in full and synthesised.
Headline finding: ML stopped being a research field and started being LHC infrastructure.

In summary: six claims about how the field changed in twelve months.

1. Tools → infrastructure. LHC experiments now publish ML-driven physics; ALEPH archive too.

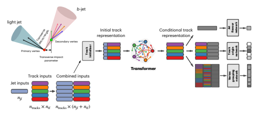

2. Two revolutions, converging. SBI methodology + foundation models. Pretrain, fine-tune, measure.

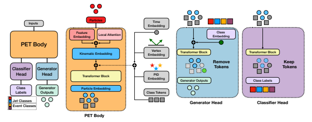

3. AI agents are the new front. 6→47 papers; four uses already differentiating; jury still out.

4. "Do we trust it?" fastest-growing. 1 in 5 recent papers is uncertainty / calibration / interpretability.

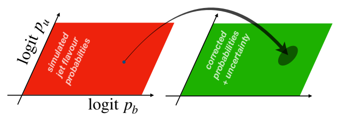

5. What's slowing down is informative. Equivariance solved into a tool; phenomenology BSM ceded to model-agnostic search.

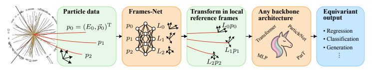

6. Theory ML crossed a capability threshold. Lattice, amplitudes, symbolic computation, structure inference.

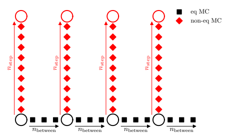

This deck is itself an exhibit of the potential — and pitfalls — of AI for the LHC physics program.

ML4LHC: 2025–2026 Stocktake | LHCP 2026 Paris | May 20, 2026 27 / 27

Demonstrating what you can do with AI at scale

[JDT, LHCP 2026, May 2026]

The rise of AI+Physics research combining physics principles with AI technologies

Machine Learning through the Lens of Physics

AI is transforming physics research, but many foundational aspects of AI can be translated into the language of physics

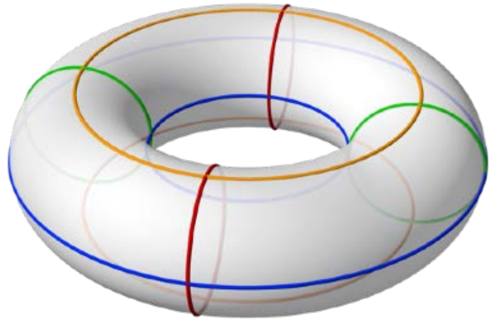

Revisiting the Modular Bootstrap

To confront thorny challenges in theoretical physics, we can reframe them as well-defined optimization problems

What Have We Actually Learned?

As we capitalize on emerging AI technologies, we still need to insist on robustness and reproducibility

Case Study: Two-Dimensional Conformal Field Theories

Highly constrained mathematical framework that we know a lot about
(though not my area of expertise…)

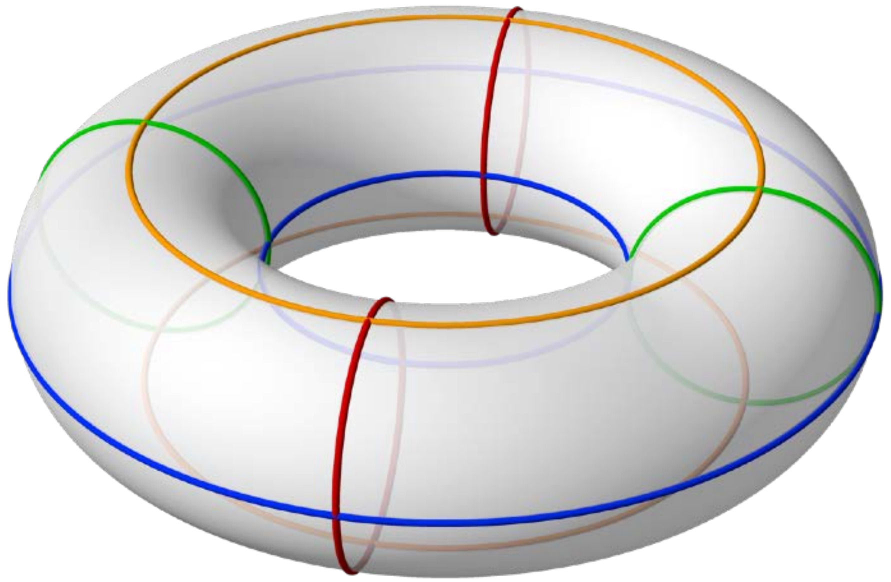

Posing a research question to Claude

Please identify a two-dimensional conformal field theory with central charge in the range of 1.00 to 1.15.

The $\mathbf{Z_5}$ **parafermion CFT** (equivalently, the $SU(2)_5/U(1)$ coset) has central charge $\mathbf{c = 8/7 \approx 1.143}$.

[Fateev, Zamolodchikov, Sov. Phys. JETP 1985]

Only four decades too late…

[h/t to Siddharth Mishra-Sharma for the Claude credits!]

Case Study: Two-Dimensional Conformal Field Theories

> Ok, great. So let me revise my question...
>
> Please identify a two-dimensional conformal field theory with central charge in the range of 1.00 to 1.14.

Ah, you've narrowed the window to exclude the Z_5 parafermion ($c = 8/7 \approx 1.1429$).

To my knowledge, there is no known unitary 2D CFT with central charge strictly in the range (1.00, 1.14). The theories cluster at the boundaries:

- **Below:** $c = 1$ theories (free boson, orbifolds)
- **Above:** Z_5 parafermion at $c = 8/7 \approx 1.143$, N=1 minimal model (m=5) at $c = 81/70 \approx 1.157$

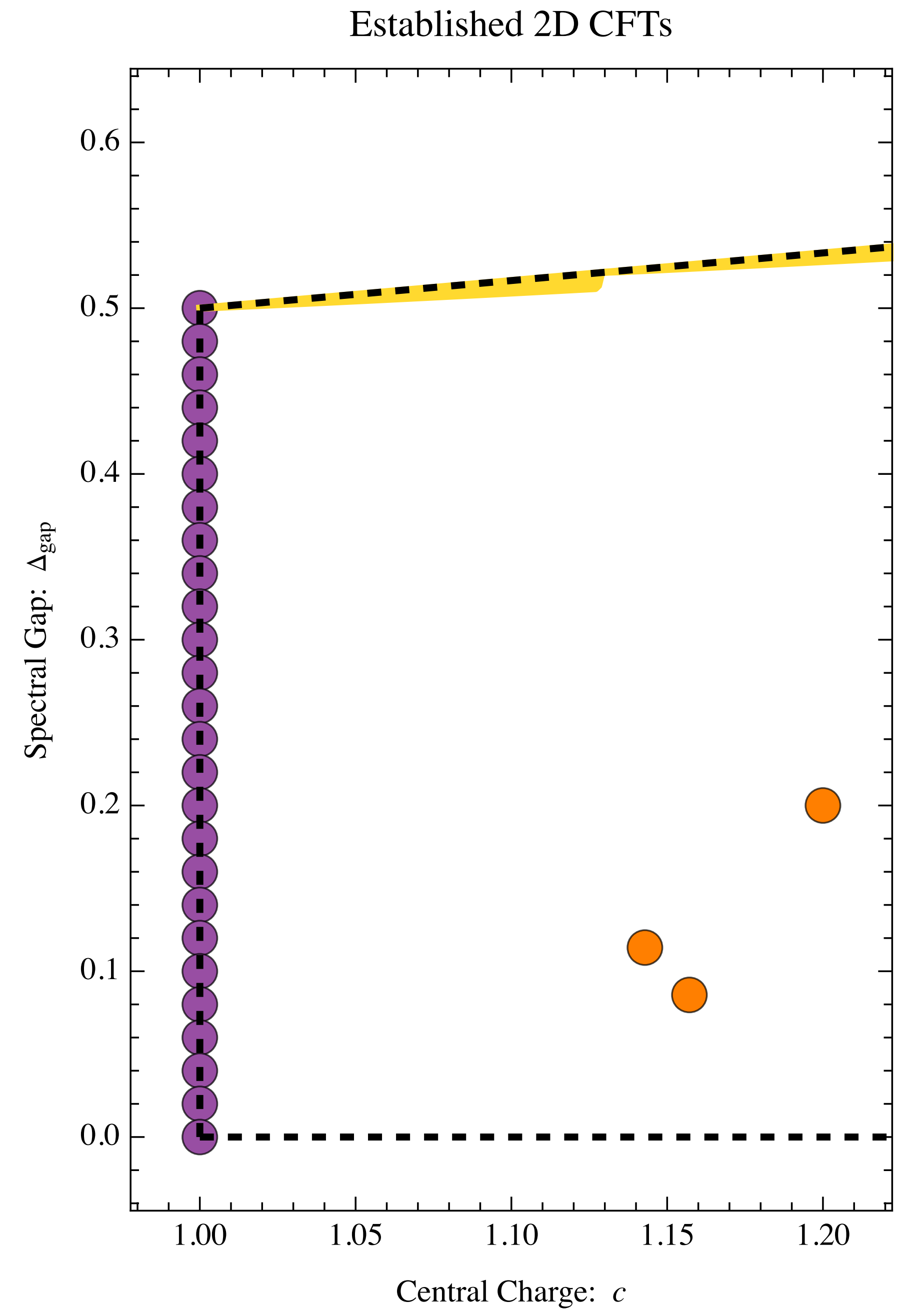

Are there any other 2D CFTs in this range?

(And could I convince the experts if I found one? Could I convince myself?!)

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
bounds from Hellerman, JHEP 2011; Friedan, Keller, JHEP 2013; Collier, Lin, Yin, JHEP 2018; Fitzpatrick, Li, JHEP 2024]



Strategies to Study/Discover CFTs

Ruling out theories (“dual”) is a lot easier than ruling them in (“primal”)

Indirect AI/ML?

Could you think about it a bit more and come up with a new theory that doesn't exist in the literature?

The honest answer: I can't construct a specific theory in (1.00, 1.14). The known RCFT constructions—cosets, W-algebra minimal models, orbifolds—all seem to have discrete spectra of central charges that skip this interval.

This could mean:

1. There's a deep reason (modular bootstrap + unitarity) forbidding unitary RCFTs here
2. There's an irrational CFT we haven't characterized
3. There's an exotic construction nobody has tried

At least it refuses to hallucinate…

Direct AI/ML?

Classification

[Chen, He, Lal, Zaz, arXiv 2020; Kuo, Seif, Lundgren, Whitsitt, Hafezi, PRR 2022]

Reinforcement Learning

[Kántor, Niarchos, Papageorgakis, PRL 2022, PRD 2022; + Richmond, PRD 2023]

Generative Modeling

[Laio, Valenzuela, Serone, PRD 2022]

Gradient-Based Methods

[Reehorst, Rychkov, Simmons-Duffin, Sirois, Su, van Rees, SciPost 2021]
[Afkhami-Jeddi, JHEP 2022]

Genetic Algorithms

[Huang, Lee, Liao, Rumbutis, arXiv 2025]

Spectral Bias Studies

[Ghosh, Kumar, Niarchos, Stergiou, arXiv 2026a, arXiv 2026b]

[Please let me know if I missed your ML-related CFT paper!]

[see reviews of dual approach to conformal bootstrap in Poland, Rychkov, Vichi, RMP 2019; Rychov, Su, RMP 2024; alternative numerical approach in Gliozzi, PRL 2013, JHEP 2016]

Warm-Up Problem: Modular Bootstrap

Necessary but not sufficient condition for unitary 2D CFT to exist

Example of more general bootstrap philosophy of defining theoretical objects by consistency conditions

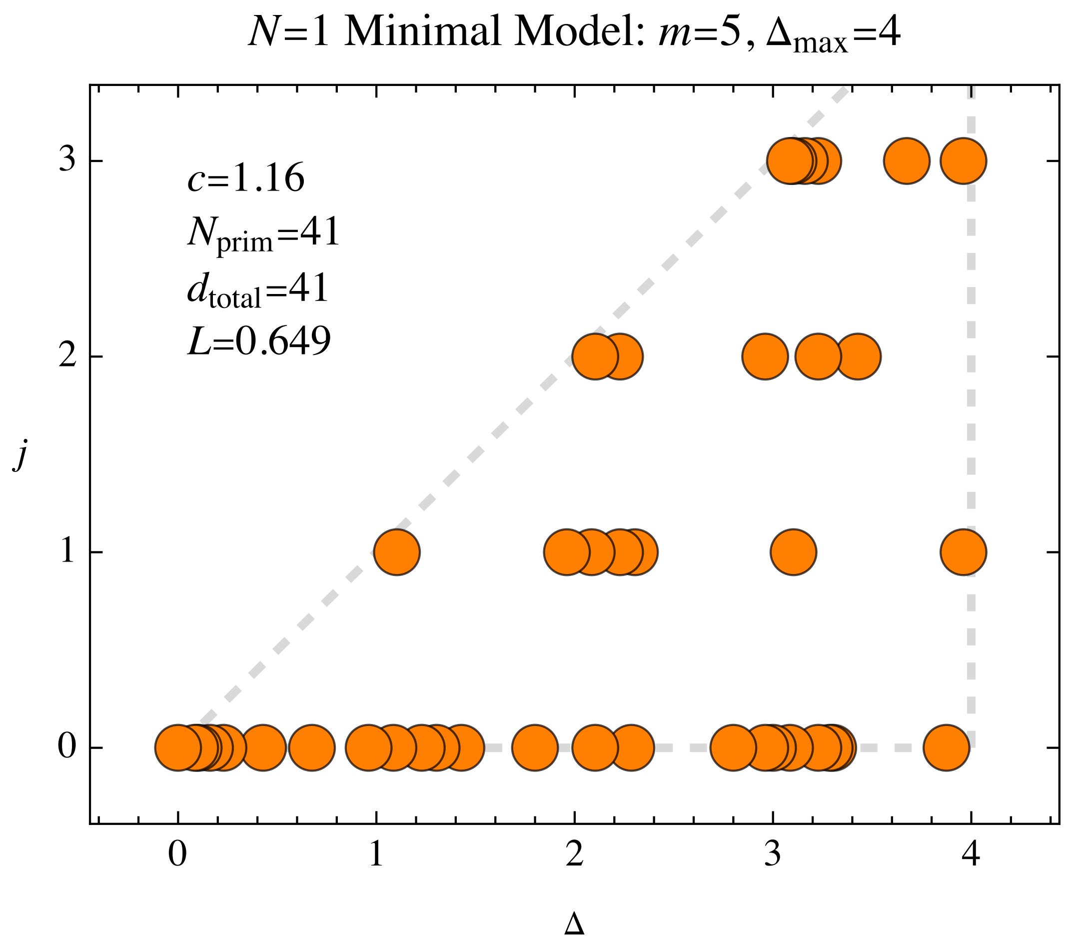

Spectrum of (Virasoro Primary) Operators:

Dimension:	Δ	$\Delta \geq j \geq 0$ assuming parity for $-j$
Spin:	j	$d > 0$
Degeneracy:	d	j, d are integers

(Euclidean) Partition Function: $Z(\tau; c, \{\Delta_a, j_a, d_a\})$

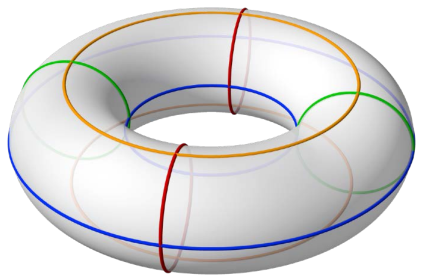

Modular Invariance:
$$\boxed{Z(\tau) = Z(-1/\tau)}$$

[Cardy, NPB 1991; Hellerman, JHEP 2011; Friedan, Keller, JHEP 2013; Collier, Lin, Yin, JHEP 2018; …]

Machine Learning without Neural Networks

Nothing stops you from optimizing any parametrized function you want, as long as its derivatives aren't too pathological

Euclidean Torus: $\tau = \tau_1 + i\,\tau_2$ $\qquad$ $\bar{\tau} = \tau_1 - i\,\tau_2$

Squared Nomes: $q = e^{2\pi i \tau}$ $\qquad$ $\bar{q} = e^{-2\pi i \bar{\tau}}$

Conformal Weights: $h_a = \dfrac{\Delta_a + j_a}{2}$ $\qquad$ $\bar{h}_a = \dfrac{\Delta_a - j_a}{2}$

Virasoro Characters: $\chi_{h=0}(\tau) = \dfrac{q^{-\frac{c-1}{24}}}{\eta(\tau)}(1-q)$ $\qquad$ $\chi_{h>0}(\tau) = \dfrac{q^{h-\frac{c-1}{24}}}{\eta(\tau)}$ $\qquad$ and barred versions

(Reduced) Partition Function: $Z_{\text{red}}(\tau) = \sqrt{\tau_2}\,|\eta(\tau)|^2\,Z(\tau)$ $\qquad$ $Z(\tau) = \sum_a d_a\,\chi_{h_a}(\tau)\,\overline{\chi}_{\bar{h}_a}(\bar{\tau})$

“Thinking Like a Machine”

Modular bootstrap as “self-supervised learning”

✓ **Training Data:**

MC Integration over τ

$$L = \int_{\mathscr{F}} \frac{d\tau\, d\bar{\tau}}{(\mathrm{Im}\,\tau)^2} \mathscr{L}$$

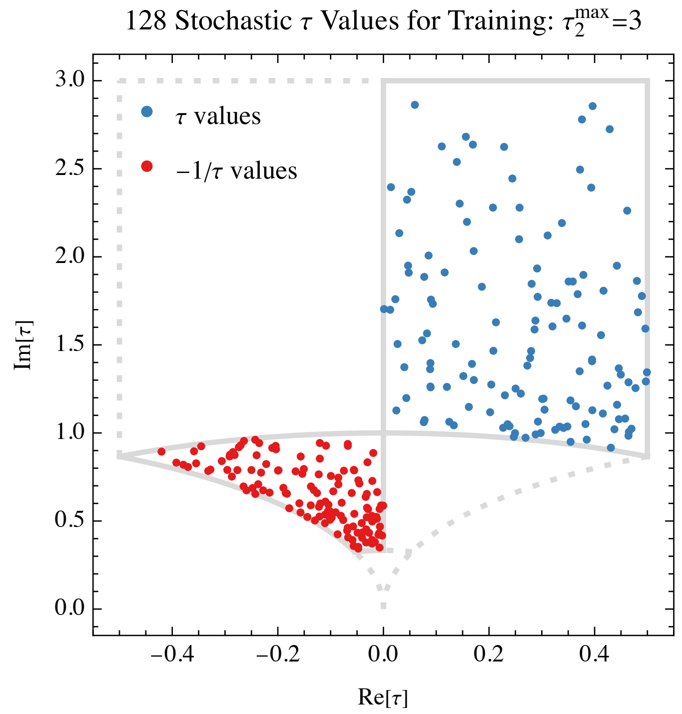

✗ **Loss Function(al):**

Modular Invariance

$$\mathscr{L} = \big(Z(\tau) - Z(-1/\tau)\big)^2$$

✓ **Learnable Function:**

Partition Function

$Z(\tau; c, \{\Delta_a, j_a, d_a\})$ with $\Delta_a < \Delta_{\max}$ as parameters

✗ **Optimizer:**

Gradient Descent

$$\Delta_a^{(i)} = \Delta_a^{(i-1)} - \eta \frac{dL}{d\Delta_a}$$

Unfortunately, this approach fails dramatically at finding anything close to physical…

[similar philosophy to physics-informed machine learning for PDEs: Karniadakis, Kevrekidis, Lu, Perdikaris, Wang, Yang, Nature Reviews 2021]
[see ML approach to S-matrix bootstrap in Dersy, Schwartz, Zhiboedov, JHEP 2024; Niarchos Papageorgakis, PRD 2024; Gumus, Leflot, Tourkine, Zhiboedov, JHEP 2025, arXiv 2026]

“Thinking Like an Experimental Physicist”

What are the Uncertainties on your Loss Function?

Use “χ^2” instead of MSE: $$\mathscr{L} = \left(\frac{Z(\tau) - Z(-1/\tau)}{\sigma(\tau)} \right)^2$$

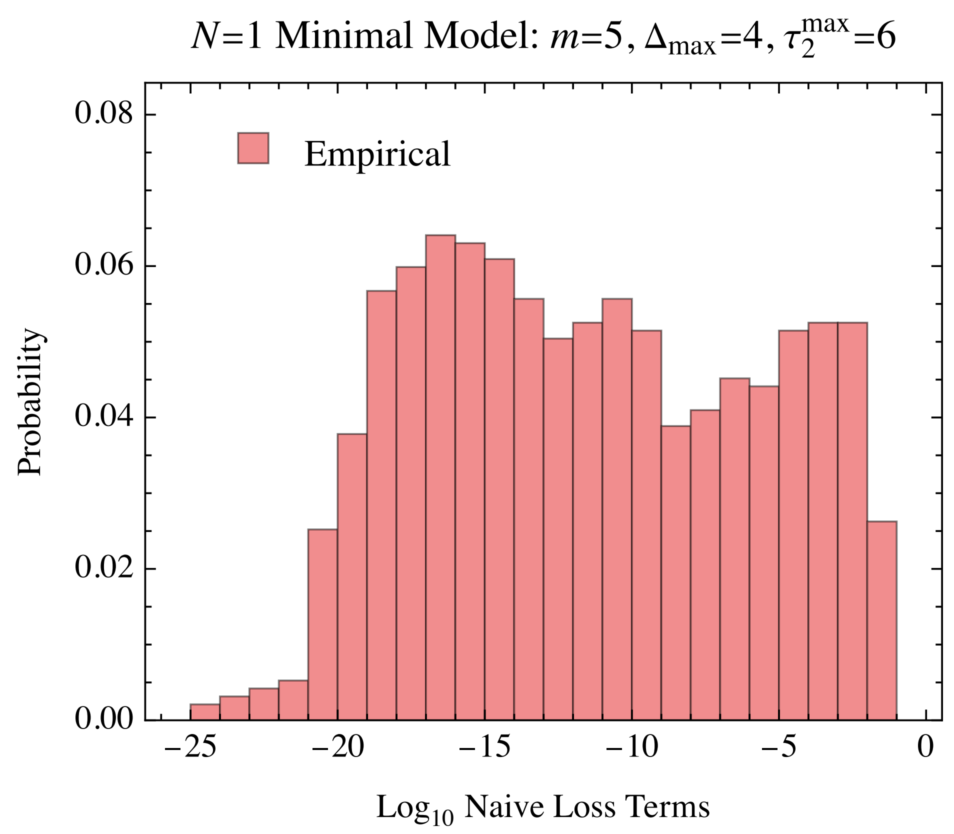

$\Rightarrow$

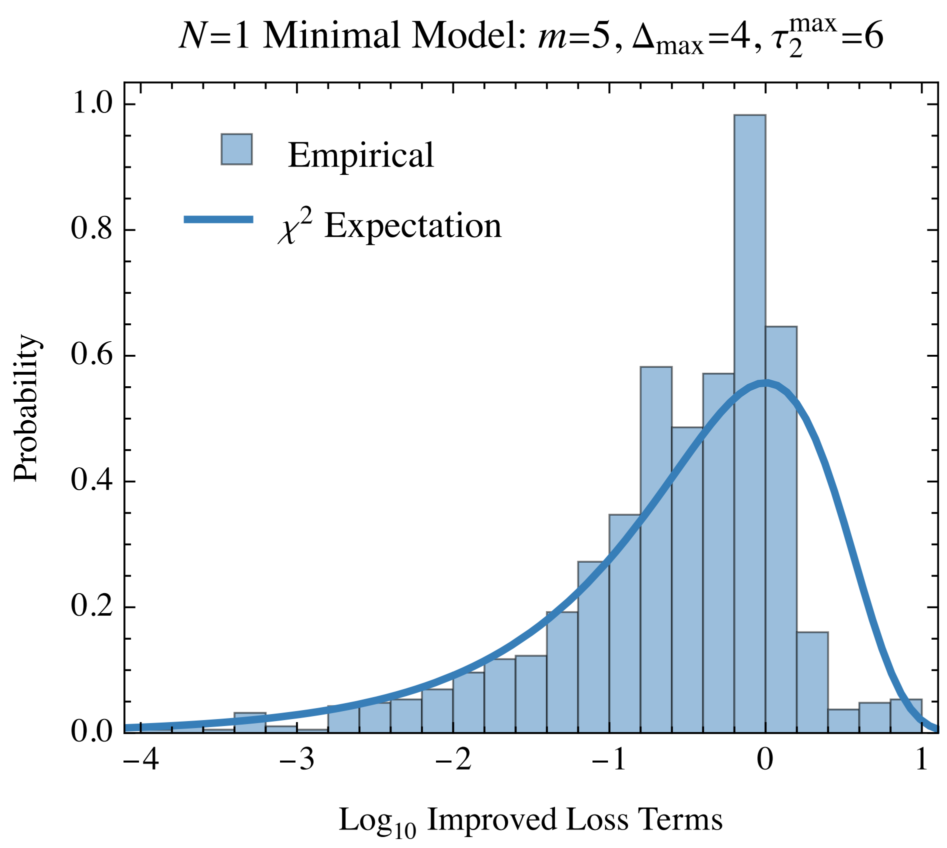

Remarkable that this distribution comes from purely theoretical data

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026]

"Thinking Like an Experimental Physicist"

Needed a robust method to quantify uncertainties

Collaborating with Claude helped me identify connections in the literature

> We previously discussed a strategy for taking a known CFT with c = 1 and deforming its partition function to mimic the Cardy scaling behavior of a CFT with c > 1, in order to estimate uncertainties associated with truncating the primary spectrum. Can you remind me of that formula and where it first appeared in the literature?

