

6D (2,0) bootstrap with the soft-actor-critic algorithmGergely Kántor^{1,*}, Vasilis Niarchos^{2,†}, Constantinos Papageorgakis^{1,‡} and Paul Richmond^{1,§}¹*Centre for Theoretical Physics, Department of Physics and Astronomy,
Queen Mary University of London, London E1 4NS, United Kingdom*²*ITCP & CCTP, Department of Physics, University of Crete, 71003 Heraklion, Greece*

(Received 4 October 2022; accepted 23 December 2022; published 10 January 2023)

We study numerically the 6D (2,0) superconformal bootstrap using the soft-actor-critic (SAC) algorithm as a stochastic optimizer. We focus on the four-point functions of scalar superconformal primaries in the energy-momentum multiplet. Starting from the supergravity limit, we perform searches for adiabatically varied central charges and derive two curves for a collection of 80 conformal field theory (CFT) data (70 of these data correspond to unprotected long multiplets and 10 to protected short multiplets). We conjecture that the two curves capture the *A*- and *D*-series (2,0) theories. Our results are competitive when compared to the existing bounds coming from standard numerical bootstrap methods, and data obtained using the OPE inversion formula. With this paper we are also releasing our PYTHON implementation of the SAC algorithm, `BOOTSTOP`. The paper discusses the main functionality features of this package.

DOI: [10.1103/PhysRevD.107.025005](https://doi.org/10.1103/PhysRevD.107.025005)**I. INTRODUCTION AND SUMMARY**

In a recent paper, [1], we proposed that stochastic optimization methods could be a useful new tool in the arsenal of the conformal bootstrap program. More specifically, in the context of truncation methods (where one looks for approximate solutions to the crossing equations with a truncated spectrum; see e.g., [2]) one can view the conformal bootstrap as a large-scale, nonconvex, continuous optimization problem. Such complex nonlinear programming problems are ubiquitous in many scientific areas and applications, and are commonly treated by stochastic algorithms.

In [1] we used an optimization method based on a popular soft-actor-critic (SAC) Reinforcement-Learning (RL) algorithm, first developed and employed in the context of robotics, [3], where it was primarily geared toward continuous control tasks. The use of RL in optimization problems is not typical. In fact, there is a vast number of stochastic algorithms and metaheuristics in the market designed to attack large scale optimization problems that do not employ machine learning (ML) techniques. Depending on the problem, some algorithms

may outperform others, but there is no single general-purpose algorithm that dominates. The development of new (hybrid) optimization algorithms and their improvement with the use of ML methods is an active research direction. In our opinion, the use of RL in this general context has promising features and deserves further study.

In this paper we present the results of a very specific exercise: We study the crossing equation of the 4-point function for the superconformal primary in the energy-momentum multiplet of interacting 6D (2,0) superconformal field theories (SCFTs). These theories, which do not have a known Lagrangian description, play a central role in M theory (they capture the low-energy dynamics of multiple M5 branes) and provide canonical examples of the AdS/CFT correspondence in string theory. Even though we study a single-correlator conformal bootstrap problem, the superconformal algebra combined with known facts from chiral-algebra techniques [4,5] leads to a crossing equation, which is expected to admit a constrained set of solutions describing known 6D (2,0) SCFTs.

Throughout the paper we employ a fixed truncation on the spectrum of superconformal primaries that can appear in the conformal block decomposition of the above crossing equation with 45 operators (10 protected up to spin 17, and 35 unprotected up to spin 12). This yields a continuous, nonconvex optimization problem in an 80-dimensional configuration space, where the algorithm aims to maximize a reward. We have chosen the reward as the inverse of a cost function that measures the violation of the truncated crossing equation on a uniform grid of 180 points in cross-ratio space. The coordinates in the

*g.kantor@qmul.ac.uk†niarchos@physics.uoc.gr‡c.papageorgakis@qmul.ac.uk§p.richmond@qmul.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.




FIG. 1. This figure depicts the chosen spectrum of 45 operators in the supergravity limit, where the operators are generalized-free with known scaling dimensions and OPE-squared coefficients. The blue dots represent nonprotected operators in $\mathcal{L}[0, 0]$ superconformal multiplets—see Sec. II. The red dots along the lowest diagonal represent protected operators in the $\mathcal{D}[0, 4]$ (at spin 0 and scaling dimension 8) and $\mathcal{B}[0, 2]$ multiplets (at odd spin).

configuration search space include 35 scaling dimensions of nonprotected operators (in leading and subleading Regge trajectories),¹ and 45 unknown OPE-squared coefficients. We do not solve for the OPE-squared coefficients at each step of the computation to reduce the dimensionality of the search space for reasons that are explained in Sec. III.

Our 80-dimensional search signifies a factor of 2 increase in search-space dimensionality compared to our previously published study cases, in examples of 2D CFTs [1]. It is also, as far as we know, one of the largest searches in truncated crossing equations to-date. However, it is still very far from our ultimate goal of truncations with thousands of operators. The latter would constitute a truly large scale optimization problem, powerful enough to start attacking with good accuracy a wide range of situations, including the bootstrap of multiple correlators. Nevertheless, the fixed truncation used in this paper allows us to demonstrate clearly a particularly interesting implementation of our approach that produces very promising results in a demanding context, where traditional Lagrangian methods are inapplicable and other conformal bootstrap methods have produced relatively few results.

For the purposes of our application, we first fixed the structure of the unknown² truncated spectrum (namely, the number of operators at each spin) in the supergravity limit

¹The scaling dimensions of the remaining 10 protected operators in short superconformal multiplets are fixed by superconformal symmetry, and therefore remain constant throughout our search.

²We emphasize that this truncation refers specifically to the crossing equation after the use of superconformal algebra relations and the known analytic results from chiral-algebra techniques [4,5]. The full crossing equation involves additional operators whose CFT data are known and have been already incorporated—see Sec. II for details.

(at $c = \infty$) where the spectrum is generalized-free (see Fig. 1). We then checked that optimizing around the analytic generalized-free configuration at $c = \infty$, our algorithm recovers a slightly modified configuration with high reward.³ That is a reassuring sign that this particular truncation is an acceptable approximation in the supergravity regime.⁴ Then, by adiabatically changing the value of c while keeping the number of operators fixed, we tracked the evolution of the supergravity spectrum of scaling dimensions and OPE-squared coefficients and produced curves for the 80 CFT data as a function of the central charge c . At the level of the single crossing equation that we are studying it is impossible to detect the discrete nature of c , so throughout our computation the latter is treated as a continuous parameter.⁵

We would like to note that in our approach the spin-partition of Fig. 1 provides crucial, additional information beyond that used to obtain the bootstrap bounds of [5]. Roughly, our method combined with a specific spin-partition determines dynamically the gap assumptions that are known to whittle down the space of allowed solutions to the crossing equations in the standard numerical bootstrap approach. Therefore, the spin-partition limits the scope of the search of solutions and singles out the configurations with minimal violation of the crossing equations. A well-informed spin-partition

³The slight deviation of the obtained configuration from the exact generalized-free result is the expected effect of the truncation.

⁴This does not guarantee, however, that the truncation remains a good approximation for all values of c . One of our goals is to probe this aspect.

⁵A recent, very interesting, study of the Ising CFT using a different adiabatic deformation in spacetime dimension appeared in [6]. This work combined the standard numerical conformal bootstrap technology with the navigator method [7] and the extremal functional method [8].

can be used to achieve a good approximation of a theory. A badly informed, or a randomly chosen, spin-partition yields a landscape of the reward function with optima that are typically harder to interpret as bona fide CFTs. Our adiabatic protocols are designed to address this difficulty.

The most prominent and encouraging features of the results obtained in this manner are the following:

- (1) Adopting a process that is outlined in Sec. III, we managed to produce two distinct sets of curves of high reward. We conjecture that they correspond to the A - and D -series of the 6D (2,0) theories. In previous applications of the numerical conformal bootstrap to 6D (2,0) SCFTs, [5,9], it had been impossible to make a distinction between the two series. This distinction was unclear even in [9], where a very interesting attempt was made to go beyond the rigorous bounds by employing an iterative procedure based on the OPE-inversion formula [10].
- (2) The overall comparison of the CFT data that we produced against the known results from the literature is very promising. Most notably, we demonstrate that the data of the lowest-lying protected and unprotected scalar operators in Figs 3–4, are within the known numerical conformal bootstrap bounds, and compare well with the results obtained using the OPE-inversion formula in [9]. For CFT data at higher spin, some data behave worse compared to the results of [5,9], but the overall emerging picture supports the message that our approach produces competitive results for a wide range of central charge values (even with a fixed truncation of only 45 operators).
- (3) In total, we obtained two sets of curves as a function of c for 80 CFT data. Most of these data refer to operators of unprotected long multiplets in subleading Regge trajectories. It is the first time that predictions have been made for these data. However, the obvious word of caution applies here: As one goes higher in scaling dimension some of these predictions are less reliable. This reflects clearly in, e.g., the statistical standard deviations of our searches.

In our opinion, these results are very encouraging and motivate further work in this direction. With improved algorithms and better implementation one should be able to achieve more reliable, higher-quality numerics. These can then be fed back to the standard numerical bootstrap methods to produce even better results. Along the same lines, for the 6D (2,0) numerical bootstrap, the results presented in this paper are merely a first step toward a more complete investigation and should not be viewed as our final best answer to this problem. As an immediate next step, an enlargement of the truncation by tens of operators is already possible with the current algorithm. We have

made our PYTHON code, `BootSTOP`, public and encourage the interested reader to explore the many possibilities. The broader question of scalability for the present algorithm, or the development of a different more efficient one is currently under investigation.

The rest of this article is organized as follows. In Sec. II we review the key features of the 6D (2,0) crossing equation that we analyze in our computations. In Sec. III we outline the details of the `BootSTOP` package, detail the improvements made compared to [1], and highlight some of the potential alternatives that we have not yet explored. Section IV contains the main results of the paper on the adiabatic evolution of the truncated spectrum as a function of the central charge c . We conclude with a summary and brief comments on future directions in Sec. V. In Appendix we present a blind run of the SAC algorithm with 10 k agents, to demonstrate how it performs in a very wide search (starting with search windows of size 10 in scaling dimensions, and 20 in OPE-squared coefficients).

II. 6D (2,0) CROSSING EQUATIONS

Our main purpose in this section is to set up the language needed to follow our numerical approach. For a more complete review of the pertinent details of 6D (2,0) SCFTs in the context of the conformal bootstrap, we refer the reader to the excellent presentation of Ref. [5].

In the conformal bootstrap analysis of 4-point functions, in any dimension, one needs to solve crossing equations of the general form

$$\sum_I \mathfrak{C}_I F_{\Delta_I}(z, \bar{z}) = 0, \quad (2.1)$$

where \mathfrak{C}_I are OPE-squared coefficients for the I th operator that runs in the conformal block decomposition and $F_{\Delta_I}(z, \bar{z})$ a conformal block (combination) for the corresponding operator, the scaling dimension of which is denoted as Δ_I .⁶ The unknowns of the problem are the CFT data $\{\Delta_I, \mathfrak{C}_I\}$. Equation (2.1) gets contributions from two channels and depends on the complex conjugate parameters (z, \bar{z}) , which express the conformal cross ratios of the four operator insertions. It has to be satisfied at all values of the z parameters where the conformal block expansion (in both channels) is convergent.

