

Monte Carlo approach to the conformal bootstrap

Alessandro Laio,^{1,2} Uriel Luviano Valenzuela^{1,3} , and Marco Serone^{1,3} 

¹SISSA, Via Bonomea 265, I-34136 Trieste, Italy

²ICTP, Strada Costiera 11, I-34151 Trieste, Italy

³INFN, Sezione di Trieste, Via Valerio 2, I-34127 Trieste, Italy



(Received 16 June 2022; accepted 18 July 2022; published 26 July 2022)

We introduce an approach to find approximate numerical solutions of truncated bootstrap equations for conformal field theories (CFTs) in arbitrary dimensions. The method is based on a stochastic search via a Metropolis algorithm guided by an action S which is the logarithm of the truncated bootstrap equations for a single scalar field correlator. While numerical conformal bootstrap methods based on semidefinite programming put rigorous exclusion bounds on CFTs, this method looks for approximate solutions, which correspond to local minima of S , when present, and can be even far from the extremality region. By this protocol we find that if no constraint on the operator scaling dimensions is imposed, S has a single minimum, corresponding to the free theory. If we fix the external operator dimension, however, we encounter minima that can be studied with our approach. Imposing a conserved stress-tensor, a \mathbf{Z}_2 symmetry and one relevant scalar, we identify two regions where local minima of S are present. When projected in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane, σ and ϵ being the external and the lightest exchanged operators, one of these regions essentially coincides with the extremality line found in previous bootstrap studies. The other region is along the generalized free theories in $d = 2$ and below that in both $d = 3$ and $d = 4$. We empirically prove that some of the minima found are associated to known theories, including the $2d$ and $3d$ Ising theories and the $2d$ Yang-Lee model.

DOI: [10.1103/PhysRevD.106.025019](https://doi.org/10.1103/PhysRevD.106.025019)

I. INTRODUCTION

Starting from the pioneering work [1], numerical bootstrap methods have become a widespread work-tool for studying conformal field theories (CFTs) in $d > 2$ space-time dimensions, see [2] for a review with an extensive list of references. The most established and developed methods make use of semidefinite programming algorithms applied to 4-point functions of local operators to rule in or rule out theories, the state of the art program being currently SDPB 2.0 [3,4]. In this approach one makes assumptions on the CFT spectrum (typically on the low-lying scaling dimension of primary operators) and is able to rigorously establish if that assumption is allowed or not in the space of all possible CFTs. By imposing suitable assumptions to a set of 4-point functions one can restrict the allowed region to a small island in parameter space, so that precise and rigorous properties of theories living at the extremality region can be determined [5–7]. Such CFTs at the edge of the extremality bounds can also be studied using the so called extremal functional method, first discussed in [8] and

later developed in [9,10]. Although not as rigorous as the previous method, the extremal functional allows a quite accurate and precise determination of the CFT data, see e.g., [11] for an application to the $3d$ Ising model. These powerful methods generally require that the theory in question lies at the boundary of the allowed region in parameter space to start with. In contrast, the “interior” of the space of allowed CFTs is harder to characterize using the current techniques.

In fact, it could be useful to have some numerical tool to explore the structure of the space of allowed CFTs and, in principle, a way to construct approximate CFT data that are consistent with crossing symmetry of 4-point correlation function of local operators.¹ Locating, even approximately, putative CFTs could be a useful first step to understand which assumptions to impose so that these theories become extremal and hence amenable to, e.g., a fully rigorous bootstrap analysis. The aim of this paper is to propose a novel numerical method to analyze the bootstrap equations and possibly find a subset of its approximate solutions. More specifically, we will try to answer to the following question: Given a finite set of n_i operators with spin l_i below a given scaling dimension Δ^* , exchanged in some 4-point correlation function, is there a set of CFT data (scaling dimensions and OPE coefficients) for the n_i

¹See [12] for recent progress in this direction.

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.

operators which best satisfy the truncated crossing equations? If so, where is the putative solution located in the allowed CFT parameter space?

In one-dimensional systems, where there is no spin, accurate solutions can be constructed [13]. For $d > 1$ this problem is much harder because one has to look for solutions by choosing among the many possible partitions of the operators in the different spin channels. In addition to that, it is unclear how to choose the initial conditions for any search and which constraints should be imposed to possibly get a finite set of solutions. Finally, it is far from clear that the search would lead to an actual theory. In particular, deterministic methods such as gradient descent, would uniquely give an answer, but most likely this answer will not correspond to an approximate solution to a physical CFT. This is easily understood as follows. Suppose we find a suitable set of constraints which guarantee the existence of a finite set of approximate solutions (we will see below that these are indeed necessary, but easy to identify). If we define a positive definite function F as, say, the modulus square of the conformal bootstrap equations for a single 4-point correlator, we see that in the infinite dimensional space of all possible CFT data, scaling dimensions Δ_i and operator product expansion (OPE) coefficients λ_{ijk} , generally $F \geq 0$. Putative physical theories can only be associated to points where $F = 0$, where crossing symmetry is realized.² Suppose now to use a deterministic numerical method that allows us, starting from an arbitrary point in the (truncated) CFT data parameter space where $F > 0$, to minimize F . The details of how to actually achieve this do not matter for the argument we want to make here. In a truncated numerical approach, of course, F will never be zero, but we will nevertheless uniquely get the minimum value F_0 , given the chosen initial conditions. It is clear that generally F_0 will be an approximation to a local minimum of F (with $F_{\min} > 0$) and not an approximation to one of the global minima of F (where $F = 0$). Thus, the solution found will not correspond to any CFT; it will be in the “swampland.”

In order to alleviate the problem of being trapped at fake CFT minima, we use a stochastic minimization via a Metropolis algorithm. The function S we choose to minimize (the “action”) is the logarithm of the truncated bootstrap equations for a single scalar field correlator, $S \sim \log F$. Configurations with higher actions are from time to time accepted with a probability which depends on a parameter T (“temperature”). In this way, thermal fluctuations allow us to “escape” from local minima in the quest for the global ones.³ The problem of choosing among the many possible partitions of the operators in the different

spin channels is solved by adding as many operators as possible to each spin channel and then allowing operators to possibly decouple along the minimization process. The complex dependence of S on the truncated CFT data requires further nontrivial numerical manipulations, which will be discussed in Sec. II, to finally determine the approximate CFT data of a putative theory.

We consider in this work the simplest bootstrap set-up, namely a single 4-point function of identical scalars in Euclidean CFTs and focus on the exploration of the CFT parameter space where the external scalar is close to the unitarity bound. We start in Sec. II by presenting the method. In Sec. II A we explain how the bootstrap equations are truncated and discretized, in Sec. II B how the stochastic minimization is defined and performed and finally in Sec. II C the various steps entering in the minimization process are reported and explained.

The results of our numerical explorations, all based on severe truncations of the bootstrap equations with $\Delta^* \lesssim 15$ and a total number of operators included $N_{\text{Ops}} \lesssim 20$, are reported in Sec. III. The first general qualitative feature that emerged from our analysis is discussed in Sec. III A and is the predominance of the free theory scalar theory. If no constraints are imposed—aside unitarity, the presence of an energy momentum tensor and a \mathbf{Z}_2 global symmetry—the minimization of S leads only to a single global minimum, associated to the free theory. It is enough to impose a single constraint, such as fixing the scaling dimension of the external operator, to get a finite set of minima. At fixed Δ_σ we first study in Sec. III B the simplest spin partition with one operator per spin. We test the method by rediscovering the free theory and give indications that there are no other CFTs with one operator per spin, at least in their lowest spin channels and lowest lying operators. Our most interesting results are contained in Sec. III C, where we study the more realistic setup of more operators per spin. Now the situation is reversed and we get generally too many minima, unless extra assumptions are imposed. Demanding the presence of only one \mathbf{Z}_2 -even relevant scalar drastically reduces the number of minima and makes the analysis feasible. A general qualitative feature that emerged from our analysis is the presence of two regions in the space of CFT data where the bootstrap equations are more easily satisfied for CFTs admitting a conserved energy-momentum tensor, a \mathbf{Z}_2 global symmetry and one \mathbf{Z}_2 -even relevant scalar. When projected in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane, σ and ϵ being the external and the lightest exchanged operator, respectively, one of these regions essentially coincides with the extremality line found in previous bootstrap studies. The other region is along the generalized free theories (GFTs) in $d = 2$ and below the GFT line in both $d = 3$ and $d = 4$. The minima found in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane are reported in Figs. 5, 9, and 12 for the $d = 2, 3, 4$ cases. The extremal minima in the $d = 2$ case are identified with the generalized minimal models of [16], one extremal minimum in $d = 3$

²Of course this is only a necessary condition, since other constraints might arise from other correlators, etc.

³Very recently numerical methods based on reinforcement learning have been used to find approximate solutions to the bootstrap equations [14,15]. Though this is an interesting line of research to pursue, it is not clear to us if and how such techniques can alleviate the problem of fake local minima.

corresponds to the $3d$ Ising model, while all the other minima have not been identified. Unfortunately the numerical accuracy we have does not allow us to reach higher values of Δ^* to establish more firmly if (and which of) these minima are numerical artifact or not. CFTs without an energy momentum tensor are considered in Sec. III D. The minima now cluster around the GFT line and GFTs themselves are well reproduced, especially in $d = 2$ where our numerical accuracy is higher. We finally briefly show in Sec. III E how our algorithm can be applied to nonunitary theories by rediscovering the $2d$ Yang-Lee model.

We discuss in the final outlook the limitations of our algorithm and possible ways to improve it in the future. Three Appendices contain some other details of our numerical procedure and further more technical results.

II. METHOD

We present in this section the numerical method we used to find approximate solutions of the bootstrap equations for a 4-point function of identical scalar primary operators in a generic CFT in d dimensions.

A. Discretization of the bootstrap equations

The crossing equations for 4 identical scalar fields σ with scaling dimension Δ_σ can be written as (notation as in [2])

$$\sum_{\Delta > 0, \ell} \lambda_{\sigma\sigma\mathcal{O}_{\Delta\ell}}^2 \mathcal{F}_{\Delta,\ell}(z) = I(z). \quad (2.1)$$

Here $\lambda_{\sigma\sigma\mathcal{O}_{\Delta\ell}}$ are OPE coefficients, Δ and ℓ are the scaling dimensions and spin of the exchanged primary operator $\mathcal{O}_{\Delta\ell}$ (even spin traceless symmetric tensors), z and \bar{z} are the usual Dolan-Osborn conformal invariant cross ratios [17,18], $I(z, \bar{z}) = |z|^{2\Delta_\sigma} - |1-z|^{2\Delta_\sigma}$ is the identity contribution to the OPE expansion, and finally

$$\mathcal{F}_{\Delta,\ell}(z) = |1-z|^{2\Delta_\sigma} g_{\Delta,\ell}(z, \bar{z}) - |z|^{2\Delta_\sigma} g_{\Delta,\ell}(1-z, 1-\bar{z}), \quad (2.2)$$

with $g_{\Delta,\ell}(z, \bar{z})$ the conformal blocks, normalized as in [1] with an extra 2^ℓ factor.

As discussed in the introduction, we do not look for rigorous disallowed regions in the space of CFT data, but rather for approximate solutions of the above equations, namely for approximate possibly allowed CFT data. To this aim, we severely truncate (2.1) to the finite set of primary operators with scaling dimension up to a given Δ^* (chosen to be integer), and assign a given number of operators n_ℓ to each spin channel ℓ up to some value ℓ_{\max} . If unitarity is imposed, Δ^* and ℓ_{\max} are related. We take

$$\ell_{\max} = \Delta^* - d - 1, \quad (2.3)$$

so that the scaling dimension of the operators with spin ℓ_{\max} have a large enough range where they can be varied, compatibly with unitarity: $\Delta^* - 3 \leq \Delta_{\ell_{\max}} \leq \Delta^*$. The total number of exchanged operators \mathcal{O}_a that we consider is

$$\sum_{\ell=0,2,\dots,\ell_{\max}} n_\ell = N_{\text{Ops}}. \quad (2.4)$$

We then choose a finite sample of N_z points in the complex z -plane in the neighborhood of $z = \bar{z} = 1/2$ and define the $N_z \times N_{\text{Ops}}$ matrix \mathcal{M} and the N_z vector I_i

$$\begin{aligned} \mathcal{M}_{i,a} &\equiv \mathcal{F}_{\Delta_a,\ell}(z_i), & I_i &\equiv I(z_i), & i &= 1, \dots, N_z, \\ a &= 1, \dots, N_{\text{Ops}}. \end{aligned} \quad (2.5)$$

We define an ‘‘action’’ as (spin dependence omitted for simplicity)

$$\exp(S(\Delta_a, \rho_a)) \equiv \frac{1}{N_z} \sum_{i=1}^{N_z} \sigma_i^{-2} \left| \sum_{a=1}^{N_{\text{Ops}}} \rho_a \mathcal{M}_{i,a} - I_i \right|^2, \quad (2.6)$$

where $\rho_a \equiv \lambda_{\sigma\sigma\mathcal{O}_a}^2$, by taking a weighted sum of the absolute square of the truncated crossing equations (2.1) at the N_z points, with

$$\sigma_i = |\mathcal{F}_{\Delta^*,0}(z_i)|. \quad (2.7)$$

The factor (2.7) is introduced to take into account that the points with the best convergent conformal block expansion are around $z = 1/2$, point where on the other hand the function \mathcal{F} vanishes. For any finite truncation, our aim would be then to minimize S in the space of CFT data.

B. Metropolis Montecarlo and stochastic minimization

Finding the $2N_{\text{Ops}}$ CFT data (ρ_a, Δ_a) that minimize S is a challenging task. First of all, in $d > 1$ dimensions, where we have spin, one should in principle look for solutions in all possible partitions of the n_ℓ operators in each spin channel. In addition to that, deterministic methods, such as gradient descent, can be highly inefficient because they converge to a minimum of S depending on the initial conditions. The actual dependence of S on the CFT data is expectedly complicated and depends on the constraints imposed, and can be characterized by the presence of several local minima. One of the main motivations of this work is in fact to shed some light on the ‘‘landscape/swampland’’ space of CFT data. As we will explain in what follows, both problems (avoiding to land on a local minimum and the need to consider all spin partitions) are significantly alleviated by the use of stochastic searches based on Monte Carlo (MC) methods. We will in particular

sample the parameter space via the Metropolis Monte Carlo method, whose structure we now briefly describe.

Given a function $f(x):U \rightarrow \mathbb{R}$, where $U \subset \mathbb{R}^n$, the minima of f in U are sampled by taking an initial point $x^{(0)} \in U$ and proposing an update

$$x^{(1)} = x^{(0)} + \delta x, \quad (2.8)$$

which is accepted with a probability given by

$$P(x^{(0)} \rightarrow x^{(1)}) = \min \left(1, \exp \left(-\frac{f(x^{(1)}) - f(x^{(0)})}{T} \right) \right). \quad (2.9)$$

The parameter T can be interpreted as a temperature and the exponential factor in (2.9) as a Boltzmann suppression factor. While a gradient descent method would uniquely find the local minimum of f connected to the point x_0 , the Metropolis algorithm, accepting from time to time moves where the function increases rather than decreases, is able to overcome barriers between different minima and sample wider regions of U . The ‘‘efficiency’’ in overcoming barriers depends on T . At fixed δx , the higher T the easier it is. Of course, if T is chosen too large, we would explore the entire region U but we might not detect the minima, because with such a diffusive dynamics the probability density $P(x)$ will tend to be uniform in x .

