

RECEIVED: April 16, 2024

REVISED: June 5, 2024

ACCEPTED: June 12, 2024

PUBLISHED: July 10, 2024

Gravitational duals from equations of state

Yago Bea ^{a,b} **Raul Jimenez** ^{b,c} **David Mateos** ^{a,b,c} **Shuheng Liu** ^d
Pavlos Protopapas ^d **Pedro Tarancón-Álvarez** ^{a,b} and **Pablo Tejerina-Pérez** ^{a,b}

^a*Departament de Física Quàntica i Astrofísica, Universitat de Barcelona,
Martí i Franquès 1, ES-08028, Barcelona, Spain*

^b*Institut de Ciències del Cosmos (ICC), Universitat de Barcelona,
Martí i Franquès 1, ES-08028, Barcelona, Spain*

^c*Institució Catalana de Recerca i Estudis Avançats (ICREA),
Passeig Lluís Companys 23, ES-08010, Barcelona, Spain*

^d*Institute for Applied Computational Science, Harvard University,
Cambridge, MA, U.S.A.*

E-mail: yago.bea.b@gmail.com, raul.jimenez@icc.ub.edu, dmateos@fqa.ub.edu,
shuheng_liu@g.harvard.edu, pprotopapas@g.harvard.edu,
pedro.tarancon@fqa.ub.edu, pablo.tejerina@icc.ub.edu

ABSTRACT: Holography relates gravitational theories in five dimensions to four-dimensional quantum field theories in flat space. Under this map, the equation of state of the field theory is encoded in the black hole solutions of the gravitational theory. Solving the five-dimensional Einstein's equations to determine the equation of state is an algorithmic, direct problem. Determining the gravitational theory that gives rise to a prescribed equation of state is a much more challenging, inverse problem. We present a novel approach to solve this problem based on physics-informed neural networks. The resulting algorithm is not only data-driven but also informed by the physics of the Einstein's equations. We successfully apply it to theories with crossovers, first- and second-order phase transitions.

KEYWORDS: AdS-CFT Correspondence, Gauge-Gravity Correspondence

ARXIV EPRINT: [2403.14763](https://arxiv.org/abs/2403.14763)

Contents

1	Introduction	1
2	Holography and the direct problem	3
3	Inverse problem	5
4	Methodology	7
4.1	Background	8
4.2	Solution bundles	9
4.3	Solving the Einstein-Klein-Gordon equations	10
4.4	Technical aspects of setup and training	13
4.5	Numerical tests	16
4.6	Computational considerations	18
5	Results	18
5.1	Loss function	18
5.2	Recovered potential and equation of state	19
5.3	Solutions to the ODEs	24
6	Discussion	24
A	High-temperature behaviour	33

1 Introduction

Holography [1–3] provides a valuable tool to study strongly coupled gauge theories with a large number of colors. In this limit, the duality maps the quantum properties of the gauge theory to the classical properties of a gravitational theory in one dimension higher. This geometrization has profound consequences. A celebrated one is the fact that the equilibrium, thermodynamic properties of the gauge theory are encoded in properties of black hole horizons in the dual gravitational solutions. Not surprisingly, near-equilibrium aspects of the gauge theory, such as transport coefficients, are then encoded in small perturbations of the corresponding black hole horizons. Yet, the most dramatic consequence of the holographic map is that even the arbitrarily-far-from-equilibrium dynamics of the gauge theory can be determined by evolving in time the dual Einstein’s equations.

The equilibrium regime can be studied from first principles with conventional methods such as the lattice formulation of gauge theories. In contrast, in the far-from-equilibrium regime holography is often the only tool with which systematic, first-principle calculations are possible. A fruitful strategy is therefore to construct a holographic model that reproduces the equilibrium properties of a gauge theory of interest, and then use the holographic side to study the far-from-equilibrium dynamics. From this perspective, one can think of

holography as model building of a unique type, namely, one with the right to access the far-from-equilibrium regime.