To shift a partition function from central charge c to a target c', use:

$$\tilde{Z}_{c'}(\tau) = (\operatorname{Im}\tau)^{(c'-c)/2} \cdot |\eta(\tau)|^{-(c'-c)} \cdot Z_c(\tau)$$

The $|\eta|$ factor shifts the vacuum energy to give the correct $q^{-c'/24}$ scaling, while the $(\operatorname{Im}\tau)^{1/2}$ factor preserves modular covariance under $S : \tau \to -1/\tau$.

A similar construction appears in **Benjamin, Collier, Fitzpatrick, Maloney, and Perlmutter, "Harmonic analysis of 2d CFT partition functions" (arXiv:2107.10744)**, where they define the primary partition function $Z_p = (y^{1/2}/|\eta|^2)^c Z$ to enable spectral decomposition on the modular fundamental domain.

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026]

“Thinking Like an Mathematical Physicist”

What is the Geometry of your Loss Landscape?

Use quasi-second order methods instead of SGD: $$\Delta_a^{(i)} = \Delta_a^{(i-1)} - g_{ab}\frac{dL}{d\Delta_b}$$

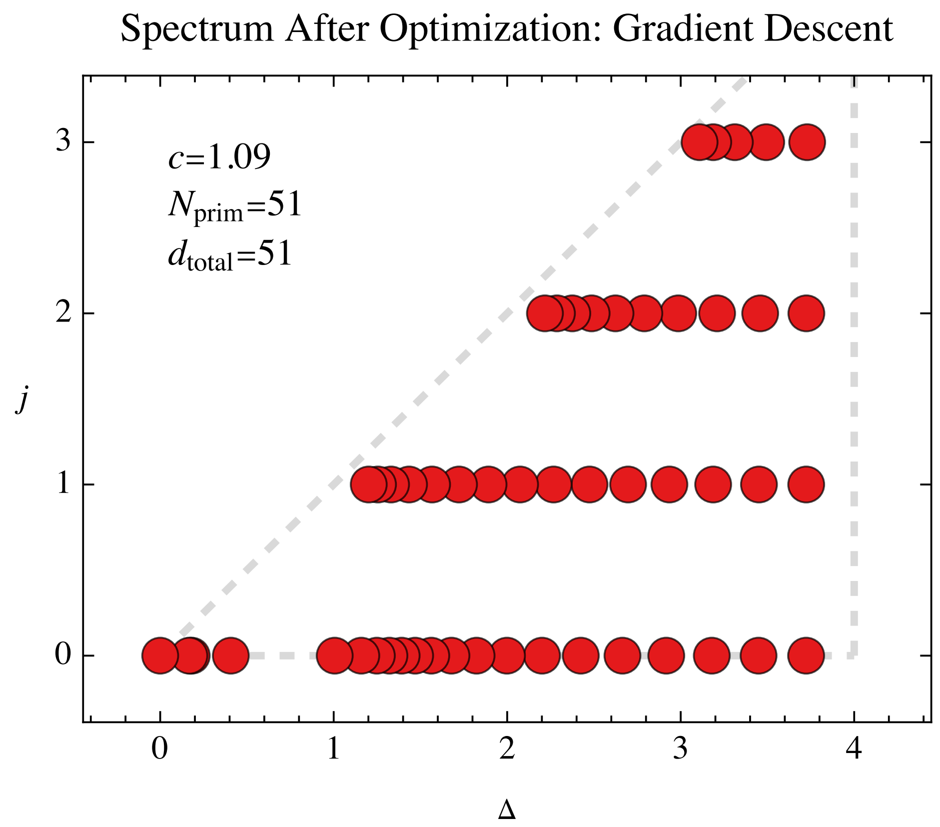

$\Longrightarrow$

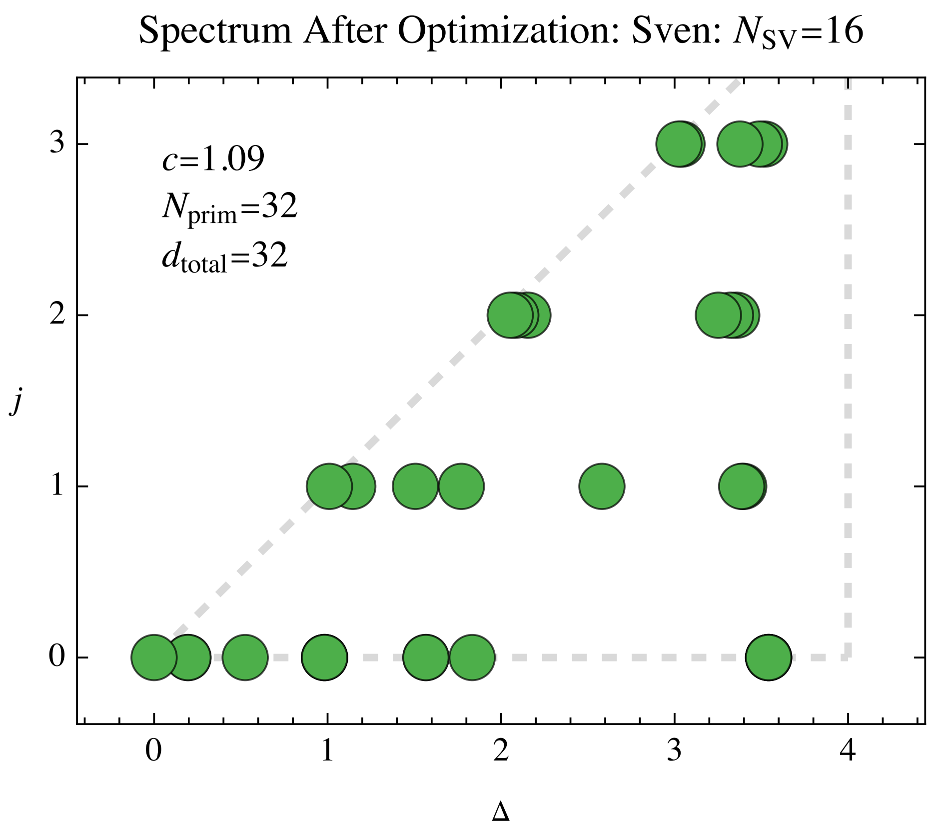

Code name: “Sven”
Singular Value dEsceNt

Now available in GitHub!

Ask me offline: basically a generalized/regularized version of Gauss-Newton

[Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Descending into the Modular Bootstrap

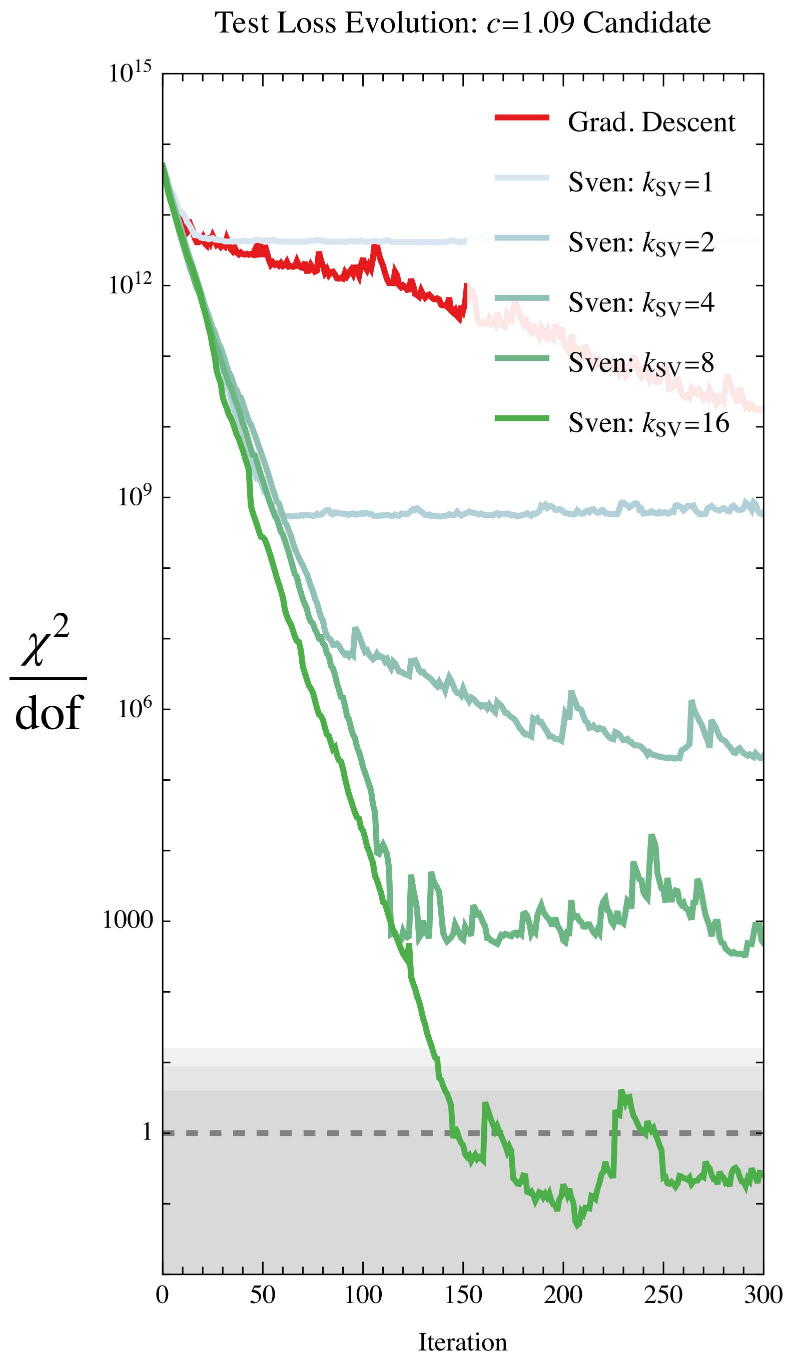

Tight constraints on individual operators…

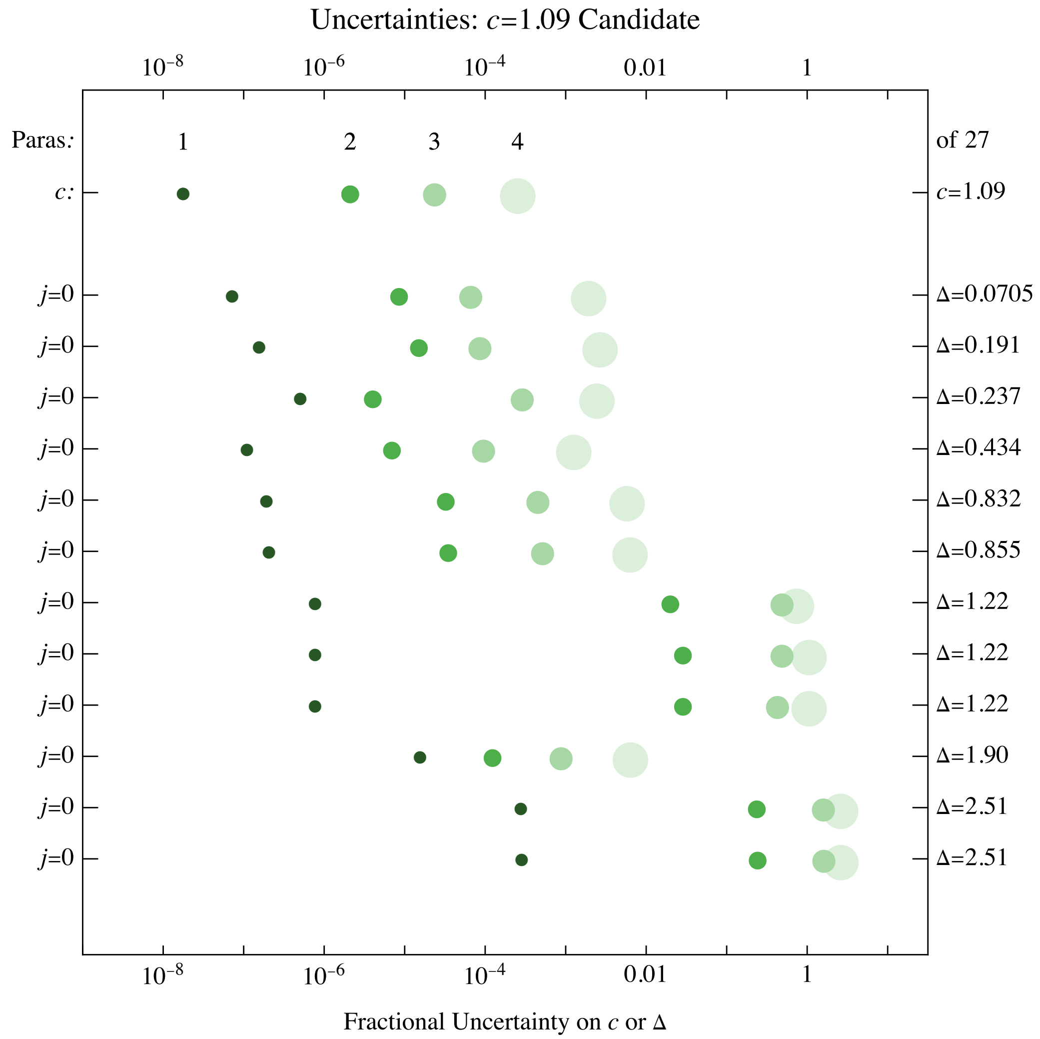

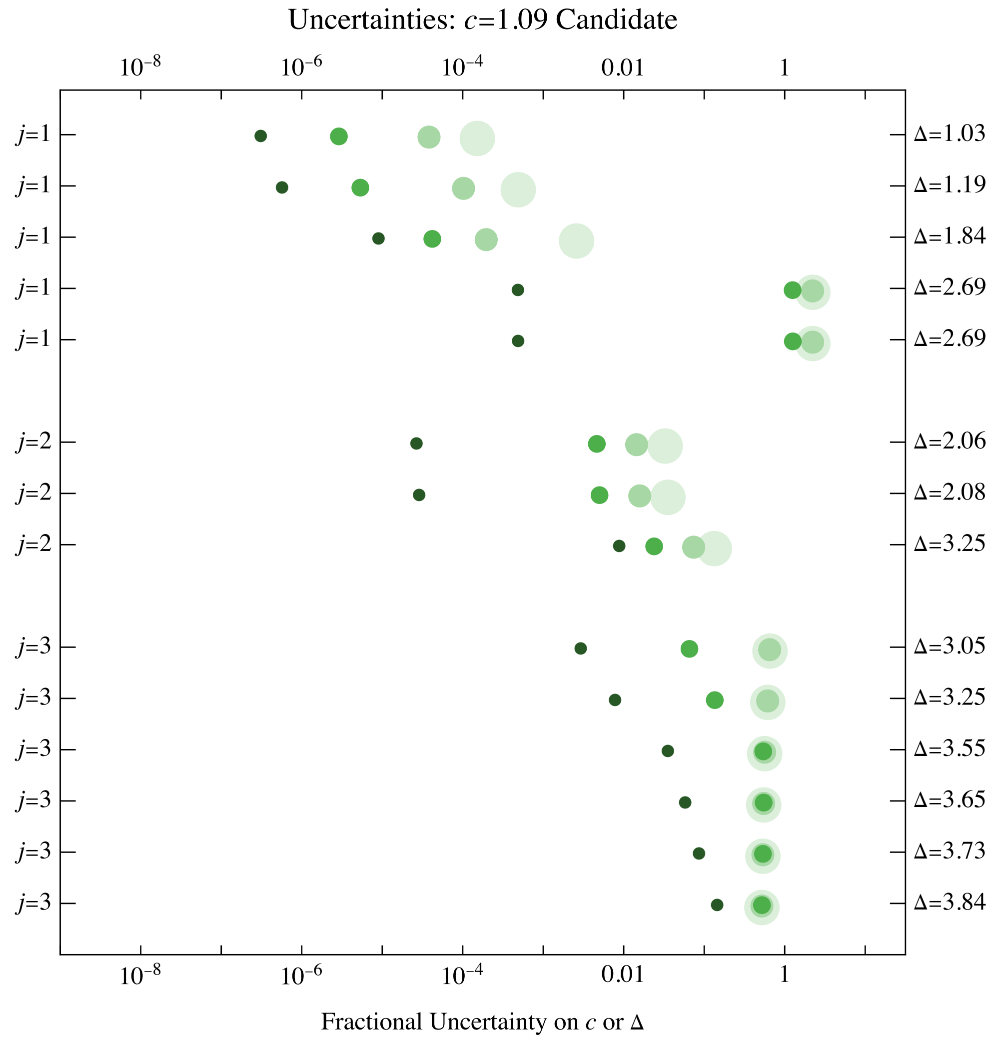

…almost certainly spanning a continuous space of solutions!

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026; Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Updating the Space of Candidate 2D CFTs

Remember: Modular invariance is necessary but not sufficient

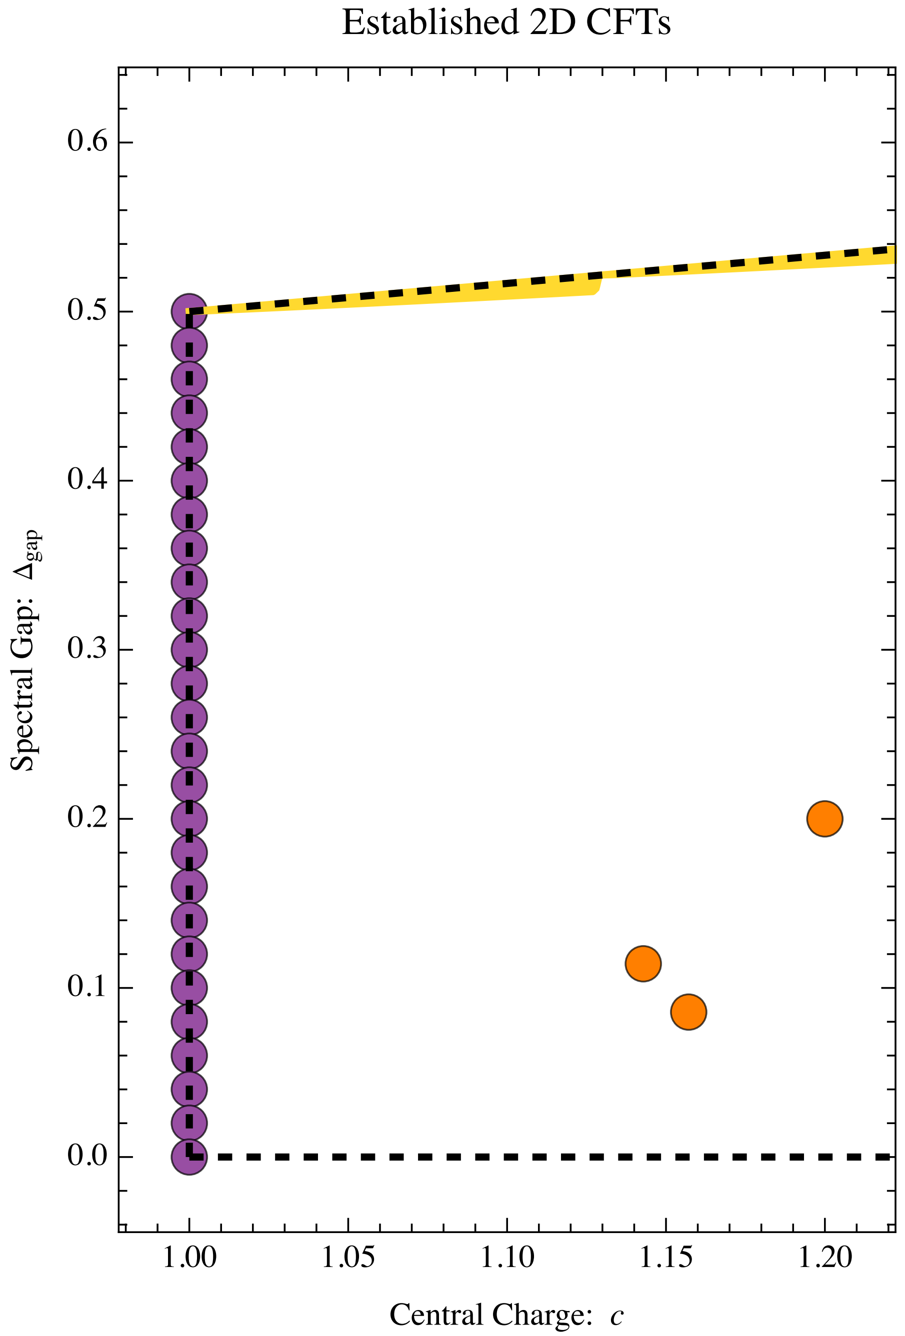

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Updating the Space of Candidate 2D CFTs

Remember: Modular invariance is necessary but not sufficient

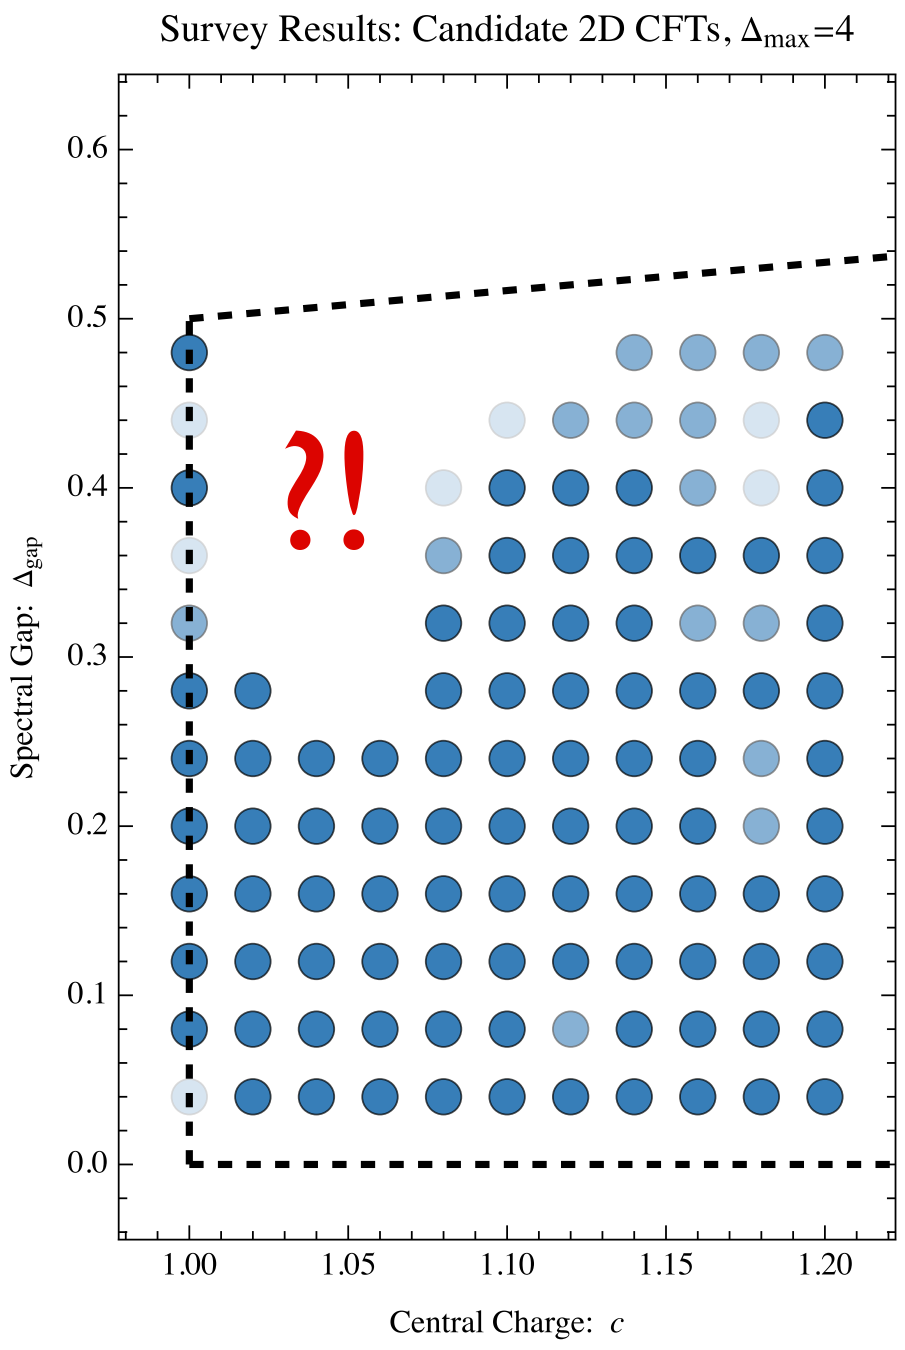

Highly suggestive of an obstruction…

I don't know how we would have gotten this result without AI/ML methods!

