Towards Solving CFTs with Reinforcement Learning

Costis Papageorgakis, QMUL

String Data 2021, 15th of December 2021

Based on [arxiv:2108.08859, arXiv:2108.09330] with Gergely Kántor and Vasilis Niarchos

Motivation

- CFTs are ubiquitous: UV/IR behaviour of QFTs, phase transitions, quantum gravity via AdS/CFT…
- But their nonperturbative solution is hard
- (Modern) Conformal Bootstrap: make assumptions about CFT spectrum and check consistency with crossing equations
  [Rattazzi, Rychkov, Tonni, Vichi ’08]
- (Numerical Modern) Conformal Bootstrap: Linear/semi-definite programming methods give sharp bounds in parameter space

Here: Introduce a complementary numerical approach making use of Reinforcement Learning techniques

Theoretical framework

Consider conformal primary scalar operators $\mathcal{O}_i$ of scaling dimension Δ_i in a CFT

OPE: $$\mathcal{O}_i(x_1)\,\mathcal{O}_j(x_2) = \sum_k C_{ij}^k \hat{f}_{ij}^k\left(x_1, x_2, \partial_{x_2}\right)\mathcal{O}_k(x_2)$$

All CFT data encoded in conformal scaling dimensions Δ_i and OPE coefficients C_{ij}^k

Associativity of OPE for $\left\langle \mathcal{O}_{i_1}(x_1)\,\mathcal{O}_{i_2}(x_2)\,\mathcal{O}_{i_3}(x_3)\,\mathcal{O}_{i_4}(x_4)\right\rangle$ leads to:

crossing equations $\Rightarrow$ constraints on CFT data

Concrete examples: 2D CFT

The crossing equation is:

$$\sum_{h,\bar{h}} {}_s\mathfrak{C}_{h,\bar{h}} g^{(1234)}_{h,\bar{h}}(z,\bar{z}) - (-1)^{\left(h_{41}+\bar{h}_{41}\right)} \frac{z^{h_1+h_2}}{(z-1)^{h_2+h_3}} \frac{\bar{z}^{\bar{h}_1+\bar{h}_2}}{(\bar{z}-1)^{\bar{h}_2+\bar{h}_3}} \sum_{h',\bar{h}'} {}_t\mathfrak{C}_{h',\bar{h}'} g^{(3214)}_{h',\bar{h}'}(1-z,1-\bar{z}) = 0$$

- OPE-coefficients squared: ${}_{s,t}\mathfrak{C}_{h,\bar{h}}$
- Positions of $\mathcal{O}_{1,2,3,4}$: $(z,\bar{z})$
- Scaling dimension: $\Delta = h + \bar{h}$ and spin: $\ell = h - \bar{h}$ (fix spin)
- 2D conformal blocks:

$$g^{1234}_{h,\bar{h}}(z,\bar{z}) = z^h \bar{z}^{\bar{h}} \, {}_2F_1\left(h - h_{12}, h + h_{34}; 2h; z\right) \times {}_2F_1\left(\bar{h} - \bar{h}_{12}, \bar{h} + \bar{h}_{34}; 2\bar{h}; \bar{z}\right)$$

$\Rightarrow$ Crossing equations are hard to solve exactly

$\Rightarrow$ Simplify to produce a numerically tractable system

- Truncate sum over intermediate operators in conformal-block expansion up to some Δ_{max} [Gliozzi '13]

- Reduce single (functional) crossing equation $E(z, \bar{z}; \overrightarrow{\Delta}, \overrightarrow{\mathfrak{C}})$ to discrete set of algebraic equations $\overrightarrow{E}(\overrightarrow{\Delta}, \overrightarrow{\mathfrak{C}})$ for specific $(z, \bar{z})$

For fixed external operators, leads to finite number of equations N_z for finite number of unknowns N_{unknown}, $(\overrightarrow{\Delta}, \overrightarrow{\mathfrak{C}})$

Evaluate truncated, reduced crossing equations $\vec{E}(\vec{\Delta}, \vec{\mathfrak{C}}) = 0$ for $N_z > N_{\text{unknown}}$ points

No exact solution but we can look for approximate (numerical) solutions minimising $\vec{E}(\vec{\Delta}, \vec{\mathfrak{C}})$ [Li '17]

Note: Choice of z-sampling motivated by:

- Theory (OPE convergence)
- Numerics (evaluation of ${}_2F_1$ hypergeometrics in Python)

$\Rightarrow$ This is a problem especially well suited to Reinforcement Learning

Attack the problem using “off the shelf” RL algorithms

⇒ Chose to work with popular soft-Actor-Critic algorithm
[Harnooja, Zhou, Abbeel, Levine ‘18]

- Involves a stochastic actor aiming to maximise both reward and entropy
- Can handle continuous state and action spaces
- It is model-free and stable when applied to several scenarios

Reward function: total violation of reduced crossing equations

$$R := -\|\vec{E}\| \qquad \text{or} \qquad R := 1/\|\vec{E}\|$$

Calibration with 2D Ising Model

- $\langle \sigma\sigma\sigma\sigma \rangle$ correlator for $N_z = 29$

$\Delta_\sigma = \frac{1}{8}$				
spin	analytic Δ	RL Δ	analytic $\mathfrak{C}$	RL $\mathfrak{C}$
0	4	3.9331603	2.44141×10^{-4}	3.657538×10^{-4}
0	1	0.9881525	0.25	0.25254947
2	2	1.9802496	0.015625	0.015717817
4	4	3.9497	2.19727×10^{-4}	2.4715587×10^{-4}
6	6	5.971367	1.36239×10^{-5}	$0.54007314 \times 10^{-5}$