To set up the exact crossing equations of the 6D (2,0) theory, we follow closely the Refs. [5,9]. We focus on the crossing equations arising from the 4-point function of superconformal primaries in the energy-momentum multiplet. We remind the reader that unitary representations of the 6D (2,0) superconformal algebra are characterized

⁶In superconformal theories, like the 6D (2,0) theory in this paper, it is convenient to work with the full superconformal blocks, which are then used to express the functions $F_{\Delta_I}(z, \bar{z})$.

by linear relations between the quantum numbers of the superconformal primaries. They can be classified as [11]⁷:

$$\begin{aligned}
 \mathcal{L}: \Delta &> h_1 + h_2 - h_3 + 2(d_1 + d_2) + 6, & h_1 &\geq h_2 \geq h_3 \\
 \mathcal{A}: \Delta &= h_1 + h_2 - h_3 + 2(d_1 + d_2) + 6, & h_1 &\geq h_2 \geq h_3 \\
 \mathcal{B}: \Delta &= h_1 + 2(d_1 + d_2) + 4, & h_1 &\geq h_2 = h_3 \\
 \mathcal{C}: \Delta &= h_1 + 2(d_1 + d_2) + 2, & h_1 &= h_2 = h_3 \\
 \mathcal{D}: \Delta &= 2(d_1 + d_2), & h_1 &= h_2 = h_3 = 0,
 \end{aligned} \tag{2.2}$$

where the quantum numbers of a state are given by $[h_1, h_2, h_3; d_1, d_2; \Delta]$. These correspond to eigenvalues of the $so(6)$ generators of rotations in three orthogonal planes in \mathbb{R}^6 , generators of a Cartan subalgebra of the R -symmetry $so(5)_R$, and the dilatation generator respectively. The conformal dimensions of the short (BPS) multiplets are fixed in terms of the other quantum numbers, while the spin of a state is related to the $so(6)$ quantum numbers through $\ell = h_1 + h_2 - h_3$. The energy-momentum multiplet is denoted $\mathcal{D}[2, 0] \equiv \mathcal{D}[0, 0, 0; 2, 0; 4]$, the superconformal primary of which is a spacetime scalar Φ^{IJ} , $I, J = 1, \dots, 5$, of conformal dimension $\Delta = 4$, transforming in the symmetric-traceless representation of $so(5)_R$.

The OPE selection rules for two energy-momentum multiplets in an interacting 6D (2,0) theory include [13,14]

$$\begin{aligned}
 \mathcal{D}[2, 0] \times \mathcal{D}[2, 0] &= \mathbf{1} + \mathcal{D}[4, 0] + \mathcal{D}[2, 0] + \mathcal{D}[0, 4] \\
 &+ \sum_{\ell=0,2,\dots} (\mathcal{B}[2, 0]_{\ell} + \mathcal{B}[0, 2]_{\ell+1}) \\
 &+ \sum_{\substack{\ell=0,2,\dots \\ \Delta > 6+\ell}} \mathcal{L}[0, 0]_{\Delta, \ell}.
 \end{aligned} \tag{2.3}$$

With this information, the corresponding exact crossing equations boil down to an expression of the form [5]

$$c(z, \bar{z}) - c(1-z, 1-\bar{z}) = 0, \tag{2.4}$$

where (z, \bar{z}) are complex coordinates on the cross-ratio plane, and

$$c(z, \bar{z}) = z\bar{z}(a^u(z, \bar{z}) + a^{\chi}(z, \bar{z})) + \mathcal{C}_h(z, \bar{z}). \tag{2.5}$$

The function \mathcal{C}_h is known exactly:

$$\mathcal{C}_h(z, \bar{z}) = \frac{1}{(z-\bar{z})^3} \frac{h(z) - h(\bar{z})}{z\bar{z}} \tag{2.6}$$

with

$$\begin{aligned}
 h(z) &= -\left(\frac{z^3}{3} - \frac{1}{z-1} - \frac{1}{(z-1)^2} - \frac{1}{3(z-1)^3} - \frac{1}{z}\right) \\
 &- \frac{8}{c} \left(z - \frac{1}{z-1} + \log(1-z)\right) - \frac{1}{6} + \frac{8}{c}
 \end{aligned} \tag{2.7}$$

and c is the central charge of the 6D (2,0) theory, whose value for the A_{N-1}/D_N -series is respectively

$$c_{A_{N-1}} = 4N^3 - 3N - 1, \quad c_{D_N} = 16N^3 - 24N^2 + 9N. \tag{2.8}$$

To complete the definition of (2.4), consider the ‘‘atomic’’ function

$$a_{\Delta, \ell}^{\text{at}}(z, \bar{z}) = \frac{4}{z^6 \bar{z}^6 (\Delta - \ell - 2)(\Delta + \ell + 2)} \mathcal{G}_{\Delta+4}^{(\ell)}(0, -2; z, \bar{z}), \tag{2.9}$$

where $\mathcal{G}_{\Delta}^{(\ell)}(\Delta_1 - \Delta_2, \Delta_3 - \Delta_4; z, \bar{z})$ are the 6D nonsuper-symmetric conformal blocks [15].⁸ Then in (2.5) the piece

$$a^{\chi}(z, \bar{z}) = \sum_{\substack{\ell=0,2,\dots \\ \ell \text{ even}}} 2^{\ell} b_{\ell} a_{\ell+4, \ell}^{\text{at}}(z, \bar{z}) \tag{2.10}$$

represents the contributions from the identity, $\mathcal{D}[2, 0]$, $\mathcal{D}[4, 0]$ and $\mathcal{B}[2, 0]_{\ell-2}$ superconformal multiplets appearing in the self-OPE of Φ^{IJ} from (2.3). The OPE coefficients b_{ℓ} can be fixed through the associated 2D chiral algebra [4,5] and read

$$\begin{aligned}
 b_{\ell} &= \frac{(\ell+1)(\ell+3)(\ell+2)^2 \frac{\ell}{2}! (\frac{\ell}{2}+2)! (\frac{\ell}{2}+3)! (\ell+5)!}{18(\ell+2)! (2\ell+5)!} \\
 &+ \frac{8(2^{-\frac{\ell}{2}-1}(\ell(\ell+7)+11)(\ell+3)! \Gamma(\frac{\ell}{2}+2))}{c(2\ell+5)!}.
 \end{aligned} \tag{2.11}$$

The remaining piece in (2.5), $a^u(z, \bar{z})$, contains the unknown data that one hopes to constrain through the crossing equations. These include the conformal dimension and OPE-squared coefficients corresponding to the long multiplets $\mathcal{L}[0, 0]_{\Delta, \ell}$, and the OPE-squared coefficients corresponding to the short $\mathcal{D}[0, 4]$ and $\mathcal{B}[0, 2]_{\ell-1}$ multiplets. In full,

$$a^u(z, \bar{z}) = \sum_{\substack{\Delta \geq \ell+6 \\ \ell \geq 0, \ell \text{ even}}} \lambda_{\Delta, \ell}^2 a_{\Delta, \ell}^{\text{at}}(z, \bar{z}). \tag{2.12}$$

The conformal dimensions of the $\mathcal{D}[0, 4]$ and $\mathcal{B}[0, 2]_{\ell-1}$ (with $\ell > 0$) multiplets are fixed via (2.2) and given by 8 and $\ell + 7$ respectively (ℓ is always an even integer).

For the purposes of truncating the 6D (2,0) crossing equations to the spin partition displayed in Fig. 1, we

⁷See also [12] for the explicit multiplet construction.

⁸These can also be explicitly found in Appendix B of [5].

specifically used in place of Eqs (2.10), (2.12), the truncated versions

$$a_{\text{tr}}^{\chi}(z, \bar{z}) = \sum_{\ell=0}^{30} 2^{\ell} b_{\ell} a_{\ell+4, \ell}^{\text{at}}(z, \bar{z}) \quad (2.13)$$

and

$$a_{\text{tr}}^{\mu}(z, \bar{z}) = (\lambda_{8,0}^{\mathcal{D}[0,4]})^2 a_{6,0}^{\text{at}}(z, \bar{z}) + \sum_{\ell=2}^{18} (\lambda_{\ell+7, \ell}^{\mathcal{B}[0,2]_{\ell-1}})^2 a_{\ell+6, \ell}^{\text{at}}(z, \bar{z}) \\ + \sum_{\substack{\ell=0 \\ \Delta > \ell+6}}^{12} (\lambda_{\Delta, \ell}^{\mathcal{L}[0,0]_{\Delta, \ell}})^2 a_{\Delta, \ell}^{\text{at}}(z, \bar{z}), \quad (2.14)$$

where $a_{\Delta, \ell}^{\text{at}}(z, \bar{z})$ is given in (2.9).

III. BootSTOP: REMARKS ON SAC IMPLEMENTATION

Throughout this paper we make use of the soft-actor-critic (SAC) reinforcement learning algorithm as a continuous optimizer, as already outlined in the context of the conformal bootstrap in [1]. We will not review the details of the SAC algorithm here, but instead refer the reader to [1] and [16] for further details and useful references. We are also making our PYTHON implementation, BootSTOP, available on GitHub at this link. The README file in that repository contains additional information. In this section we summarize the main components of the package, its functionalities, and highlight the key improvements included in BootSTOP compared to our earlier work [1]. These improvements extend beyond the current application of BootSTOP to the 6D (2,0) theory, and we plan on upgrading its functionality in the near future to incorporate CFTs in diverse dimensions, as well as a choice of different continuous optimizers, in addition to SAC.

When the sum over the intermediate operators I in (2.1) is truncated, the crossing equations cannot be satisfied exactly. Instead, one attempts to minimize an associated semi-positive-definite cost function \mathcal{C} . The choice of this function is not unique. In this paper, we have chosen to discretize the z -plane using a uniform grid of 180 points (following the multipoint scheme of [17]); these are detailed in a file called `data_z_sample.py`. We have also chosen to define \mathcal{C} as the Euclidean quadratic norm of a vector, the components of which are given by the left-hand side (lhs) of (2.1),

$$\sum_I^{\text{cutoff}} \mathfrak{G}_I F_{\Delta_i}(z_i, \bar{z}_i), \quad (3.1)$$

evaluated at each point z_i on the grid. It is straightforward to implement other possibilities, e.g., use Taylor-expansion coefficients around the $z = \frac{1}{2}$ point (which is common

practice in the numerical conformal bootstrap), or define a different cost function. Once the cost function is specified, we define the reward function as $\mathbf{R} = \mathcal{C}^{-1}$, which the algorithm aims to maximize.

A. PYTHON libraries, numerical accuracy, and speed-up techniques

BootSTOP has been implemented using the PyTorch package for PYTHON 3.7. One of the main factors affecting the speed of the code is the evaluation of the conformal blocks, which takes place at every cycle of the algorithm. The conformal blocks contain several hypergeometric functions, the numerical evaluation of which is subtle and costs time. Two commonly used mathematical libraries in PYTHON are “MPMATH” and “SciPy.” The former is slower, but offers greater accuracy (including the possibility of arbitrary numerical precision). The latter is considerably faster, but exhibits known inaccuracies in the evaluation of the hypergeometric functions at specific arguments. We have chosen to employ the SciPy library in BootSTOP and have avoided the evaluation inaccuracies by carefully selecting the points on the z -grid.

An additional and considerable speed-up (typically by a factor of 10) can be obtained by computing the conformal blocks only once outside the main loop. This was carried out by performing a discretization in the space of the unknown scaling dimensions with 60 k lattice points, evaluating the conformal blocks and storing them in a separate set of CSV files.⁹ When BootSTOP starts running, it pre-loads all the conformal blocks before the initialization of the main loop. Once the main loop starts, the agent explores continuously the space of scaling dimensions, but the code rounds them to the closest value on the pre-evaluated scaling-dimension lattice to obtain information about the reward. This procedure limits numerical accuracy but offers a significant speed-up compared to the original implementation in [1].