In our MC implementation, the variables x which are changed at each iteration procedure are the scaling dimensions Δ_a (and possibly the external scaling dimension Δ_σ), while the function f is identified with the action S in (2.6). The procedure works as follows. Starting from some initial scaling dimensions $(\Delta_\sigma^{(0)}, \Delta_a^{(0)})$, we analytically minimize S with respect to the OPE coefficients ρ_a , given that they enter quadratically in the action. We get

$$\begin{aligned} \rho_a^{(0)} &= \sum_b A_{ab}^{-1} \sum_i \text{Re}(\mathcal{M}_{b,i}) \sigma_i^{-2} I_i, \\ A_{ab} &\equiv \sum_i (\mathcal{M}_{ai}^* \mathcal{M}_{bi} + \mathcal{M}_{ai} \mathcal{M}_{bi}^*) \sigma_i^{-2}, \end{aligned} \quad (2.10)$$

where all the terms in the right-hand side of (2.10) are evaluated at $(\Delta_\sigma^{(0)}, \Delta_a^{(0)})$. We then plug (2.10) into (2.6) to get an action $S^{(0)} = S(\Delta_a^{(0)}, \rho_a^{(0)})$ which depends on operator scaling dimensions only (spin and external operator dependence omitted for simplicity). At this stage, the scaling dimensions $(\Delta_\sigma^{(0)}, \Delta_a^{(0)})$ are changed as in (2.8), the OPE coefficients recomputed as above and then the move is accepted or not according to (2.9). This procedure is then iterated N_{steps} times. In this way we explore different the local minima of S , and then study the collective trends of these minima as a function of Δ^* . Of course, for the CFT data $(\Delta_\sigma, \Delta_a, \rho_a)$ associated to a physical theory we should have

$$\lim_{\Delta^*, N_{\text{ops}} \rightarrow \infty} S(\Delta_a, \rho_a) = -\infty. \quad (2.11)$$

It is then important to establish that, in a finite truncation, the value of S at the minima decreases as Δ^* increases.

This is in a nutshell how our protocol works. In the next subsection we will describe in detail the various steps involved.

C. The search protocol in detail

We choose to sample the N_z points in the z -plane from a uniform grid with the constraint

$$\lambda(z) < \lambda_0, \quad (2.12)$$

with $\lambda(z)$ defined as in (3.11) of [19]. This condition defines a compact and convex region around $z = \bar{z} = 1/2$, the point of best convergence of the bootstrap equations. The parameter λ_0 is, together with Δ^* and T , the most relevant ‘‘hyperparameter’’ of our computational pipeline, and has to be chosen carefully in order for the method to work.

The number of MC iterations N_{steps} needed to get sensible results can be quite large ($\sim 10^8$). In order to be able to perform exhaustive searches in parameter space we developed a numerical approach which allows estimating *efficiently* the action S defined in (2.6). Indeed, conformal blocks are computationally expensive functions of Δ , ℓ , and z , which would make an extensive sampling impossible with moderate computational resources without an appropriate numerical optimization. We tabulate the blocks for a fixed z -point sample and each ℓ , and then use cubic splines [20] to interpolate their values. In this way we get a speedup of $O(10^4)$. The interpolation grid is chosen in such a way that the numerical error induced by this procedure is negligible compared to the truncation error (see Fig. 17 and the discussion in Appendix B 2).

It is important here to note that this speed up comes with the tradeoff of limiting our working precision to that of double type in Fortran (i.e., 64-bit). This might come as a surprise to our more technical bootstrap readers, but an important result of this paper is indeed that approximate solutions to the bootstrap equations can be found without resorting to 200 digits of precision. Nevertheless, limited precision is a bottleneck to go to larger values of Δ^* and to improve the quality of our approximate solutions.

We now describe each step of the protocol more thoroughly. Further details can be found in Appendix B.

1. Determining the hyper parameters

Once a physical system of interest (understood as a global symmetry structure and the relevant bounds on the operators’ scaling dimensions) has been chosen, some exploratory work must be done in order to make our search protocol as efficient as possible. In Appendix B we

list the choices for each search discussed in Sec. III. We will here concentrate on motivating our choice for T , since we have found empirically that the other hyperparameters are less determinant to the success of the search.

Since we will explore different truncations, the order of magnitude of δS , see Eq. (2.9), can be very different from spectrum to spectrum⁴ and thus the temperature T must be adjusted accordingly. Although we will use an adaptive step that guarantees that in the stationary state roughly half of the moves will be rejected, choosing a temperature that is too high will make $\delta\Delta$ too coarse and the MC will “miss” narrow features of the landscape. On the other hand, if T is too small our method will move too slowly around the landscape or might even “freeze” in the basin of a local minimum.

In the following we will empirically choose T in such a way that the projection of the trajectory to each coordinate axis (which in our case correspond to the operators’ scaling dimensions) covers the whole allowed range at least 10 times, but still low enough so that the dynamics does not become diffusive, namely the local minima can still be detected as metastable states.

For further detail, we show in Fig. 18 an example of archetypal trajectories. This figure and further considerations about the choice of T are presented and analyzed in B 3.

2. Finite temperature MCs

For each ansatz $\Delta_a^{(0)}$ ⁵ we sample the CFT parameter space at the temperature T chosen according to the criteria discussed in the previous step. This is the most important and computationally expensive step, but one must bear in mind that it gives all the necessary information to understand the landscape of possible CFTs at several different truncations.

The variation of the scaling dimensions $\delta\Delta_a$ is taken randomly with uniform distribution over an hyperparallellepiped centered around the origin. This was observed to be more efficient than moving one operator at a time since the landscape shows highly correlated features and thus attempting moves in which all the operators change simultaneously improves the sampling efficiency. The user can choose the length of each side but the overall scale is adaptive and changes at each MC step. This is done in order to guarantee a 50% acceptance rate in the stationary state. In practice, when a move is rejected (resp. accepted) this overall scale is decreased (resp. increased) by a fixed factor. Each operator scaling dimension is confined within a finite region within two boundary values decided by the user. The MC then explores the landscape for $\sim 10^8$ steps. Our choice of the temperature guarantees that after this number of

iterations the initial conditions are irrelevant. For each of these spin partitions we store each 1000th configuration visited. We will refer to these as frames.

3. Separation by sectors

In a given MC run not all the operators in the ansatz will have sizable OPE coefficients at each iteration step. In this work we assume that if the OPE coefficient associated to a primary operator is smaller than a given tolerance (we found 10^{-8} to be a sensible choice for the ℓ_{\max} used in this paper) this operator is decoupled and can be neglected. If unitarity is imposed, OPE coefficients that turn out to be negative from (2.10) are set to zero in that iteration. These two effects lead to a classification of the different points visited by the MC in terms of the effective spectrum contributing to the crossing equations. In order to describe these different sectors we introduce the following shorthand notation to indicate each one of them: $(m_0)_{-}(m_2)_{-}\dots_{-}(m_{\ell_{\max}})$. For example, a point with 3 scalars, 2 spin-2 operators, 2 spin 4 and one spin-6 would belong to spin sector 3_2_2_1. Crucially,

$$\sum_{\ell} m_{\ell} \leq N_{\text{Ops}}. \quad (2.13)$$

This implies that there is no need to consider different partitions of operators in each spin channel. All partitions with $m_{\ell} \leq n_{\ell}$ will be automatically searched for by MC. Moreover, the same sector can be visited in two different MC runs with different n_{ℓ} . Thus, this step effectively recombines all the information from the different spin partitions studied in the previous step.

4. Identification of putative local minima

Now that we have sampled the landscape we must identify the points that are more likely to be close to a local minimum of the action. For a low-dimensional parameter space we could visualize the action as a contour plot and identify the points by inspection. It is clear, though, that for more than 3 scaling dimensions, this approach cannot be followed.

Our proposal is thus the following. For each set of frames obtained in the previous step (namely, those corresponding to spin sectors with more than 120k frames and those describing the trajectory of every search)⁶ we sort their elements according to the number of nearest neighbors with higher action, in the philosophy of density peak clustering [21]. We consider as putative local minima the 20 configurations with highest rank. We have observed empirically that in each

⁴We have found empirically that the most important parameter in this case is ℓ_{\max} , not necessarily Δ^* on its own.

⁵For simplicity of notation, we do not explicitly report $\Delta_{\sigma}^{(0)}$ as input data.

⁶It might seem redundant to analyze sector by sector frames that are already contained in the trajectories that result from the MC search described in point 2. Since exhaustivity is paramount and the definition of local minimum depends on the neighboring points, it makes sense to consider the same frame in two different settings in order not to miss any interesting features.

sector the number of actual minima is always much smaller than this number, and that many of the smaller rank frames actually converge to other minima, so this is a conservative estimate. To clarify, the point with the highest rank will be the global minimum of that set of frames because every other point will have higher action.

5. Local minimization

For each putative local minimum we launch Newton—Raphson (NR) minimizations. This deterministic algorithm is justified because the points found in the previous step are likely contained in the basin of attraction of a local minimum. More details about the implementation of this algorithm can be found in Appendix A.

All the endpoints from the NR minimizations are collected and those that violate the bounds imposed on the spectrum are discarded.⁷ We then assign to the corresponding sector the minima where one or more operators decouple (in the sense described above). In the case of two or more minimizations converging to the same point (i.e., the relative difference for any scaling dimension is smaller than 20%) we keep only the point with the lowest action. At the end of this step we have a list of approximate solutions to crossing that are local minima of S .

6. Convergence and boundary check

Since we usually impose boundaries on the subleading operators in order to better control the type of CFTs that we might find, it is important to rule out the possibility of some of the local minima found in the previous step actually being induced by the presence of this boundary.

To this end, a final NR minimization is performed from all the minima found in the previous step, relaxing all the boundaries except strict unitarity (if imposed). We consider as true minima the ones that reach convergence according to our NR implementation and are contained within the region of interest: this means a minimum can be at the boundary of unitarity, but not at one of the boundaries introduced in step 1.

7. Identifying akin minima at different truncations

Minima belonging to different spin sectors but with similar CFT data in their common operators should be identified. In this case the minimum with more primaries should have a value of S lower than the minimum with less primaries. This is a necessary (but by no means sufficient) condition to associate such minima to approximate solutions of the crossing equations of one and the same CFT. In order to find akin minima and quantify their similarity we organize all the minima in a directed graph where the minima are the nodes and $A \rightarrow B$ if: (i) the sector to which

A belongs is contained in the one of B , namely if B has at most `maxMismatch` more operators than A and (ii) the operators which are in common between the two sectors differ in their scaling dimensions by at most `relTol` in relative error. The final result of our protocol are the leaves of this directed graph, namely the nodes with no outgoing links. One can also study the branches (i.e., subsets of this directed graph) in order to understand the asymptotic behavior of certain values, such as the action S or the Δ of some leading operators.

III. RESULTS

We report in this section the results obtained with our method. We have considered a single 4-point function with four identical scalar operators in $d = 2, 3, 4$ dimensions. Except for a short discussion of the Yang Lee CFT (see Sec. III E), we have considered CFTs where the fusion rules are compatible with a \mathbf{Z}_2 discrete symmetry under which the external operator is odd.⁸ We will use a Ising-like notation and denote the external operator by σ and the lowest dimensional exchanged scalar (besides the identity) by ϵ . As will be later explained, demanding that ϵ be the only \mathbf{Z}_2 -even relevant (in the renormalization group sense) primary operator that can appear in the $\sigma \times \sigma$ OPE is necessary to avoid the proliferation of minima. We consider both local and nonlocal theories, demanding or not the existence of a spin-two energy-momentum tensor operator at the unitarity bound $\Delta_T = d$. Given its relevance in our analysis, we have also considered the space of nonlocal theories where the scaling dimension of the first spin 2 operator is fixed and equals that of generalized free theories (GFT), $\Delta_T = 2\Delta_\sigma + 2$. We then get further insights by allowing Δ_T to vary. We focus mostly on unitary theories, though our method does not rely on it. As a proof of concept, we study the nonunitary landscape around the 2d Yang-Lee model.

As already anticipated in the introduction, when Δ_σ is allowed to vary, the landscape of S becomes very simple: the global minimum is at the lower edge of the unitarity region, which in $d = 3, 4$ coincides with the free theory (FT) of a scalar field. Moreover, for generic spectra we find that it is in fact the only reachable minimum. For this reason in most of our analysis we will study “slices” of the landscape by keeping fixed Δ_σ in each MC run.

The section is structured as follows. We start in Sec. III A by showing that the FT is the global (and in many cases the only) minimum of S . From Sec. III B we look for local minima by keeping the external operator fixed. We first consider the simplest possible scenario, namely very sparse CFTs where just one operator per spin appears below Δ^* .

⁷Spectra where a whole spin sector is missing, such as 2_1.0_1.1, have been ignored for simplicity.

⁸Note that in general such a \mathbf{Z}_2 symmetry is not required to be a genuine faithful symmetry of the whole CFT. We will encounter an instance of this phenomenon when discussing 2d minimal models in Sec. III C.

We then discuss in Sec. III C more general CFTs with multiple operators per spin. In Sec. III D we discuss nonlocal theories with a special focus on the GFT. Finally, we briefly discuss the 2d Yang-Lee model as a proof of concept of the validity of the method for nonunitary theories in Sec. III E.

A. The free theory is the global minimum of S

Before attempting to identify the local minima of S it is useful to have a qualitative idea of how S behaves as a function of the CFT data. This analysis can be made by launching MC runs with a given temperature and spin partition of operators, and simply looking for the values of the Δ 's corresponding to the lowest value of S explored during the run.

The most important property that comes out from this analysis, common to every scenario studied in this paper, is that the *global* minimum of S corresponds to the FT (or the point $\Delta_\sigma = \Delta_\epsilon = 0$ for $d = 2$). This has been observed in all our searches, and in particular for all the spin partition of operators which we considered. Of course, this does not imply the nonsensical result that the FT is the only consistent CFT. Rather, it implies that among all the physical CFTs allowed within the given assumptions at fixed truncation, within our numerical accuracy and in absence of further constraints, only the FT shows up as an isolated minimum of S .