As we will see, implementing this strategy requires solving a challenging inverse problem. Early (semi)analytical attempts include [4–6] (see [7] for a review). In this paper we will solve the problem using Physics-Informed Neural Networks (PINNs). With the exception of [8–10], previous applications of NNs to holography [11–31] have focused on the reconstruction of a specific *solution* of the dual gravitational theory. In contrast, our algorithm reconstructs the gravitational *theory* itself. Although we will illustrate the method in a simple setup, we expect it to be generally applicable to inverse problems involving highly non-linear partial differential equations. We envisage many potential applications that will be explored elsewhere.

To illustrate our method, we will consider a four-dimensional Conformal Field Theory (CFT) deformed by a relevant operator $\mathcal{O}$. For concreteness, in this paper we will assume that $\mathcal{O}$ has conformal dimension three, but this is easily generalized. The holographic dual consists of five-dimensional gravity coupled to a scalar field ϕ with an appropriate potential $V(\phi)$ [32, 33]. The existence of an ultraviolet (UV) fixed point in the gauge theory translates into anti-de Sitter space (AdS) boundary conditions on the gravity side. Despite their simplicity, this type of models capture the essence of many fundamental gauge theory properties such as thermal phase transitions [34], confinement [35], etc.

All the properties of the gauge theory, in particular its thermodynamics, are encoded in the scalar potential $V(\phi)$. Given the potential, in order to determine the thermodynamics one needs to solve a *direct* problem. This consists of finding all the static, regular black hole solutions of the Einstein-Klein-Gordon (EKG) equations. This set of solutions can be conveniently parameterized by the value of ϕ at the horizon, ϕ_H . For each solution one computes the area density and the surface gravity of the horizon. From them one extracts the Bekenstein-Hawking entropy density $S(\phi_H)$ and the Hawking temperature $T(\phi_H)$. These quantities are then identified with the entropy density and the temperature of the corresponding thermal state in the gauge theory. The result is an equation of state $S(T)$. Any other thermodynamic property such as the energy density, the pressure, etc. can be obtained from $S(T)$ via thermodynamic identities.

Since the scalar potential enters the EKG equations that must be solved to determine the thermodynamics, it is clear that the function $S(T)$ is a functional of the function $V(\phi)$. The determination of $S(T)$ from $V(\phi)$ is a direct problem because it can be done algorithmically, as we have just described. The *inverse* problem consists of finding a potential $V(\phi)$ that reproduces a given $S(T)$. The fact that the EKG equations are highly non-linear differential equations makes this an extremely challenging problem. We will solve it using PINNs, which are uniquely suited for inverse problems. Our approach differs from previous literature in two respects. First, in contrast to other possible uses of the term “physics-informed”, our loss function is given directly by the residue of the differential equations. As we will see, the role of the equation of state is to provide the necessary physical input to solve these equations, namely the boundary conditions for the EKG equations. Second, we do not adopt any a priori parameterization of the potential in terms of a finite number of variables, as in e.g. [10], but instead reconstruct it as a free-form function by encoding it as a second NN. We will illustrate the applicability of the method by reconstructing gravitational theories that give rise to equations of state with crossovers, first- and second-order phase transitions.

2 Holography and the direct problem

As explained above, we will focus on an EKG theory in five dimensions. The action takes the form

$$S = \frac{2}{\kappa_5^2} \int dx^5 \sqrt{-g} \left[\frac{1}{4} R - \frac{1}{2} (\nabla \phi)^2 - V(\phi) \right], \quad (2.1)$$

where κ_5 is the gravitational coupling and $V(\phi)$ is the scalar potential.