Another possibility in the implementation of the SAC algorithm is the following. When using a quadratic Euclidean norm cost function it is in principle straightforward to take advantage of the fact that the OPE-squared coefficients appear quadratically: one could solve the corresponding extremization equations to obtain the OPE-squared coefficients at each step as a function of the unknown scaling dimensions. Inserting the answer back into the cost function, results into a new cost function that depends only on the scaling dimensions. The obvious benefit of this method is that it allows one to reduce the dimensionality of the search space significantly. However, there are also considerable disadvantages.

⁹For the runs presented in this paper this amounts to a lattice separation of 5×10^{-4} in the scaling dimensions that enter the hypergeometric functions in the conformal blocks.

First, by employing this intermediate extremization, one changes the optimization problem and finds different solutions at finite truncation. As the size of the truncation grows and one approximates the exact solutions better and better, one expects the global minima of the cost functions in both approaches to converge to each other. It is not, however, obvious how the search is affected by the intermediate extremization at finite truncation (especially, if there are multiple extrema of interest, which are not global minima). Second, the solution of the extremization equations for the OPE-squared coefficients involves the inversion of a matrix, the size of which is set by the number of unknown scaling dimensions. This inversion becomes increasingly computationally expensive as the dimension of the search space grows. Moreover, it can become quite unstable if some scaling dimensions are almost degenerate. Third, this approach is harder to implement with non-quadratic cost functions, which can be useful for some problems. For all these reasons, we have chosen to treat the OPE-squared coefficients as independent unknowns in `BootSTOP`, over which we optimize. Nevertheless, the option of intermediate extremization can be incorporated into our code with the obvious modifications.

B. Summary of parameters and `BootSTOP` functionalities

We now dive a little deeper into the options built into `BootSTOP`. The code has several sets of parameters which we broadly describe as: neural network hyperparameters, learning loop, automation and environment parameters. In this section we briefly summarize the functionality of each set, and highlight them in bold for easy reference. The reader can find the values that were used in our searches on GitHub.

Neural network hyperparameters: These control the behavior of the neural networks employed in the SAC algorithm. The relevant details on the architecture of the SAC algorithm can be found in [1, 16] and references therein.

Learning loop and automation parameters: These control how the code handles the quenching of the size of the search windows and the reinitialization of the SAC algorithm during a run. When a run is initialized, the agent moves through the environment and continuously updates the memory buffer, optimizing the SAC neural networks and the stochastic policies of the associated Markov decision process. After a number of iterations the reward stops improving and the result saturates. Before quenching the size of the search windows, it is useful to repeat the search anew, with the same search windows but retaining knowledge of the previous highest reward configuration. During such a re-initialization, the memory buffer is flushed and the agent starts learning from scratch. Typically, this leads to immediate improvement. The parameter controlling the maximum time spent without improving before re-initialization is called **`faff_max`**. Another parameter, **`pc_max`**, controls how

many re-initializations without improvement in the reward are made before the sizes of the search windows are quenched. Once **`pc_max`** is saturated, the search windows are quenched by some percentage, which is controlled by the **`window_rate`** parameter. The parameter **`max_window_exp`** specifies how many window quenches are carried out before the run ends. We note that this automation was missing in [1], where all reported results were the outcome of single runs on a laptop computer, and with the user manually performing the re-initializations and search-window quenches. `BootSTOP` on the other hand, does not require user supervision and can be run in parallel on a computing cluster.

Environment parameters: This set comprises parameters that specify the environment and how the agent interacts with it. In the former class are parameters like the central charge, which enters explicitly the 6D crossing equations that we are analyzing in this paper. The latter class gives the user control over several functionalities. There are two arrays of Booleans, called **`guessing_run_list_deltas`** and **`guessing_run_list_opes`**, which specify for each CFT datum if the run starts in “guessing mode.” In this mode, every subsequent value in the search is generated stochastically within a fixed range of values specified by the search windows. In nonguessing mode, instead, every next value is generated stochastically around the previous best result. The latter gives the algorithm the capacity to move dynamically over an arbitrarily large, in principle, region even if the search windows are small. The guessing mode is ideal for blind searches with little prior knowledge of where to look. This mode was applied in the blind-search example of Appendix. The nonguessing mode is suitable for searches focused in a particular region of configuration space. This mode was employed in the adiabatic searches of Sec. IV. In `BootSTOP` a search that starts in guessing mode reverts automatically to nonguessing mode after the first **`pc_max`** is reached.

The initial sizes of the search windows are specified by the arrays **`guess_sizes_deltas`** and **`guess_sizes_opes`**. The user can also specify the lower bounds of each CFT datum using two arrays: **`shifts_deltas`** and **`shifts_opecoeffs`**. This is useful as an option that enforces the unitarity bounds in unitary CFTs. Additionally, `BootSTOP` can impose a minimum separation of scaling dimension between operators of the same spin through use of the flag **`same_spin_hierarchy`**. In this context, the uniform gap is set with the parameter **`dyn_shift`**.

C. Statistics, parallel runs, and the search for basins of attraction

The SAC algorithm is in principle able to navigate continuous search spaces and escape local minima through a carefully balanced explore/exploit strategy. However, for high-dimensional, continuous search spaces, where there can be many local minima and the number of saddles

increases exponentially [18], this approach is not feasible in finite time and can be affected by the choice of parameters. For this reason, and in order to improve performance, we took advantage of the automation introduced in the previous subsection, and ran `BootSTOP` in parallel on a computing cluster.

In stochastic searches the collection of statistics with a large number of independent searches is very important, because every run produces a different result. `BootSTOP` can easily run in parallel on computing clusters (with each run being completely independent). For instance, the results in Appendix were obtained using 10 k parallel runs. For the adiabatic results in Sec. IV we had to perform a rather large number of sequential runs at different values of c , which imposed practical restrictions on our running times on the computing cluster. As a result, we ended up performing a parallel run of 200 agents for each value of c , which still produced useful statistics.¹⁰

It is well known that global-search algorithms can perform well at locating the basin of the optimal solution, but are less efficient at identifying the unique, most optimal solution within that basin. The problem becomes harder in high-dimensional searches even when stochastic algorithms are combined with local-search gradient-descent methods. For that reason, we would like to put emphasis on a strategy that focuses primarily on the identification of basins of attraction, and secondarily on the exploration of the small-scale structure within such basins, which can typically be very spiky.

The results of Appendix illustrate how such a strategy would work for the 6D (2,0) SCFT at $c = 25$ in a blind search without any additional, prior, theory-dependent information. In that search, which begins in guessing mode, large arbitrary search windows were chosen for the CFT data (10 for scaling dimensions and 20 for OPE-squared coefficients). The SAC algorithm proves quite effective in identifying basins of attraction in this manner, albeit with significant width. We refer the reader to Appendix for a more detailed discussion of this example and how this run performed against independent expectations. We note that the initial search, and the corresponding identification of the basin, can improve by suitable choices of parameters. Subsequent runs in nonguessing modes around the initial statistical average will further improve the reward and decrease the statistical spread of the results. This, however, can be a delicate and time-consuming process without any guarantees that the global minimum will be identified. In general, the quality of the results arising from this process depend on the specifics of the problem and the corresponding complexity of the search. There are also situations, where the global

minimum may not be the only minimum of physical interest. The 6D bootstrap problem that we analyze in this paper poses such a situation.

An alternative approach is based on a guided, *adiabatic* implementation of SAC starting from the known results for the CFT data in a convenient limit of parameters (in this paper, this is the supergravity limit) and gradually changing the parameters to explore the theory in a general parametric region. In our 6D context, this process involves the adiabatic decrease of the value of the central charge, which appears explicitly in the crossing equation. This procedure can be applied in many different situations. It avoids the pitfalls of having to navigate the full high-dimensional potential, leading to much narrower basins of attraction. We implemented this approach using `BootSTOP` in Sec. IV to obtain very promising results in the 6D (2,0) theory exhibiting two separate basins of attraction.

D. Protocols for adiabatic searches

As already noted, in Sec. IV we present results for adiabatic searches starting from the supergravity limit at large c and gradually flowing toward smaller values up to $c = 25$, where one expects to find the A_1 (2,0) theory. This theory is expected to be the (2,0) SCFT with the smallest possible value of central charge. Anticipating two distinct sets of solutions (one corresponding to the A_{N-1} -series and the other to the D_N -series 6D (2,0) SCFTs), we performed the adiabatic search using two distinct protocols. In both protocols, and for each value of c , we executed two sequential 8 hr runs with 200 parallel agents in nonguessing mode with large `faff_max`.¹¹ This process is especially subtle in the supergravity limit where the two series are expected to be very close. Nevertheless, both protocols worked surprisingly well producing two distinct curves of comparable reward. In Fig. 2 we present the corresponding values of the reward function in each of the two curves. The details of each protocol are as follows.

Protocol 1: The first protocol is the most natural one in an adiabatic search: after changing the value of c by a small amount, the next run is performed around the statistical average of the previous one.¹² During the first run for each new value of c , the `guess_sizes_deltas` parameter was set uniformly at the value 0.01, while the `guess_sizes_opes` was set uniformly at the value 0.001. In the second run at each value of c , these parameters were set at the 1σ standard deviation of the statistics collected during the first run. We conjecture in Sec. IV that this protocol produces a curve for the D_N -series (2,0) SCFTs.

Protocol 2: The second protocol, which is less obvious and more subtle to implement, is motivated by the

¹⁰These searches were implemented on the Queen Mary University of London High Performance Compute cluster Apocrita [19]. Each run utilized 1 core and 5 GB of memory.

¹¹We used the parameters: `faff_max` = 10000, `pc_max` = 5 and `window_rate` = 0.7.

¹²The configuration around which a run is performed can be defined in `BootSTOP` by suitably setting the array `global_best`.




FIG. 2. A plot of the rewards obtained in the adiabatic searches of Sec. IV with the two distinct protocols described in Sec. III D The blue data are conjectured in Sec. IV to capture the A -series (2,0) SCFTs, while the red data the D -series (2,0) SCFTs.

following observations. The OPE-squared coefficient of the single scalar superconformal primary in the $\mathcal{D}[0, 4]$ multiplet, $\lambda_{\mathcal{D}[0,4]}^2$, has the distinguishing characteristic that it vanishes exactly in the A_1 theory. Therefore, one might suspect that the value of $\lambda_{\mathcal{D}[0,4]}^2$ is smaller in the A_{N-1} -series compared to the D_N -series, and that the same happens in the corresponding basins of attraction that represent these solutions in our optimization problem.¹³ This was, indeed, also the expectation in [20].¹⁴ As a result, the second protocol was set up in a manner that would potentially allow us to identify the basin with the smaller value of $\lambda_{\mathcal{D}[0,4]}^2$. In Sec. IV we will provide further supporting evidence for the conjecture that the second protocol captures the A_{N-1} -series (2,0) SCFTs. The two runs of Protocol 2 involved the same guess sizes as in Protocol 1, but all runs were performed around a point with comparatively high reward and smaller value of $\lambda_{\mathcal{D}[0,4]}^2$. Prior to each run, this point was selected after the inspection of the full set of configurations obtained during the previous 200 independent runs. It was chosen as the highest-reward configuration in this set that fit the above criteria on $\lambda_{\mathcal{D}[0,4]}^2$. In many cases, a suitable configuration immediately stood out in the dataset.