In order to test the (non)existence of other minima, we performed a very large number of MC minimizations at $T = 10^{-4}$. This temperature, being of the order of the numerical error of the action,⁹ guarantees that the MC can only diffuse toward configurations of lower action, since the temperature is too low to cross any relevant barrier. Independently of the choice of n_ℓ and of the operator scaling dimensions $(\Delta_\sigma^{(0)}, \Delta_\epsilon^{(0)})$, all the low temperature MC runs we performed converge to the FT point within 10^8 steps, although longer times might be needed for higher N_{Ops} .¹⁰

We report as an example in Fig. 1 the low- T evolution for a $3d$ CFT spectrum with $\ell_{\text{max}} = 6$, $N_{\text{Ops}} = 8$ and an initial spin partition $3_2.2_2.1$. Here the first spin two operator is fixed at $\Delta_T = 3$ and no gap on the number of relevant scalars has been imposed.¹¹ The starting point of each colored line (indicated with a black bullet) corresponds to a MC with initial values of $(\Delta_\sigma^{(0)}, \Delta_\epsilon^{(0)})$ at that point. The black arrows on the colored lines indicate the

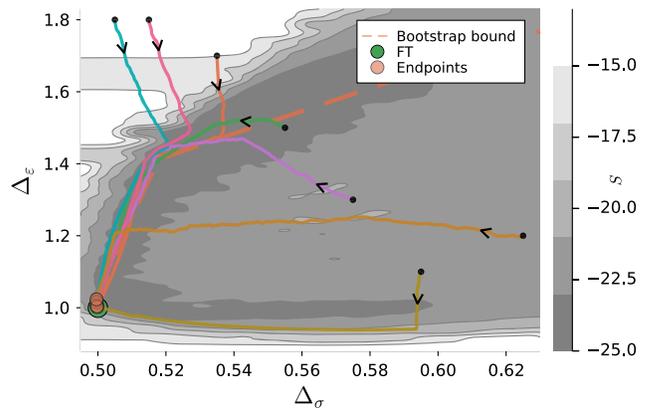


FIG. 1. Convergence of different trajectories to the $d = 3$ free scalar CFT. The contour lines represent the minimum value of the action in each bin for a spectrum of the kind $3_2.2_2.1$. The dashed line is the upper bound for Δ_ϵ from [22]. The black arrows on the colored lines indicate the direction of the MC evolution.

direction of the MC evolution. As can be seen, the final point of all the MC runs (highlighted by red circles) have $\Delta_\sigma \approx 1/2$, $\Delta_\epsilon \approx 1$. The contour plot shown in the background is indicative of the value of S for the same spin partition but sampled at a much higher temperature (roughly 100 times larger). A two-dimensional representation of S is then obtained by binning the points visited on the $(\Delta_\sigma, \Delta_\epsilon)$ -plane and plotting the minimum value of S in each bin.¹² They are reported as a visual aid. As can be seen from the figure, the initial values of $(\Delta_\sigma^{(0)}, \Delta_\epsilon^{(0)})$ can be taken in the excluded region, in which case they quickly approach the extremality line. The discrepancy between the rigorous bootstrap extremal line and the trajectories of our MC solutions is due to the severity of our truncation. During the MC evolution the OPE coefficients of the extra scalars and higher spin operators decrease in such a way that at the final point we effectively have the FT spectrum with one operator per spin. We illustrate this phenomenon by plotting in Fig. 2 the scaling dimensions of σ , ϵ and the spin-4 current as well as the squared OPE coefficients of ϵ , T and ϵ' (the second exchanged scalar) associated to the green line starting at $(0.55, 1.5)$ in Fig. 1. A similar situation occurs for other choices of initial partitions and by keeping Δ_T unconstrained. The same also holds in $d = 2$ and $d = 4$ dimensions. This is the first important result of our analysis: for most of the spin partitions studied in this work, the landscape of S has a single minimum, which can be reached by trivial gradient descent from (generically) any initial configuration. Note that most trajectories in Fig. 1 reach the FT point following two broad paths in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane: (i) the extremality line from above, (ii) a curve, almost constant in ϵ , from below. The presence of regions in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane where the

⁹See Appendix B 2.

¹⁰There is an important exception described in detail in Sec. III B. In $d = 3$ for spectra of the kind $1_1.1_1$ and $1_1.1_1.1$ there are actually two local minima, though the FT still is the global one.

¹¹Note that in $d = 3$ (the case shown in Fig. 1), if a scalar gap is imposed (for example, demanding that there can only be one exchanged relevant scalar), depending on the initial CFT data it can happen that the FT cannot be reached and a local minimum is induced.

¹²Understood as a cell of fixed width in Δ_σ and Δ_ϵ .

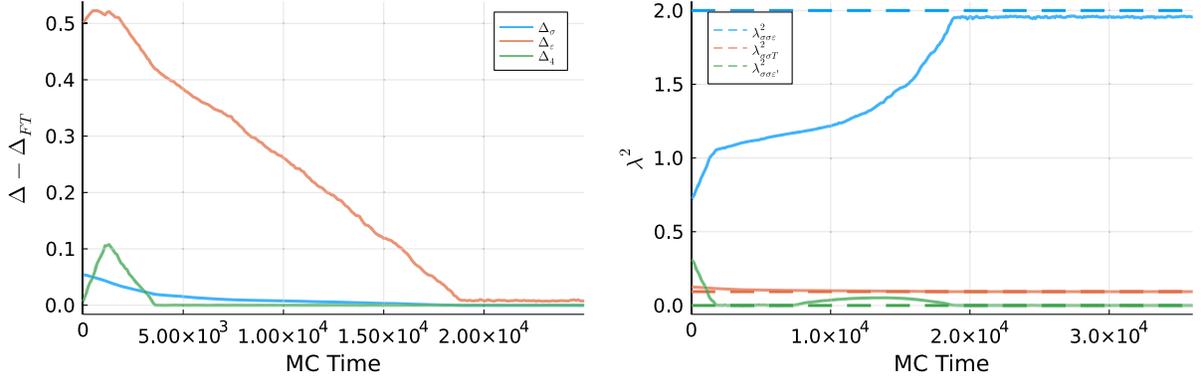


FIG. 2. Convergence in t_{MC} for representative scaling dimensions (left) and squared OPE coefficients (right).

bootstrap equations are more easily satisfied is going to be a common theme in our results.

As we are going to show next, if one fixes at least another Δ one can find local minima along these slices. In the following subsections we will apply the protocol described in Sec. II C for fixed values of Δ_σ and we will show that some of the local minima that are present correspond to known theories (for example the Ising model). However, it is important to remark that none of the minima that we will describe are stable if Δ_σ is allowed to vary, not even the minimum corresponding to the Ising models in $d = 2$ and $d = 3$, in agreement with the results presented in this subsection.

B. One operator per spin: $d = 3, 4$

We now describe the results of a systematic search of local minima at fixed Δ_σ by the protocol described in Sec. II C. It is instructive to begin by exploring the simplest possible scenario where

$$n_\ell = 1, \quad \ell = 0, 2, \dots, \ell_{\max}. \quad (3.1)$$

Although this assumption looks very restrictive, some basic features of the shape of the landscape of CFTs is already visible in this set-up. Step 3 of the protocol is clearly unnecessary when $n_\ell = 1$, since all the configurations in each run will belong to just one spin sector. We focus on $d = 3$ and $d = 4$ CFTs and ask the following question: aside from the FT, are there other viable approximate solutions to the crossing equations in which (3.1) is satisfied? This question might seem academic, given that interacting CFTs have an infinite number of Regge trajectories [23]. However, we are severely truncating the spectrum, so the question we address here is if we can have CFTs where there is a substantial gap between the first Regge trajectory and the others.

An extensive high- T search in $d = 3$ with $\ell_{\max} = 6$ and free Δ_σ —the contour lines in the background of Fig. 3(a)—shows that there are two (very elliptic) disjoint basins of attraction. The one at the bottom of the panel contains the

FT. The upper one, elongated along the extremality bootstrap line, approximately contains the Ising point in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. Note that the latter basin of attraction has a higher action with respect to the former one.

We searched local minima of S at different fixed values of $\Delta_\sigma = 1/2 + n/20$, $n = 0, 1, 2, 3$, with input spin sectors 1_1_1, 1_1_1_1, 1_1_1_1_1 and 1_1_1_1_1_1. The minima that we found are represented in Fig. 3(a). In order to better appreciate the value of S at the minima i , the latter are indicated with a circle whose radius r_i is given by

$$r_i \propto S_{\max} - S_i + \alpha, \quad (3.2)$$

where S_{\max} is the value of S in the minimum with larger action, and α is an offset so that this minimum can still be visible in the figure.

For $n = 0, 1$ two classes of minima are found, one for each of the two basins of attraction. For $n = 2, 3$, instead, we find only minima along the lower basin, with $\Delta_\epsilon < 1$. It should be clear that there is nothing special to the above chosen values of Δ_σ . For each fixed value of Δ_σ we expect to find two local minima if $\Delta_\sigma \lesssim 0.6$ and one for higher values.

We then use the protocol described in Sec. II 3 with `relTol` = 0.1 and `maxMismatch` = 2 to identify the different branches, which correspond to sets of minima with different input sectors that can however be considered to represent the same “theory” at different truncations. In this way we find that, among these minima, some are isolated and do not form any branch (orange diamonds in panel (b) of Fig. 3), while others form branches (green, pink and blue bullets). The only branch in which the action decreases as a function of Δ^* is the one associated to the free theory, the blue branch in panel (b) of Fig. 3.

In $d = 4$ a similar high- T search with $\ell_{\max} = 6$ and free Δ_σ gives rise to the contour lines in the background of panel (c) of Fig. 3. A search of the minima at different fixed values of $\Delta_\sigma = 1 + n/10$, $n = 0, 1, 2, 3$, with input sectors 1_1_1, 1_1_1_1, 1_1_1_1_1 and 1_1_1_1_1_1 with `relTol` = 0.1 and `maxMismatch` = 2 allows us to identify, also in this

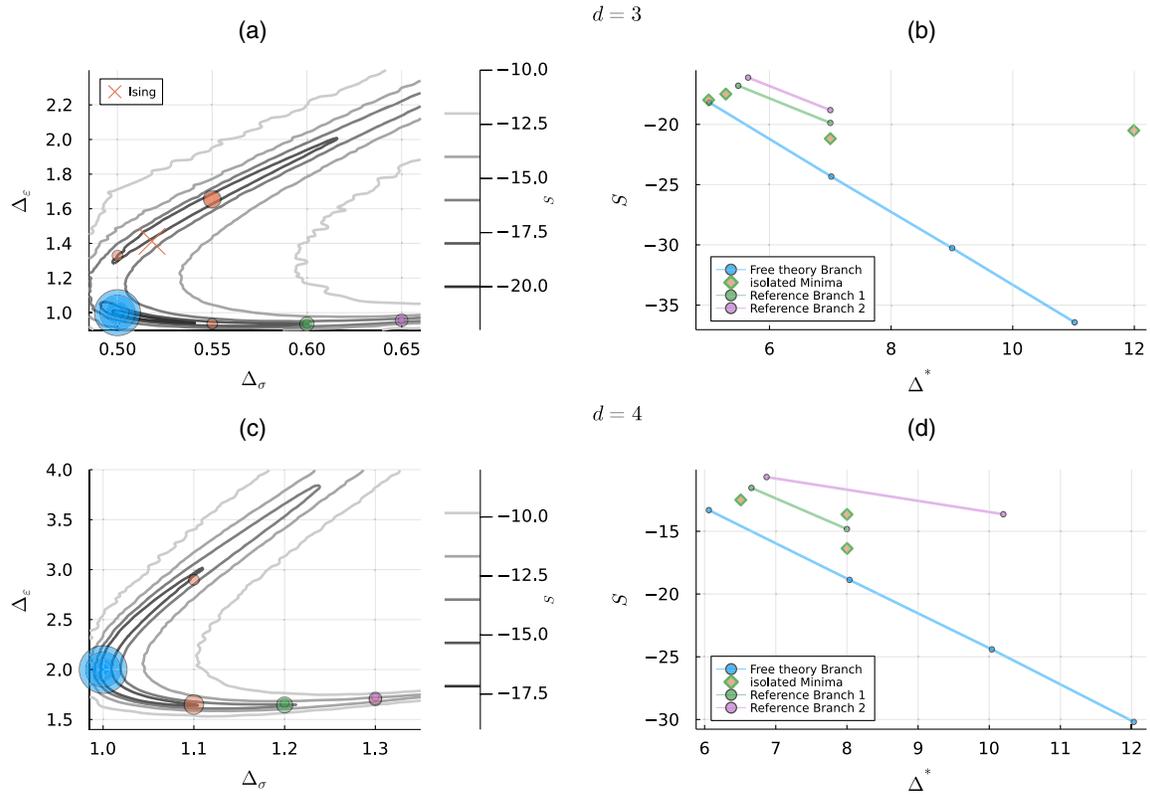


FIG. 3. Overview of the different minima found in $d = 3$ (top) and $d = 4$ (bottom) with spectra containing one operator per spin. The color code is the same across the panels. Left: minima at fixed Δ_σ . The radii indicate the value of the action S . The larger the radius the smaller S . The contour lines are taken from an high- T search with $\ell_{\max} = 6$. Right: value of S for each point along the colored branches in the left panel, identified by its Δ^* .

case, some local minima at fixed Δ_σ . Like in the $d = 3$ case, some minima are isolated and do not from any branch (orange circles in panel (d) of Fig. 3), while others form branches (green, pink and blue circles). The only branch in which S consistently decreases with Δ^* is the one associated to the FT, the blue branch in Fig. 3(d).

As we will show in the following, the main structure of the S landscape with two elongated basins of attractions (valleys) will appear repeatedly in the analysis of more complex situations. Note the striking difference between the $d = 3$ and $d = 4$ cases: in the former the basins are two, in the latter the basin is only one. One could argue that this difference is due to the presence of another actual theory in $d = 3$ close to the FT point, the Ising theory, and none in $4d$ within the above assumptions. Indeed, our results in Secs. III B and III C further point to this explanation.

We also find that if the unitarity bounds are relaxed (both on Δ and on the OPE coefficients squared), the aforementioned disjoint valleys in $d = 3$ join in a nonunitary point of the parameter space. This is shown in Fig. 4, where we compare the lowest-lying points in an unitary MC search to those of a nonunitary setting.

The main take-away of this test is that the evolution with respect to Δ^* of the action in branches formed by akin minima seems to be a discriminant factor between spurious

minima and physically meaningful ones (the FT in this case); namely, the spurious minima do not create long branches that go down into small values of the action S , while the FT is present for every truncation with very high consistency. We find that in the FT minimum high spin

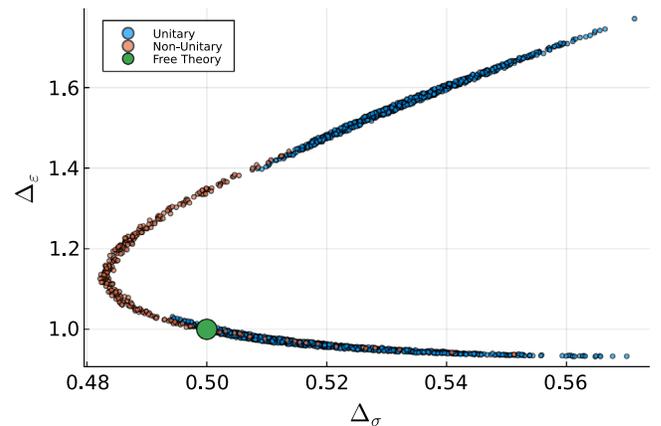


FIG. 4. Locus of the lowest-action points visited by a wide-search MC in $d = 3$ with spectrum 1-1-1-1. When unitarity is imposed we find two disjoint basins. Allowing for nonunitary configurations gives rise to a single connected basin.

TABLE I. Scaling dimensions of the first operators appearing in the OPE of σ in the free scalar theory. In these settings $\Delta_\sigma = (d-2)/2$ and $\Delta_T = d$.

	$\Delta_0 (\Delta_\epsilon)$	Δ_4	Δ_6	Δ_8	Δ_{10}
$d = 3$	1.0000	5.0000	6.9999	8.9999	11.0178
$d = 4$	1.9999	6.0000	7.9999	9.9993	12.0308

TABLE II. Scaling dimensions of σ and of the first operators appearing in the OPE of σ in the free scalar theory. Besides unitarity the only constraint imposed is $\Delta_T = d$.