The thermodynamics of the gauge theory is determined by finding all the static, planar black hole solutions of (2.1). Because of their translation invariance along the gauge theory directions, these solutions are usually referred to as “black brane” solutions. In order to find them, we consider an asymptotically AdS spacetime and we choose coordinates t, x, y, z for the Minkowski metric at the boundary. We then write the bulk metric in Eddington-Finkelstein coordinates as

$$ds^2 = -A dt^2 + \Sigma^2 (dx^2 + dy^2 + dz^2) - \frac{2}{u^2} dt du, \quad (2.2)$$

where u is the null holographic coordinate. The AdS boundary is located at $u = 0$. This region corresponds to the UV of the gauge theory. The metric functions A and Σ , as well as the scalar field ϕ , are functions of u . Substituting the ansatz (2.2) into the equations of motion derived from (2.1) we obtain four ordinary differential equations (ODEs)

$$\Sigma'' + \frac{2}{u} \Sigma' + \frac{2}{3} \Sigma \phi'^2 = 0, \quad (2.3a)$$

$$A'' + \frac{8}{3u^4} V(\phi) + A' \left(\frac{2}{u} + \frac{3\Sigma'}{\Sigma} \right) = 0, \quad (2.3b)$$

$$\phi'' - \frac{1}{u^4 A} \frac{\partial V(\phi)}{\partial \phi} + \phi' \left(\frac{2}{u} + \frac{A'}{A} + 3 \frac{\Sigma'}{\Sigma} \right) = 0, \quad (2.3c)$$

$$A' + 2A \frac{\Sigma'}{\Sigma} - \frac{2\Sigma}{3\Sigma'} \left(A \phi'^2 - \frac{2}{u^4} V(\phi) \right) = 0, \quad (2.3d)$$

where a prime indicates differentiation with respect to u . The last equation is first-order in derivatives because it is a constraint associated to our gauge choice for the u -coordinate. As a consequence, the four equations (2.3) are not all independent: two of the second-order equations plus the constraint imply the third second-order equation. This redundancy will be useful for our PINN.

The equations of motion must be solved subject to appropriate boundary conditions. Near $u = 0$, these ensure that the geometry is asymptotically AdS and take the form

$$A(u) = \frac{1}{u^2} + \dots, \quad (2.4a)$$

$$\Sigma(u) = \frac{1}{u} + \dots, \quad (2.4b)$$

$$\phi(u) = \Lambda u + \dots, \quad (2.4c)$$

where the dots indicate subleading terms in the limit $u \rightarrow 0$. In the expressions for the metric functions we have fixed some integration constants by setting to unity the coefficients of the

leading terms. As anticipated above, in the boundary condition for the scalar field we have assumed that the dual scalar operator, $\mathcal{O}$, has conformal dimension

$$\Delta = 3. \quad (2.5)$$

This means that the constant Λ has mass dimension 1.

For a static black brane, the horizon is characterized by a simple zero in the g_{tt} component of the metric, namely in $A(u)$, at some position $u = u_H$. Performing a Taylor expansion around this point we obtain

$$A(u) = A_1(u_H - u) + \dots, \quad (2.6a)$$

$$\Sigma(u) = \Sigma_0 + \Sigma_1(u_H - u) + \dots, \quad (2.6b)$$

$$\phi(u) = \phi_H + \phi_1(u_H - u) + \dots. \quad (2.6c)$$

Reparametrization invariance allows us to set the horizon of the black brane at $u_H = 1$, which we will do hereafter. In these equations we have set to zero some integration constants in order to ensure regularity at the horizon. This choice, together with the choice of integration constants in (2.4), implies that all the coefficients in (2.6) except for ϕ_H are fixed by the requirement that a solution of the form (2.6) at the horizon matches a solution of the form (2.4) at the boundary. It follows that there is a one-to-one map between the value of the scalar field at the horizon and physically different black brane solutions. This is consistent with the expectation that the different equilibrium states of the gauge theory are characterized by a single parameter, namely the temperature T . This and the entropy density can be computed from the surface gravity and the area density of the horizon as

$$\frac{T}{\Lambda} = \frac{1}{4\pi} \frac{A_1}{\Lambda}, \quad (2.7a)$$

$$\frac{S}{\Lambda^3} = \frac{2\pi}{\kappa^2} \left(\frac{\Sigma_0}{\Lambda} \right)^3. \quad (2.7b)$$

In these equations we have assumed (2.5) again. The scale Λ is the only intrinsic scale in the gauge theory, so we will work in units in which

$$\Lambda = 1. \quad (2.8)$$

Similarly, on the gravity side it will be convenient to work in units such that $\kappa^2 = 2$.