Direct AI/ML: Reframe modular bootstrap in optimization language

Indirect AI/ML: Reason through conceptual issue of estimating uncertainties

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Updating the Space of Candidate 2D CFTs

Remember: Modular invariance is necessary but not sufficient

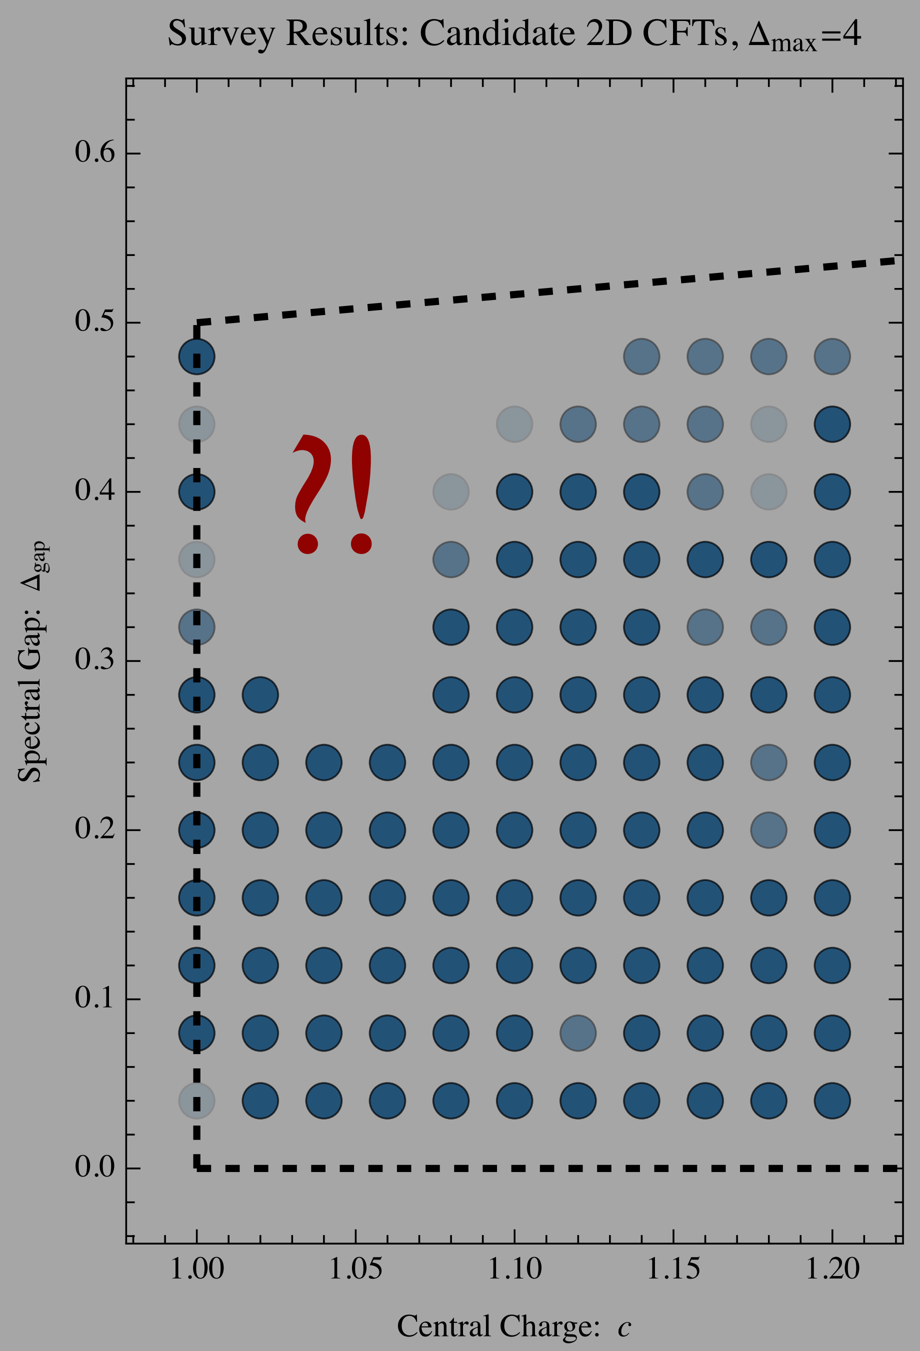

But I also don't know how we would have gotten this result without physics input!

Identify tractable problem of relevance to experts in the field

Develop new optimization algorithms based on physics intuition

Pursue statistical robustness even in a purely theoretical setting

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026;
Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

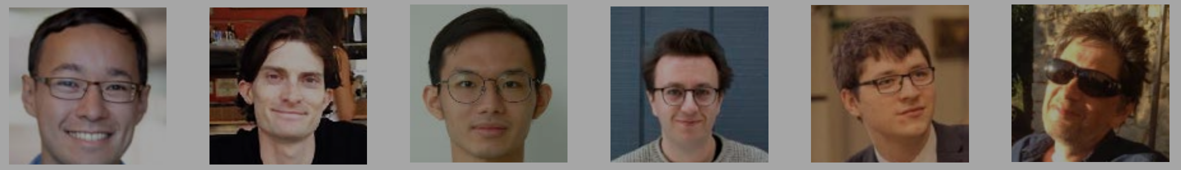

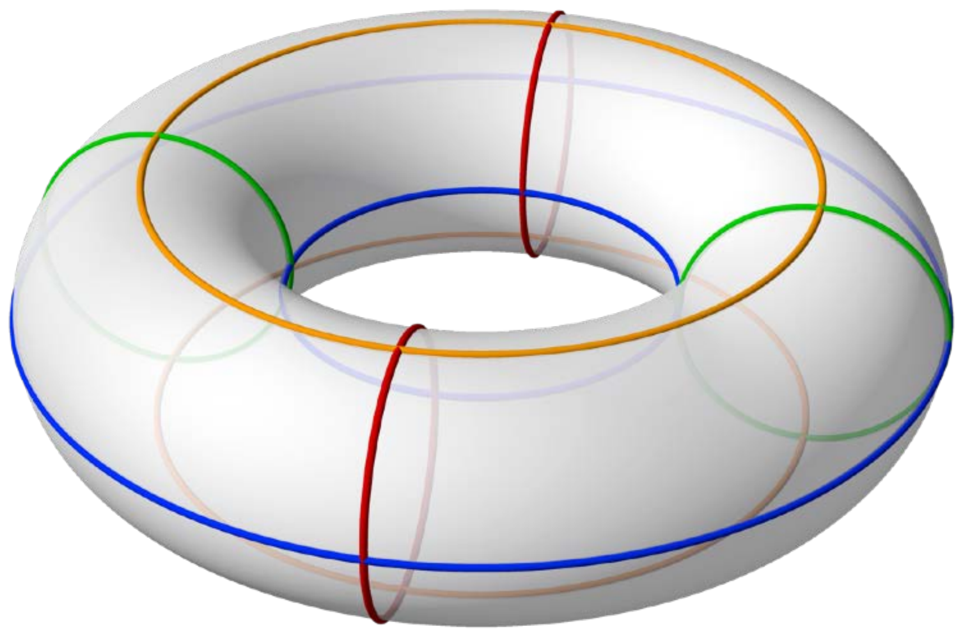

Advancing theoretical physics through a methodological merging of physical reasoning and computational algorithms

Machine Learning through the Lens of Physics

AI is transforming physics research, but many foundational aspects of AI can be translated into the language of physics

Revisiting the Modular Bootstrap

To confront thorny challenges in theoretical physics, we can reframe them as well-defined optimization problems

What Have We Actually Learned?

As we capitalize on emerging AI technologies, we still need to insist on robustness and reproducibility

Theoretical Research via Numerical Optimization

If your problem can be stated as…

$$\boxed{\textit{Find parameters } \theta \textit{ such that } f(x;\theta) = 0 \textit{ for all } x}$$

…then it can be attacked with ML-style optimization

The lesson from our study is that you have to know
how close to "0" you need to be to declare success

This is essentially a problem of uncertainty quantification and
I don't know any shortcut other that detailed robustness studies

"What Have We Actually Learned?"

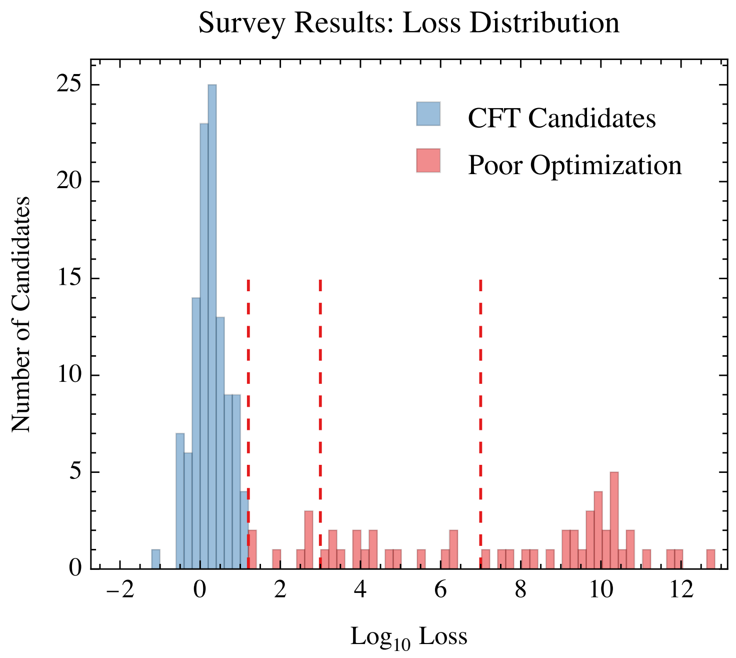

With very high confidence:

We found approximate solutions to the truncated modular bootstrap equations

However, we cannot claim to have identified new CFTs

In fact, these examples are almost certainly not CFTs

The modular bootstrap does not constrain OPE coefficients

Moreover, we cannot even claim to have excluded possible CFTs!

In the primal approach, the absence of evidence is not evidence of absence

Maybe Sven just isn't a powerful enough optimizer for this problem

“What Have We Actually Learned?”

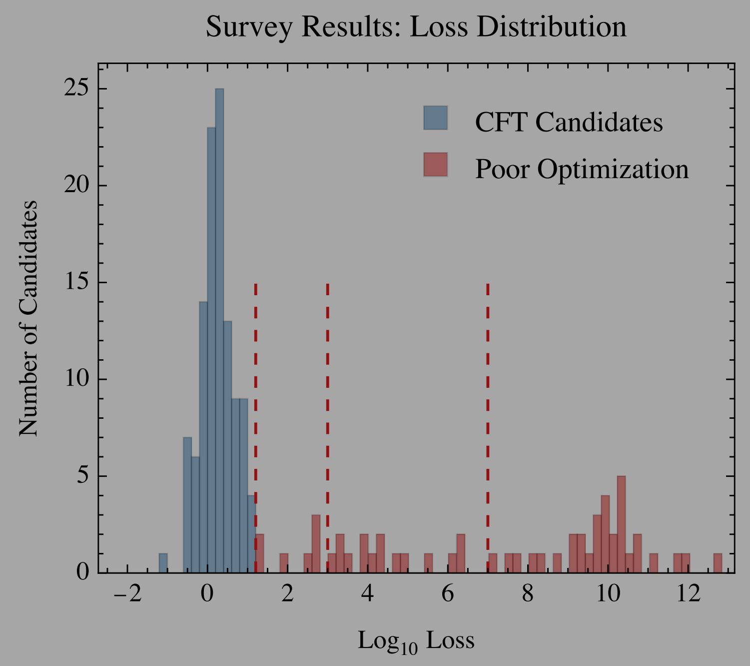

With very high confidence:

We found approximate solutions to the truncated modular bootstrap equations

What we can do is perform numerous studies to increase our confidence that this apparent obstruction is robust and reproducible

I do not claim that these are proofs of anything, but there were numerous opportunities for these checks to fail

Visual Inspection: (Non-)Candidate 2D CFT Spectra at c = 1.08

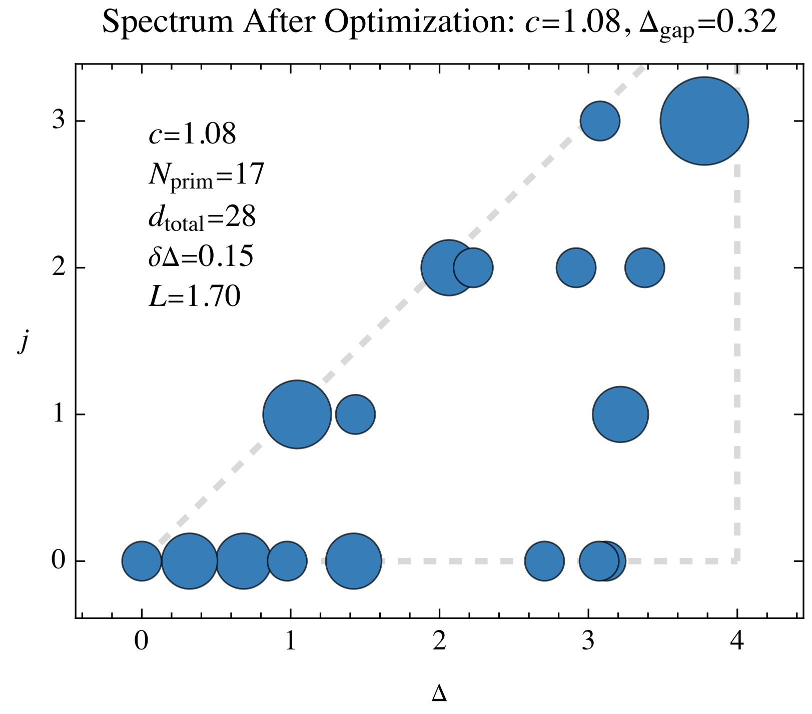

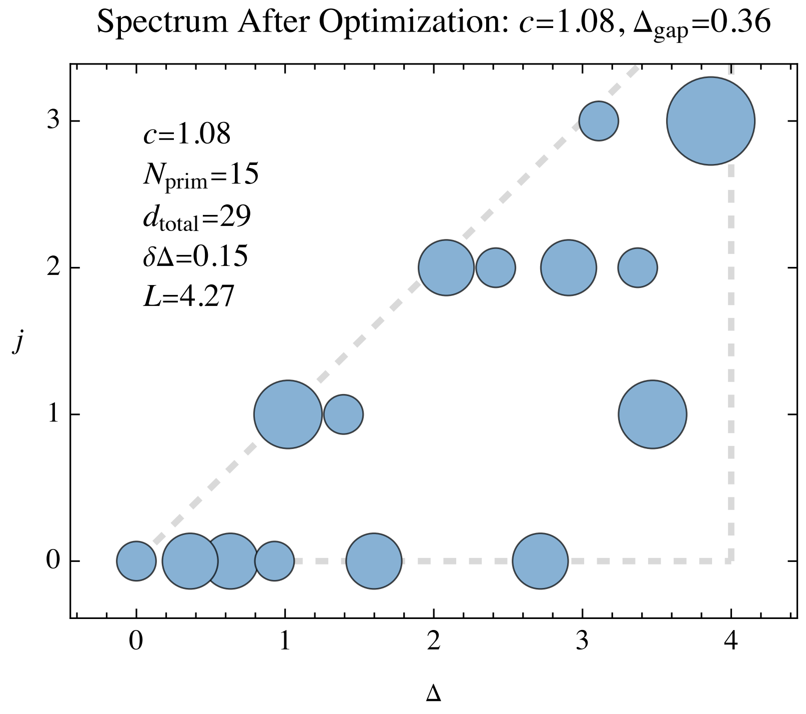

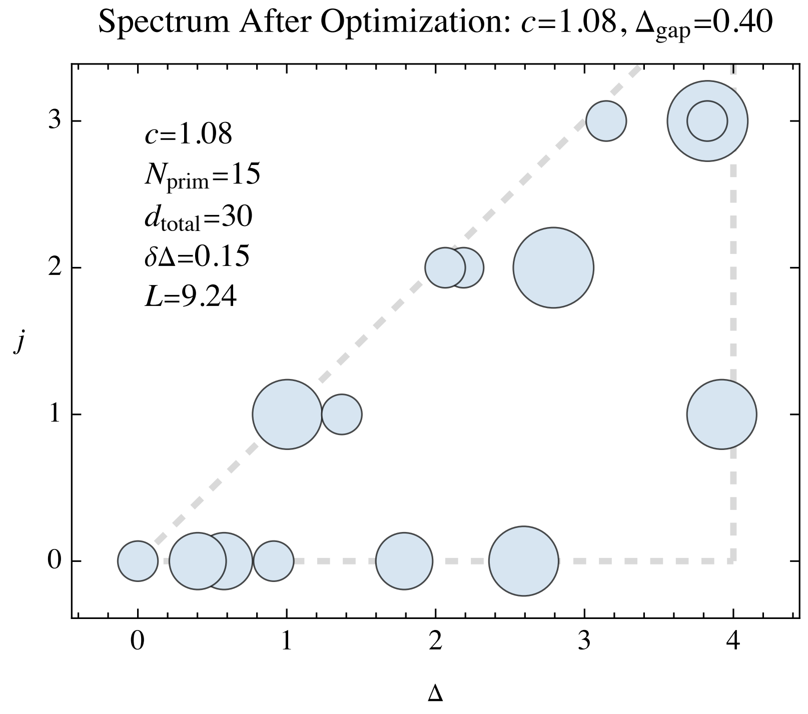

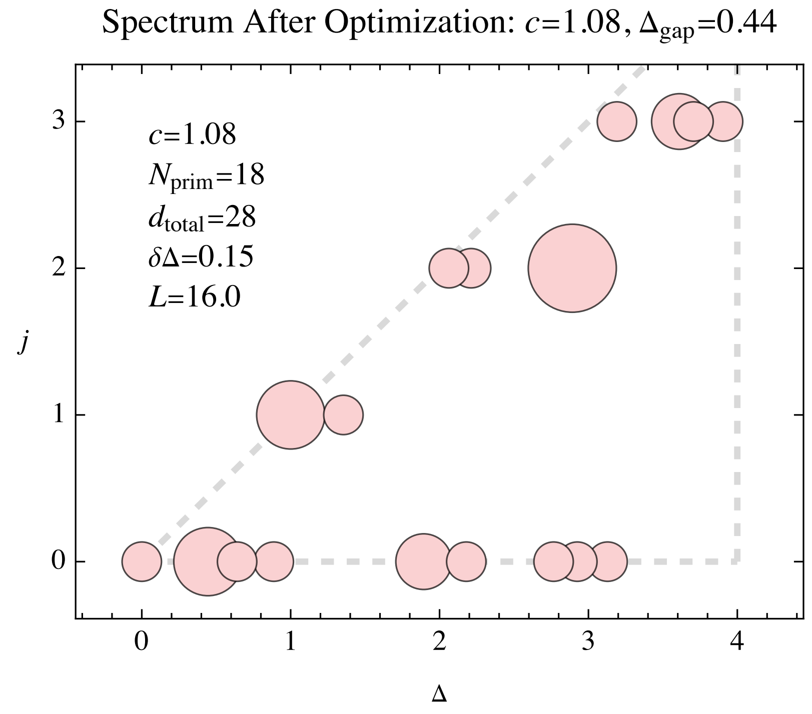

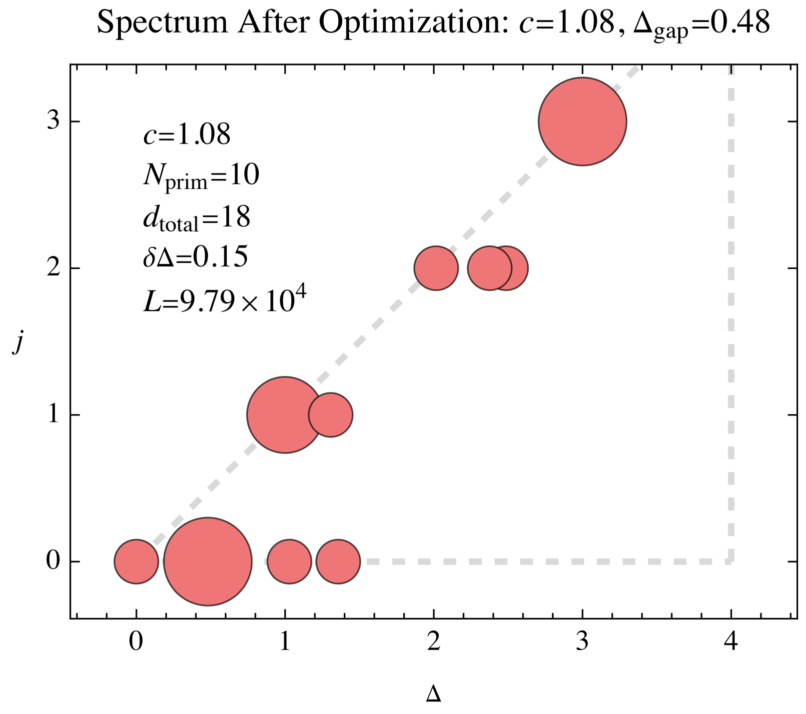

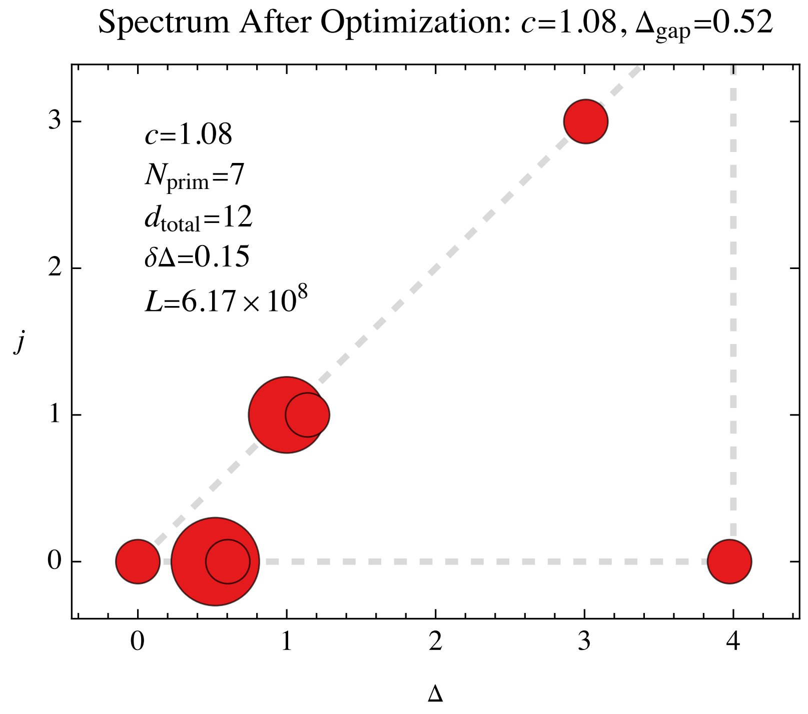

Checking the Definition of Uncertainties

Strategy: Take an known CFT partition function with central charge c_0
Deform it to make it have effective central charge c
Compute the uncertainty from truncating spectrum at $\Delta_{\max}$

$$\sigma_{\text{trunc}}(\tau; c; \Delta_{\max})^2 = \left(Z_{\text{deformed}}(i\,|\tau|; c) - Z^{\text{trunc}}_{\text{deformed}}(i\,|\tau|; c; \Delta_{\max})\right)^2 + (\tau \to \tau_S)$$