	Δ_σ	$\Delta_0 (\Delta_\epsilon)$	Δ_4	Δ_6	Δ_8	Δ_{10}
$d = 3$	0.50009	0.9997	5.0009	6.9974	8.9973	11.1505
$d = 4$	1.00000	2.0035	6.0000	7.9995	9.9997	12.0368

operators are determined quite accurately ($\sim 1\%$ error) except for the highest- Δ at each truncation, as can be seen in Table I.

Although it is clear that the minima at $\Delta_\sigma = (d-2)/2$ are deeper than those away from it, we performed a MC search at $\ell_{\max} = 10$ without any constraint on Δ_σ besides unitarity. This allows us to understand the precise shape of the landscape around the FT point. After 2×10^8 steps of an extensive search ($T \sim 1$), we chose the global minimum and performed a $T = 10^{-4}$ minimization in order to refine the precise location of this minimum. In Table II we report these values. For completeness we also show in Table III the OPE coefficients determined in this case.

We finally note that for higher Δ^* , namely higher N_{Ops} , the FT minimum is much deeper than the other minima. This means that the latter are very poor solutions to the crossing equations. In other words, we did not find other consistent approximate solutions to the crossing equations, besides the FT, where (3.1) is satisfied up to $\ell_{\max} = 10$.

C. More operators per spin

We now use our protocol to study CFTs with more than one operator per spin with the assumptions discussed at the beginning of the section (presence of an effective \mathbf{Z}_2 symmetry and only one relevant \mathbf{Z}_2 -even scalar).

We here consider viable solutions to the crossing equations the endpoint of branches¹³ that arrive up to a given ℓ_{\max} and include at least 2 different minima with strictly decreasing S . For brevity, we will refer to these end points simply as “end-minima”. For example, in the case of only one operator per spin, the “end-minima” are the blue, pink and green circles at the end of the branches in the right panels of Fig. 3. The rationale behind concentrating on

¹³Our parameters for determining the connection between minima are `relTol=0.2` and `maxMismatch=2`.

end-minima is that they summarize the information from minima at higher ℓ_{\max} and thus allow for a cleaner visualization. Since this approach might exclude possible branches that start only at the largest Δ^* here considered, we will also include in our discussion isolated minima at the largest ℓ_{\max} considered. They will be also called end-minima in a slight abuse of notation.

Our whole search space includes spectra with $\ell_{\max} = 4, 6, 8, 10$. For $d = 2$ we limit the analysis to $\ell_{\max} = 8$, since at $\ell_{\max} = 10$ the end-minima that we find are so many that they obscure the visualization. Here we only observe that, qualitatively, all the minima that we find at $\ell_{\max} = 10$ differ only by the least relevant operators.

The role of the stress-energy tensor deserves some consideration. For the study of local theories it is necessary to impose the presence of a spin-2 operator with $\Delta_T = d$. In practice, we have also found it necessary to impose a gap on the scaling dimension $\Delta_{T'}$ of the second spin-2 operator so that it does not come too close to d and create the effect of a nonlocal theory by “supplanting” the stress energy tensor. In Sec. III D we will study the effect of allowing for nonlocal theories but hereafter it should be understood that unless specified otherwise we fix $\Delta_T = d$ and impose the gap $\Delta_{T'} \geq d + 1$.

We report the end-minima found in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane in Figs. 5, 9, and 12 for the $d = 2, 3, 4$ cases, respectively. We discuss these results separately in the sections below, but a few considerations apply to all cases. The contour lines of the lowest-lying configuration sampled by a high- T MC is shown in the background as a visual aid. The orange dashed line corresponds to rigorous bootstrap bounds, while the green continuous line is the line of GFTs. We use the radius of the points to represent the action as before [see (3.2)] and use transparent markers so that coincident minima appear darker. To understand the fact that several minima are superimposed in this representation it is important to keep in mind that there can be up to 10 other dimensions not shown in the plot. Indeed, in many cases there are end-minima which have almost identical values of Δ_σ and Δ_ϵ but differ in the scaling dimensions of the other operators.

1. $d=2$

One can wonder if our method can work in $2d$ given that the landscape of CFTs is notoriously very rich. On the other hand, $2d$ is the ideal playground to test the method, since entire classes of CFTs are exactly soluble.

We have considered various MCs for 8 different fixed values of Δ_σ :

$$\Delta_\sigma = \frac{1}{8}, \frac{3}{20}, \frac{7}{40}, \frac{1}{5}, \frac{1}{4}, \frac{2}{7}, \frac{3}{10}, \frac{\sqrt{5}}{10}. \quad (3.3)$$

When Δ_σ is too small the numerical analysis becomes quite noisy. For this reason, we did not explore values of

TABLE III. Comparison between the exact and the numerically determined values of the OPE coefficients squared for the first operators in the free scalar theory. Besides unitarity the only constraint imposed is $\Delta_T = d$.

	λ_ϵ^2	λ_7^2	λ_4^2	λ_6^2	λ_8^2	λ_{10}^2
$d = 3$	2.0022	0.093796	0.0042723	0.00022128	1.1604×10^{-5}	8.3384×10^{-7}
Exact	2.0	0.09375	0.0042725	0.0002203	1.1986×10^{-5}	6.7214×10^{-7}
$d = 4$	1.9927	0.33343	0.02853	0.0021857	0.00014554	1.4538×10^{-5}
Exact	2.0	0.33333	0.028571	0.0021645	0.0001554	1.0825×10^{-5}

$\Delta_\sigma < 1/8$. 7 of the values in (3.3) have been chosen to be “small” rational numbers, while one was chosen irrational, e.g., numerically equivalent to a “large” rational number. Rational scaling dimensions can be associated to exactly known rational conformal field theories (RCFT). As we will see, in order to limit the number of allowed theories and make the analysis feasible, the assumption of having only one relevant \mathbf{Z}_2 -even scalar turns out to be particularly important in $2d$. Before presenting our results, it is useful to report some (by no means exhaustive) instances of known CFTs admitting scalars with the above scaling dimensions.

A glimpse at the 2d CFT landscape The landscape of the known CFTs in $2d$ is amazingly rich. In contrast to the situation for $d > 2$, in particular, $2d$ CFTs can have a continuum spectrum (e.g., Liouville field theory). It is clear that we cannot detect such theories.¹⁴ Exactly marginal deformations are ubiquitous in $2d$. These theories will not give rise to minima of S , but to a valley of minima, each corresponding to a CFT as we vary the marginal couplings. Since we look for isolated minima in our search, such theories will also be difficult to detect. The CFTs we can hope to find should have a discrete, relatively sparse, spectrum and should be isolated. Their central charge c is expected to be of order one. Prototypical theories of this kind are RCFTs. Even if we focus on these theories only, we need to further restrict our search space to avoid a proliferation of theories. We will do so by assuming, as mentioned, that only one relevant scalar ϵ can appear in the $\sigma \times \sigma$ OPE. The importance of this assumption can be appreciated by focusing on specific exactly soluble CFTs. We will briefly consider in what follows the space of diagonal unitary minimal models ($c < 1$), the S^1/\mathbf{Z}_2 orbifold ($c = 1$) and $SU(2)_k$ Wess-Zumino-Witten (WZW) models ($c > 1$).

Recall that the scaling dimensions of the (Virasoro) primary fields $\phi_{r,s}$ in diagonal unitary minimal models $\mathcal{M}(m+1, m)$ is (conventions as in [25])¹⁵:

$$\Delta_{r,s} = h_{r,s} + \bar{h}_{r,s}, \quad (3.4)$$

where

$$h_{r,s} = \bar{h}_{r,s} = \frac{((m+1)r - ms)^2 - 1}{4m(m+1)}, \quad 1 \leq r \leq m-1, \\ 1 \leq s \leq m, \quad m \in \mathbf{Z}_{\geq 3}. \quad (3.5)$$

The range in s in (3.5) can be restricted to $1 \leq s \leq r$ using the field equivalence

$$\phi_{r,s} = \phi_{m-r, m+1-s}. \quad (3.6)$$

Minimal models feature a faithful discrete \mathbf{Z}_2 symmetry for any m (see e.g., [26] for a pedagogical exposition).¹⁶ Using the equivalence class (3.6), for any r and s we can choose the representative field $\phi_{r,s}$ with $r+s$ an even integer. On these representatives, \mathbf{Z}_2 -even fields are those with r and s odd, while \mathbf{Z}_2 -odd fields have r and s even (the identity operator is $\phi_{1,1}$). It is a straightforward exercise to show that there are one or more unitary minimal models which contain a \mathbf{Z}_2 -odd field σ with scaling dimension equal to each of the first 7 values of Δ_σ reported in (3.3).¹⁷ For instance, for $\Delta_\sigma = 3/20$ within minimal models with $m \leq 100$ we find two which contain a Virasoro primary with that dimension: $\phi_{8,8} \subset \mathcal{M}(15, 14)$ and $\phi_{14,14} \subset \mathcal{M}(26, 25)$. The fusion rules $\sigma \times \sigma$ contain respectively 42 and 132 Virasoro primaries. Aside from the identity operator, 11 and 21 of these \mathbf{Z}_2 -even fields have $\Delta < 2$, respectively. The smallest Virasoro primaries in the two cases are $\phi_{5,5}$ and $\phi_{3,3}$, with scaling dimensions $\Delta_\epsilon = 2/105, 2/325$, respectively. Models in such kind of field correlators cannot be detected by our method, because (i) their OPE contain too many light operators and (ii) operators with a too small scaling dimension makes the numerical analysis unfeasible.

¹⁴Unless the continuum starts beyond Δ^* in all spin channels, in which case the CFT will simply appear discrete.

¹⁵In the $2d$ CFT literature, Virasoro and $SL(2, \mathbb{C})$ primaries are usually denoted primaries and quasiprimaries (respectively). In numerical bootstrap analysis we usually care about $SL(2, \mathbb{C})$ primaries, which are the analogs of the primary fields in $d > 2$ CFTs. Most of the considerations that follow can be made at the level of Virasoro primaries with no need to decompose them in $SL(2, \mathbb{C})$ primaries.

¹⁶The presence of such a \mathbf{Z}_2 symmetry is clear from the equivalence of minimal models with Wess-Zumino-Witten cosets of the form $[SU(2)_{m-2} \times SU(2)_1]/SU(2)_{m-1}$, where \mathbf{Z}_2 is the combination of the two \mathbf{Z}_2 center symmetries of $SU(2)_{m-2}$ and $SU(2)_1$ which is not gauged by the center of $SU(2)_{m-1}$.

¹⁷Of course, minimal models associated to the irrational value of Δ_σ reported in (3.3) can also be found provided we take m large enough and approximate $\sqrt{5}/10$ with an appropriately chosen rational number.

Imposing that we have only one relevant \mathbf{Z}_2 -even scalar removes most of these field identifications. The only identification left, for any $m \geq 3$, is the one where $\sigma = \phi_{m-1,m-1}(\sim\phi_{1,2})$ and $\epsilon = \phi_{1,3}$, already considered in [27], which have

$$\Delta_\sigma = \frac{1}{2} - \frac{3}{2(m+1)}, \quad \Delta_\epsilon = 2 - \frac{4}{m+1}, \quad (3.7)$$

and the simple fusion rules

$$\sigma \times \sigma = 1 + \epsilon. \quad (3.8)$$

Only 4 among the 7 rational values of Δ_σ in (3.3) equal the scaling dimensions of $\phi_{m-1,m-1}$: ordinary ($m=3$, $\Delta_\sigma = 1/8$), tricritical ($m=4$, $\Delta_\sigma = 1/5$), tetracritical ($m=5$, $\Delta_\sigma = 1/4$) and pentacritical ($m=6$, $\Delta_\sigma = 2/7$) Ising theories. Note that for even m we have $\mathbf{Z}_2(\sigma) = +\sigma$. However, as far as we consider the four-point function of σ only, given the fusion rules (3.8), we can effectively assign an unfaithful \mathbf{Z}'_2 charge to σ such that $\mathbf{Z}'_2(\sigma) = -\sigma$ and identify such \mathbf{Z}'_2 as the global discrete symmetry of the four-point correlator.

Another notable example which shows the importance of restricting to one relevant scalar is given by the theory of a free compact scalar on a S^1/\mathbf{Z}_2 orbifold. This CFT has two Virasoro primaries with $\Delta_\sigma = 1/8$, which are the twisted ground states of the orbifold. They are odd under the quantum \mathbf{Z}_2 global symmetry that emerges after the gauging and hence we can consider any of them as our field σ . As well-known, the S^1/\mathbf{Z}_2 orbifold is a $c=1$ theory admitting an exactly marginal deformation, given by the radius R of S^1 . At $R^2=1$ the theory is equivalent to a decoupled pair of Ising theories.¹⁸ Deforming the radius gives rise to a family of theories related to the Ashkin-Teller model, namely the theory obtained by coupling two Ising models [28]. The OPE of two twist fields give rise to an infinite number of Kaluza-Klein and winding vertex operators (Virasoro primaries) with scaling dimensions [29]

$$\Delta_{m,n} = \frac{m^2}{R^2} + \frac{n^2 R^2}{4}, \quad m \in \mathbf{Z}, \quad n \in 2\mathbf{Z}. \quad (3.9)$$

At $R^2=1$, the two Ising theories are decoupled and this theory is indistinguishable from a single Ising copy. When $R \neq 1$ the two couples to each other and we get a continuum of $c=1$ CFTs parametrized by R . For any R such theories contain more than 1 relevant \mathbf{Z}_2 -even scalar, as we see from (3.9) by setting either m or n to zero. In this case, restricting to only one relevant \mathbf{Z}_2 -even scalar is crucial to avoid a continuum of theories, which would be a challenge for our numerical methods.

¹⁸In our normalizations, like those of [25], the self-dual radius is $R^2=2$.

The external operator σ could also be part of a CFT with a continuous global symmetry G . CFTs with continuous global symmetries should of course be analyzed in a covariant way exploiting the symmetry, see e.g., [30]. However, if $G \supset \mathbf{Z}_2$ the field σ could be considered some (real) component of a multiplet in a representation of G for which $\mathbf{Z}_2(\sigma) = -\sigma$. For example, σ could be the real component of a spin j multiplet, with half-integer j , of a $SU(2)_k$ WZW model. More specifically, we could have

$$\sigma \sim \phi_{j,m}(z)\bar{\phi}_{j,-m}(\bar{z}) + \phi_{j,-m}(z)\bar{\phi}_{j,m}(\bar{z}), \quad (3.10)$$

where $\phi_{j,m}(z)$ and $\bar{\phi}_{j,\bar{m}}(\bar{z})$ are the holomorphic and anti-holomorphic highest-weight state components of the spin j field. The symmetry \mathbf{Z}_2 could be identified with the center of either the holomorphic or of the antiholomorphic $SU(2)$ gauge group. With the identification (3.10), we have

$$\Delta_\sigma = \frac{2j(j+1)}{k+2}, \quad j \in \frac{\mathbf{Z}}{2}. \quad (3.11)$$

At $\Delta_\sigma = 1/5, 2/7$ there are no solutions with integer k and half-integer j of (3.11). For the remaining rational values in (3.3) we have an infinite number of solutions with increasing values of j and k . Solutions with $j > 1/2$ have more than one relevant \mathbf{Z}_2 -even scalar exchanged in the $\sigma \times \sigma$ OPE. On the other hand, for $j = 1/2$ we get only one relevant \mathbf{Z}_2 -even scalar with¹⁹

$$\Delta_\epsilon = \frac{4}{k+2}. \quad (3.12)$$

This occurs for four of the values reported in (3.11):

$$\left(k=3, \Delta_\sigma = \frac{3}{10} \right), \quad \left(k=4, \Delta_\sigma = \frac{1}{4} \right), \\ \left(k=8, \Delta_\sigma = \frac{3}{20} \right), \quad \left(k=10, \Delta_\sigma = \frac{1}{8} \right). \quad (3.13)$$

This quick detour on some well-known $2d$ CFTs shows the importance of reducing the space of viable theories by demanding only one relevant scalar. Needless to say, the space of CFTs is vastly larger than the three classes we discussed, so it is likely that other CFTs with σ operators with scaling dimensions as in (3.3) and only one relevant \mathbf{Z}_2 -even scalar exist.