We can now state what solving the direct problem means. The only input is some scalar potential $V(\phi)$. Given this function, we must solve the EKG equations (2.3). A standard way to do this is to numerically integrate the equations from the horizon to the boundary. For each value of ϕ_H there is a unique choice of A_1 and Σ_0 such that the solution near the boundary takes the form (2.4). This requirement determines A_1 and Σ_0 in terms of ϕ_H . In figures 10 and 11 we show the functions $A(u)$, $\Sigma(u)$ and $\phi(u)$ for three different black brane solutions for two different gauge theories. Substituting the values of $A_1(\phi_H)$ and $\Sigma_0(\phi_H)$ in (2.7) we obtain $T(\phi_H)$ and $S(\phi_H)$ and, as a consequence, the thermodynamic curve $S(T)$. Note that, while there is always a single state for each value of ϕ_H , in theories with phase transitions there can be more than one state for a given value of T . In other words, while the map from ϕ_H to equilibrium states is single-valued, the map from T can be multi-valued.

3 Inverse problem

It is clear from section 2 that solving the direct problem is algorithmic. Solving the inverse problem is not. In the inverse problem the only input is the thermodynamic curve $S(T)$. Through equations (2.7), each pair of values (T, S) determines a pair (A_1, Σ_0) . In turn, each of these pairs provides boundary conditions at the horizon through equations (2.6). Given this knowledge, the inverse problem consists of finding a unique potential $V(\phi)$ such that the EKG equations with this same potential admit black brane solutions for all these boundary conditions.

Our goal is thus to construct a PINN able to solve this inverse problem. In order to test its accuracy, we will ask the PINN to find the potentials for a family of equations of state $S_{\text{theory}}(T)$ constructed in [36] and labeled by one parameter dubbed ϕ_M . As an illustration, figure 1 shows the equations of state for $\phi_M = 5, 1.08$ and 1 . The reason why we choose these examples is that they cover the three possible cases of a crossover, a second-order or a first-order phase transition.

In the high-temperature limit, the entropy density behaves as

$$S_{\text{theory}}(T) = \pi^4 T^3 + \frac{3\pi^4}{64} T + \dots \quad (3.1)$$

Comparing with eq. (A.5), we see that this corresponds to a CFT deformed by an operator $\mathcal{O}$ of dimension (2.5), in agreement with the boundary condition (2.4c) that we imposed on the dual scalar field ϕ . Solving the EKG equations (2.3) perturbatively near the boundary, it can be shown that, in our conventions, the form of the entropy density (3.1) implies the following conditions on the scalar potential at $\phi = 0$:

$$C_0 \equiv V(0) + 3 = 0, \quad C_1 \equiv V'(0) = 0, \quad C_2 \equiv V''(0) + 3 = 0. \quad (3.2)$$

The fact that these conditions can be imposed at $\phi = 0$ can always be achieved by a shift in ϕ . The first condition fixes the coefficient of the T^3 term in S_{theory} . The last two conditions state that $\phi = 0$ is a maximum of the potential. On the gravity side, this corresponds to the AdS boundary condition on the metric, and on the gauge theory side it corresponds to the presence of an UV fixed point, which in turn fixes the leading T^3 scaling in S_{theory} . Finally, the last condition imposes that the dimension of the gauge theory operator dual to ϕ is precisely (2.5), in agreement with the fact that the leading correction in the entropy density is linear in T . Since all this information is contained in the equation of state, which is assumed to be given, we will provide our PINN with these boundary conditions on the potential. All the results reported in section 5 were obtained under these assumptions. Nevertheless, we have also performed tests in which we do not provide the PINN with the information in (3.2). We will comment on these in section 6.