Claude's Deformation: $Z_{\text{deformed}}(\tau; c) = \left(\sqrt{\tau_2}\,|\eta(\tau)|^2\right)^{c_0 - c} Z_{\text{exact}}(\tau; c_0)$

Smear out vacuum to reproduce Cardy scaling; never a valid CFT

Human Rebuttal! $Z_{\text{deformed}}(\tau; c) = \Big(Z_{\text{exact}}(\tau; c_0)\Big)^{\alpha}$ $\qquad \alpha = \dfrac{c}{c_0}$

Take α copies of a CFT; valid theory for integer α

Humans win on numerical accuracy! Though AI wins on computational efficiency…
And it makes absolutely no difference to the story

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026]

Checking the Impact of Integrality

“Allowing for non-integer degeneracies, your optimizer should be able to saturate the Hellerman-Collier-Lin-Yin bound”

Main Result: $\Delta_{\max} = 4$

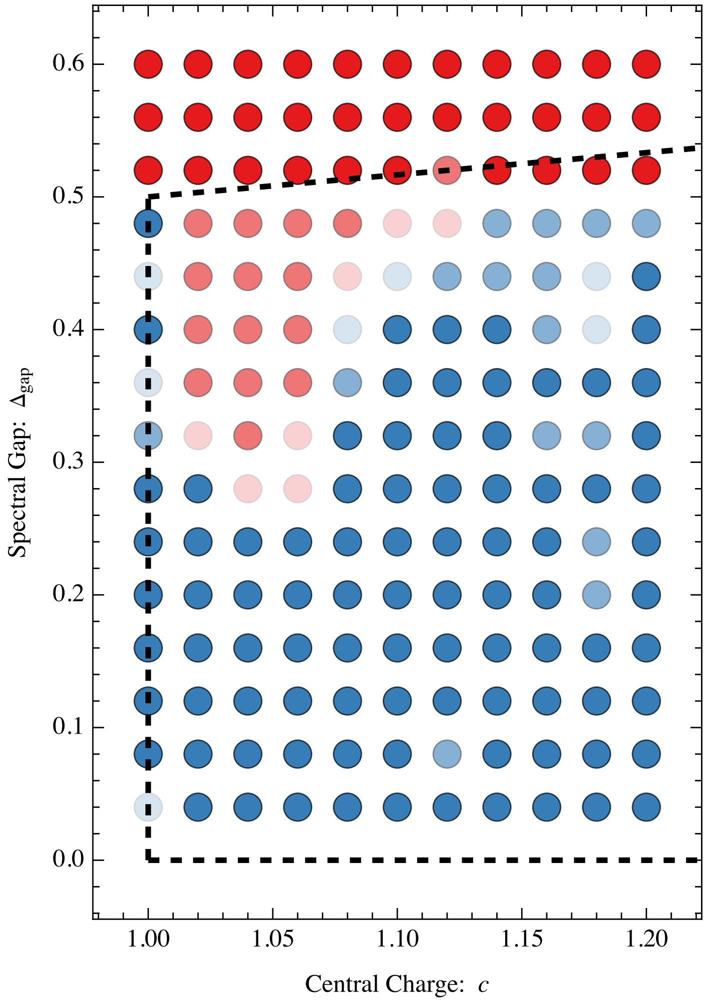

$\Rightarrow$

Floating Degeneracy

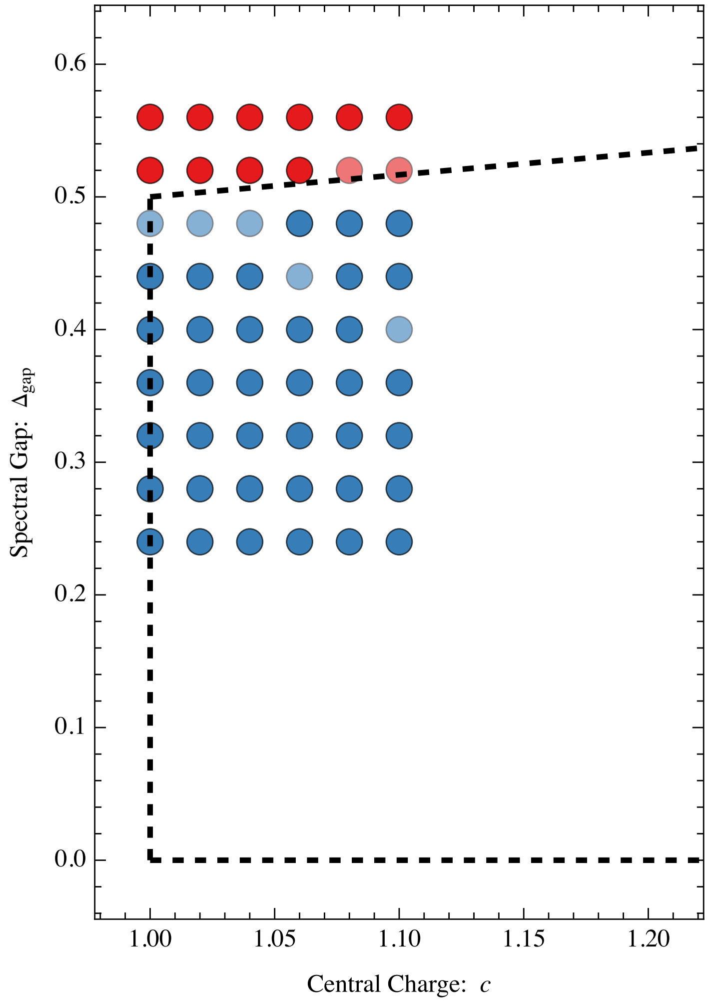

$\leftarrow \Delta_{\text{gap}} \leq \frac{c}{6} + \frac{1}{3}$

With additional freedom, Sven does find solutions that fill the allowed space

Integrality appears essential for the obstruction

[bounds from Hellerman, JHEP 2011; Collier, Lin, Yin, JHEP 2018; Benjamin, Fitzpatrick, Li, JDT, arXiv 2026; Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Checking the Impact of the Truncation Scale

"If you lower the truncation scale, the uncertainties should be larger, so it should be easier to find solutions to the bootstrap equations"

Main Result: $\Delta_{\max} = 4$

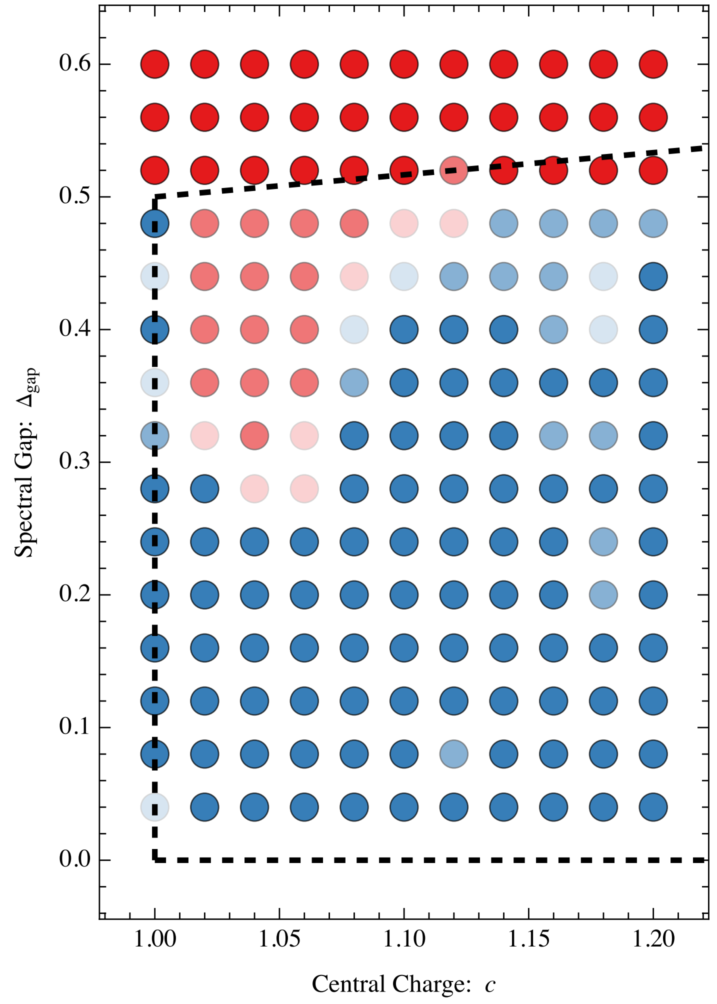

$\Rightarrow$

Lowering: $\Delta_{\max} = 3$

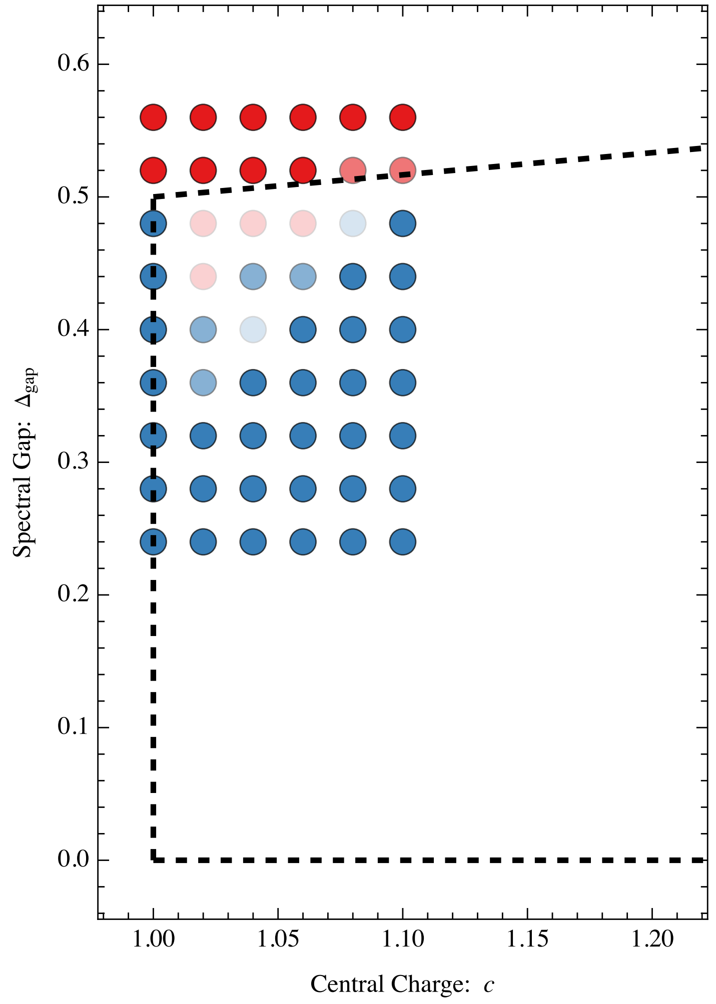

Loosening the constraints does yield more apparent solutions

High dimension (but low twist) operators appear essential for the obstruction

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026; Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Pushing Sven to the Max

"If you raise the truncation scale, the uncertainties should be smaller, so it should be harder to find solutions to the bootstrap equations"

Main Result: $\Delta_{\max} = 4$

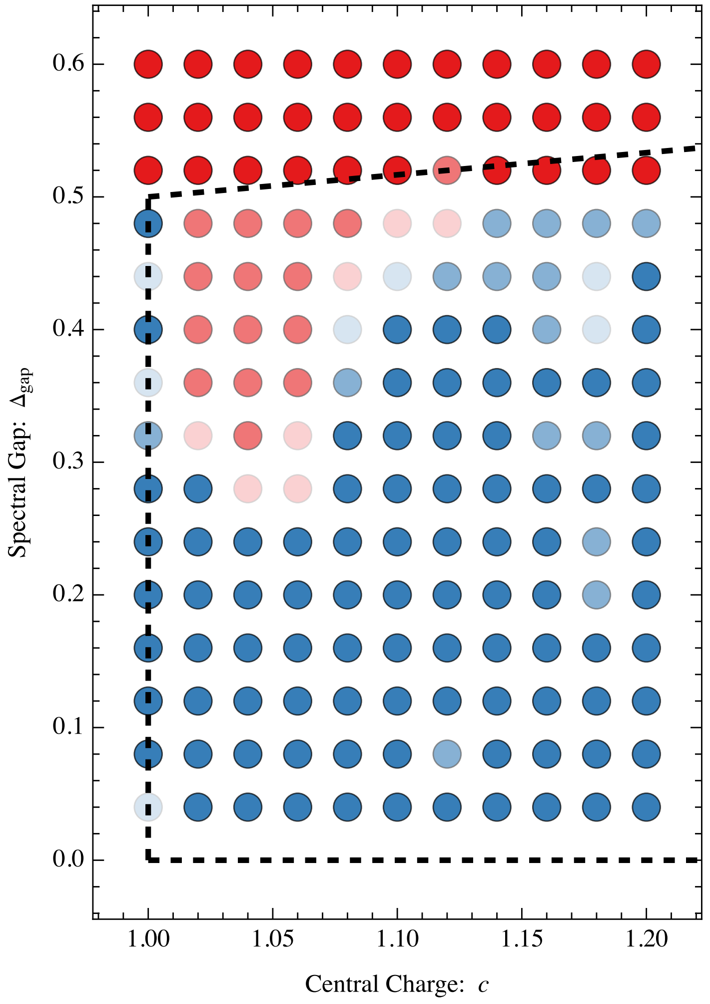

$\Rightarrow$

Raising: $\Delta_{\max} = 5$

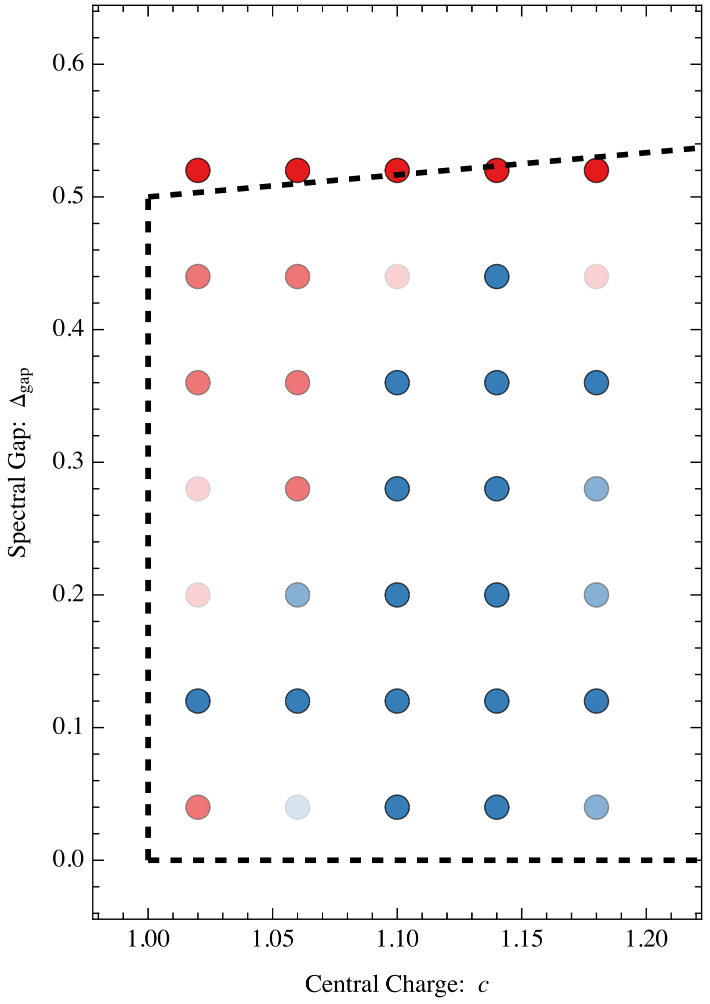

Tightening the constraints does reduce the space of apparent solutions

Further evidence that high dimension (but low twist) operators are essential for the obstruction

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026; Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

Essential Closure Test

"If you find optima with a higher truncation scale, but then truncate at Δ_{max}, your answers should agree with the directly optimizing with Δ_{max}"

Main Result: $\Delta_{\text{max}} = 4$

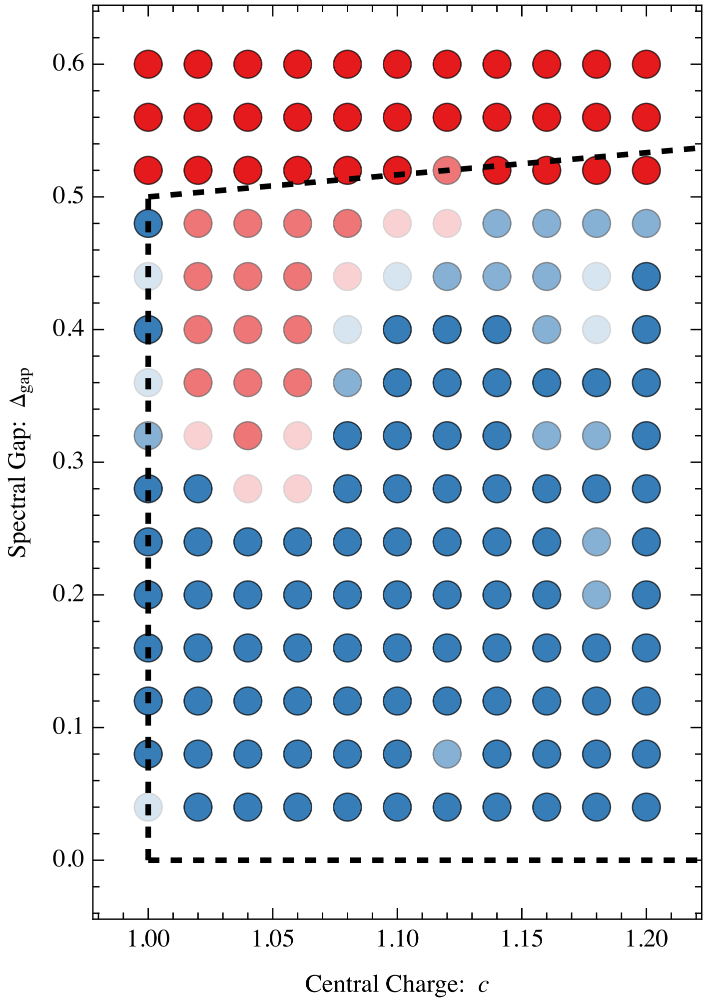

$\Rightarrow$

Truncating $\Delta_{\text{max}} = 5$ to 4

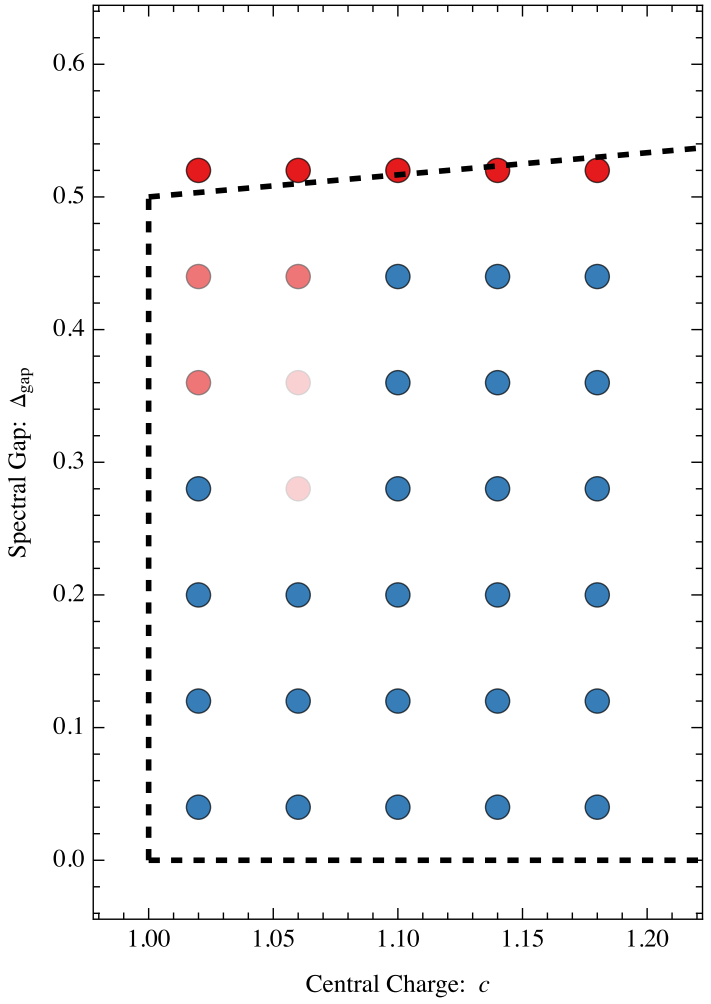

Same space of solutions!
Closure test passed!

More difficult optimization yields consistent solution to less constrained problem

[Benjamin, Fitzpatrick, Li, JDT, arXiv 2026; Bright-Thonney, Harvey, Lukas, JDT, arXiv 2026]

The Ultimate Test of Reproducibility

Humans wrote every line of Mathematica code with zero AI assistance (despite trying)
Could a human+AI team reproduce the results in PyTorch? Could they go even further??