Results in $d=2$ We report in Figs. 5 and 6 the location of the end-minima found with $\ell_{\max} \leq 8$ in the $(\Delta_\sigma, \Delta_\epsilon)$ and (Δ_σ, c) -planes, respectively. The end-minima belong to different sectors, all of them contained in 4_4_3_2_1.

¹⁹We also have marginal scalars with $\Delta=2$ associated to the $\mathcal{J}\bar{\mathcal{J}}$ operators arising from the identity character.

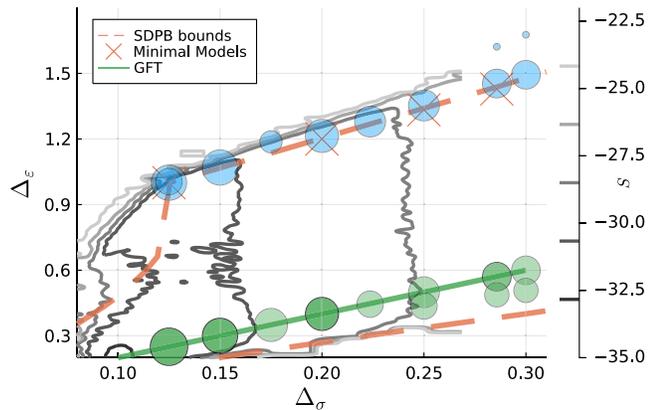


FIG. 5. $d = 2$. Location of the end-minima with $\ell_{\max} \leq 8$ in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. The end-minima belong to different sectors all of them contained in 4.4_3_2_1. The dashed orange line corresponds to the bounds obtained in [24] assuming only one \mathbf{Z}_2 -even relevant scalar. All the end minima correspond to local theories $\Delta_T = 2$, $\Delta'_T \geq 3$.

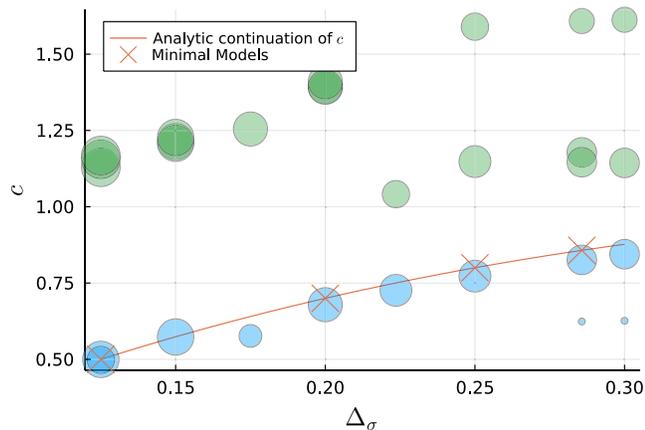


FIG. 6. Central charge c as a function of Δ_σ for end-minima with $\ell_{\max} \leq 8$ in $d = 2$. The end-minima belong to different sectors all of them contained in 4.4_3_2_1. Comparison of our results to the analytic expectation (3.15) (red line) for the generalized minimal models. All the end minima correspond to local theories $\Delta_T = 2$, $\Delta'_T \geq 3$.

Most of the end-minima are aligned along two approximately straight trajectories in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. The blue end-minima are along the upper extremal bootstrap bounds, while the green ones are aligned along the GFT line, despite the constraint $\Delta_T = 2$.

Let us first discuss the blue end-minima. Two less significant blue end-minima appear above the extremal line in the disallowed region, indication of being fake end-minima not associated to CFTs. The rest are in good agreement with the bootstrap bounds.

It is reassuring that, among such end-minima, we reproduce the first four minimal models (red crosses in Fig. 5), with the field identification as in (3.7) discussed

before. The remaining four points along the line correspond instead to the so called generalized minimal models [16], i.e., nonunitary minimal models obtained by analytically continuing the CFT data appearing in the 4-point correlation function from integer values of m (minimal models) to real values of m (generalized minimal models). Such theories are visible in the conformal bootstrap, despite the enforcement of unitarity, because the OPE coefficients squared of the exchanged quasiprimaries have been conjectured in [16] (see also [31]) and then proved in [24], to be positive for any m .

Given (3.7), the blue end-minima should be aligned along the line

$$\Delta_\epsilon = \frac{2}{3} + \frac{8}{3} \Delta_\sigma. \quad (3.14)$$

In terms of the central charge, we have

$$c = \frac{9\Delta_\sigma}{\Delta_\sigma + 1} - 4\Delta_\sigma. \quad (3.15)$$

We compare in Fig. 6 the central charge (3.15) with the one extracted from the OPE coefficient $\lambda_{\sigma\sigma T}$ by means of the relation

$$c = \frac{d}{d-1} \frac{\Delta_\sigma^2}{4\lambda_{\sigma\sigma T}^2}, \quad (3.16)$$

valid for general d dimensions. We see that the agreement is very good, despite the severity of our truncation, and confirm also the identification of such end-minima with generalized minimal models.

As a further check we compare in Fig. 7 the OPE coefficient squared $\lambda_{\sigma\sigma\epsilon}^2$ with the one analytically found for such models in Appendix B of [16]. The good agreement

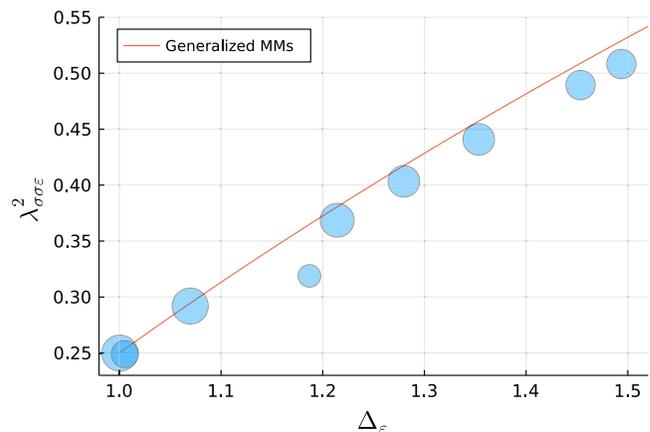


FIG. 7. $d = 2$. Comparison between the OPE squared coefficient $\lambda_{\sigma\sigma\epsilon}^2$ numerically determined from the blue end-minima in Fig. 5 and its analytic expression for generalized minimal models (red line) [16], as a function of Δ_ϵ .

found is yet another indication of the accuracy of our determination and of the nature of the blue end-minima. The only exception is the minimum at $\Delta_\sigma = 7/40$, which falls slightly above the SDPB bounds in Fig. 5 and its values of c and $\lambda_{\sigma\sigma}^2$ in Figs. 6 and 7 are slightly misaligned with those of the generalized minimal models. We expect this to be due to some numerical artifact and that a more thorough study will find a lower action minimum that better agrees with the theoretical values.

Let us now discuss the green end-minima of Fig. 5. In contrast to the blue ones, we do not have a convincing explanation for the nature of these end-minima as local CFTs, assuming they are not numerical artifacts. Most of them lie on the GFT line $\Delta_\epsilon = 2\Delta_\sigma$. A notable class of local CFTs that have in their spectrum operators with scaling dimensions related in this way are $\mathcal{N} = 2$ CFTs. It would be interesting to understand if the green points along the GFT line (or some of them) can be identified as $\mathcal{N} = 2$ supersymmetric minimal models, provided an appropriate identification of σ is made. This analysis would probably require also to understand whether generalized $\mathcal{N} = 2$ minimal models, in the spirit of [16], exist and can be constructed.

Despite the fact that we enforced the presence of an energy momentum tensor, the green end-minima along the GFT line could simply be an approximation of GFTs themselves. This observation is supported by the fact that if we remove the gap assumption $\Delta_{T'} \geq 3$ most of the green minima are unstable. In particular, $\lambda_{\sigma\sigma T}^2 \sim 0$ and $\Delta_{T'} \approx 2\Delta_\sigma + 2$ effectively replaces T as the first spin-2 operator. Similarly, if we do not assume the existence of an energy momentum tensor, most minima turn out to be GFTs. As we will see in the next subsections, the analog of the green points in Fig. 5 will not occur in $d > 2$ when we enforce the presence of an energy-momentum tensor.

It is useful to also report the spectrum of $SL(2, \mathbb{C})$ -primary operators which appear at the end-minimum at a given truncation. We report as an example in Fig. 8 the full

spectrum of some end-minima found at $\Delta_\sigma = 1/8, 1/5$ (two in each case). The operators associated to the same end-minimum are connected with lines. In the Ising case, where the identification of the end-minima with the Ising model is quite convincing, we see how the agreement of the whole low-lying spectrum is quite accurate, with small deviations occurring for higher-spin operators in the first Regge trajectory (the spin-8 orange diamond) or for operators in the second Regge trajectory (the second spin-0 green diamond, the second spin-4 orange diamond).

In the GFT-like case, as already mentioned, we do not have a convincing explanation for the nature of such theories. Despite having enforced the presence of a spin 2 operator at $\Delta = 2$, we can see from the right panel of Fig. 8 that the spectrum resembles roughly the one of GFTs. The presence of more operators with respect to the Ising case makes the accuracy of the numerics less accurate. In particular, it is not possible to disentangle if the significant deviations from the GFT spectrum at higher Regge trajectories are, in addition to numerical noise, due trivially to the fact that we have imposed the presence of a nonexistent operator (the stress-tensor), or because the theory under question is in fact distinct from a GFT. See Appendix C for some further details, in particular for examples of selected branches, the endpoints of which correspond to the blue and green end-minima discussed above.

It is well possible that other theories with a relatively sparse spectrum satisfying our assumptions (or effectively doing it, like the generalized minimal models) might have not been found by our search protocol, given also the limitations imposed by machine precision computations and the small values of Δ^* and N_{Ops} we sampled. Possible theories of this kind are for instance the $SU(2)_k$ WZW models (3.13) with the field identification (3.10).

2. $d=3, 4$

In $d = 3$ we run our protocol for the following values of Δ_σ :

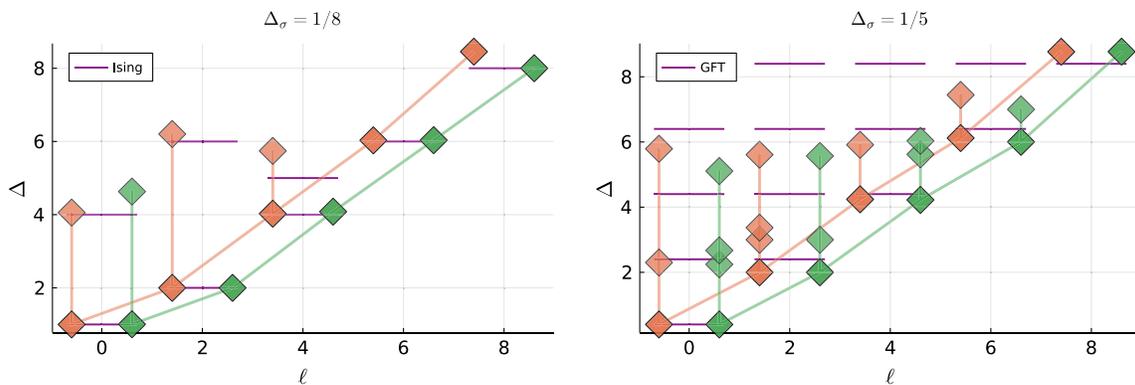


FIG. 8. Full spectrum of selected end-minima in $d = 2$ with $\Delta_{T'} = 2$ and $\Delta_{T'} \geq 3$. Left: Ising-like end-minima. Right: GFT-like end-minima. The colors and the lines connect operators within the same end-minimum. For comparison we also report the exact scaling dimension of the 2d Ising and GFT with purple lines.

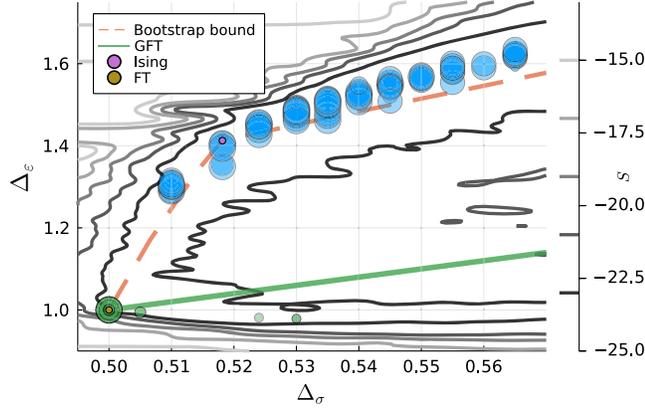


FIG. 9. $d = 3$. Location of the end-minima with $\ell_{\max} \leq 10$ in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. The end-minima belong to different sectors all of them contained in 4.4.4_3.2.1. The dashed line is the upper bound for Δ_ϵ from [22].

$$\Delta_\sigma = 0.5, 0.505, 0.51, 0.5181489, 0.524, 0.53, \\ 0.535, 0.54, 0.545, 0.55, 0.555, 0.56, 0.565. \quad (3.17)$$

We plot in Fig. 9 the end-minima found in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane, with $\Delta_T = 3$, $\Delta_{T'} \geq 4$, $\Delta_{\ell'} \geq 3$. All the end-minima lie at the extremality bound, like in the $d = 2$ case. In contrast to the latter, no end-minima along the GFT appear, while we find a couple of faint end-minima with $\Delta_\epsilon < 1$, below the GFT line.²⁰

Reassuringly, we get end-minima in correspondence of the 3d Ising model, which sits at the kink of the dashed orange lines. End-minima are found for all the Δ_σ in (3.17) with the exception of $\Delta_\sigma = 0.505$. The lack of an end-minimum at such value is due to the imposed gap $\Delta_{\ell'} \geq 3$. Indeed, by repeating the MC analysis of Sec. III A with Δ_σ unconstrained assuming $\Delta_{\ell'} \geq 3$, the “flow” toward the FT passing through the valley connecting the Ising point with the FT in Fig. 1 gets interrupted around $\Delta_\sigma = 0.510$. This is an indication that possible extremal CFTs with $0.5 \leq \Delta_\sigma \leq 0.510$ do have more than one \mathbf{Z}_2 -even scalar.²¹

We plot in Fig. 10 the central charge of the end-minima, computed using (3.16), as a function of Δ_σ . The dashed orange lines delineate the rigorous lower bound of [31]. Most of the end-minima in Figs. 9 and 10 lie in the forbidden region, slightly in the former case and more evidently in the latter. This difference should not come as a surprise, since it is known that the bounds on the central charge are more sensitive than those on Δ_ϵ to higher dimensional operators. Compare e.g., Figs. 3 and 4 of [19]

²⁰At low ℓ_{\max} we get an entire line of minima, but these do not create consistent branches. As we will see later when discussing nonlocal theories, end-minima along the GFT line are reassuringly found when we fix $\Delta_T = 2\Delta_\sigma + 2$ (the GFT value), as expected.