¹³Recall that for the purposes of the single-correlator conformal bootstrap exercise that we are performing, c is viewed as a continuous parameter that can be extended to possibly unphysical values. For each c one is, therefore, looking for both a potential A - and D -series extremum. In accordance with these expectations, in independent blind searches with the SAC algorithm at low values of c , we identified frequently two distinct types of basins—one at a higher value of $\lambda_{\mathcal{D}[0,4]}^2$ and another at a lower value.

¹⁴We note that in [21] it was conjectured that the bootstrap bounds for the physical values of c are saturated by AdS_7 maximal supergravity.

IV. RESULTS: ADIABATIC RUNS FROM THE SUPERGRAVITY LIMIT

We will now present and analyze the CFT data produced by the SAC algorithm when used to perform adiabatic deformations of the truncation given in Fig. 1. For this task we employed the two distinct protocols outlined in Sec. III D while slowly varying the central charge from the supergravity limit, $c = \infty$, down to $c = 25$. This section contains the main results of the paper.

We begin with a discussion of the CFT data analyzed in Refs [5,9] using standard numerical conformal bootstrap techniques and the OPE inversion formula. The comparison with [5,9] is useful as a performance check and toward building a consistent interpretation of our results. We conclude with results on CFT data that have not appeared previously in the literature and are new predictions.

A. Lowest protected and unprotected operators

In this subsection we focus on the short $\mathcal{D}[0, 4]$ multiplet, and the first scalar in the long $\mathcal{L}[0, 0]$ multiplet. Let us first discuss the results on the $\mathcal{D}[0, 4]$ multiplet outlined in Fig. 3.

The scaling dimension of the $\mathcal{D}[0, 4]$ multiplet is fixed by superconformal symmetry at the value $\Delta = 8$, but the OPE-squared coefficient between this multiplet and the energy-momentum multiplet depends nontrivially on the central charge c . The left plot in Fig. 3 depicts the results obtained with Protocol 2, while the right plot depicts the results obtained with Protocol 1. There are two sets of data points in each plot. The green points represent statistical averages of the 200 parallel runs, and the purple dots the result with the highest reward. The error bars are determined by the 1σ standard deviation of the statistics obtained on the second 8 hr run at each value of c . In most cases, the result of highest reward is inside the error bars. We observe that in both cases the errors are relatively small, which reflects the




FIG. 3. Plots of the A - and D -series curves for the OPE-squared coefficient of the lowest protected operator in the $D[0, 4]$ multiplet with fixed scaling dimension $\Delta = 8$. The dashed red curve represents the numerical conformal bootstrap upper bound obtained in [5] and the brown curve on the left plot the long-inverted curve of [9]. Regression fits to our data (blue and orange curves) are obtained with the LOWESS method.

fact that the 200 parallel agents concentrated naturally inside a small region for the specific datum during the second run. The blue and orange curves are local regression curves produced using the locally weighted scatterplot smoothing (LOWESS) method. The dashed red curve represents the numerical bootstrap upper bounds obtained in Ref. [5]. Anything above this red curve is expected to be inconsistent in an exact (2,0) SCFT.

Comparing with the brown curve obtained using the OPE-inversion formula in [9], we notice a strong similarity with the Protocol-2 curve and the corresponding data points. Both stay closely below the numerical bootstrap bound and the long-inverted value at $c = 25$ from [9] is also very close to our result, slightly above 0.8.¹⁵ In summary, similar to [9], we have produced a curve below the numerical bootstrap bound, which has not achieved at this level of approximation the expected analytic result of $\lambda_{D[0,4]}^2 = 0$ at $c = 25$. The agreement with [9] is our first strong indication that the SAC algorithm with Protocol 2 has identified a sensible approximation of a family of (2,0) SCFTs. Since the values of the blue curve are consistently below the ones of the orange curve (in the right plot of Fig. 3) we are tempted to conjecture (using intuition from the known supergravity results [13,20] and the related discussion in [20]) that the blue curve represents the A -series (2,0) theories and the orange curve the D -series (2,0) theories. But would such an interpretation be consistent with the existing expectations about the behavior of the D -series curve?

At first sight, the data points on the right plot (fitted by the orange LOWESS curve) appear to behave rather badly compared to the numerical bootstrap bounds. They remain close to the bounds up to roughly $c^{-1/3} \sim 0.11$ and then

gradually start diverging significantly from the bound-curve. The error bars are also increasing toward the small c region and for the last few points close to $c = 25$ the highest-reward values are significantly outside the 1σ error bars. In other words, several features of these data demonstrate that the quality of these results deteriorates at lower values of c . The same qualitative feature is also present in the values of the best reward as a function of c , as depicted in Fig. 2. In that figure the Protocol 1 curve is the red curve. There is an apparent peak around $c^{-1/3} \sim 0.15$ beyond which the reward gradually falls. Interestingly, the smallest central charge for a nontrivial (2,0) SCFT in the D -series is $c = 676$, the central charge of the D_4 theory. The D_3 theory (at $c = 243$) is dual to the A_3 theory, while the D_2 theory (at $c = 50$) is equivalent to the $A_1 \times A_1$ theory. The central charge of the D_4 theory has $c^{-1/3} \sim 0.11$, which is tantalizingly close to the region of 0.11 where we notice the deviation of the orange curve from the bootstrap bound curve. Hence, if we were to terminate the orange curve at $c = 676$ (as an exact analysis of the D -series theories would suggest), the remaining set of data would be satisfyingly close to the numerical bootstrap bounds. Similar observations can be made in all the other CFT data that we have collected, reinforcing the conjecture that the orange curve is consistent with the D -series interpretation.

Consider next the results on the CFT data for the lowest $\mathcal{L}[0, 0]$ multiplet, which are displayed in Fig. 4. The two top-row plots depict the Protocol 2 (conjectured A -series) results for the scaling dimension and OPE-squared coefficient, while the two bottom-row plots the corresponding Protocol 1 (conjectured D -series) results. Repeating the cautionary comment of Footnote 27 in Ref. [9], we point out that the dashed-red numerical bootstrap bounds for the OPE-squared coefficient on the right plots were obtained under the assumption that the corresponding conformal

¹⁵The long-inverted-corrected result at $c = 25$ quoted in [9] is a bit lower, approximately at 0.65.

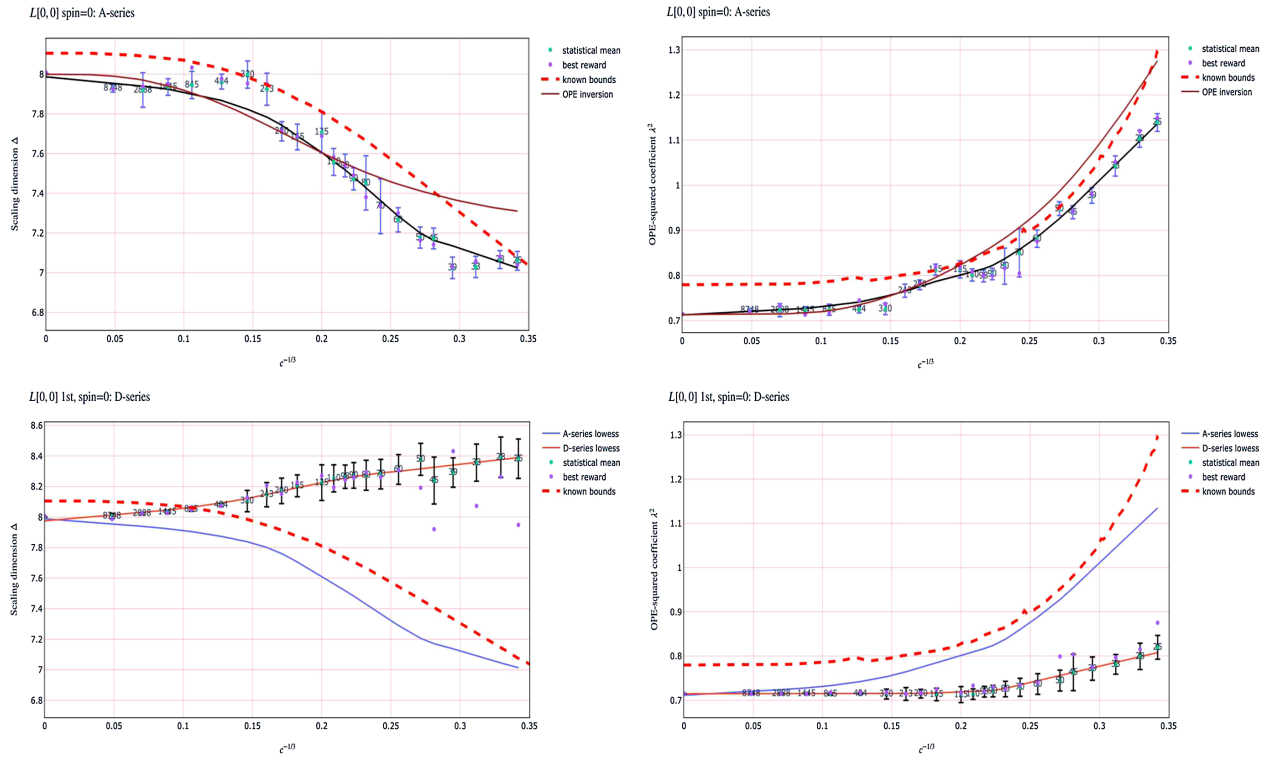


FIG. 4. Plots of the *A*- and *D*-series curves for the scaling dimension and OPE-squared coefficient of the lowest nonprotected operator in the $\mathcal{L}[0, 0]$ multiplet. The plots on the left depict the scaling-dimension results and the plots on the right the OPE-squared coefficient results. The brown curves on the top two plots exhibit the long-inverted curves of [9].

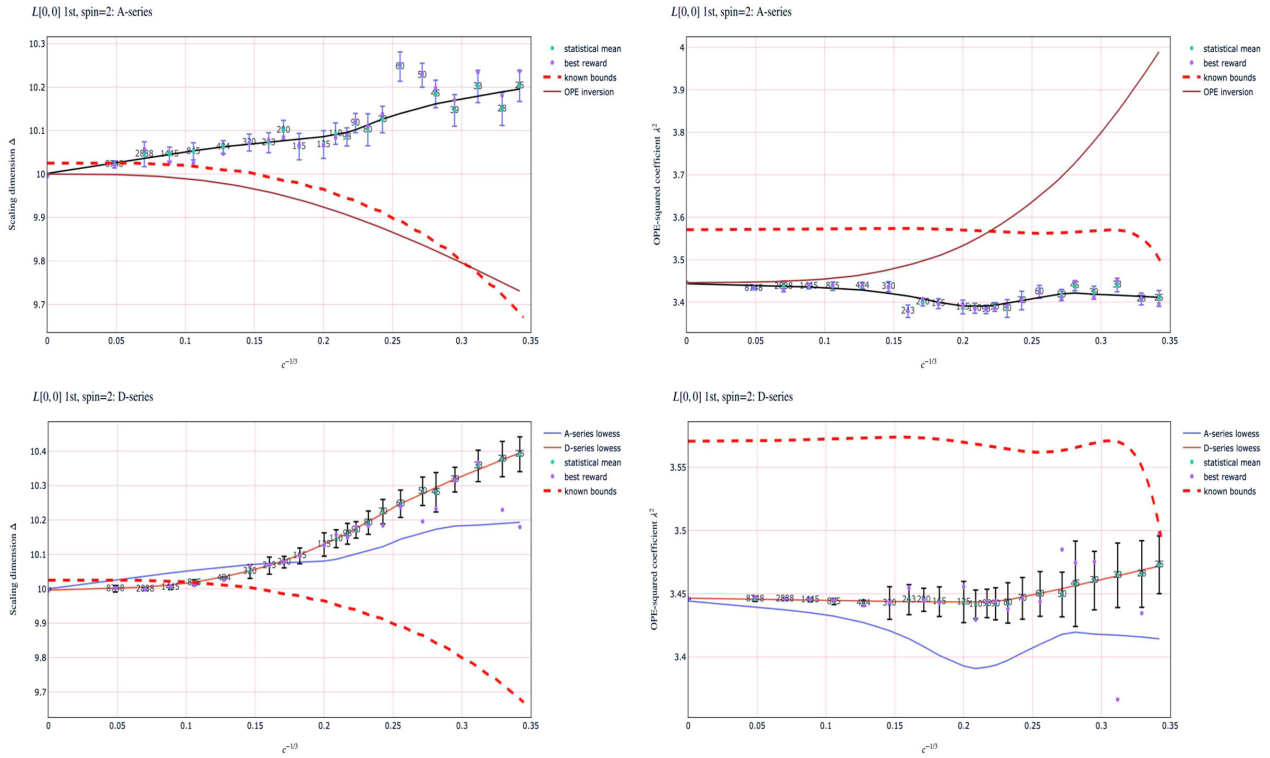


FIG. 5. Plots of the *A*- and *D*-series curves for the scaling dimension and OPE-squared coefficient of the leading nonprotected spin-2 operators in the $\mathcal{L}[0, 0]$ multiplet. The plots on the left depict the scaling-dimension results and the plots on the right the OPE-squared coefficient results. The brown curves on the top two plots exhibit the long-inverted curves in [9].

dimensions saturate their own bounds, which is not directly applicable here. As in [9], we depict these bounds under the expectation that they are a decent estimate to the actual bounds.