- Stays close to analytic values for Δ and reasonable predictions for $\mathfrak{C}$

“Discovery” of $c = 1$ compact free boson CFT

Start with minimal assumptions and see what our algorithm can do:

- Consider theory with exactly marginal coupling
- Take spin-less operators V_p with $\Delta_{V_p} := p^2$
- Assume these are charged under $U(1)$ symmetry $(V_p, \overline{V}_p)$
- Consider correlator $\langle V_p V_p \overline{V}_p \overline{V}_p \rangle$ and associated crossing equations

- s-channel: OPE of $V_p V_p$ and $\overline{V}_p \overline{V}_p$
- t-channel: OPE of $\overline{V}_p V_p$ and $V_p \overline{V}_p$
- Input that truncated crossing equations involve spin-partition:

Spin	0	1	2	3	4	5
s-channel	2	1	1	1	1	1
t-channel	2	3	2	2	1	1

Run RL algorithm with $\Delta_{max} = 5.5$ and $N_z = 49$ to get for $\Delta_{V_p} = 0.1$

Marginal deformation

U(1) current

EM tensor

Channel	spin	ML Δ	ML $\mathfrak{C}$
s	0	0.40006787	1.0057276
	0	4.336432	$0.43016876 \times 10^{-5}$
	1	5.307818	$-2.2633198 \times 10^{-4}$
	2	2.4060674	5.486169×10^{-3}
	3	3.446559	$-0.4480493 \times 10^{-5}$
	4	4.410344	$0.27796367 \times 10^{-3}$
	5	5.3354797	-9.976282×10^{-5}
t	0	2.001293	0.0056684865
	0	4.0166564	4.8836926×10^{-4}
	1	1.040068	-0.085237
	1	3.0494268	-2.271628×10^{-2}
	1	4.9848695	-9.268466×10^{-4}
	2	2.00707	0.0018059064
	2	4.045016	7.282457×10^{-4}
	3	3.0331514	-2.894943×10^{-4}
	3	4.9544168	$-3.3044117 \times 10^{-3}$
	4	3.9395354	6.668457×10^{-4}
	5	5.0390368	$-4.3607014 \times 10^{-4}$

$\Rightarrow c = 1$ compact free boson

- $V_p(z, \bar{z}) = e^{ip(X(z)+\bar{X}(\bar{z}))}$ are vertex operators with unit momentum and no winding
- $R = 1/p$

Channel	spin	analytic Δ	ML Δ	analytic $\mathfrak{C}$	ML $\mathfrak{C}$
s	0	0.4	0.40006787	1	1.0057276
	0	4.4	4.336432	1.27551×10^{-5}	$0.43016876 \times 10^{-5}$
	1	5.4	5.307818	0	$-2.2633198 \times 10^{-4}$
	2	2.4	2.4060674	3.57143×10^{-3}	5.486169×10^{-3}
	3	3.4	3.446559	0	$-0.4480493 \times 10^{-5}$
	4	4.4	4.410344	1.96039×10^{-3}	$0.27796367 \times 10^{-3}$
	5	5.4	5.3354797	0	-9.976282×10^{-5}
t	0	2	2.001293	0.01	0.0056684865
	0	4	4.0166564	2.5×10^{-5}	4.8836926×10^{-4}
	1	1	1.040068	-0.1	-0.085237
	1	3	3.0494268	-5×10^{-4}	-2.271628×10^{-2}
	1	5	4.9848695	-8.33333×10^{-7}	-9.268466×10^{-4}
	2	2	2.00707	0.005	0.0018059064
	2	4	4.045016	1.66667×10^{-5}	7.282457×10^{-4}
	3	3	3.0331514	-1.66667×10^{-4}	-2.894943×10^{-4}
	3	5	4.9544168	-4.16667×10^{-7}	$-3.3044117 \times 10^{-3}$
	4	4	3.9395354	4.16667×10^{-6}	6.668457×10^{-4}
	5	5	5.0390368	-8.33333×10^{-8}	$-4.3607014 \times 10^{-4}$

⇒ RL algorithm has succeeded:

The numbers improve as we include operators of increasing Δ

⇒ But keep in mind sources of error for $\mathfrak{C}$:

- Analytic errors from truncation of conformal-block expansion
- Choice of z-sampling in reducing the crossing-symmetry constraints
- Statistical errors
- Choice of reward function R

Conclusions

- Introduced a new RL approach to conformal bootstrap
- Our code is fully automated, using as input a spin-partition and predicts CFT data $(\vec{\Delta}, \vec{\mathfrak{C}})$ for specific theories
- Calibrated on 2D Minimal Models
- Used to "detect" the $c = 1$ free compact boson CFT. Excellent results for Δ. Results for $\mathfrak{C}$ can be improved for precision

Outlook

- Improve efficiency of RL algorithm (implementation, understanding errors)
- Parallelise computations to account for statistical errors
- Scale up searches to $O(100)$ unknowns
- Make systematic use of theory constraints (unitarity, global symmetries, supersymmetry) and multiple correlators to probe higher-dimensional CFTs
- Goal: Algorithm determines full theory by gradually adding in operators