By feeding the different $S_{\text{theory}}(T)$ to the PINN we obtain a family of reconstructed potentials $V_{\text{PINN}}(\phi)$. We emphasize that, other than the boundary conditions (2.4), (2.6) and (3.2), we do not supply the PINN with *any* information about the functional form of the potential or of the metric functions. We then solve the direct problem with each of these $V_{\text{PINN}}(\phi)$ and obtain a set of thermodynamic curves $S_{\text{PINN}}(T)$. In section 5 we will compare $S_{\text{PINN}}(T)$ with $S_{\text{theory}}(T)$. Needless to say, the success of the method requires that these two curves be as close as possible.




Figure 1. Equation of state for a gauge theory with a crossover ($\phi_M = 5$), a second-order phase transition ($\phi_M = 1.08$), and a first-order phase transition ($\phi_M = 1$) [36]. In the latter two cases, the critical temperatures are indicated by solid vertical lines. The entropy density is shown in units of the intrinsic scale in the theory (top) and in units of the temperature (bottom).

The family of equations of state $S_{\text{theory}}(T)$ were constructed in [36] by solving the direct problem for the family of potentials given by

$$V_{\text{theory}}(\phi) = -\frac{4}{3}\mathcal{W}(\phi)^2 + \frac{1}{2}\mathcal{W}'(\phi)^2, \quad (3.3)$$

where

$$\mathcal{W}(\phi) = -\frac{3}{2} - \frac{\phi^2}{2} - \frac{\phi^4}{4\phi_M^2} + \frac{\phi^6}{10}. \quad (3.4)$$




Figure 2. Potentials that give rise to the equations of state in figure 1 [36]. All potentials share the same maximum at $\phi = 0$, but they have minima at different positions indicated by the solid vertical lines.

Some of these are plotted in figure 2. The maxima of these potentials corresponds to the UV fixed point in the gauge theory. These potentials also exhibit a minimum at some $\phi > 0$, corresponding to the presence of an IR fixed point in the dual gauge theory. We will see that the PINN is able to discover this feature.

Knowledge of the true potentials means that, in section 5, we will be able to perform an additional, intermediate test: we will be able to compare the potential reconstructed by the PINN with the true potential. This will illustrate a crucial property of this inverse problem, namely the fact that the potential must be reconstructed with very high precision in order for the method to be successful.

4 Methodology

In this section we describe our procedure to solve the EKG equations and reconstruct the potential non-parametrically, i.e. as a free-form function, using our implementation of PINNs to solve ODEs. First, we review some general aspects about the use of NNs to solve differential equations.

4.1 Background

A neural network (NN) is a machine-learning algorithm that produces a structure of interconnected nodes, or neurons, disposed in a layered structure, which can be trained to solve a certain task through minimization of a loss function. In a feedforward NN, which is the type of architecture that we consider in this work, a neuron takes an input and affine-transforms it, passes the resulting transformed data through an activation function, and produces an output. Each neuron in a layer is connected to all neurons in the following layer. Every connection carries a weight w_i that multiplies the output of the neuron of the previous layer passing via some non-linear function to the i -th neuron in the next layer. These weights are the main parameters of the NN followed by the activation function. The final output of the NN will depend on the entire set of all weights, which we denote by W_{net} . A NN can be trained so that its output gives an approximate solution to a certain problem. How well the NN's output resembles the exact solution is encoded in the loss function. By changing the parameters W_{net} in a way that minimizes this loss function,¹ the NN will, in principle, approximate increasingly well the solution to the problem it is being trained on. It is worth noting that, according to the universal approximation theorem, NNs with a single hidden layer containing a finite number of neurons can approximate any continuous function to arbitrary accuracy. In the case of PINNs, the NN trains on the differential equations it is trying to solve. As a consequence, once trained, the PINN itself becomes the approximate solution.