The plots in Fig. 4 reveal a picture consistent with the observations of the previous paragraphs. The A -series data are consistently below the numerical bootstrap bounds for all values of c both for the scaling dimension and OPE-squared coefficient. In fact, when compared to the brown curve obtained using the OPE inversion formula in [9], we see that our results are behaving slightly better in relation to the known numerical bootstrap bounds. For the D -series data we observe that they exhibit the same qualitative features that were noted in the context of the $\mathcal{D}[0,4]$ multiplet above. Their quality deteriorates beyond the region of $c^{-1/3} \sim 0.15$, and a violation of the numerical bootstrap bounds for the scaling dimension starts in the vicinity of $c^{-1/3} \sim 0.11$. Hence, if one were to terminate the curve at $c = 676$, the agreement with the bounds would be very good both for the scaling dimension and OPE-squared coefficient.

We believe that all these results, collectively, are very suggestive for the efficiency of the SAC algorithm in this particular context, and the conjectured interpretation of the results.

B. Comparison with the remaining previously analyzed CFT data

We continue with the presentation of results for higher-spin operators for which certain predictions have already been made in [5,9]. More specifically, we discuss scaling dimensions and OPE-squared coefficients for long $\mathcal{L}[0,0]$ multiplets at spin $\ell = 2, 4, 6$ and OPE-squared coefficients for short $\mathcal{B}[0,2]$ multiplets at spin $\ell = 1, 3, 5$. For the depicted numerical bootstrap bounds of the OPE-squared coefficients of the long multiplets the comment of Footnote 27 in [9] still applies.

1. $\mathcal{L}[0,0]$ for spin $\ell = 2, 4, 6$

Based on the above interpretation, the A - and D -series data of the leading spin-2 operators in the $\mathcal{L}[0,0]$ multiplet appear in Fig. 5. The A -series curve for the scaling dimension (top-left plot) violates the known numerical bootstrap bounds for most values of c . It behaves significantly worse than the corresponding OPE inversion curve of Ref. [9], which also violates the numerical bootstrap bounds, but with a violation that occurs at much smaller values of c . For the OPE-squared coefficient (right-top plot) the situation is reversed. The (indicative) bootstrap bounds are respected throughout our A -series curve, but are violated significantly by the corresponding OPE inversion

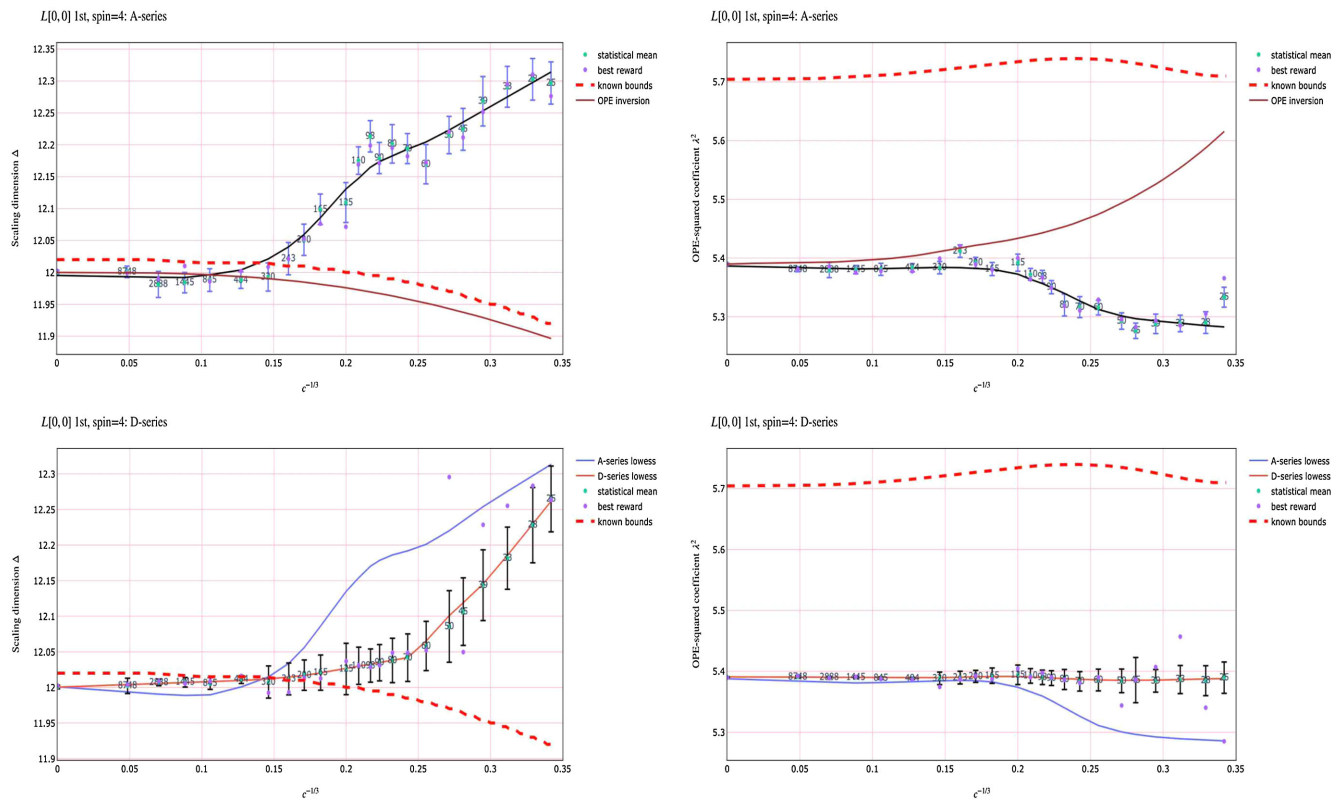


FIG. 6. Plots of the A - and D -series curves for the scaling dimension and OPE-squared coefficient of the leading nonprotected spin-4 operators in the $\mathcal{L}[0,0]$ multiplet. The plots on the left depict the scaling-dimension results and the plots on the right the OPE-squared coefficient results. The brown curves on the top two plots exhibit the long-inverted curves in [9].

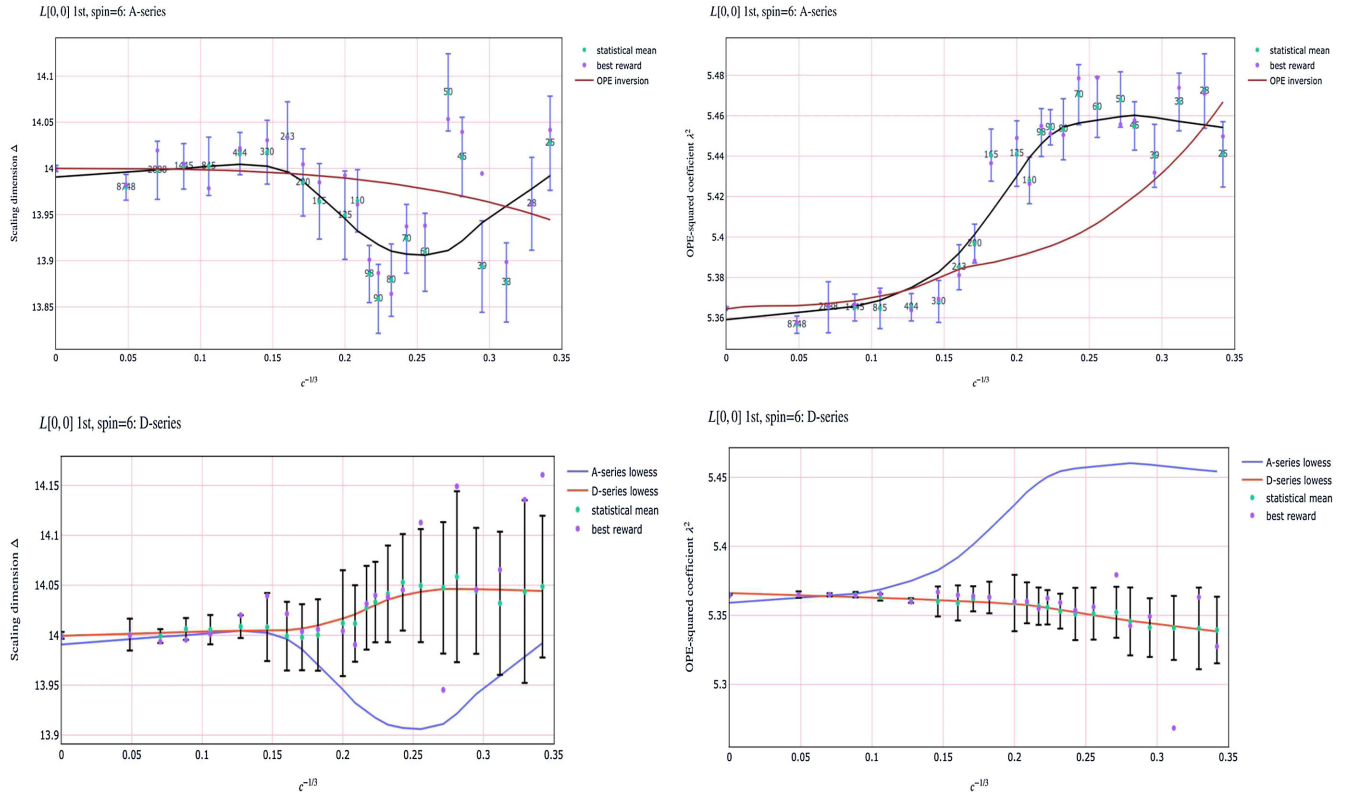


FIG. 7. Plots of the A - and D -series curves for the scaling dimension and OPE-squared coefficient of the leading nonprotected spin-6 operators in the $\mathcal{L}[0,0]$ multiplet. The plots on the left depict the scaling-dimension results and the plots on the right the OPE-squared coefficient results. The brown curves on the top two plots exhibit the long-inverted curves in [9].

curve of Ref. [9]. For the D -series scaling-dimension curve (bottom-left plot) we observe the same effect as in previous data: a significant violation occurs only for values of $c^{-1/3}$ above ~ 0.11 . The OPE-squared curve (bottom-right plot) remains everywhere below the bound.