²¹Note that this region is devoid of not only end-minima, but of minima at all.

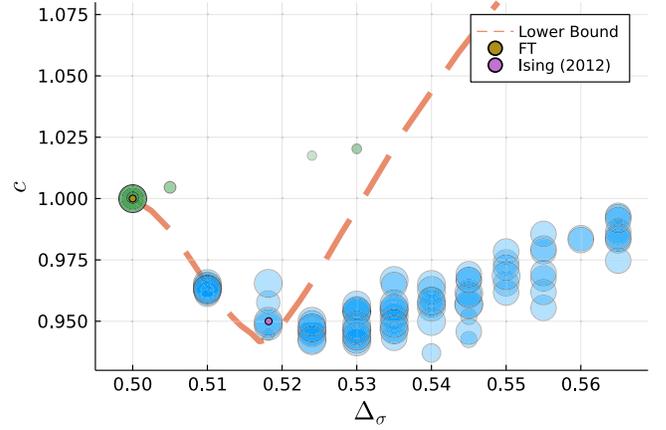


FIG. 10. Central charge c as a function of Δ_σ for end-minima with $\ell_{\max} \leq 10$. The end-minima belong to different sectors all of them contained in 4.4.4_3.2.1. The rigorous lower bounds are taken from [31]. The color code of the end-minima is the same as in Fig. 9.

to appreciate the different convergent properties of the two bounds as the degree of truncation is varied.

The spectra associated to the branches with $\Delta_\sigma = 0.5181489, 0.53$ is reported in Fig. 11. In the left panel we also report as reference values the Ising spectrum found in [11] using the extremal functional method (purple lines). The operators associated to the same end-minimum are connected with lines. In the Ising case we see how the accuracy is significantly worse with respect to the $2d$ case. We have a reasonably good agreement of the low-lying spectrum in the first Regge trajectory, but the remaining operators have a significant indeterminacy. Similarly, in the right panel of Fig. 11 the operators in the first Regge trajectory agree among the two orange and green branches, but significantly differ at higher levels.

In $d = 4$ we run our protocol for the following values of Δ_σ :

$$\Delta_\sigma = 1, 1.05, 1.1, 1.15, 1.2, 1.3. \quad (3.18)$$

Due to the limited numerical accuracy of our protocol in $d = 4$ we just report in Fig. 12 the end-minima found in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane with $\Delta_T = 4$, $\Delta_{T'} \geq 5$, $\Delta_{\ell'} \geq 4$, with $\ell_{\max} \leq 10$. Once again we have end-minima along the extremality bound but it is interesting to see how our minima satisfy the older bounds (with less derivatives) of [1] but are slightly excluded by the bounds in [32]. Interestingly enough, for each value of Δ_σ sampled we also find end-minima below the GFT line.

D. Nonlocal theories

In this subsection we analyze the structure of the minima found when relaxing the condition $\Delta_T = d$. The most notable and exactly calculable nonlocal theory is the GFT. At a given Δ^* the number of operators in the GFT

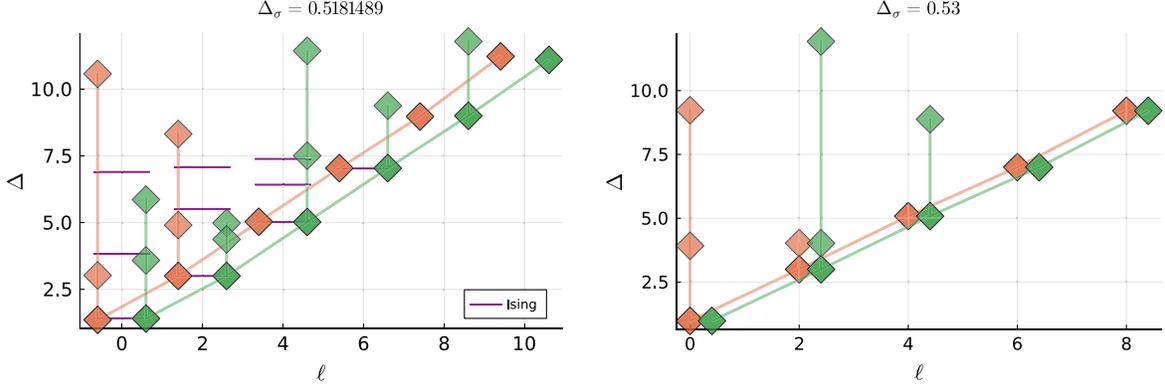


FIG. 11. Left: full spectrum of the minima found at $\Delta_\sigma = 0.5181489$, $\Delta_T = 3$ that resemble the critical Ising model in $d = 3$. Ising reference values taken from [11]. Right: full spectrum of the minima found at $\Delta_\sigma = 0.53$, $\Delta_T = 3$ that have Δ_ϵ below the GFT line. The colors and the lines connect operators within the same end-minimum.

spectrum is larger than the maximal truncation we considered in this paper. Although in principle this rules out the possibility of finding exactly the GFT as end-minima in our approach, the limitation is only theoretical, considering that our numerical accuracy does not anyhow allow to precisely determine operators beyond the first Regge trajectories, as we have seen. In fact, we will find end-minima which seem quite compatible with the leading operators of the GFT at different values of Δ_σ .

The landscape of end-minima is now much richer and cover most of the allowed region. Interestingly enough, along the Δ_T direction we find that there is a clear tendency of the minima to cluster around $\Delta_T \approx 2\Delta_\sigma + 2$, see Fig. 13(a) for $d = 2$ (a similar phenomenon occurs in $d = 3$ and $d = 4$), where we report the distribution of the end-minima in the $(\Delta_\sigma, \Delta_T)$ plane. As can be seen, we have many more end-minima with respect to those found when $\Delta_T = d$ and $\Delta_T > d + 1$ were imposed. Interestingly

enough, most end-minima have $d \leq \Delta_T \leq d + 2\Delta_\sigma$. Note that only for a limited range in Δ_σ , end-minima with $\Delta_T = d$ are found. This is particularly evident in $d = 2$, where an end-minimum at $\Delta_T = 2$ is found only at $\Delta_\sigma = 1/8$. This signals the fact that the extremal end-minima associated to the (generalized) minimal models with $m > 3$ found in Sec. III C 1 become unstable when we relax Δ_T and move toward the GFT region. In panel (a) of Fig. 13 two specific end-minima points have been singled out at $\Delta_\sigma = 1/8$ (green circle) and at $\Delta_\sigma = 3/10$ (red circle). They have been selected as those resembling most closely the GFT conformal data using (3.19) as criterion. The truncated spectrum at those points is reported in panel (b) of Fig. 13. Note the good agreement beyond the first Regge trajectory, in particular in the scalar sector where the first three scalars are well reproduced in both theories.

We have also explored the effect of fixing the scaling dimension of the first spin 2 operator at the exact GFT value, still assuming $\Delta_T \geq d + 1$. This is useful to understand whether minima at fixed Δ_T are “harder” to reach than similar points when Δ_T is allowed to vary. We show these results in Fig. 14 for the $d = 3$ case, where the blue circles are obtained at $\Delta_T = 3$ and are the same as in Fig. 9, while the red ones are those arising at $\Delta_T = 2\Delta_\sigma + 2$. The red end-minima align along the GFT line, as expected, but they spread a wider area, which however does not include the extremality region. Similar considerations apply also in $d = 2$ and $d = 4$.

We saw in Fig. 13(a) that some of the minima found by our protocol at $\Delta_\sigma = 1/8$ do resemble local theories, although they are clearly disfavored in front of the GFT-like ones. We would like to determine more precisely how well we reproduce local theories when no constraint on Δ_T is imposed. This can be attained by taking the 2d Ising theory and the GFT as reference points and defining a distance measure between our minima and these exactly known CFTs. We will here take the scaling dimensions of the first scalar and first spin-2 operators and their associated

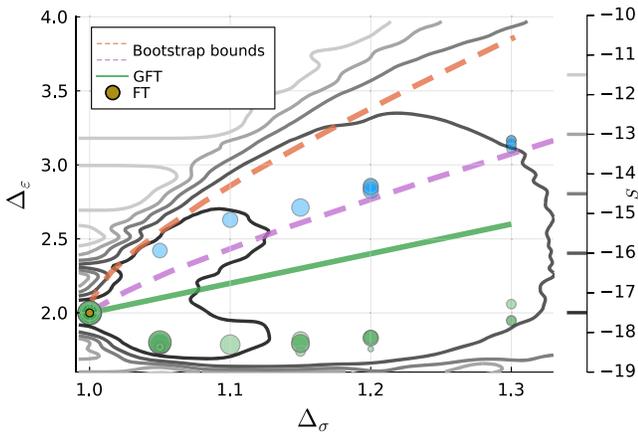


FIG. 12. $d = 4$. Location of the end-minima with $\ell_{\max} \leq 10$ in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. The end-minima belong to different sectors all of them contained in 4_4_4_3_2_1. The dashed orange and purple lines correspond to the bounds obtained in [1,32], respectively.

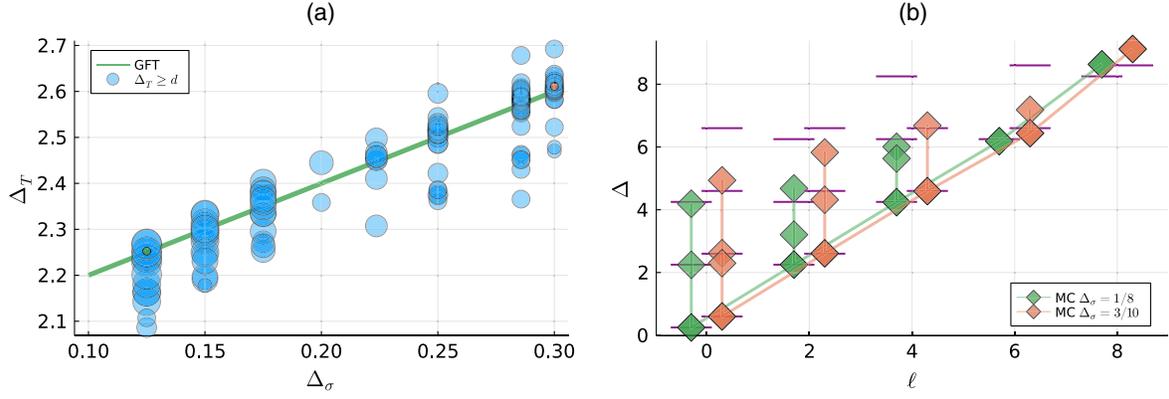


FIG. 13. (a) End-minima in $d = 2$ with no condition on Δ_T besides unitarity. (b) Spectra for the closest match to the GFT for $\Delta_\sigma = 1/8, 3/10$. The colors and the lines connect operators within the same end-minimum. Purple lines represent the theoretical values of the GFT at that Δ_σ .

OPEs coefficients, and compute the average relative error with respect to those of the known theory:

$$Q(\mathbf{v}_{\min}, \mathbf{v}_{\text{exact}}) = \frac{1}{4} \sqrt{\sum_{i=1}^4 \left(\frac{v_{\min}^i - v_{\text{exact}}^i}{v_{\text{exact}}^i} \right)^2}, \quad (3.19)$$

where $\mathbf{v} = (\Delta_e, \Delta_T, \lambda_{\sigma\sigma e}^2, \lambda_{\sigma\sigma T}^2)$.

In Fig. 15 we show how the blue circles reach configurations that are not too far away from the 2d Ising theory (roughly 10% mean relative error) but also cover configurations which reproduce the leading conformal data of the GFT even better than the MC which had $\Delta_T = 2 + 2\Delta_\sigma$ to begin with. Finally, Fig. 15 supports the interpretation of the green minima in Fig. 5 to be “images” of the GFT, since we can appreciate that they correspond to solutions which reproduce the conformal data of this theory with an error of less than 10%.

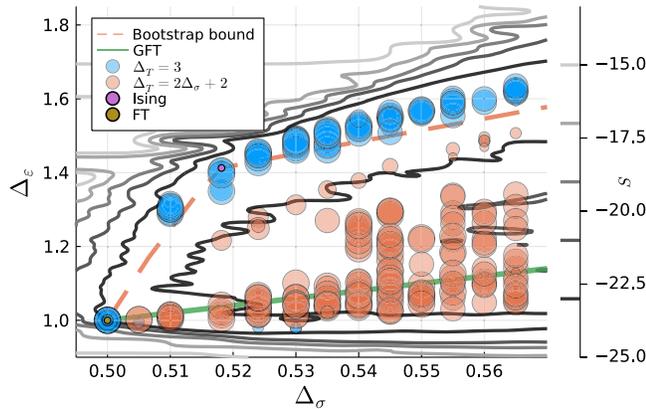


FIG. 14. $d = 3$. Same as Fig. 9 but we also show with red circles the minima found by fixing Δ_T to the GFT value. Note that we show here in blue all the end-minima that in Fig. 9 were reported in blue and green.

E. A taste of nonunitarity: The 2d Yang-Lee model

The MC search protocol of Sec. II C does not rely on unitarity. Although we mostly focus on unitary models in this paper, we would like to briefly show here a proof of concept of its use in the nonunitary realm, by checking that the $d = 2$ Yang-Lee (YL) model can easily be found using our method. Nonunitary theories have received less attention from the bootstrap community due to their non tractability with linear or positive semidefinite programming methods. The best approach so far to solve nonunitary theories with the conformal bootstrap is given by [33,34], which also severely truncates the crossing equations (as we do). As is well-known, the $d = 2$ YL model can be described by the nonunitary minimal model $\mathcal{M}(5, 2)$, which contain only two Virasoro primaries related by the following fusion rule:

$$\Phi \times \Phi \sim 1 + \Phi. \quad (3.20)$$

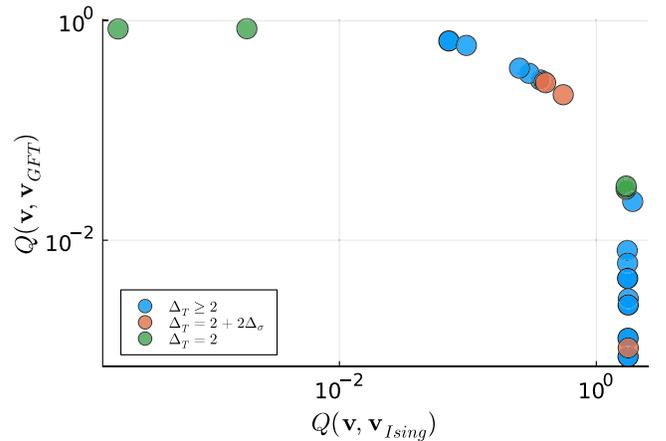


FIG. 15. Distance of different end-minima to the $d = 2$ Ising and the GFT with $\Delta_\sigma = 1/8$. The quantity Q is computed according to (3.19) for end-minima with $\Delta_T = 2, 2 + 2\Delta_\sigma$ and those in which Δ_T is determined dynamically.