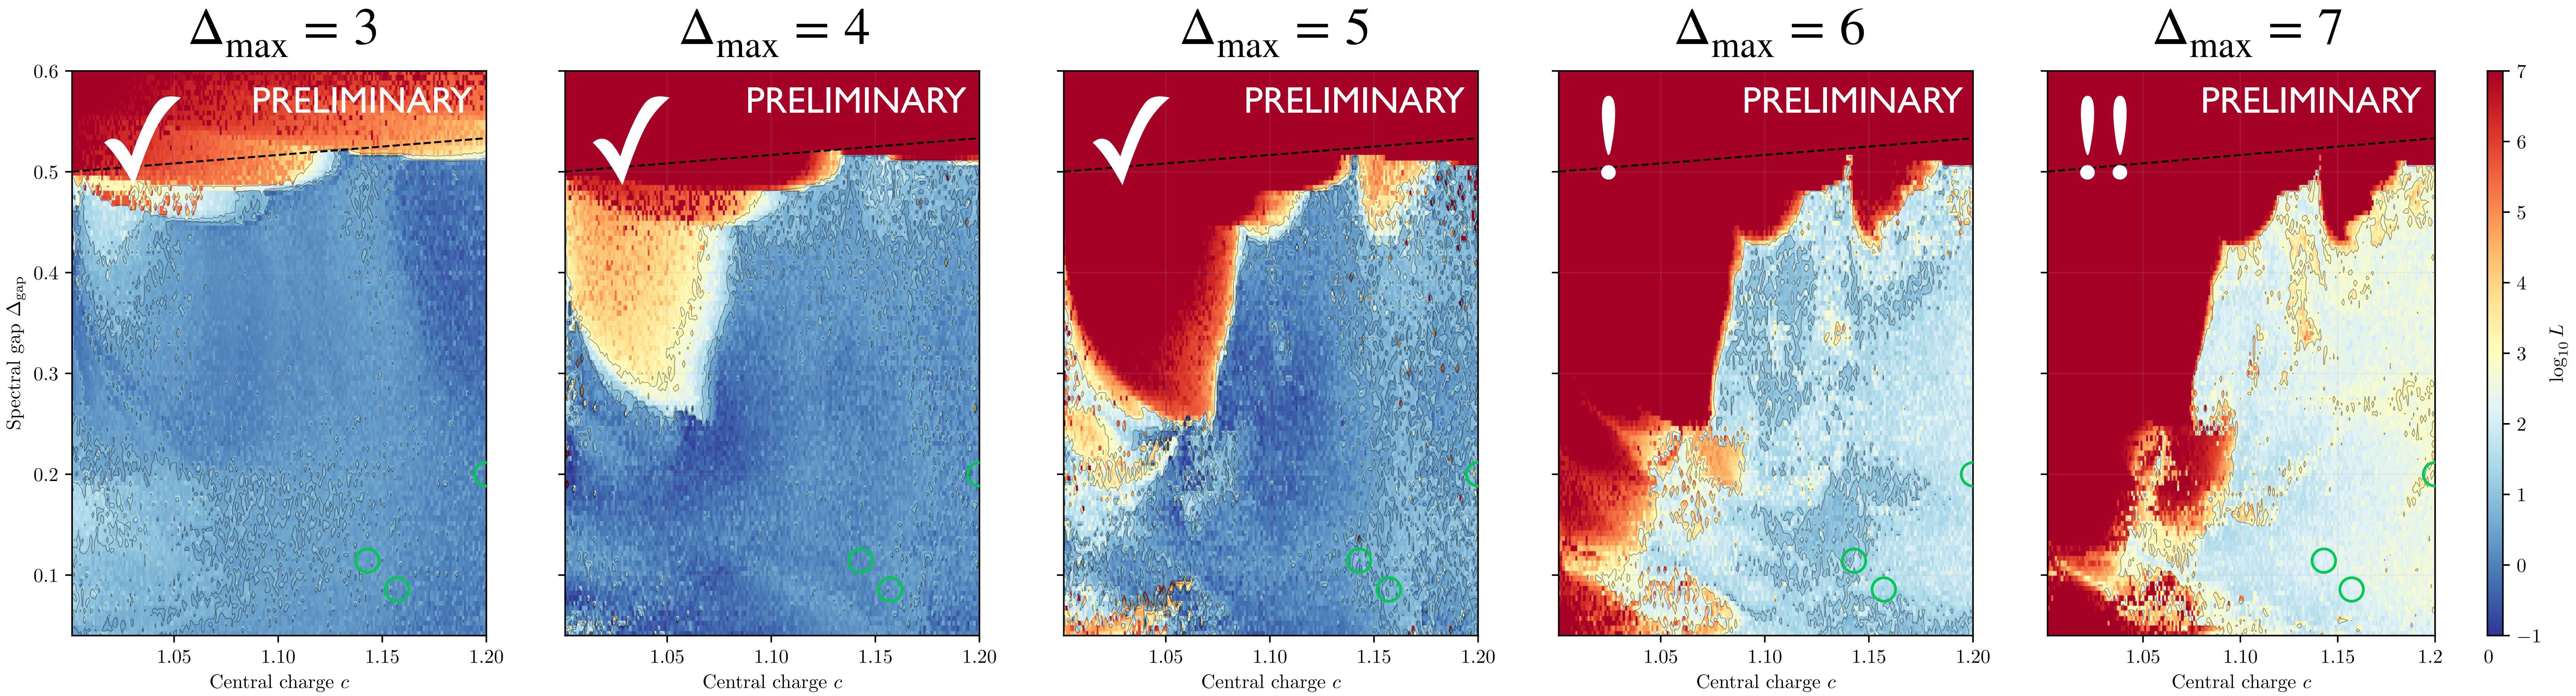

Ultra-preliminary results suggestive of an obstruction to further obstructions!

(Recent results show that yellowish regions turn more blue with better optimization hyperparameters, but red regions persist)

[Benjamin, Fitzpatrick, Li, JDT + Langhoff, work in progress]

Revisiting the Dual Constraints

"If you really pushed the dual bootstrap numerics, is there any hope to extend the known excluded region?"

Number of operators in an interval…

$$\pm\mathcal{F}_s(\Delta_a,\Delta_b) = \vec{\alpha}\cdot\vec{F}_{\text{vac}} + \sum_{\{\Delta,\ell\}\in\mathcal{S}} (n_{\Delta,\ell})\vec{\beta}\cdot\begin{pmatrix}\theta(\Delta-\Delta_b)\theta(\Delta_a-\Delta)\delta_{s,\ell}\\ \vec{F}_{\Delta,\ell}\end{pmatrix} \geq \vec{\alpha}\cdot\vec{F}_{\text{vac}}$$

…via standard SDPB optimization

Look for intervals whose F-bracket excludes integers!

Previous study:

$$\Delta_{s=0}\in\left[\Delta_{\text{gap}},\Delta_{\text{gap}}+\frac{1}{20}\right]$$

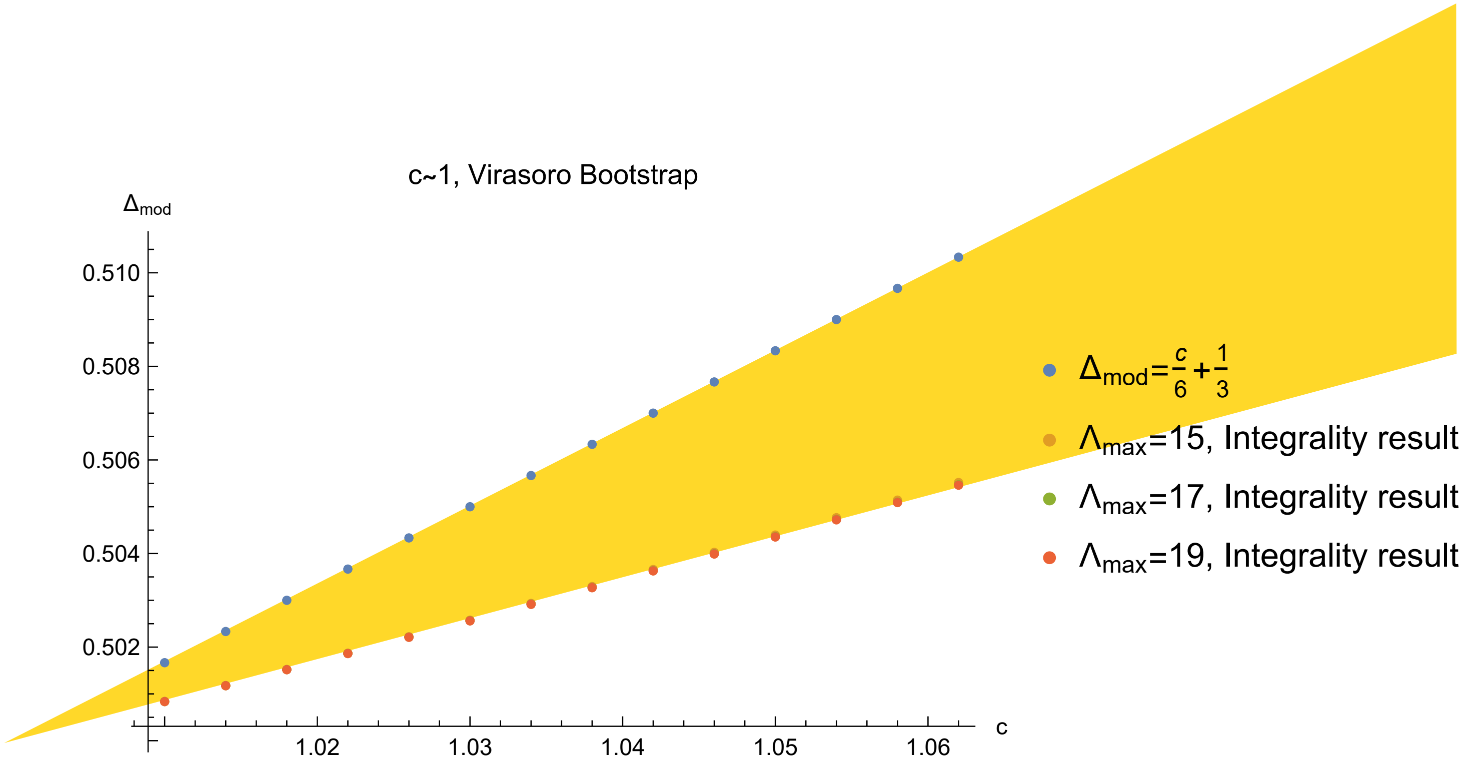

$\Rightarrow$

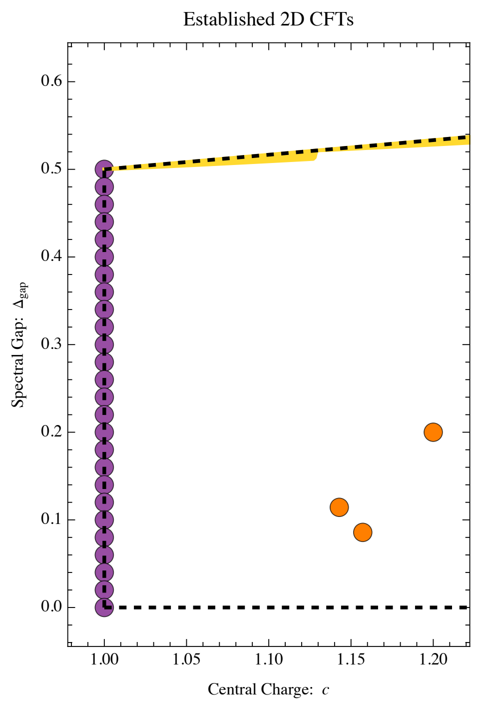

[Fitzpatrick, Li, JHEP 2024;
using Simmons-Duffin, JHEP 2015; Landry, Simmons-Duffin, arXiv 2019]


Towards a Proof of the Obstruction

Handing a JHEP article and a Mathematica notebook to my unflappable research assistant…

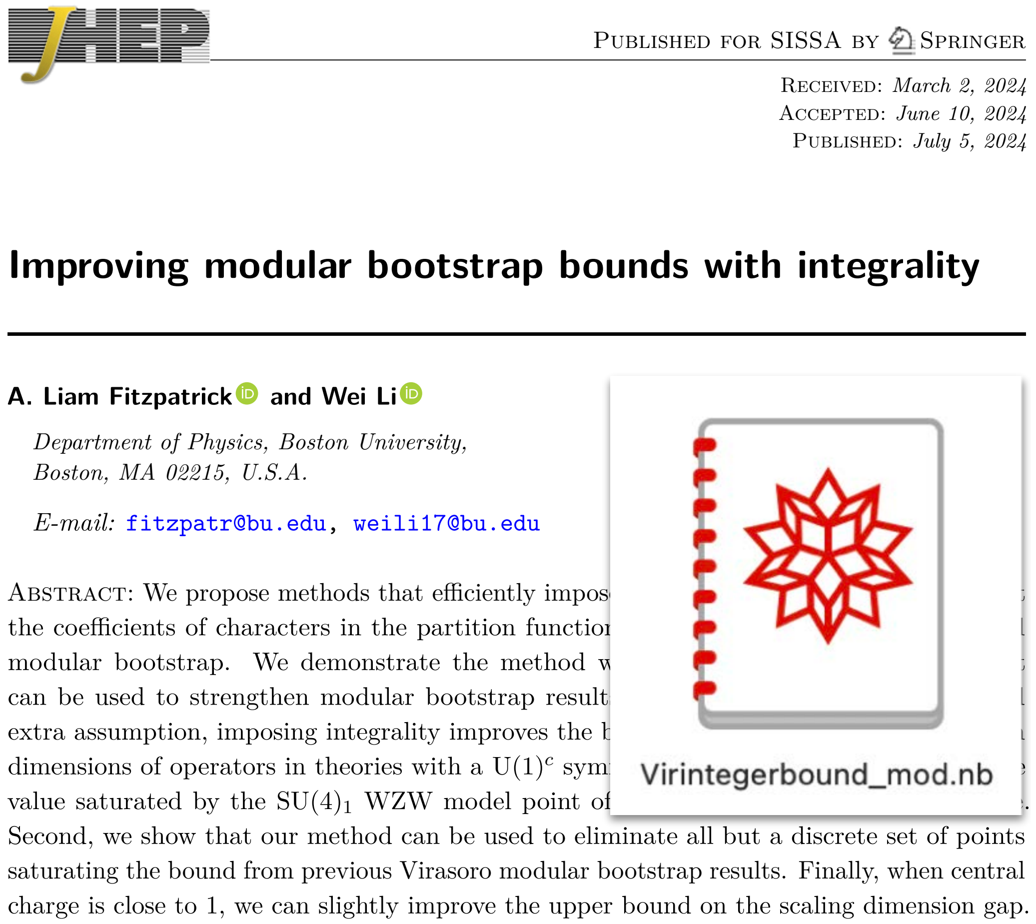

JHEP

PUBLISHED FOR SISSA BY SPRINGER

RECEIVED: March 2, 2024
ACCEPTED: June 10, 2024
PUBLISHED: July 5, 2024

Improving modular bootstrap bounds with integrality

A. Liam Fitzpatrick and Wei Li

Department of Physics, Boston University, Boston, MA 02215, U.S.A.

E-mail: fitzpatr@bu.edu, weili17@bu.edu

ABSTRACT: We propose methods that efficiently impos
the coefficients of characters in the partition functio
modular bootstrap. We demonstrate the method w
can be used to strengthen modular bootstrap result
extra assumption, imposing integrality improves the b
dimensions of operators in theories with a U(1)^c sym
value saturated by the SU(4)$_1$ WZW model point of
Second, we show that our method can be used to eliminate all but a discrete set of points saturating the bound from previous Virasoro modular bootstrap results. Finally, when central charge is close to 1, we can slightly improve the upper bound on the scaling dimension gap.

Virintegerbound_mod.nb

⇒

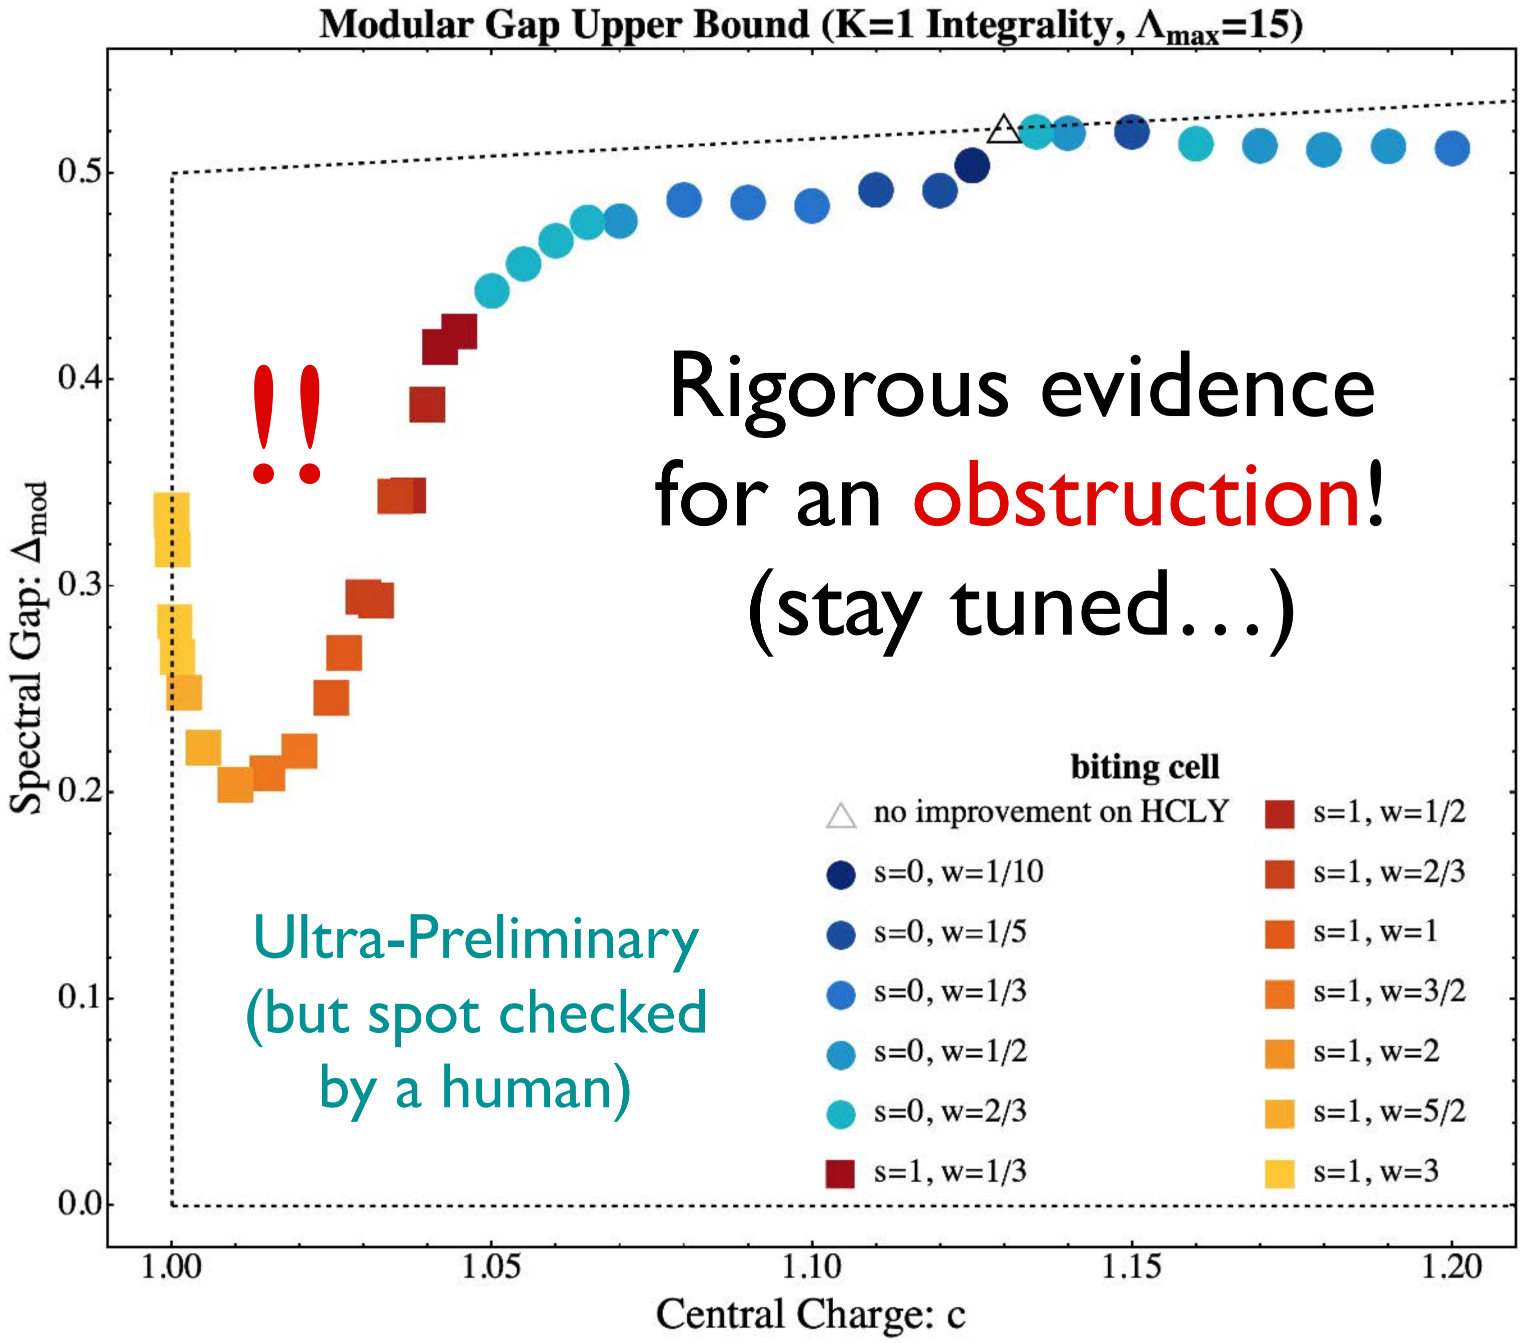

[Benjamin, Fitzpatrick, Langhoff, Li, JDT, work in progress;
using Fitzpatrick, Li, JHEP 2024; Simmons-Duffin, JHEP 2015; Landry, Simmons-Duffin, arXiv 2019]




Exciting opportunities for discovery from a combination of physics intelligence and artificial intelligence

Descending into the Modular Bootstrap

Machine Learning through the Lens of Physics

AI is transforming physics research, but many foundational aspects of AI can be translated into the language of physics

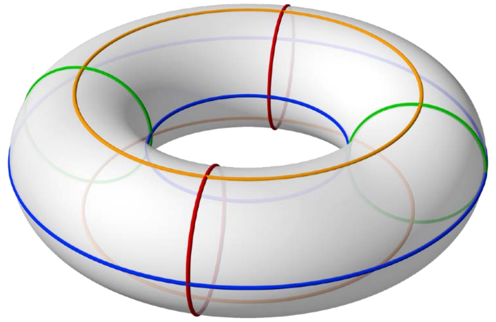

Revisiting the Modular Bootstrap

To confront thorny challenges in theoretical physics, we can reframe them as well-defined optimization problems

What Have We Actually Learned?

As we capitalize on emerging AI technologies, we still need to insist on robustness and reproducibility

Backup Slides

The Future of Artificial Intelligence and the Mathematical and Physical Sciences (AI+MPS)

Community Paper from the NSF Future of AI+MPS Workshop
Cambridge, Massachusetts — March 24–26, 2025

Cross-Disciplinary Opportunities:

- Advocate for Diverse Funding Streams
- Pursue the Science of AI →
  - AI Innovations from Science
  - Understanding AI Behaviors
  - Robust and Reproducible AI
- Establish Scalable AI Infrastructures
- Facilitate Interdisciplinary Collaborations
- Cultivate Key AI Techniques for Science
- Leverage AI for Conducting Research
- Educate and Train an AI+MPS Workforce
- Empower AI Innovation

[Ferguson, LaFleur, Ruthotto, JDT, Ting, Tiwary, Villar, et al., MLST 2026; see MIT News March 2026]

AI is Changing What it Means To…

Representative examples, far from exhaustive!