The corresponding data of the leading spin-4 operators in the $\mathcal{L}[0,0]$ multiplet exhibit comparable features (see Fig. 6). The A -series curve for the scaling dimension (left-top plot) violates the bound significantly for low enough values of c , but the D -series scaling dimensions do not, if the curve is terminated around $c^{-1/3} \sim 0.11$. In comparison, the corresponding OPE-inversion scaling-dimension curve in [9] (see brown curve in the left-top plot) does not violate the bootstrap bounds. Finally, both of our OPE-squared curves are well within the (indicative) allowed regions at slightly lower values compared to the corresponding brown OPE-inversion curve (see brown curve in the right-top plot).

For the leading spin-6 operators in the $\mathcal{L}[0,0]$ multiplet we are not aware of any published numerical bootstrap bounds. Our A - and D -series results appear in Fig. 7. Compared with the corresponding curves derived using the OPE-inversion formula, our A -series data evolve in the same neighborhood of values (with visible differences in the shape of the curves). The scaling dimensions are in the

neighborhood of 13.9 and 14 and the OPE-squared coefficients in the neighborhood 5.35 and 5.45. Once again, we observe that the error bars of the D -series curves exhibit a visible increase around the point $c^{-1/3} \sim 0.15$.

2. $\mathcal{B}[0,2]$ for spin $\ell = 1, 3, 5$

To complete the comparison with the data obtained in Ref. [9], we exhibit in Fig. 8 the A and D versions of the curves for the OPE-squared coefficients of the $\mathcal{B}[0,2]$ multiplets at spin $\ell = 1, 3, 5$. The A -series curves appear on the left column while the D -series curves on the right column. Violations of the numerical bootstrap bounds are observed in the spin-1 and spin-3 cases for the A -series. For spin 5 our results are comparable with the results of [9], with the most significant deviation in the rough region of $c^{-1/3}$ between 0.2 and 0.25. Once again, the D -series curves confirm the pattern observed in previous data.

3. Summary of the comparison with known bootstrap results

In summary, so far in this section we presented two versions of 12 CFT data. In 9 of them we could compare with known numerical bootstrap bounds. When terminated in the vicinity of $c = 676$, the D -series curves did not

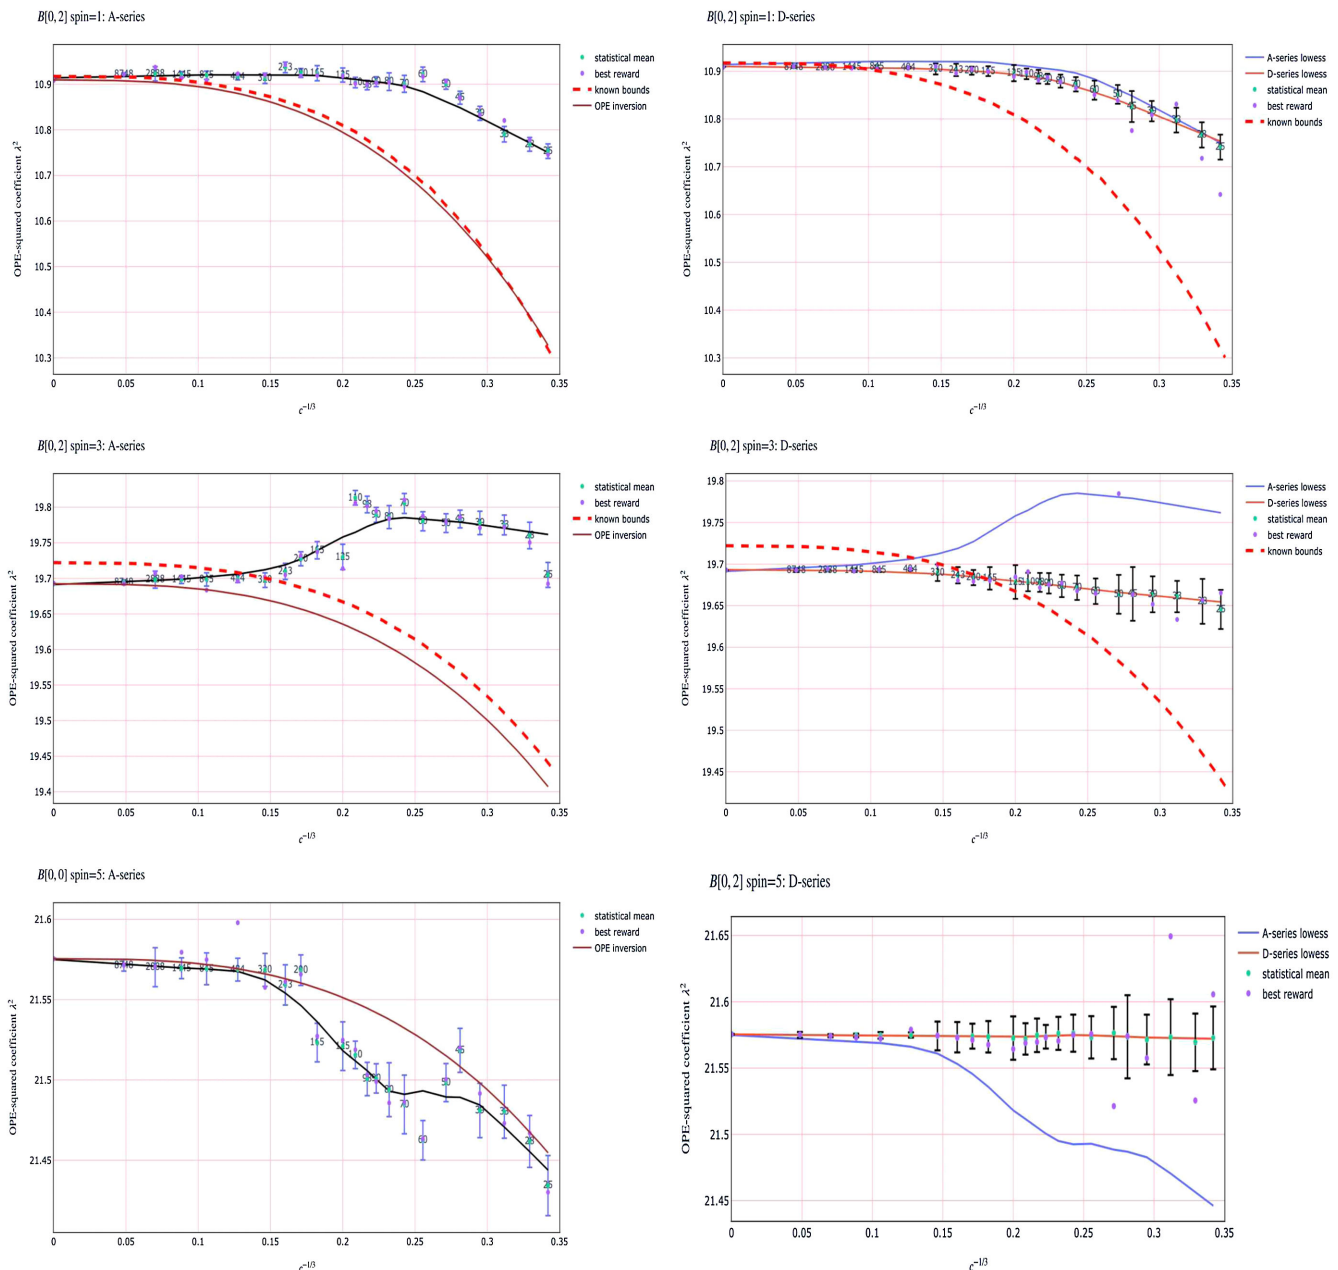


FIG. 8. Plots of the A - and D -series curves for the OPE-squared coefficients of the protected operators in the $\mathcal{B}[0, 2]$ multiplets at spin $\ell = 1, 3, 5$. The brown curves on the left three plots exhibit the long-inverted curves in [9].

exhibit any significant violations of these bounds. The A -series curves extended all the way down to $c = 25$ and violated the bounds for sufficiently low values of c in 4 cases: the spin-2 and spin-4 scaling dimensions of $\mathcal{L}[0, 0]$ multiplets in Figs. 5, 6 and the spin-1 and spin-3 OPE-squared coefficients of the $\mathcal{B}[0, 2]$ multiplets in the corresponding plots of Fig. 8. These violations reflect the limitations of the specific truncation of 45 operators that we used throughout the computation. It would be very interesting to increase significantly the number of operators and determine the extent to which these violations persist. We have not explored this very important aspect of the

problem in the present work, because we want to treat it separately elsewhere within a more focused investigation of systematic errors.

C. More results beyond the leading Regge trajectory

In total, the above application of the SAC algorithm has provided an A and a D -series version of 80 CFT data. Most of these data are new predictions that have not appeared previously in the literature. We expect the quality of these predictions to be worse for operators that are closer to the