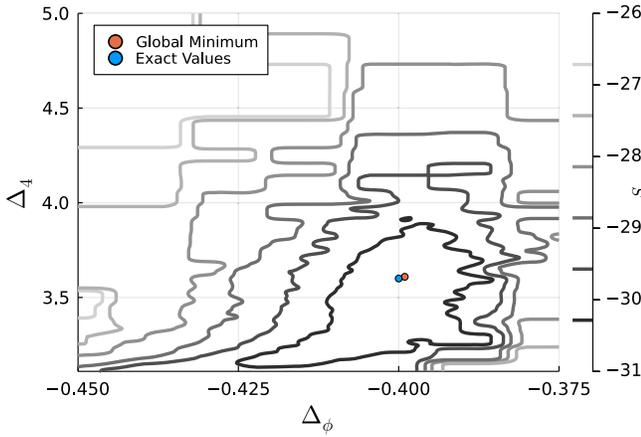


FIG. 16. Projection of the value of the action S on the (Δ_ϕ, Δ_4) -plane. The lowest-action point visited by the MC at $T = 0.5$ and the exact values are shown as red and blue circles, respectively.

This means in practice that the external dimension Δ_ϕ will also be the scaling dimension of the first exchanged scalar and that we must not impose positivity of the OPE coefficients.

We launched a Metropolis Monte Carlo search with $T = 0.5$ and a spectrum 3.1_2.1 ($\Delta^* = 9$) to sample a parameter space close to the one containing the exact CFT data of the $2d$ YL model. We find that the global minimum is compatible with the exact values, as can be seen in Fig. 16 for Δ_ϕ and Δ_4 (the first spin 4 operator). At the global minimum the other operators show also similar agreements with the exact values.

We note in passing that in $d = 3$, if we relax the positivity constraint on the OPE coefficients, a search with a 1_1_1_1 spectrum shows another basin of attraction besides those discussed in Sec. III B. The basin is in a region well below the unitarity bounds for Δ_σ and lies roughly along the line $\Delta_\sigma = \Delta_\epsilon$. Besides being an interesting feature of the landscape in itself, we consider this to be an indication that the $d = 3$ YL model should also appear as a minimum if investigated with our method.

IV. OUTLOOK

Motivated by the pressing need of a tool to explore the inner part of the allowed space of CFTs, we have introduced in this paper a method based on a stochastic minimization via a Metropolis algorithm to find approximate numerical solutions with a sufficiently sparse spectrum of operators to truncated bootstrap equations for CFTs in arbitrary dimensions. Given an initial value of scaling dimensions for a reduced set of primary operators up to some scaling dimension Δ^* , the action (2.6) is minimized and an approximate set of CFT data found.

While for all physical CFTs we should eventually have $F = e^S \rightarrow 0$ in the limit when all operators are taken into

account, at fixed operator truncation the rate in which this limit is reached does matter. For instance, we have seen that for arbitrary severe truncations of the operator spectrum, the free theory minimum dominates over all and no other theory can be found, see e.g., Fig. 1. In order to find other theories, extra conditions, such as fixing the external operator dimension, have to be imposed. It is clear that, at fixed truncation, not all CFTs are equally accessible, even assuming the same sparsity of spectrum. We expect that a richer set of solutions can be found by generalizing our method to multiple correlators and by imposing further constraints that would halt the MC evolution toward specific attractors (e.g., the extremality line or the GFT line in $d = 2$). Though we mostly focused on unitary theories, the method does not rely on it and can efficiently be used in the nonunitary realm. This is particularly important considering that for nonunitary theories we still do not have rigorous numerical protocols. It would be interesting to see how our numerical method performs with respect to Gliozzi's method [33,34] and its subsequent versions [35,36]. The aim of this work is very similar to that of [14,15], where approximate solutions to crossing have been achieved using a similar logic but a different technique (reinforcement learning instead of MC techniques). It would be interesting to systematically compare the two approaches.

In this work the bootstrap equations have been evaluated in a given sample of points in cross-ratio space. This choice has been primarily dictated by the flexibility, when developing the algorithm, of changing the number of points and the possibility (used in earlier versions of the code) of sampling stochastically these points. The actual performance of the code sensitively depends on how points in cross-ratio space are chosen. It would be interesting to further improve on this aspect. In particular, an improvement could result by choosing derivatives evaluated at the crossing symmetric point $z = \bar{z} = 1/2$, as currently done in functional bootstrap studies.

The FORTRAN-code used in this work is limited to machine-like precision. While this allowed us to perform extensive MC searches at little cost, limited precision does not allow us to systematically increase the value of Δ^* and N_{Ops} . For example we do not have enough accuracy to understand the nature of the green end-minima we found in $2d$ (Fig. 5), whether they are unidentified local CFTs of some kind or some sort of numerical artifacts resembling GFT theories. For $d > 2$ the accuracy is even less and the nature of the found end-minima is yet to be understood. Despite the limited accuracy of the algorithm, we think the results of this paper have shown that the method works. Upgrading the current code to arbitrary precision will be the key for a fully working program and is the most important thing to do in the near future. We believe this is feasible, given the ample room for improvement of the code performance in terms of speed.

In summary, we hope that the nonrigorous approach initiated in this work, properly extended and improved, might be a useful tool to navigate in the space of allowed CFT and guide subsequent analysis performed using more rigorous numerical algorithms.

ACKNOWLEDGMENTS

U. L. V. and M. S. thank the organizers of the Bootstrap 2019 conference held in PI for the hospitality and all the participants for interesting discussions. M. S. also thanks F. Benini, S. Cecotti, L. Eberhardt and G. Mussardo for discussions on $2d$ CFTs. Work partially supported by INFN Iniziativa Specifica ST&FI.

APPENDIX A: NEWTON RHAPSON MINIMIZATION

The Newton-Rhapson algorithm is a way of finding the local minima of \mathcal{C}^2 -functions with a positive definite Hessian H . We will now discuss the basic idea behind this method and the implementation used in this work. For a detailed explanation the reader is suggested to consult [20].

In short, the method uses the Hessian matrix to approximate the function to be minimized up to quadratic order and solves iteratively for the minimum of this hyper paraboloid. In this work a modified version of this method is used (known as the Levenberg-Marquardt algorithm [37,38]), in which a damping factor β interpolates between gradient descent and the bare Newton-Rhapson method, avoiding the overshooting of the minimum when very close to it. This is necessary because due to numerical noise some eigenvalues of the Hessian might not be positive, which violates the conditions for the convergence Newton-Rhapson method.

More concretely, let $f(x):\mathbb{R}^n \rightarrow \mathbb{R}$ be a twice-differentiable function. Given an initial guess x_0 we find the next point by using the following formula:

$$x_1 = x_0 - (H(x_0) + \beta \mathbb{1})^{-1} \cdot \nabla f(x_0). \quad (\text{A1})$$

At each step, if $f(x_1) < f(x_0)$ then β is decreased by a factor of 3. Otherwise it is increased by a factor of 2. This is known in the literature as “delayed gratification” [39] and gives stable results in the minimization problems explored in this work.

The factor β also gives the termination condition for the algorithm: when $\beta > 10^6$ it is clear that a stationary point has been reached.

APPENDIX B: IMPLEMENTATION DETAILS

The protocol detailed in Sec. II C was implemented in a series of automatized bash scripts that used Julia for the data processing and FORTRAN90 (compiled with `ifort v2021.2`) for the numeric evaluation of the action and the different minimization algorithms. LAPACK was used for the linear algebra manipulations. They were run on Intel XeonTM processors in an institutional cluster. The CPU time to run 10^8 steps was of roughly 12 hours for representative spin partitions. The parameters used in the searches are shown in Table IV.

In order to provide the technical details we will discuss in the following subsections the different elements of our pipeline and their implementation.

1. Numerical evaluation of S

We start by discussing the approximation and numerical evaluation of (2.6). We describe each of the approximations used in order to render computationally feasible sampling $O(10^8)$ points in ~ 10 hours on a single CPU.

a. Discretization in z

The grid is parametrized as $z = x_0 + an + y_0 + ibm$ with $m, n \in \mathbb{N}$ and $x_0 = 1/2 + 1/1000$, $y_0 = 1/1000$. In this work we used $a = 1/100$ and $b = 1/50$ but in the computation of S only a randomly chosen sample of `nZ` points was considered. This was done by ordering them with respect to their value of λ and then sampling this list.

TABLE IV. Summary of the relevant parameters for the wide searches (first step of the protocol).

Parameter	Value(s)	Description
<code>nZ</code>	200	Number of points taken in the z -plane.
<code>lambda0</code>	0.42	Size of the region in which the points are sampled.
<code>NT</code>	2×10^8	Number of steps (MC time).
<code>positiveOPes</code>	<code>true</code>	Enforcement of real OPE coefficients.
<code>externalInOPE</code>	<code>false</code>	Fixing the first exchanged scalar to be equal to the external operator.
<code>Temp</code>	0.3–0.9	Temperature (T) of the Metropolis Monte Carlo algorithm.
<code>MCstep</code>	0.001	Overall scale factor for the MC-step.
<code>wall</code>	10^4	Constant for the quadratic penalty that enforces the bounds on Δ_i .
<code>frameRate</code>	1000	The MC saves one out of <code>frameRate</code> steps as a frame.
<code>nops</code>	3–20	Number of operators (besides the identity).
<code>boundaries</code>	$[\Delta_{\text{unit}}, \Delta^*]$	Range to which each operator is constrained.

b. Truncation of the spectrum

This was implemented by passing to the FORTRAN routine the list of the spins of each operator in the truncation under study. It is in this step that we input important parameters such as the spin partition, Δ^* and ℓ_{\max} .

c. Evaluation of the conformal blocks

The conformal blocks were tabulated using 160 digits of internal precision in MATHEMATICA and stored as plain text files for each ℓ . In $d = 2, 4$ we used the closed expressions in terms of hypergeometric functions [17]. For $d = 3$ the recursion relations implemented as in [40] were used. For each point in z we computed the CBs at 5000 points evenly spaced in $[\Delta_{\text{unit}}(\ell) - 99/100, \Delta^*]$ with $\Delta^* = d + 15$. We note in passing that this gives a lower bound on the error in our determination of the scaling dimensions (Δ) of the exchanged operators. In the FORTRAN code, the interpolation is done with cubic splines.

d. Determination of the OPE coefficients

Once the matrix of the conformal blocks for each z and operator (Δ, ℓ) has been computed, the determination of the OPE coefficients is then a weighted least squares problem which we solve using LAPACK (see (2.7)).

If `positiveOPES` is true, this is not the end of the story because the OPE coefficients that solve the quadratic problem might as well be negative. To enforce unitarity, we iteratively decouple the offending operators until every OPE coefficient is positive. It is important to underline that this is done at each point evaluated in our algorithm, and thus operators do not decouple for the whole test but only for those particular configurations.

e. Computation of S

As described in Sec. II, we obtain S by taking the logarithm of the sum of the squared residuals. If any operator has a Δ outside its allowed boundaries we impose a quadratic penalty that is then added to the action. This guarantees the analyticity of our potential and thus avoids noxious boundary effects.

More concretely, suppose an operator bound to be in the interval $[\Delta_0, \Delta_1]$ has $\Delta < \Delta_0$. In that case, the effective action will be

$$S_{\text{eff}} = S + \text{wall}(\Delta - \Delta_0)^2,$$

where the constant `wall` is set to 10^4 .

2. Numerical error

After all the approximations discussed above, reasonable doubt could remain about the accuracy of S computed with our method. Thus, we consider it necessary to quantify the magnitude of the numerical errors in our implementation.

We accomplish this by computing S at representative points with our framework and by comparing them to the values of S obtained in *Mathematica* with arbitrary precision.

From this analysis we conclude that the action computed with *Mathematica* is systematically lower than the approximated one using our FORTRAN implementation. However, the discrepancy is such that at the values of Δ^* studied in this paper we can still trust the results. Naturally, the errors become larger for bigger N_{Ops} and Δ^* , but even in the cases with the worst agreement the relative errors are smaller than 10^{-3} . Moreover, the largest discrepancy sets a natural lower bound for T , since we must allow for fluctuations of this height in order not to wrongly identify artifacts as minima. For example, in $d = 3$ the minimum temperature for $\ell_{\max} = 6$ is $T \sim 10^{-4}$, whereas for $\ell_{\max} = 8$ is $T \sim 10^{-3}$. We note that this resolution is likely enough for any physical minimum, given that in the $\Delta^* \rightarrow \infty$ limit any physical theory should have $S \rightarrow \infty$.

As an example, we show in Fig. 17 the relative error computed for 10^4 gaussian perturbations of a local minimum in sector 2_2_1_1_1 for $d = 3$, $\lambda_0 = 0.42$ and $nZ = 200$. The fact that the error is mostly negative can be easily understood from the delicate numerics of the bootstrap equations. It is known that in order to solve (2.1) to a high degree of precision, very fine-tuned cancellations between the terms must occur. When using our double-precision interpolation of the blocks, there will be points where these errors will spoil the cancellations and thus the value of S can only be bigger than the “exact” one.

Another point that deserves discussion is the determination of the OPE coefficients (ρ in eq. (2.6)). Solving for ρ amounts to finding the minimum of a quadratic form. We find that the curvature of this form is very anisotropic, with hierarchies of as much as 10 orders of magnitude. This is in fact one of the main bottlenecks to adding too many

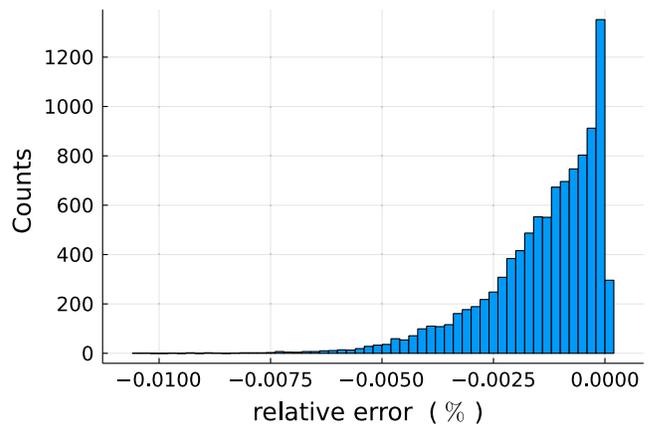


FIG. 17. Distribution of the relative error defined as $S_{\text{Mathematica}} - S_{\text{FORTRAN}}$. The action was evaluated at 10^4 points normally distributed around an apparent local minimum in $d = 3$ with spin sector 2_2_1_1_1, $\lambda_0 = 0.42$ and $nZ = 200$. Note the long tail for negative values.

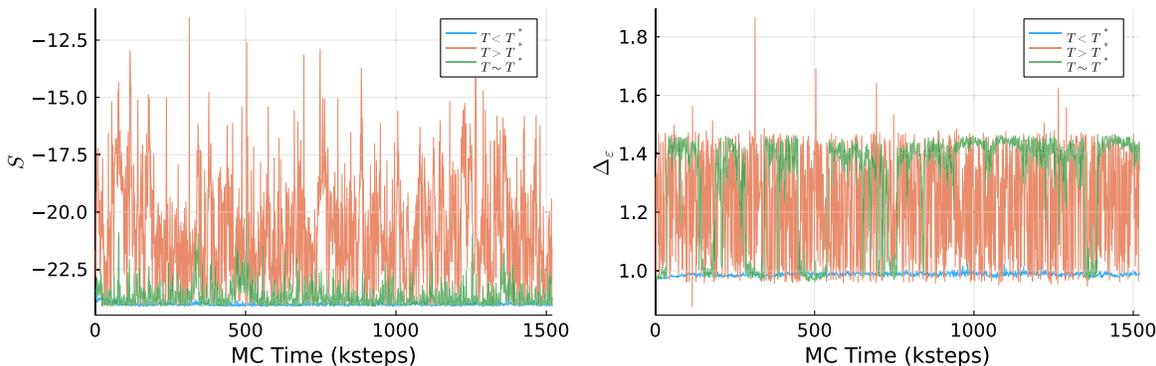


FIG. 18. Sample trajectories for MC runs at different temperatures. Left: value of the action S of the points visited. Right: scaling dimension of the lightest exchanged scalar operator. This particular test was done in $d = 3$ by taking $\Delta_\sigma = 0.519$ fixed, with a spin partition 3_2_2_1 and $\Delta^* = 10$. The trajectories are shown for $T = 0.05, 0.3, 0.9$. In this case, $T^* = 0.3$.

operators: the curvature of the subleading operators at some points becomes too small to be resolved with double precision and the OPE coefficient of these operators cannot be reliably determined.