Design/Build/Operate Experiments

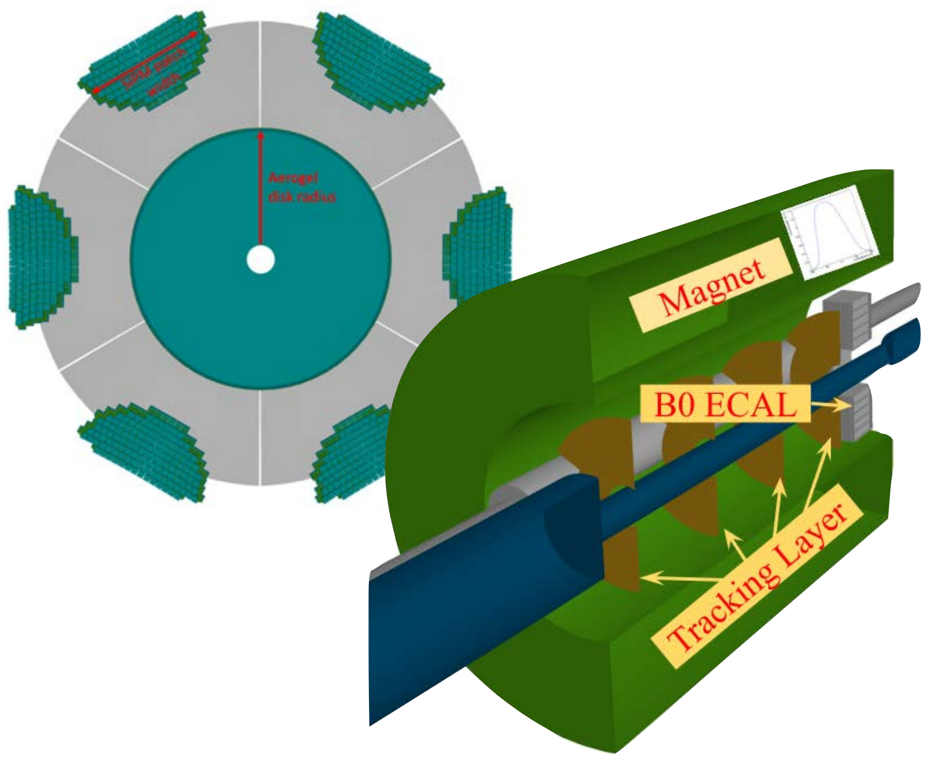

e.g. ePIC detector optimization for EIC

[AID(2)E Collaboration, JINST 2024]

Perform/Report Measurements

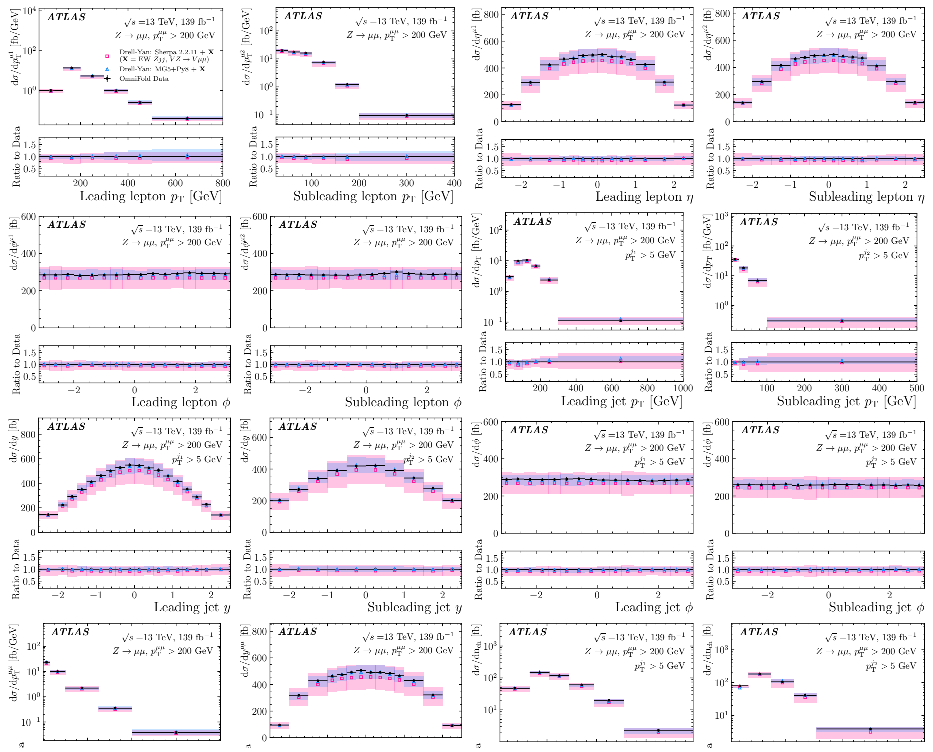

e.g. ATLAS 24-dimensional unbinned unfolding

[ATLAS, PRL 2024]

Search for New Physics

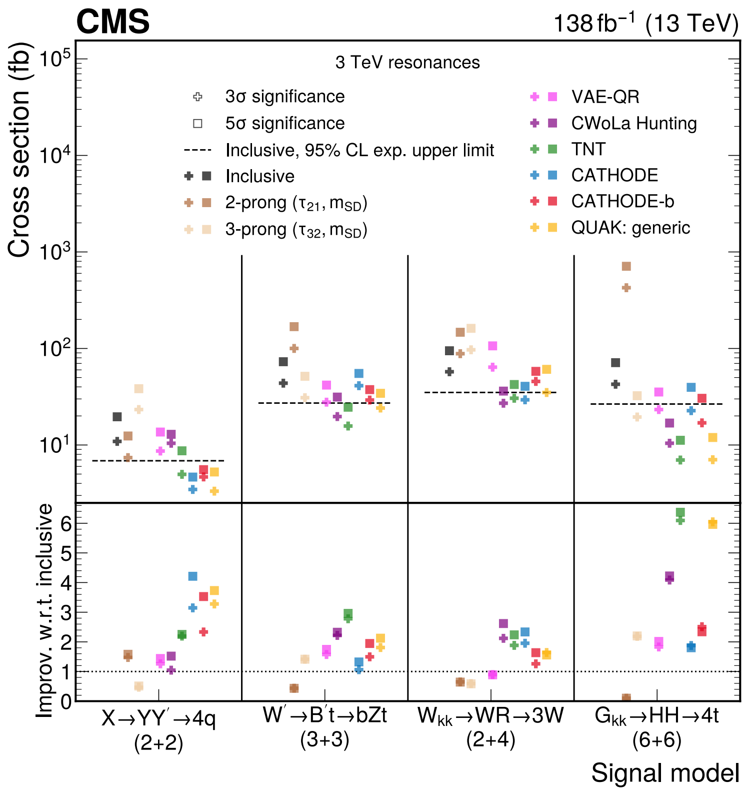

e.g. CMS anomalous dijet resonance search

[CMS, RPP 2025]

[from my GGI lectures, January 2026; more examples at HEP ML Living Review]

AI is Changing What it Means To…

Representative examples, far from exhaustive!
Experimental implications in backup slides

Make Theory Predictions

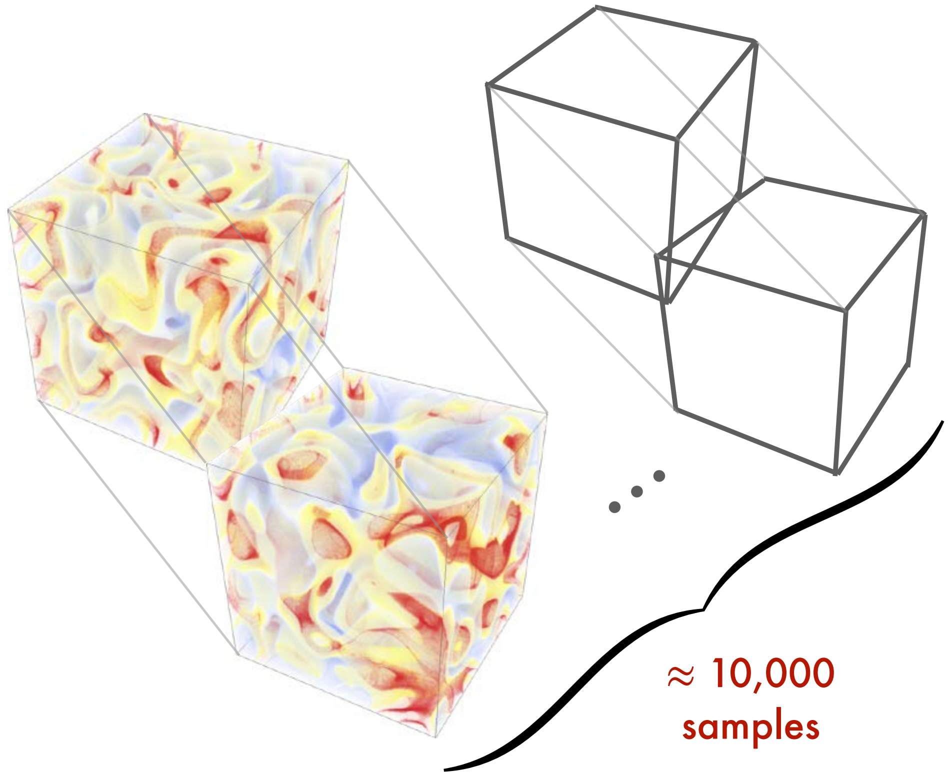

e.g. generative modeling
for lattice field theory

[Cranmer, Kanwar, Racanière, Rezende, Shanahan, Nature Reviews 2023]

Be a Physicist

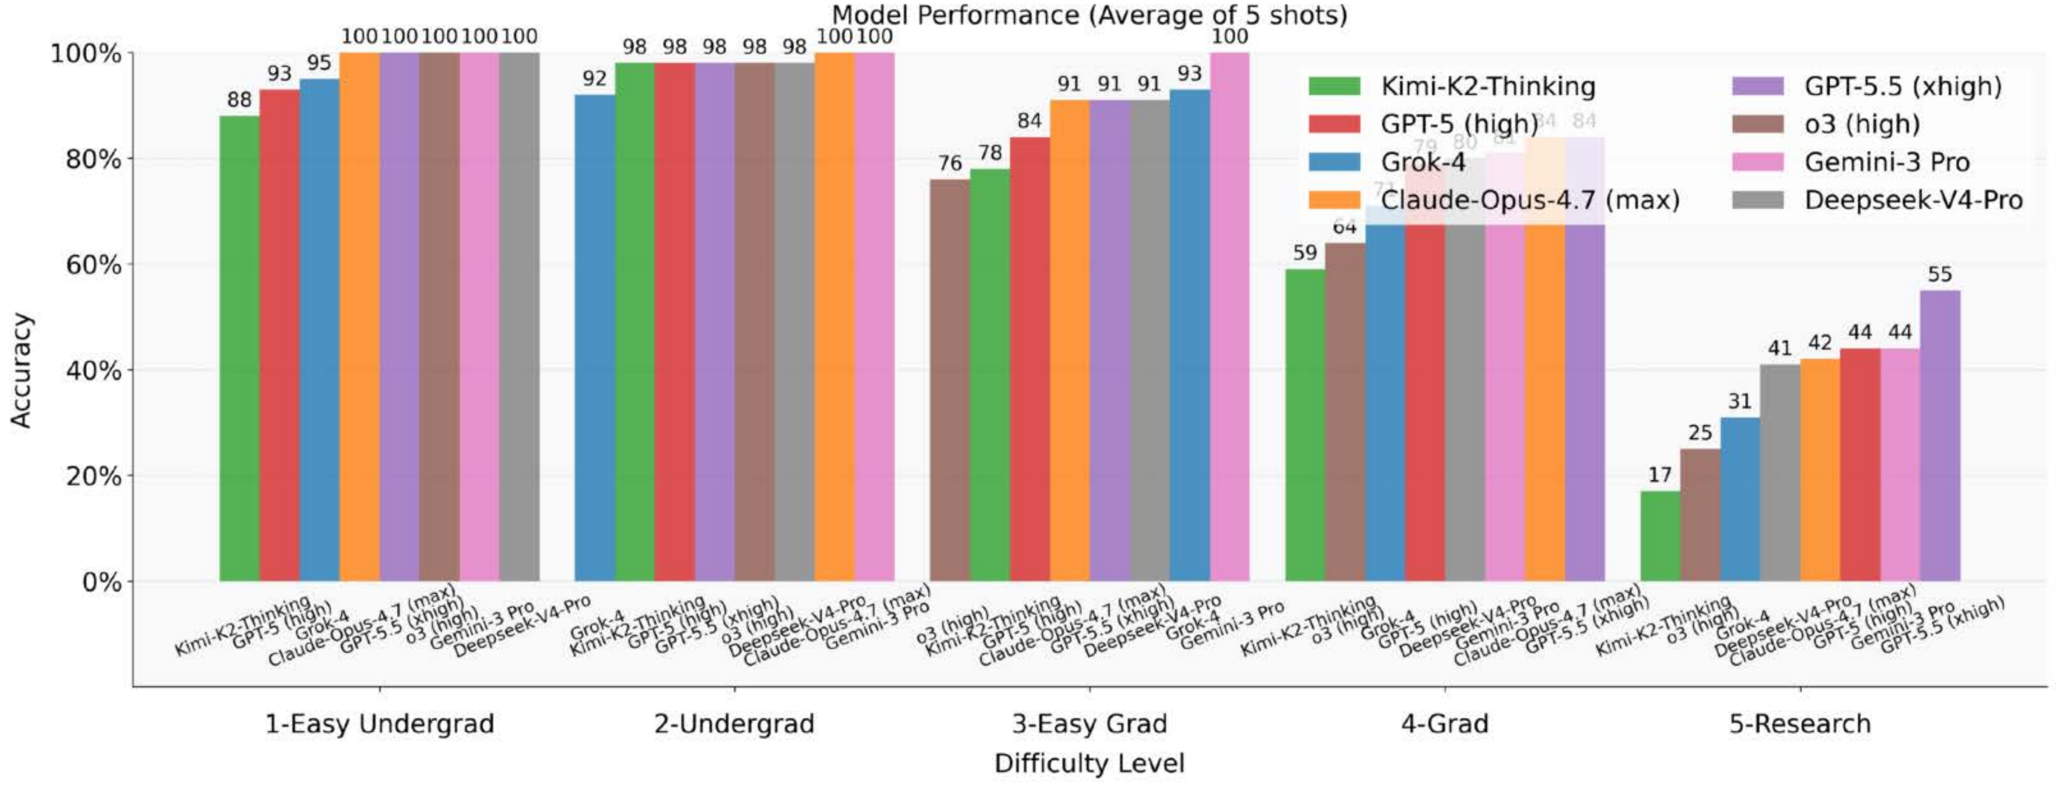

e.g. evaluating AI reasoning
abilities on TPBench

[Chung, Gao, Kvasiuk, Li, Münchmeyer, Rudolph, Sala, Tadepalli, MLST 2025]

[from my GGI lectures, January 2026; more examples at HEP ML Living Review and Ferguson, LaFleur, Ruthotto, JDT, Ting, Tiwary, Villar, et al., MLST 2026]

AI is Changing What it Means To…

Representative examples, far from exhaustive!
Experimental implications in backup slides

Make Theory Predictions

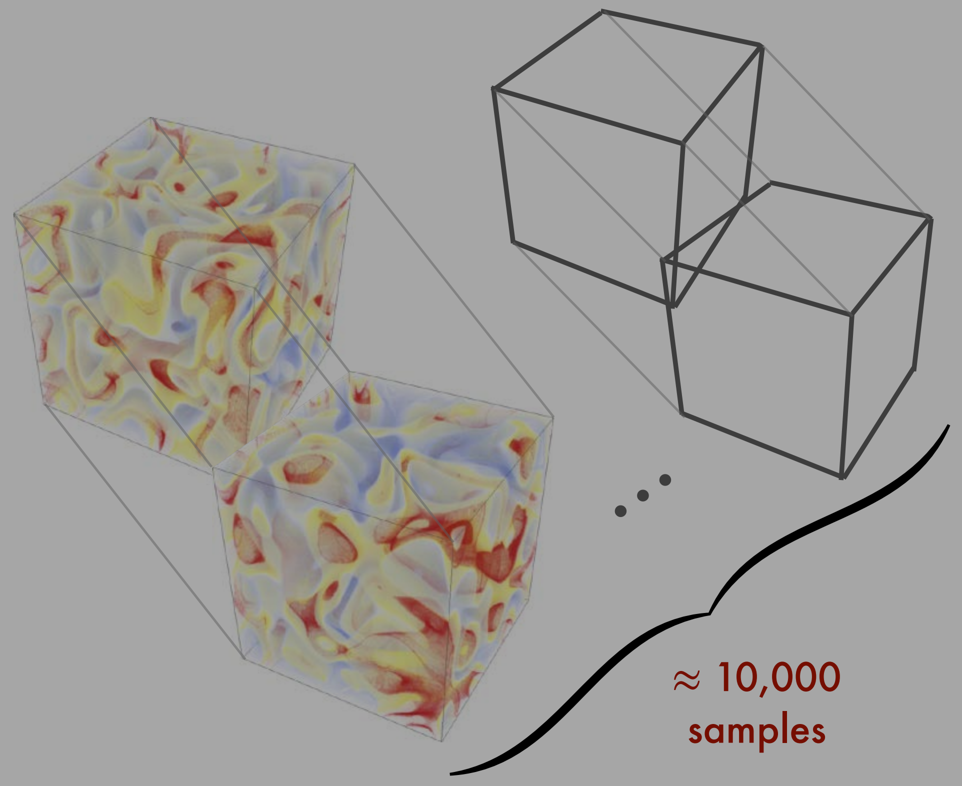

e.g. generative modeling
for lattice field theory

[Cranmer, Kanwar, Racanière, Rezende, Shanahan, Nature Reviews 2023]

Be a Physicist

Resummation of the C-Parameter Sudakov Shoulder Using Effective Field Theory

Matthew D. Schwartz[1,2]

[1] *Department of Physics, Harvard University, Cambridge, MA 02138, USA*

[2] *Institute for Artificial Intelligence and Fundamental Interactions (IAIFI)*

schwartz@g.harvard.edu

AI Research Assistant: Claude Opus 4.5 (Anthropic)

January 7, 2026

[Schwartz, arXiv 2026; see also Shih, arXiv 2026a, arXiv 2026b]

[from my GGI lectures, January 2026; more examples at HEP ML Living Review and Ferguson, LaFleur, Ruthotto, JDT, Ting, Tiwary, Villar, et al., MLST 2026]

Direct AI/ML: QCD Theory Meets Information Theory

Goal:	Synthesize Multiple QCD Calculations
Training Data:	Precision Moments of EFPs
Loss Function(al):	Information-Theoretic Similarity
Learnable Function:	Custom Boltzmann Factor
Optimizer:	Non-Linear Solver (L-BFGS)

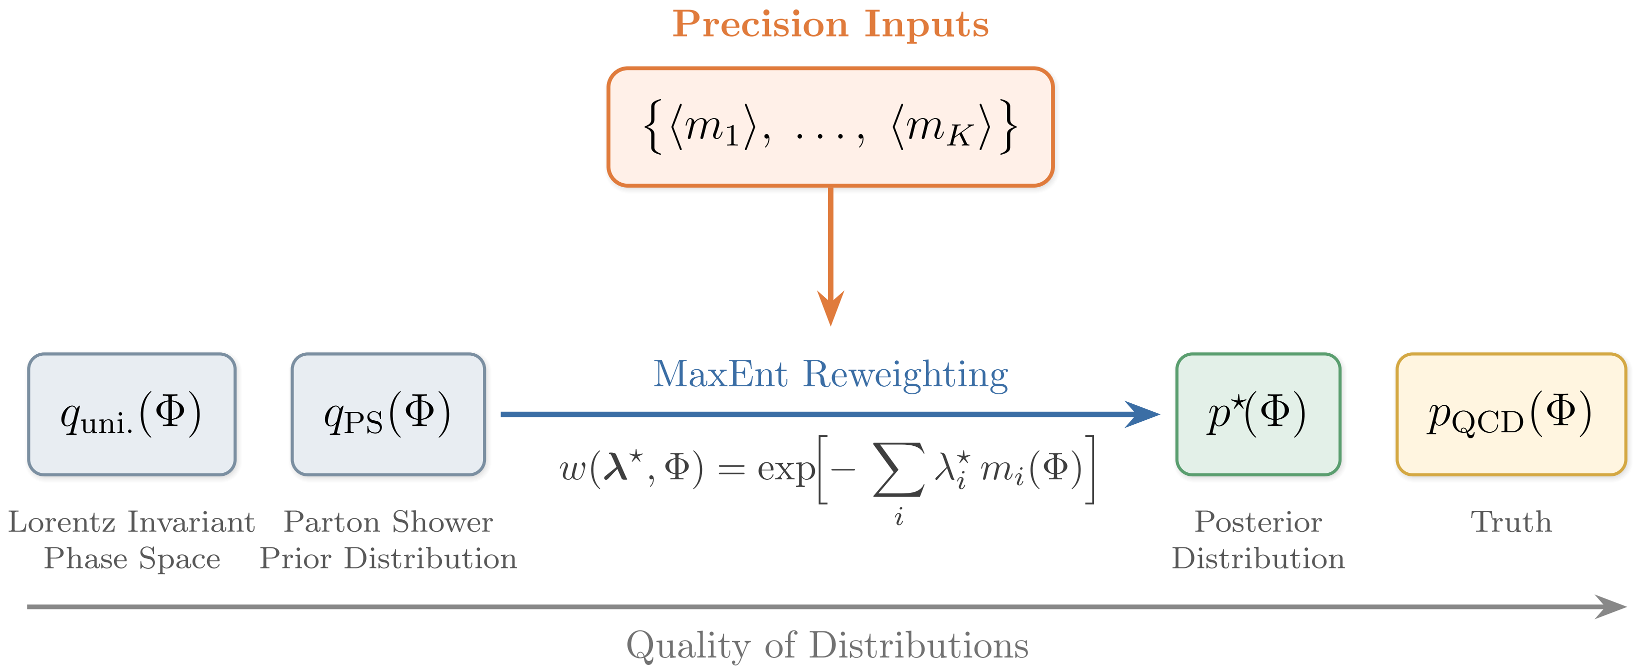

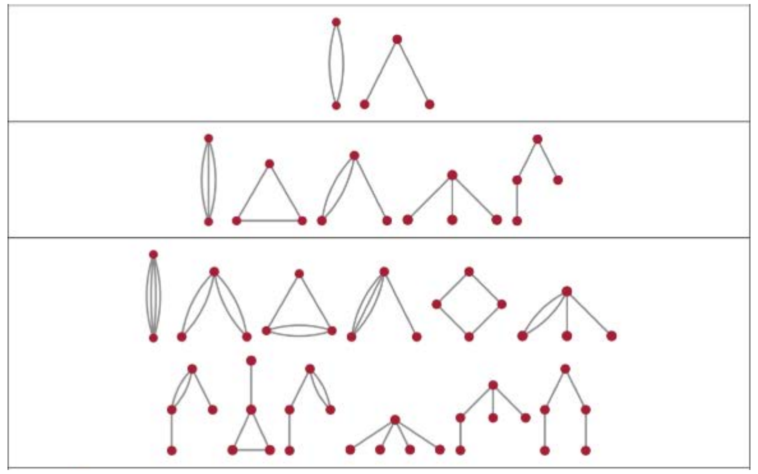

EFPs = Energy Flow Polynomials

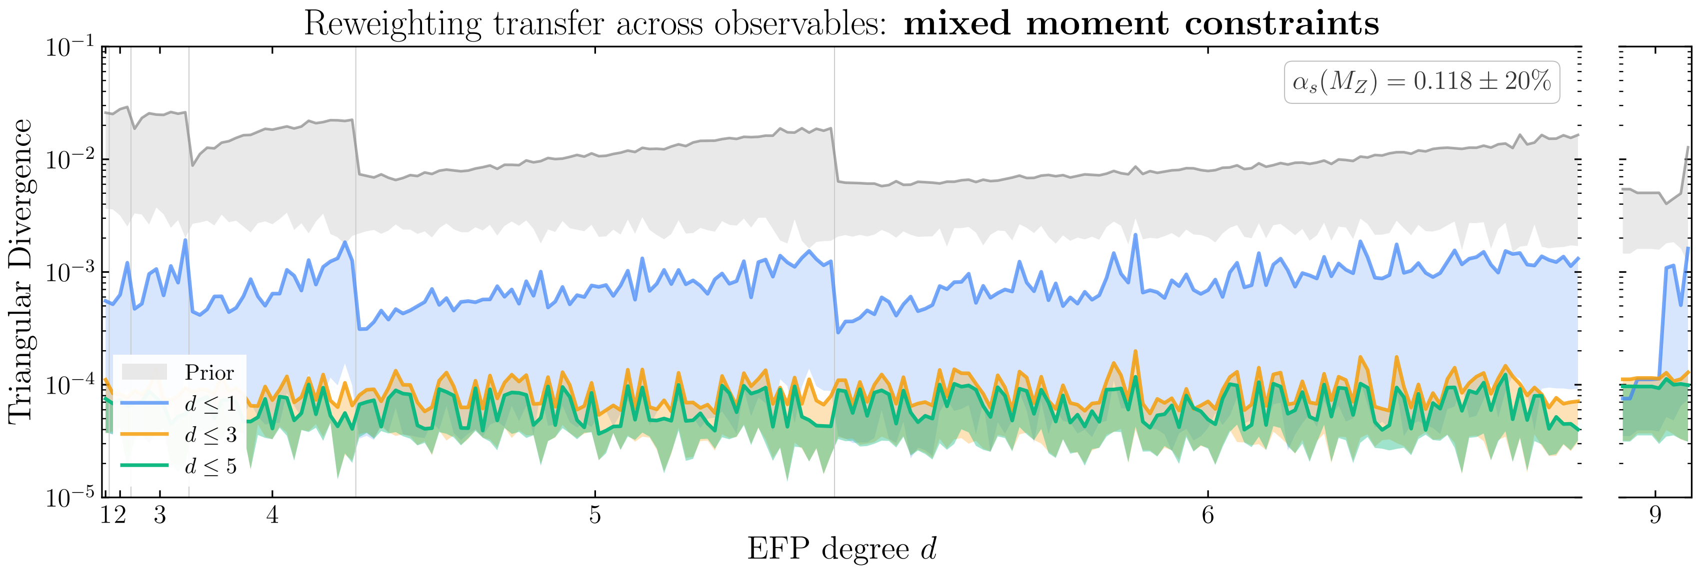

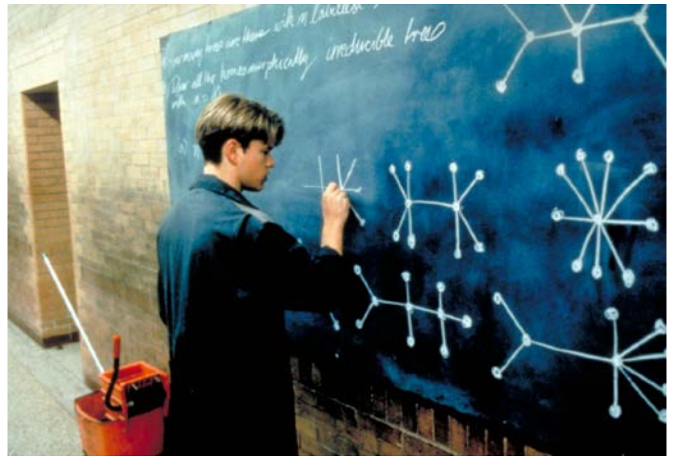

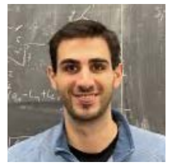  

[Assi, Lee, JDT, arXiv 2026; Assi, Lee, Höche, JDT, PRL 2025; using Komiske, Metodiev, JDT, JHEP 2018]

From AI Curmudgeon… to AI Evangelist!