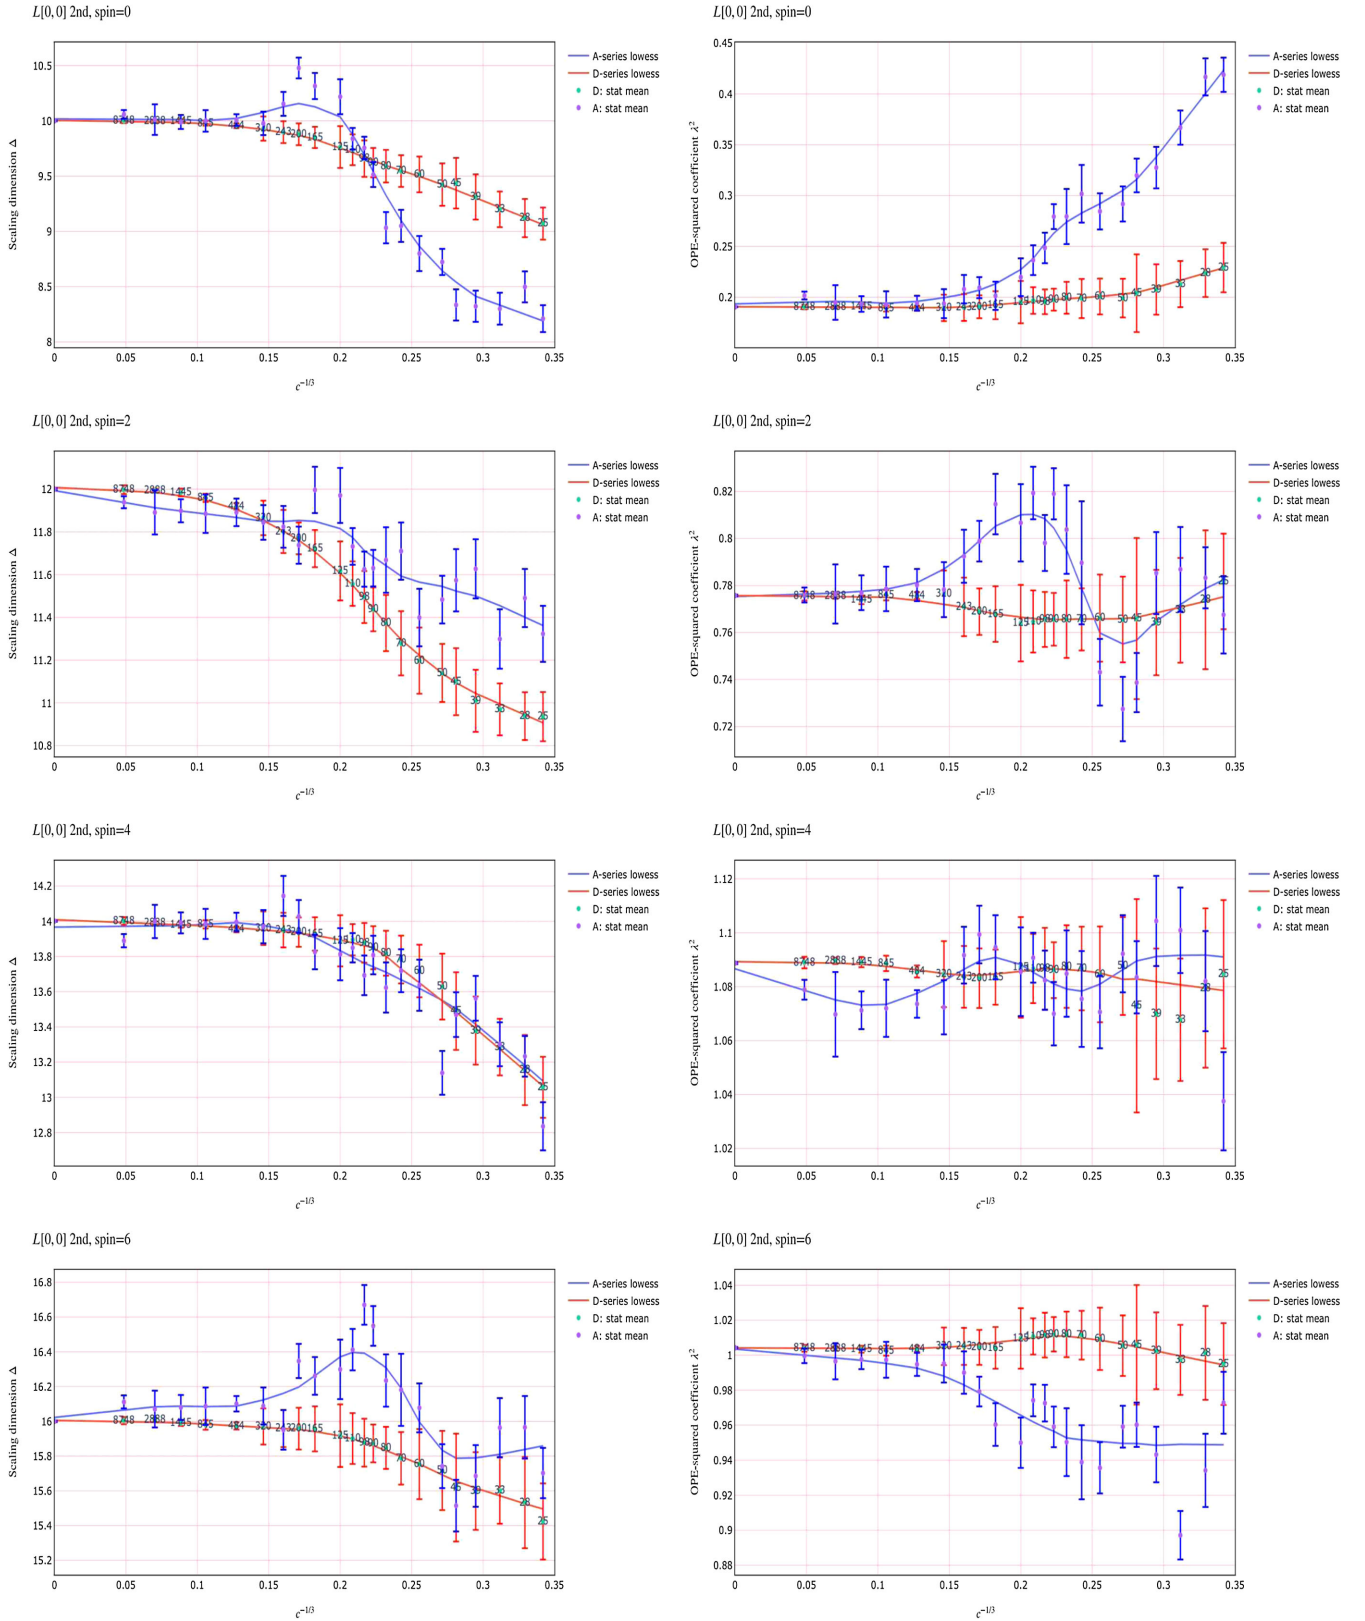


FIG. 9. Plots of the A- and D-series curves for the scaling dimensions and OPE-squared coefficients of the first subleading nonprotected operators at spin $\ell = 0, 2, 4, 6$.

upper bound (in scaling dimension) of the truncation. In our case, the truncation contained long operators up to spin 12 and short operators up to spin 17. In the previous subsections we observed that the results for the lowest-lying operators are consistent with known bounds, but some higher spin data in leading Regge trajectories exhibit violations of known bounds. We expect such potential violations for other higher-dimension data in leading and subleading Regge trajectories. Nevertheless, for future reference we present in Fig. 9 data for the first subleading long operators at spin $\ell = 0, 2, 4, 6$. PDF files with all the plots presented in this paper, and more plots that were not included here, can be found at this GitHub folder.

V. CONCLUSIONS AND OUTLOOK

In this paper we employed the SAC algorithm as a stochastic optimizer to obtain 80 CFT data for protected and unprotected operators in interacting 6D (2,0) theories. These data appear in the truncated superconformal block expansion of the 4-point function of superconformal primaries in the energy-momentum multiplet. The computations were carried out using our `BootSTOP` package. Our best results were achieved via an adiabatic strategy, whereupon starting from the analytic supergravity values for the CFT data, the central charge was gradually changed to $c = 25$, the lowest expected value corresponding to a physical theory. We were able to identify two curves for each datum. There is evidence that suggests that each of these curves corresponds to the *A*- and *D*-series (2,0) SCFTs. On the whole, these results are competitive when compared to numerical bootstrap bounds [5] and more recent approaches using the OPE inversion formula [9], even with a truncation of only 45 unknown operators (a number which is still low compared to our ultimate target for large-scale optimization searches). We expect that the accuracy of the reported results will be improved upon increasing the size of the truncation. The rate and limit of improvement are important open questions, which are part of our general on-going investigation. The addition of tens of extra operators can be explored immediately with the current version of `BootSTOP`. We encourage the interested reader to do so.

Although `BootSTOP` is currently using SAC exclusively to optimize the 6D (2,0) crossing equations, we intend to enlarge its functionality in the very near future. This will include more traditional stochastic-optimization algorithms, such as Metropolis–Hastings Monte–Carlo that was recently used in a similar context [22], and will allow it to serve as a benchmarking tool. Similarly, we will incorporate conformal blocks for a number of dimensions, so that `BootSTOP` can be used by the inquisitive reader to attack a number of interesting CFTs.

This research utilized Queen Mary’s Apocrita HPC facility, supported by QMUL Research-IT [19].

`BootSTOP` can be publicly accessed on GitHub at [23]. The specific values that we used in our adiabatic runs can be found at [24].

ACKNOWLEDGMENTS

We would like to thank J. Halverson, P. Kravchuk, F. Ruehle, K. Skenderis, A. Stergiou and especially M. Lemos for stimulating discussions. We would also like to thank the authors of [9] for sharing the data contained in their plots. C.P. is supported in part by a Royal Society University Research Fellowship No. URFAR\180009 and the STFC Consolidated Grant No. ST/P000754/1. P.R. is funded through the Royal Society Research Fellows Enhancement Grant No. RGF\EA\181049. C.P. and V.N. would like to acknowledge the Mainz Institute for Theoretical Physics for hospitality during the workshop “A Deep-Learning Era of Particle Theory.” V.N. would also like to acknowledge Queen Mary University of London for hospitality during the completion of this work.

APPENDIX: EXAMPLE OF A WIDE SEARCH FROM SCRATCH AT $c = 25$

In this appendix we showcase the results of a blind search without any additional, theory-dependent prior information, at $c = 25$. The parameters used for this search were: **guess_sizes_deltas** = 10 for all scaling dimensions, **guess_sizes_opes** = 20 for all OPE-squared coefficients, while choosing **faff.max** = 12000, **pc.max** = 5 and **window.rate** = 0.5, to allow the SAC algorithm enough time to explore the search space. Each search started in guessing mode, the functionality of which was detailed in Sec. III B. We ran 10k agents on the computing cluster, with each run terminating after 24 hrs.

A small sample of our results is collected in Fig. 10, where we display the CFT data for the $\mathcal{D}[0, 4]$ and spin-1 $\mathcal{B}[0, 2]$ protected operators, as well as for the spin-0 and spin-2 $\mathcal{L}[0, 0]$ unprotected operators in the leading Regge trajectory. The plots contain distributions of CFT-data values, for the top 100 k reward improvements achieved by any agent among the 10 k runs. Note that this is not the final, best reward for each run, so there could be imbalanced contributions from a small number of particularly lucky agents.¹⁶ For the rewards of data thus selected within a given bin, a corresponding weighted average is computed, where the weight was chosen as the square of the reward normalized by the overall best reward. The resulting heat map exhibits the formation of basins of attraction, with the highest concentrations of weighted-average reward denoted as deeper red on the plots. Although the basins of attraction

¹⁶We experimented with many blind-search runs in different contexts and with different parameters. Another way to plot the results is with heat maps that record the 5 best rewards of each agent. Typically, these maps exhibit very similar qualitative features to the heat maps presented here.