3. Ideal temperature

As mentioned earlier in Sec. II, the key step of our protocol is the wide search performed using the Metropolis Monte Carlo algorithm. In order for this step to be efficient the temperature must be chosen wisely. To this end we performed test searches for different values of d , N_{Ops} and Δ^* , where several different temperatures were used. Then, we determined by inspection the ideal one in each setting, according to the criterion described in Sec. II C.

In Fig. 18 we show a concrete example of how to identify the ideal temperature. We show “cold,” “hot” and just right temperatures in blue, red and green (respectively). We can see clearly that the blue line “freezes” into the first local

minimum found by the MC and then stays there for the whole test, whereas the red trajectory travels back and forth randomly. We consider the green trajectory to be representative of an efficient search because it oscillates around local minima for some time before making a transition into another one. While this could mislead the reader into thinking that the time spent oscillating around a local minimum is in some sense “wasted” one must bear in mind that during those steps, the MC effectively samples the rest of the scaling dimensions, thus refining the solution to crossing in that neighborhood.

APPENDIX C: MINIMA AND END-MINIMA

In this appendix we illustrate the relation between minima and end-minima. As discussed in the main text, end-minima are defined as those special minima which are the endpoints of branches. The latter are a set of akin minima, supposedly associated to the same CFT.

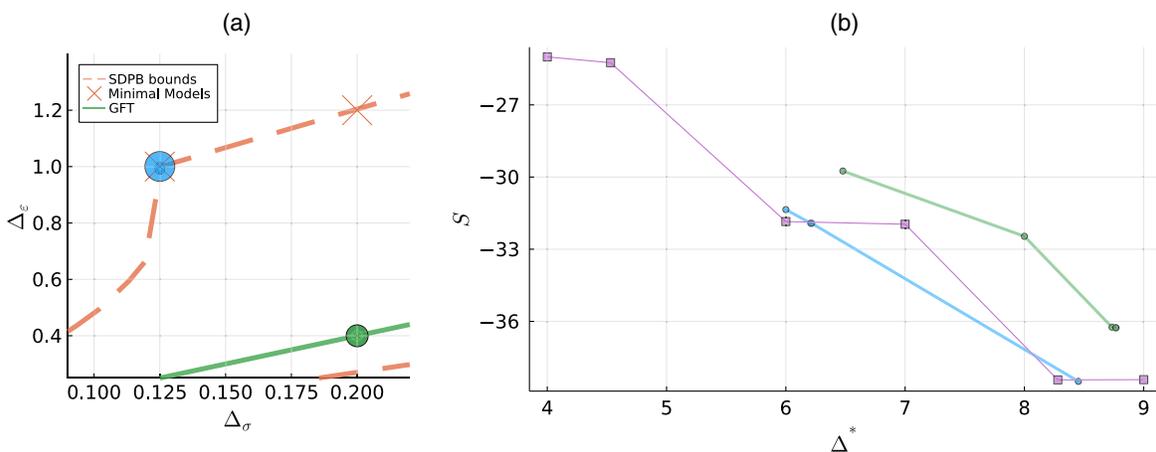


FIG. 19. Representative branches in $d = 2$ local CFTs. (a) Location of individual minima (colored circles) with $\ell_{\text{max}} \leq 8$ in the $(\Delta_\sigma, \Delta_\epsilon)$ -plane. The minima belong to different sectors all of them contained in 4_4_3_2_1. The dashed orange line corresponds to the bounds obtained in [24] assuming only one \mathbf{Z}_2 -even relevant scalar. (b) Value of the action as a function of Δ^* for each branch. Color codes are the same in the two panels. Namely, each minimum in (a) corresponds to a related small circle in (b) of the same color. The purple Ising ref. branch is the one obtained by our protocol when starting from the exact scaling dimensions for the 2d Ising model.

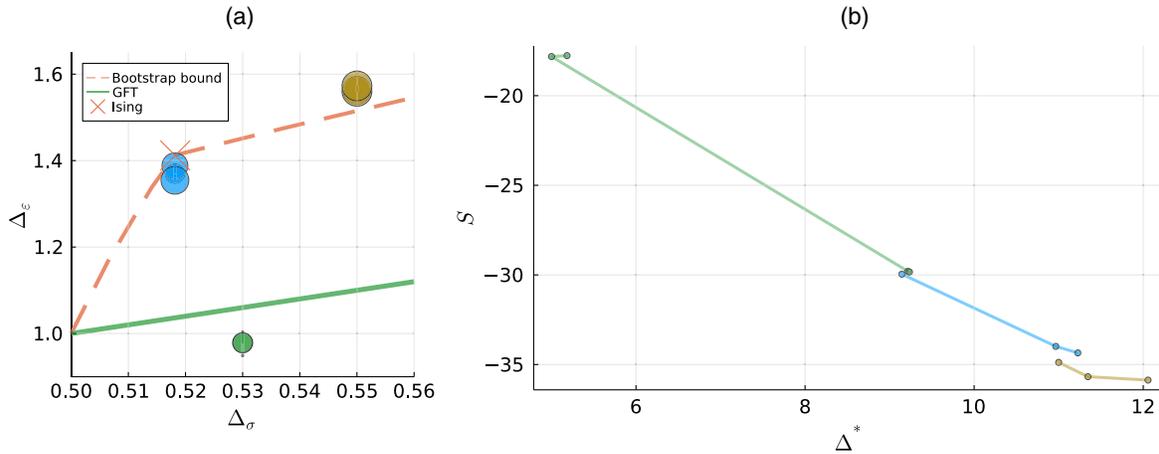


FIG. 20. Representative branches in $d = 3$, $\Delta_T = 3$. (a) Location of individual minima (colored circles) with $\ell_{\max} \leq 10$ in the $(\Delta_\sigma, \Delta_c)$ -plane. The minima belong to different sectors all of them contained in 4_4_4_3_2_1. Several minima overlap and appear as a single one. The dashed orange line is the bound reported in [31]. (b) Value of the action as a function of Δ^* for each branch. Color codes are the same in the two panels. Each minimum in (a) corresponds to a related small circle in (b) of the same color.

We report in the left panel of Fig. 19 the individual minima found in $d = 2$ for some selected branches at $\Delta_\sigma = 1/8$ (Ising) and at $\Delta_\sigma = 1/5$ (GFT-like). Note the difference with respect to Fig. 5. There we reported only the end-minima associated to the end-points of *all* the branches. Here, we pick up a given branch and show all the minima contained in it. Since most of these minima overlap in Fig. 19 and similarly several branches overlap in Fig. 5 (detectable from the darker color of the circles), the difference between the two figures can be overlooked at first glance. In the right panel of Fig. 19 we report the value of the action S for each of these minima making the branches manifest. We also note in passing that the spectrum at the end-minimum for each value of Δ_σ is shown in Fig. 8 in orange (the green spectrum corresponds to a similar spectrum not belonging to the same branch).

In order to have an estimate of how the action is supposed to decrease with Δ^* , we also report the branch obtained by starting low T -MCs from several truncations of the exact Ising scaling dimensions (purple squares in the right panel of Fig. 19). This implies that there are several minima in the vicinity of the exact Ising values and that our

protocol is able to find at least a subset of them. It is also reassuring to see that the decrease in S of the branches found and of the Ising benchmark branch are fully compatible. The good behavior of the GFT-like branches is the main reason why we believe such theories are not numerical artifacts.

As further example we show in the left panel of Fig. 20 the individual minima for three selected end-minima in $d = 3$. The blue one corresponds to the Ising model, the green one to the end-minimum below the GFT green line and the olive one to one of the end-minima at $\Delta_\sigma = 0.55$. In the right panel we show the value of the action S for each of these minima (the color code indicates the different branches).

We point out in passing that a consistent branch should not necessarily be monotonically decreasing as Δ^* increases, because it can happen that minima with less operators have a slightly higher Δ^* than a related minima with a denser spectrum. For instance, the green branch in Fig. 20(b) would seem to “turn around” at the beginning, but the important fact is that we make sure that minima with more operators have strictly smaller S .

- [1] R. Rattazzi, V.S. Rychkov, E. Tonni, and A. Vichi, Bounding scalar operator dimensions in 4D CFT, *J. High Energy Phys.* **12** (2008) 031.
- [2] D. Poland, S. Rychkov, and A. Vichi, The conformal bootstrap: Theory, numerical techniques, and applications, *Rev. Mod. Phys.* **91**, 015002 (2019).
- [3] D. Simmons-Duffin, A semidefinite program solver for the conformal bootstrap, *J. High Energy Phys.* **06** (2015) 174.

- [4] W. Landry and D. Simmons-Duffin, Scaling the semi-definite program solver SDPB, [arXiv:1909.09745](https://arxiv.org/abs/1909.09745).
- [5] F. Kos, D. Poland, and D. Simmons-Duffin, Bootstrapping mixed correlators in the 3D Ising model, *J. High Energy Phys.* **11** (2014) 109.
- [6] F. Kos, D. Poland, D. Simmons-Duffin, and A. Vichi, Bootstrapping the $O(N)$ archipelago, *J. High Energy Phys.* **11** (2015) 106.

- [7] F. Kos, D. Poland, D. Simmons-Duffin, and A. Vichi, Precision islands in the Ising and $O(N)$ models, *J. High Energy Phys.* **08** (2016) 036.
- [8] D. Poland and D. Simmons-Duffin, Bounds on 4D conformal and superconformal field theories, *J. High Energy Phys.* **05** (2011) 017.
- [9] S. El-Showk and M. F. Paulos, Bootstrapping Conformal Field Theories with the Extremal Functional Method, *Phys. Rev. Lett.* **111**, 241601 (2013).
- [10] S. El-Showk and M. F. Paulos, Extremal bootstrapping: Go with the flow, *J. High Energy Phys.* **03** (2018) 148.
- [11] D. Simmons-Duffin, The lightcone bootstrap and the spectrum of the 3D Ising CFT, *J. High Energy Phys.* **03** (2017) 086.
- [12] 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).
- [13] N. Afkhami-Jeddi, T. Hartman, and A. Tajdini, Fast conformal bootstrap and constraints on 3D gravity, *J. High Energy Phys.* **05** (2019) 087.
- [14] G. Kántor, V. Niarchos, and C. Papageorgakis, Solving Conformal Field Theories with Artificial Intelligence, *Phys. Rev. Lett.* **128**, 041601 (2022).
- [15] G. Kántor, V. Niarchos, and C. Papageorgakis, Conformal bootstrap with reinforcement learning, *Phys. Rev. D* **105**, 025018 (2022).
- [16] P. Liendo, L. Rastelli, and B. C. van Rees, The bootstrap program for boundary CFT_d, *J. High Energy Phys.* **07** (2013) 113.
- [17] F. A. Dolan and H. Osborn, Conformal four point functions and the operator product expansion, *Nucl. Phys.* **B599**, 459 (2001).
- [18] F. A. Dolan and H. Osborn, Conformal partial waves and the operator product expansion, *Nucl. Phys.* **B678**, 491 (2004).
- [19] A. Castedo Echeverri, B. von Harling, and M. Serone, The effective bootstrap, *J. High Energy Phys.* **09** (2016) 097.
- [20] B. P. Flannery, Saul A. Teukolsky, William H. Press, and William T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, England, 2007).
- [21] A. Rodriguez and A. Laio, Clustering by fast search and find of density peaks, *Science* **344**, 1492 (2014).
- [22] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, Solving the 3D Ising model with the conformal bootstrap, *Phys. Rev. D* **86**, 025022 (2012).
- [23] S. Caron-Huot, Analyticity in spin in conformal theories, *J. High Energy Phys.* **09** (2017) 078.
- [24] C. Behan, Unitary subsector of generalized minimal models, *Phys. Rev. D* **97**, 094020 (2018).
- [25] P. Di Francesco, P. Mathieu, and D. Sénéchal, *Conformal Field Theory*, Graduate Texts in Contemporary Physics (Springer, New York, 1997).
- [26] P. H. Ginsparg, Applied conformal field theory, in *Les Houches Summer School in Theoretical Physics: Fields, Strings, Critical Phenomena* (North-Holland, Amsterdam, 1988), [arXiv:hep-th/9108028](https://arxiv.org/abs/hep-th/9108028).
- [27] V. S. Rychkov and A. Vichi, Universal constraints on conformal operator dimensions, *Phys. Rev. D* **80**, 045006 (2009).
- [28] A. B. Zamolodchikov, Two-dimensional conformal symmetry and critical four-spin correlation functions in the Ashkin-Teller model, *Sov. J. Exp. Theor. Phys.* **63**, 1061 (1986).
- [29] L. J. Dixon, D. Friedan, E. J. Martinec, and S. H. Shenker, The conformal field theory of orbifolds, *Nucl. Phys.* **B282**, 13 (1987).
- [30] R. Rattazzi, S. Rychkov, and A. Vichi, Bounds in 4D conformal field theories with global symmetry, *J. Phys. A* **44**, 035402 (2011).
- [31] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, Solving the 3D Ising model with the conformal bootstrap II. c-minimization and precise critical exponents, *J. Stat. Phys.* **157**, 869 (2014).
- [32] D. Poland, D. Simmons-Duffin, and A. Vichi, Carving out the space of 4D CFTs, *J. High Energy Phys.* **05** (2012) 110.
- [33] F. Gliozzi, More Constraining Conformal Bootstrap, *Phys. Rev. Lett.* **111**, 161602 (2013).
- [34] F. Gliozzi and A. Rago, Critical exponents of the 3D Ising and related models from conformal bootstrap, *J. High Energy Phys.* **10** (2014) 042.
- [35] I. Esterlis, A. L. Fitzpatrick, and D. Ramirez, Closure of the operator product expansion in the non-unitary bootstrap, *J. High Energy Phys.* **11** (2016) 030.
- [36] W. Li, New method for the conformal bootstrap with OPE truncations, [arXiv:1711.09075](https://arxiv.org/abs/1711.09075).
- [37] K. Levenberg, A method for the solution of certain nonlinear problems in least squares, *Q. Appl. Math.* **2**, 164 (1944).
- [38] D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *J. Soc. Ind. Appl. Math.* **11**, 431 (1963).
- [39] M. K. Transtrum, B. B. Machta, and J. P. Sethna, The geometry of nonlinear least squares with applications to sloppy models and optimization, *Phys. Rev. E* **83**, 036701 (2011).
- [40] F. Kos, D. Poland, and D. Simmons-Duffin, Bootstrapping the $O(N)$ vector models, *J. High Energy Phys.* **06** (2014) 091.