MIT News
ON CAMPUS AND AROUND THE WORLD

Laboratory for Nuclear Science
August 26, 2020

National Science Foundation announces MIT-led Institute for Artificial Intelligence and Fundamental Interactions

IAIFI will advance physics knowledge — from the smallest building blocks of nature to the largest structures in the universe — and galvanize AI research innovation.

The New York Times

By Dennis Overbye Nov. 23, 2020

Can a Computer Devise a Theory of Everything?

"In five to 10 years from now, I'm going to want to do exactly what you're getting at: Here's the data, here's a very rough tool kit; find the equation I could put on a T-shirt, the equation that replaces the Standard Model of particle physics. What's the equation that replaces Einstein's general relativity?"

N.B. This was November 2020. ChatGPT was released November 2022.

From AI Curmudgeon… to AI Evangelist!

MIT News
ON CAMPUS AND AROUND THE WORLD

Laboratory for Nuclear Science
August 26, 2020

National Science Foundation announces MIT-led Institute for Artificial Intelligence and Fundamental Interactions

IAIFI will advance physics knowledge — from the smallest building blocks of nature to the largest structures in the universe — and galvanize AI research innovation.

The New York Times

By Dennis Overbye   Nov. 23, 2020

Can a Computer Devise a Theory of Everything?

This dream seems very far away to me now, since the search for BSM physics appears limited by experiment not theory…

Nevertheless, AI tutors already enable Matrix-style knowledge downloads and AI systems can crack theory problems amenable to brute force attack