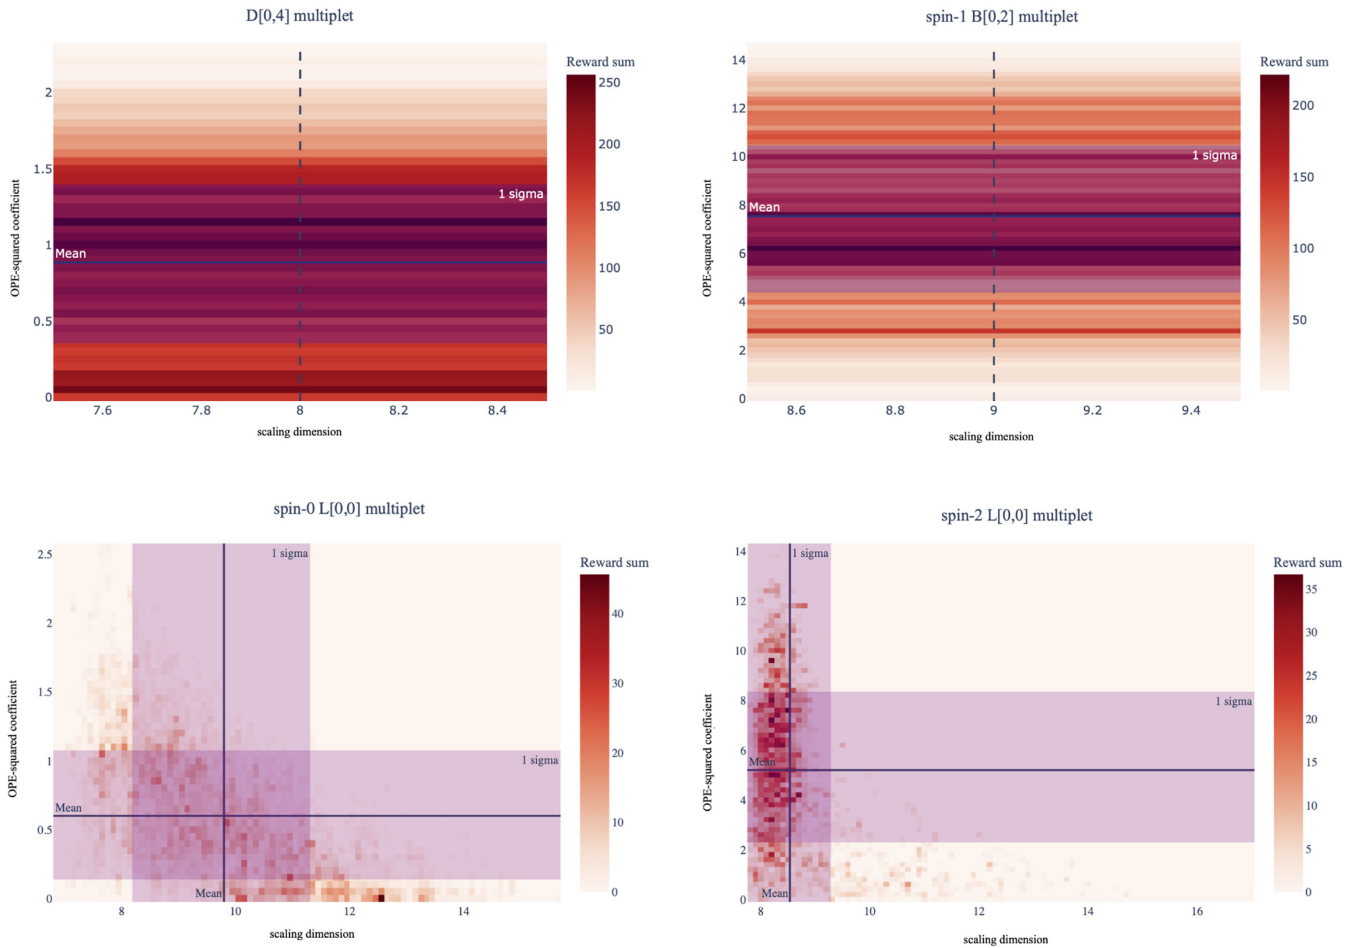


FIG. 10. Sample heat-map plots for the blind-search results of 10 k agents for $\mathcal{D}[0, 4]$, $\ell = 1$ $\mathcal{B}[0, 2]$ and $\ell = 0, 2$ $\mathcal{L}[0, 0]$ multiplets in the leading Regge trajectory. In the two top plots the scaling dimension is fixed at the value denoted by the vertical line. We plot the heat map with uniform color intensity on the horizontal axis for visual effect. The shaded region, which is more clearly visible in the two bottom plots, represents the 1σ standard deviation spread of the results.

exist in the full 80-dimensional search space, we find it natural to portray a two-dimensional projection involving the scaling dimension and OPE-squared coefficient for each operator. It is also useful to include the statistical mean and 1σ standard deviation.

Overall, the SAC algorithm proves quite effective in identifying basins of attraction. We observe that the values expected from bootstrap bounds, the OPE inversion formula or our adiabatic searches are in general within the blind-search basins.

For example, for the OPE-squared coefficient of the $\mathcal{D}[0, 4]$ multiplet, we observe that the basin in the top-left of plot of Fig. 10 contains both the A -series value at $\lambda_{\mathcal{D}[0,4]}^2 \sim 0.8$ and the D -series value at $\lambda_{\mathcal{D}[0,4]}^2 \sim 1.4$ obtained in the adiabatic search of Sec. IV. For the spin-1 $\mathcal{B}[0, 2]$ multiplet the conformal bootstrap bound gives a value $\lambda_{\mathcal{B}[0,2]}^2 \sim 10.3$ and our adiabatic searches $\lambda_{\mathcal{B}[0,2]}^2 \sim 10.75$. Both values are contained within the upper part of the region of interest identified by the basin in the top-right plot of Fig. 10.

The plots for the unprotected $\mathcal{L}[0, 0]$ operators exhibit a similar clustering. The expected values for the scaling dimensions and OPE-squared coefficients for spin 0 are $\Delta_{\mathcal{L}[0,0]_0} \sim 7$ and $\lambda_{\mathcal{L}[0,0]_0}^2 \sim 1.1$ from the A -series adiabatic search and the numerical bootstrap bounds, and $\Delta_{\mathcal{L}[0,0]_0} \sim 8.4$ and $\lambda_{\mathcal{L}[0,0]_0}^2 \sim 0.8$ from the D -series adiabatic search. The 1σ regions of the bottom-left plot in Fig. 10 favors the D -series result. The clustering is more pronounced in the spin-2 case. The expected values from the bootstrap bounds ($\Delta_{\mathcal{L}[0,0]_2} \sim 9.6$ and $\lambda_{\mathcal{L}[0,0]_2}^2 \sim 3.5$) are within the 1σ region. The scaling dimension expected from the adiabatic searches is higher at $\Delta_{\mathcal{L}[0,0]_2} \sim 10.2$ and lies in the upper part of the distribution, slightly above the 1σ boundary.

The rest of our results exhibit similar qualitative features.¹⁷ In these results, one can also observe that there are

¹⁷Plots of the remaining CFT data can be made available upon request.

cases like the OPE-squared coefficient of the spin-5 $\mathcal{B}[0, 2]$ multiplet, where although the original region of search in the guessing mode was between 0 and 20, the subsequent search in the nonguessing mode dynamically evolved to explore values above 20, where one anticipates to find the expected result.

The blind-search results should be understood as the first step in a longer process. One should perform additional iterations around the initial statistical averages obtained here, with the search windows adjusted, for example, to the 1σ ranges for each unknown. This typically improves the reward, makes the basins of attraction immediately more pronounced and further reduces the statistical spread. It can also be accompanied by a

corresponding tuning of other parameters. Unfortunately, as we also stressed in Sec. III C, due to the generically complicated nature of high-dimensional search spaces there is no guarantee that the global minimum will be reached in this fashion, without an additional delicate and time-consuming fine tuning of parameters. Moreover, there can be situations where there are multiple minima corresponding to physical theories that this approach will not be able to distinguish. For these reasons, we find the adiabatic method of Sec. IV far more promising as a tool toward reaching the final more accurate answer. Blind searches with large windows will be less accurate in general, but can still be useful at providing further intuition about the problem at hand.

-
- [1] G. Kántor, V. Niarchos, and C. Papageorgakis, Solving Conformal Field Theories with Artificial Intelligence, *Phys. Rev. Lett.* **128**, 041601 (2022); Conformal bootstrap with reinforcement learning, *Phys. Rev. D* **105**, 025018 (2022).
- [2] F. Gliozzi, More Constraining Conformal Bootstrap, *Phys. Rev. Lett.* **111**, 161602 (2013); F. Gliozzi and A. Rago, Critical exponents of the 3d Ising and related models from conformal bootstrap, *J. High Energy Phys.* **10** (2014) 042; F. Gliozzi, P. Liendo, M. Meineri, and A. Rago, Boundary and interface CFTs from the conformal bootstrap, *J. High Energy Phys.* **05** (2015) 036; F. Gliozzi, Truncatable bootstrap equations in algebraic form and critical surface exponents, *J. High Energy Phys.* **10** (2016) 037; W. Li, New method for the conformal bootstrap with OPE truncations, [arXiv:1711.09075](https://arxiv.org/abs/1711.09075).
- [3] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, [arXiv:1801.01290](https://arxiv.org/abs/1801.01290).
- [4] C. Beem, L. Rastelli, and B. C. van Rees, \mathcal{W} symmetry in six dimensions, *J. High Energy Phys.* **05** (2015) 017.
- [5] C. Beem, M. Lemos, L. Rastelli, and B. C. van Rees, The (2, 0) superconformal bootstrap, *Phys. Rev. D* **93**, 025016 (2016).
- [6] J. Henriksson, S. R. Kousvos, and M. Reehorst, Spectrum continuity and level repulsion: The Ising CFT from infinitesimal to finite ϵ , [arXiv:2207.10118](https://arxiv.org/abs/2207.10118).
- [7] M. Reehorst, S. Rychkov, D. Simmons-Duffin, B. Sirois, N. Su, and B. Van Rees, Navigator function for the conformal bootstrap, *SciPost Phys.* **11**, 072 (2021).
- [8] S. El-Showk and M. F. Paulos, Bootstrapping Conformal Field Theories with the Extremal Functional Method, *Phys. Rev. Lett.* **111**, 241601 (2013); Extremal bootstrapping: Go with the flow, *J. High Energy Phys.* **03** (2018) 148.
- [9] M. Lemos, B. C. van Rees, and X. Zhao, Regge trajectories for the (2, 0) theories, *J. High Energy Phys.* **01** (2022) 022.
- [10] S. Caron-Huot, Analyticity in spin in conformal theories, *J. High Energy Phys.* **09** (2017) 078.
- [11] S. Minwalla, Restrictions imposed by superconformal invariance on quantum field theories, *Adv. Theor. Math. Phys.* **2**, 783 (1998); V. K. Dobrev, Positive energy unitary irreducible representations of $D = 6$ conformal supersymmetry, *J. Phys. A* **35**, 7079 (2002).
- [12] M. Buican, J. Hayling, and C. Papageorgakis, Aspects of superconformal multiplets in $D > 4$, *J. High Energy Phys.* **11** (2016) 091; C. Cordova, T. T. Dumitrescu, and K. Intriligator, Multiplets of superconformal symmetry in diverse dimensions, *J. High Energy Phys.* **03** (2019) 163.
- [13] P. J. Heslop, Aspects of superconformal field theories in six dimensions, *J. High Energy Phys.* **07** (2004) 056.
- [14] B. Eden, S. Ferrara, and E. Sokatchev, (2,0) superconformal OPEs in $D = 6$, selection rules and nonrenormalization theorems, *J. High Energy Phys.* **11** (2001) 020.
- [15] F. A. Dolan and H. Osborn, Conformal partial waves and the operator product expansion, *Nucl. Phys.* **B678**, 491 (2004); Conformal partial waves: Further mathematical results, [arXiv:1108.6194](https://arxiv.org/abs/1108.6194).
- [16] G. Kantor, Exact and numerical methods in (super)conformal field theories, <https://qmro.qmul.ac.uk/xmlui/handle/123456789/79600>.
- [17] A. Castedo Echeverri, B. von Harling, and M. Serone, The effective bootstrap, *J. High Energy Phys.* **09** (2016) 097.
- [18] Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, Identifying and attacking the saddle point problem in high-dimensional nonconvex optimization, in *Advances in Neural Information Processing Systems*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger (Curran Associates, Inc., 2014).
- [19] T. King, S. Butcher, and L. Zalewski, Apocrita—High performance computing cluster for Queen Mary University of London, *J. High Energy Phys.* **11** (2016) 091.

-
- [20] L. F. Alday, S. M. Chester, and H. Raj, 6d (2,0) and M-theory at 1-loop, *J. High Energy Phys.* **01** (2021) 133.
- [21] L. F. Alday and S. M. Chester, Pure Anti-de Sitter Supergravity and the Conformal Bootstrap, *Phys. Rev. Lett.* **129**, 211601 (2022).
- [22] A. Laio, U. L. Valenzuela, and M. Serone, Monte Carlo approach to the conformal bootstrap, *Phys. Rev. D* **106**, 025019 (2022).
- [23] <https://github.com/vniarchos/BootSTOP>.
- [24] https://github.com/vniarchos/BootSTOP/blob/main/Applications/6D/2022_08_results/data_z_sample.py.