Don't be an ostrich! Try AI in your research and form an opinion

Known 2D CFT Spectra: Free Boson Circle Branch

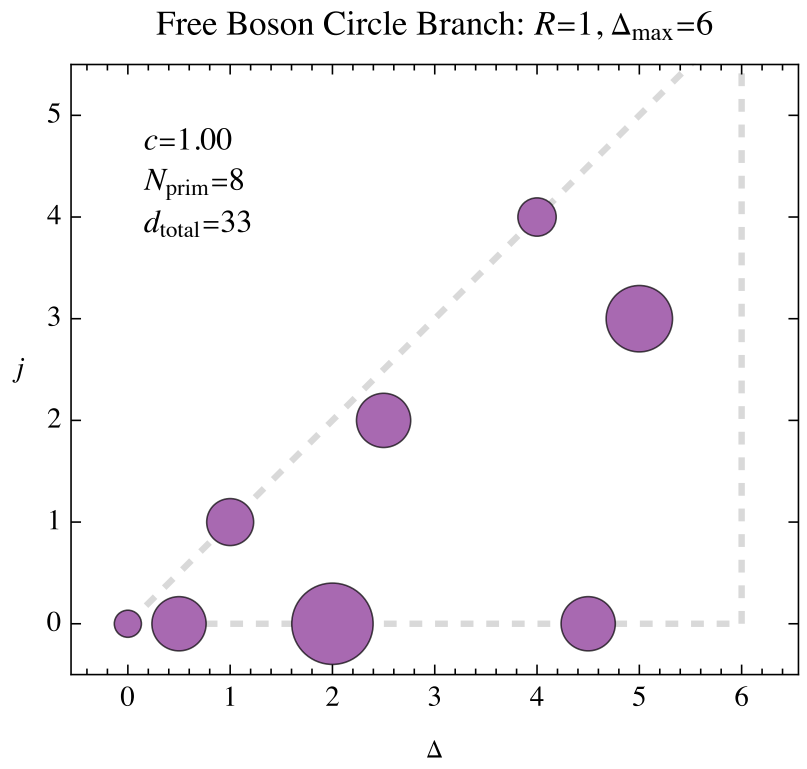

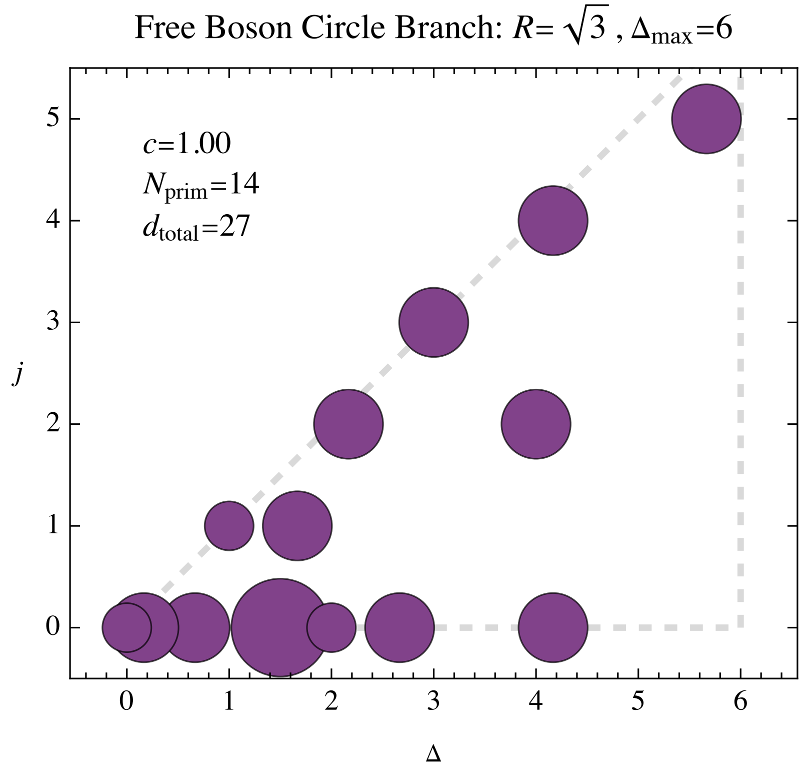

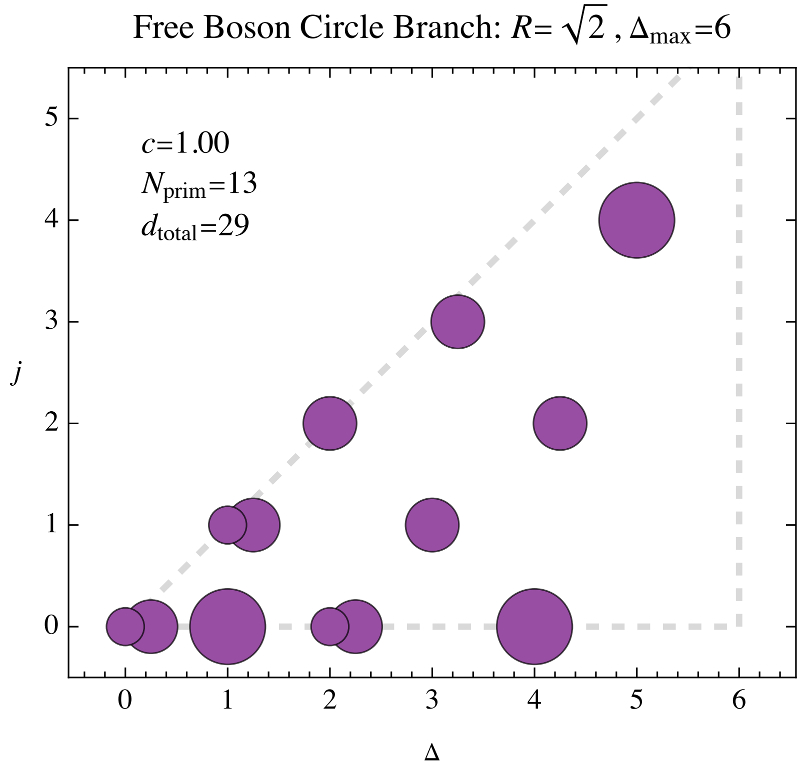

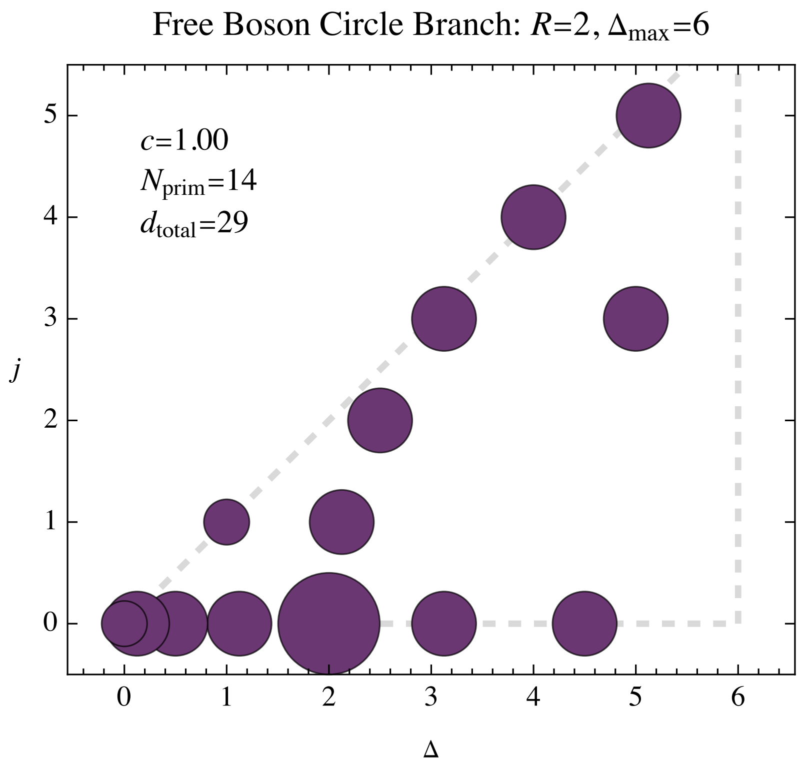

Known 2D CFT Spectra: N=1 Minimal Models

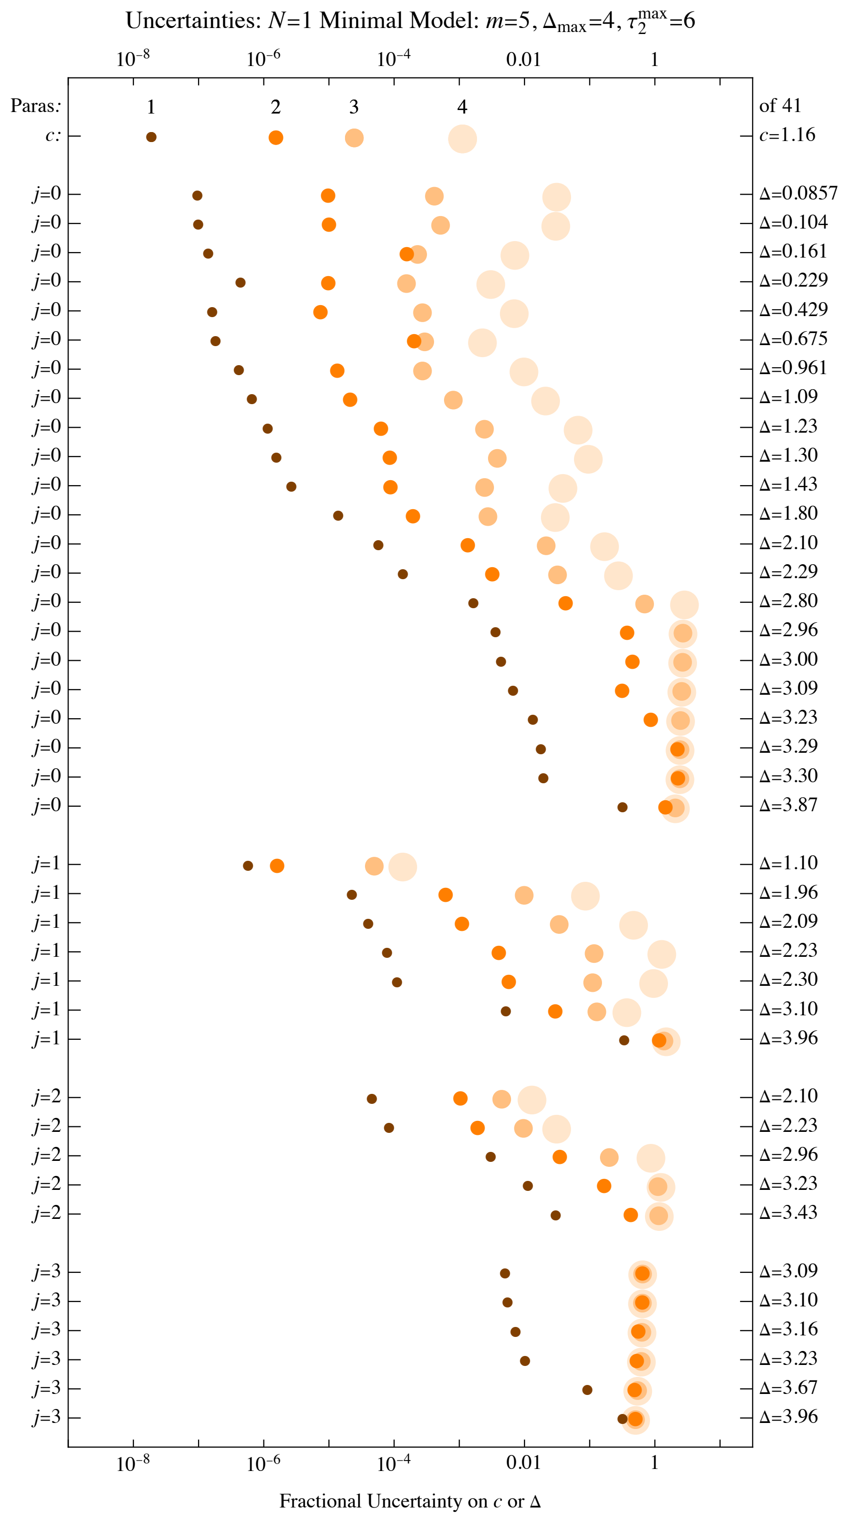

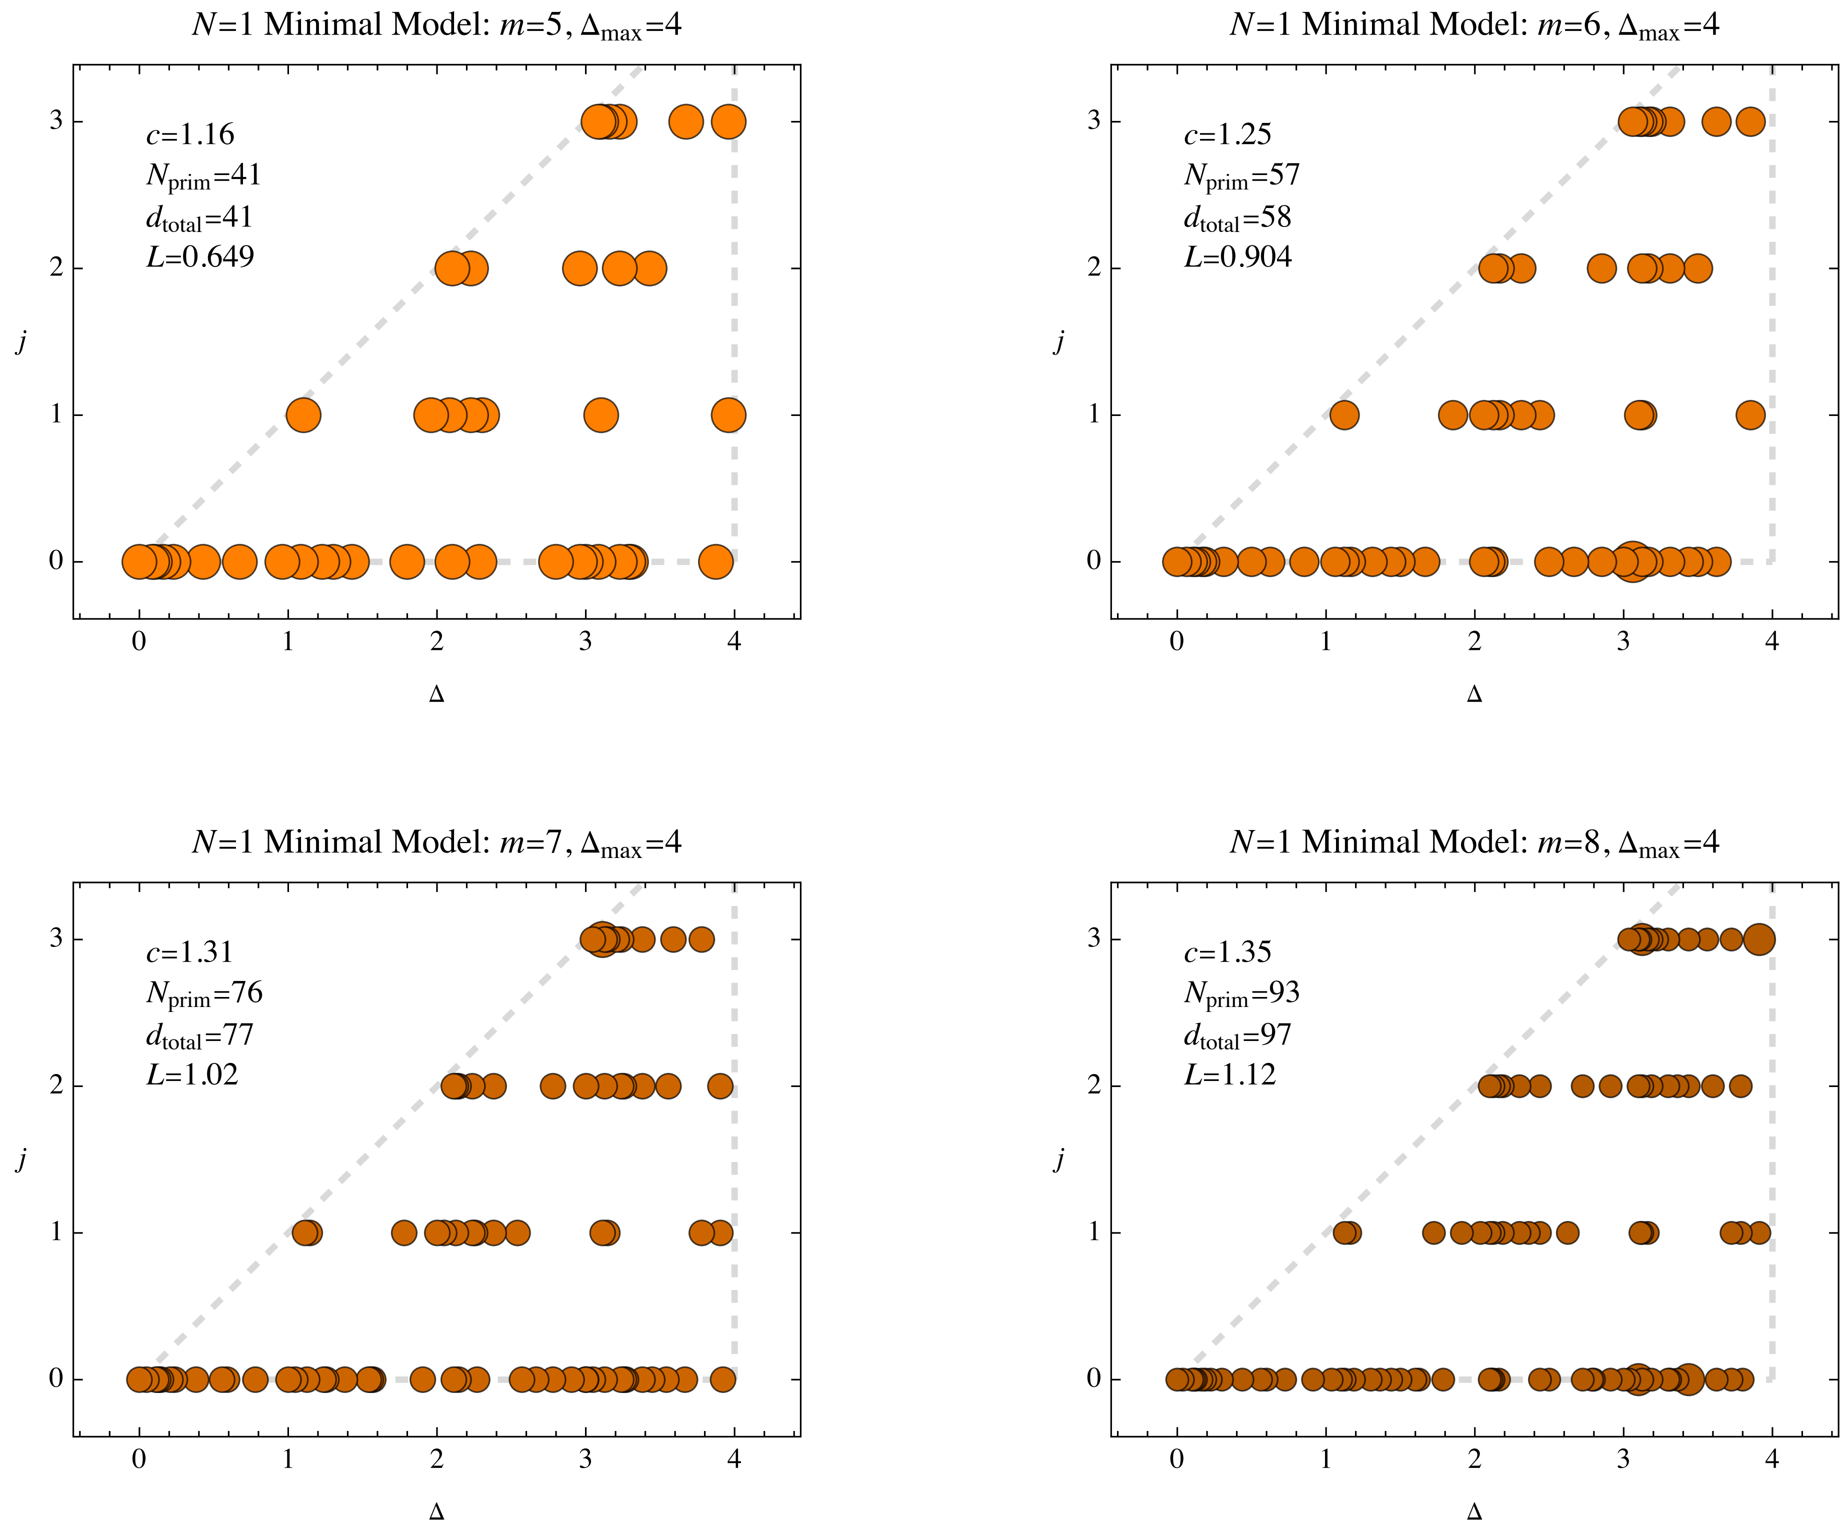

Candidate 2D CFT Spectra: Floating Δ_{gap}

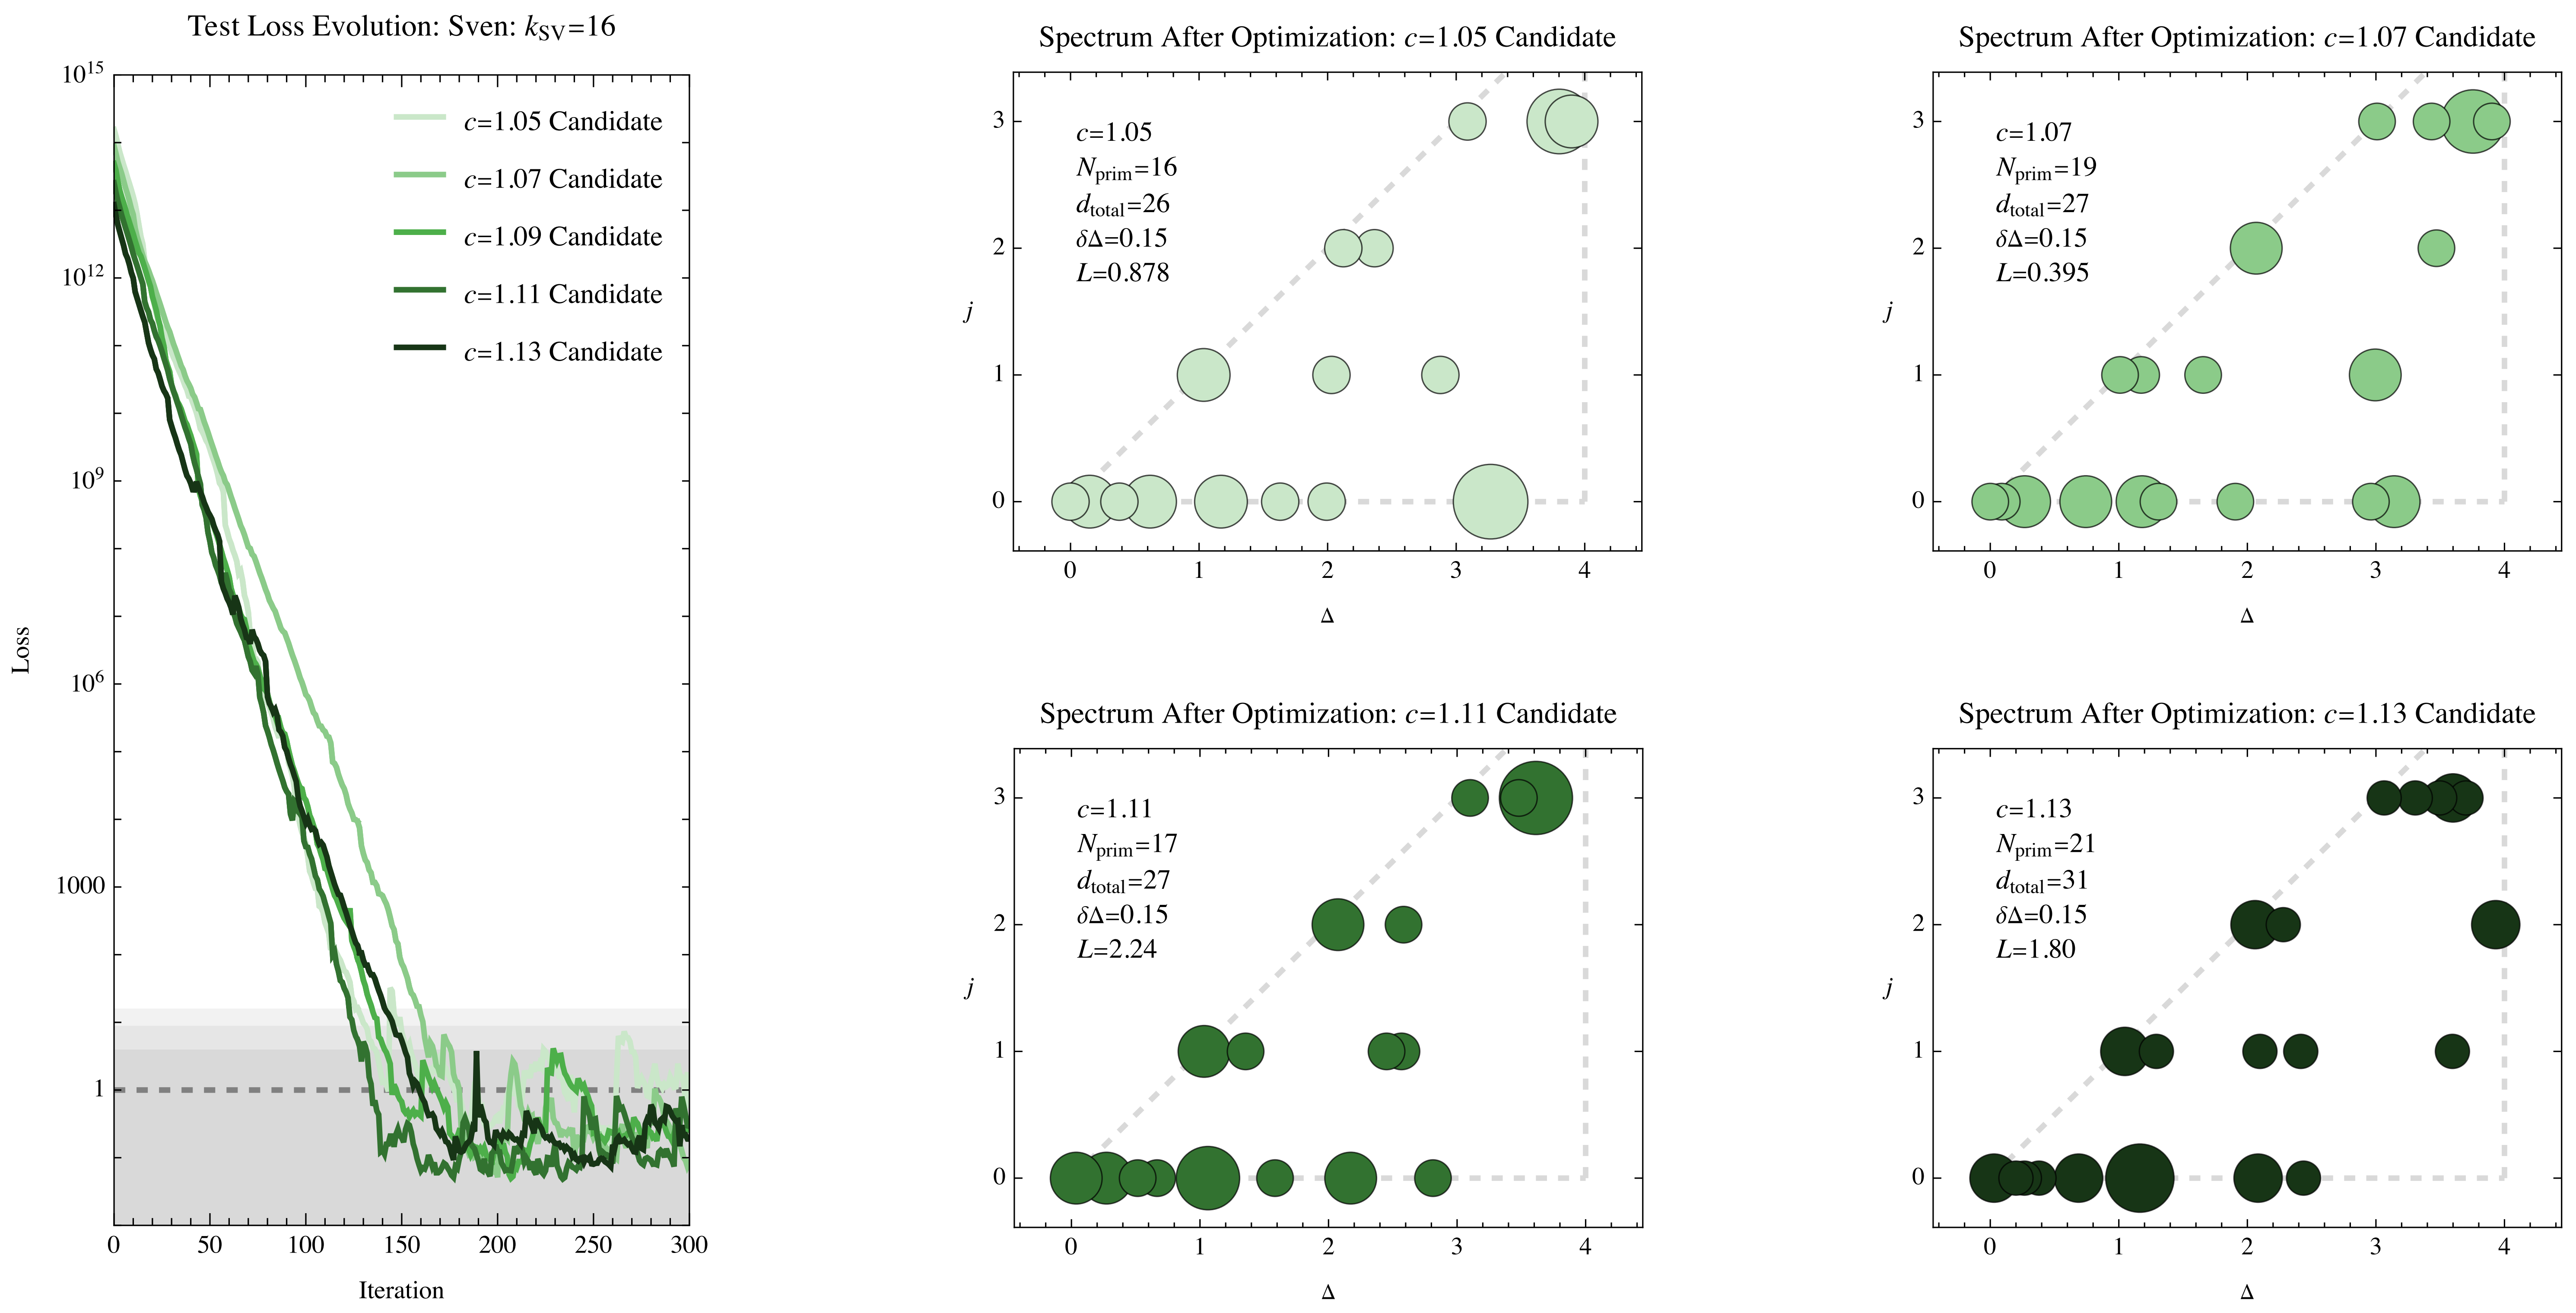

Sven: Sweeping Number of Singular Values

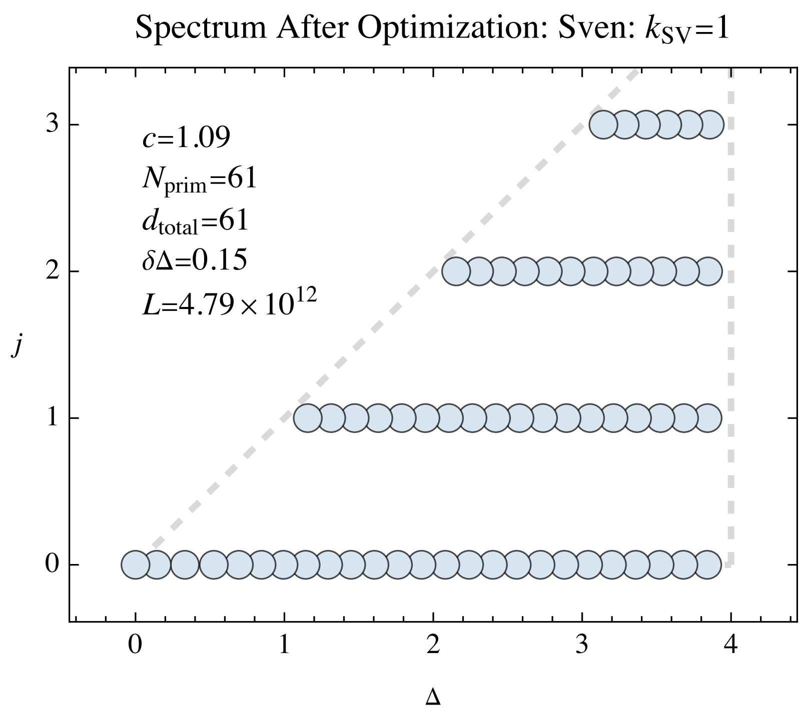

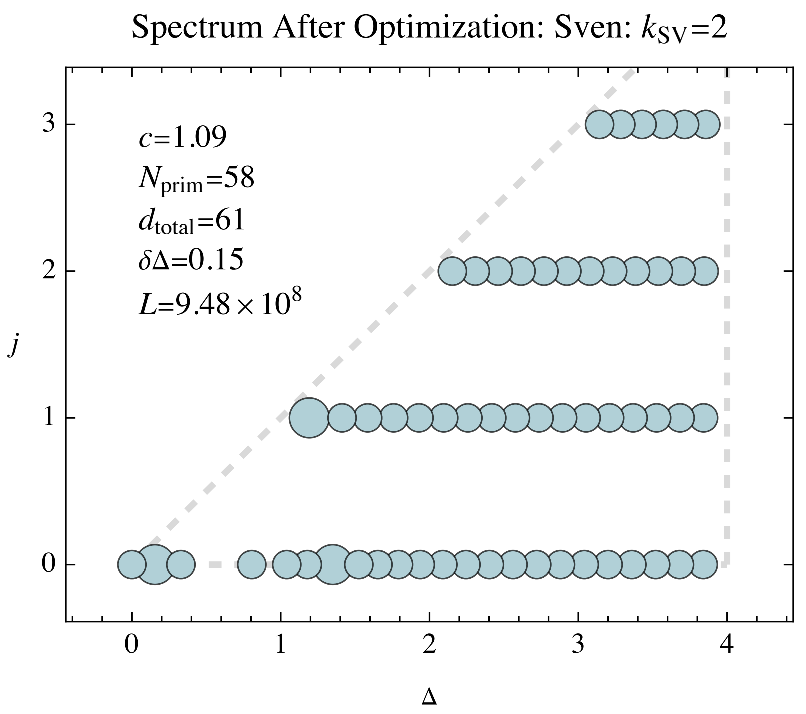

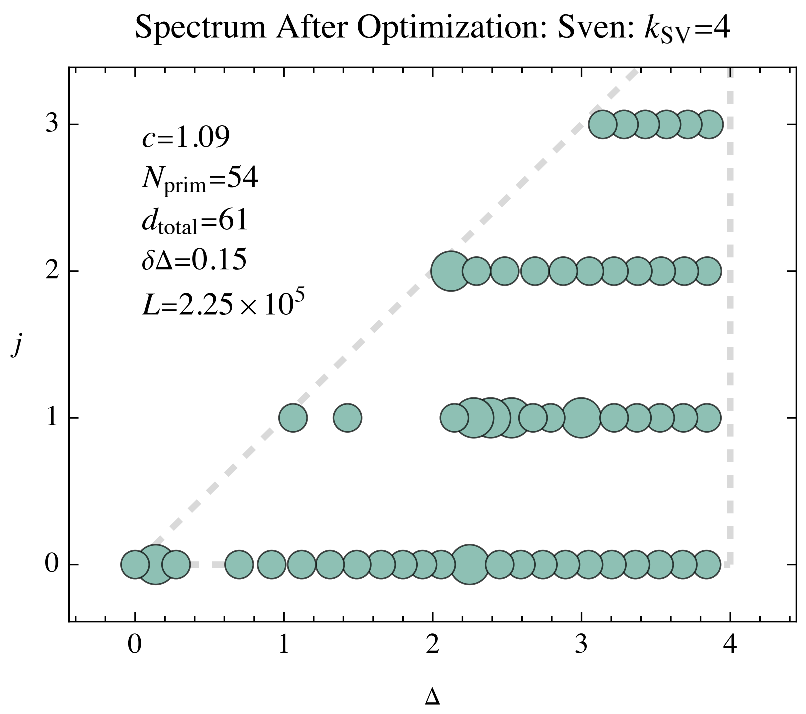

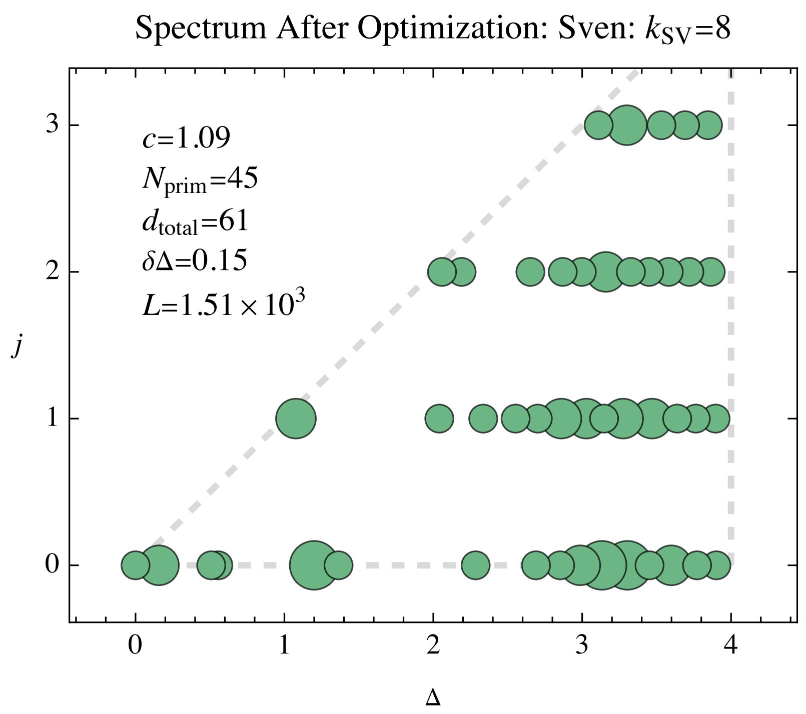

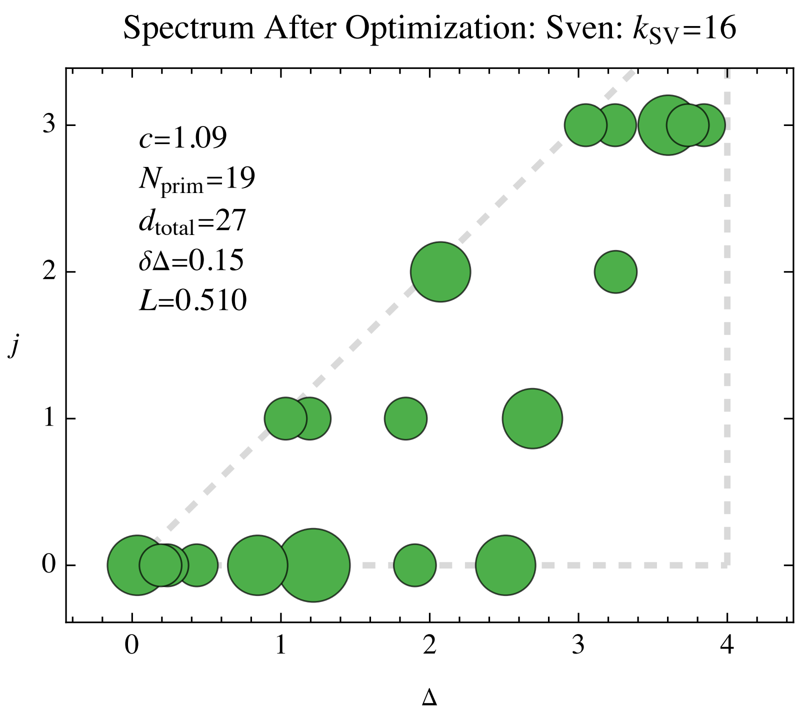

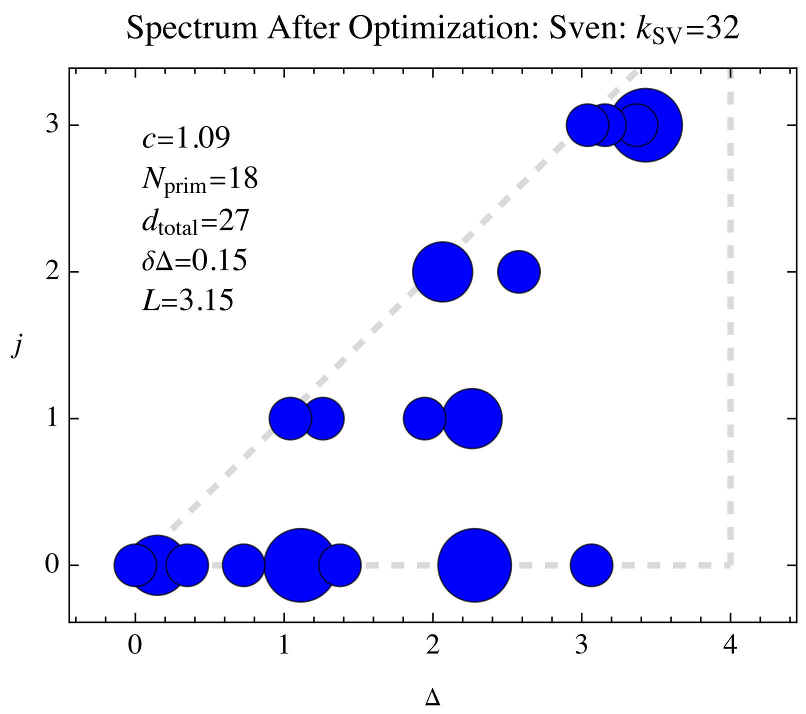

Robustness of Loss Values

N1MMs

Claude's Way:

	$m = 5$	$m = 6$	$m = 7$	$m = 8$	$m = 9$
$R = 1$	0.851	1.16	1.29	1.40	1.65
$\boldsymbol{R = \sqrt{2}}$	**0.649**	**0.904**	**1.02**	**1.12**	**1.32**
$R = \sqrt{3}$	1.58	1.89	2.03	2.13	2.48
$R = 2$	0.777	1.03	1.14	1.24	1.46

Human's Way:

	$m = 5$	$m = 6$	$m = 7$	$m = 8$	$m = 9$
$R = 1$	0.424	0.449	0.431	0.428	0.469
$\boldsymbol{R = \sqrt{2}}$	**0.371**	**0.399**	**0.388**	**0.389**	**0.429**
$R = \sqrt{3}$	1.14	1.04	0.982	0.957	1.05
$R = 2$	0.422	0.453	0.449	0.455	0.506

Floating Δ_{gap}

Claude's Way:

	$c = 1.05$	$c = 1.07$	$c = 1.09$	$c = 1.11$	$c = 1.13$
$R = 1$	1.19	0.528	0.680	3.02	2.39
$\boldsymbol{R = \sqrt{2}}$	**0.878**	**0.395**	**0.510**	**2.24**	**1.80**
$R = \sqrt{3}$	3.39	1.30	1.61	6.64	5.11
$R = 2$	1.14	0.513	0.649	2.78	2.18

Human's Way:

	$c = 1.05$	$c = 1.07$	$c = 1.09$	$c = 1.11$	$c = 1.13$
$R = 1$	0.870	0.358	0.416	1.73	1.31
$\boldsymbol{R = \sqrt{2}}$	**0.725**	**0.305**	**0.360**	**1.50**	**1.15**
$R = \sqrt{3}$	3.42	1.27	1.43	5.71	4.46
$R = 2$	0.866	0.364	0.421	1.71	1.31