Our implementation to solve the inverse problem is relatively new and original for this kind of problems, but it exploits previously known PINNs. These are deep neural networks (DNNs) that can approximate the solution of differential equations (DEs) and initial and boundary values. The early results of Dissanayake and Phan-Thien [37] and then by Lagaris [38], solving DEs using NNs, were followed by a rapid growth of publications on PINNs in recent years. PINNs have many advantages compared to traditional numerical solvers. They provide closed-form solutions, eliminating the need for iterative solvers and reducing computational overhead. PINNs are mesh-free, allowing for on-demand solution computation post-training, which enhances efficiency in solving complex problems. Their ability to leverage transfer learning enables rapid discovery of new solutions by transferring knowledge from related tasks, enhancing adaptability and accelerating convergence. PINNs are invertible, which makes them suitable for inverse problems where obtaining the original input from the output is essential. Additionally, PINNs can be parametrized to include parameters of the differential equation as input.

Many dynamical systems are described by ODEs which relate the rates and values of state variables and external driving functions. While some simple ODEs have closed-form solutions, the vast majority must be solved approximately using discretization of the domain or via spectral methods [39]. The former approximating methods are more general, with Runge-Kutta and multi-step methods as typical examples. These methods seek to numerically integrate the ODEs starting from a value of the independent variable at a boundary and stepping away. These conventional methods are generally efficient in determining the state of a system for a sequence of values. However, if we are only interested in the state at a

¹That is, flowing towards (or against) the steepest gradient of the loss function represented in a high dimensional parameter space.

specific later value, substantial computational effort must still be spent determining all the states at steps leading up to the state of interest. This ordering also limits parallelizability of the conventional single- and multi-step methods because, until the preceding state is known, processors tasked with finding a segment of the system's evolution over a later interval cannot start calculating the correct piece of the trajectory.

The idea of approximating DE solutions with NN was first developed by Dissanayake and Phan-Thien, where training was carried out by minimizing a loss based on the network's satisfaction with the boundary conditions and DEs themselves [37]. Lagaris et al. [40] showed that the form of the network could be chosen to satisfy the boundary conditions by construction, and that automatic differentiation could be used to compute the derivatives appearing in the loss function [38]. After training, the approximate value of the solution at any point within the training range can be computed without having to compute the previous states first. This method has been extended to systems with irregular boundaries [41, 42], applied to solving partial differential equations (PDE) arising in fluid mechanics [43] and Hamiltonian systems [44], and software packages have been developed to facilitate its application [45–47]. In the Lagaris approach, the NN learns a single solution to the ODE. For different sets of initial conditions or for different sets of parameters in the differential equation, the network has to be re-trained on the new task.

One of us (PP) has developed an extension of the Lagaris method [38] where the neural network is instead taught a variety of solutions to a parameterized differential equation. This increases the reusability of the trained network and can speed up tasks that require knowing many solutions, such as for Bayesian parameter inference, propagating uncertainty distributions in dynamical systems, or inverse problems. It is straightforward to extend this approach to problems containing various types of boundary conditions, to PDEs and higher derivatives. Here, we focus on boundary-condition problems in a system of first-order ODEs. We will show that our method has promise when applied to a variety of tasks requiring quick, parallel evaluation of multiple solutions to an ODE, and where it is useful to be able to differentiate the state at a particular value of the evolution variable with respect to the initial conditions or ODE parameters.

4.2 Solution bundles

When working with ODEs it is common to require multiple solutions corresponding to different initial and boundary values. In dynamical systems, each of these solutions represents a trajectory, tracing out an alternate time evolution of the state of the system. Similar situations apply when we have multiple boundary values. In addition, when an ODE depends on a parameter, for example, a physical constant whose value has an associated uncertainty, it can be useful to have different solutions for different values of the parameter.

In our case, we have a range of boundary conditions: the range of values of $S(T)$. Finding the free-form function $V(\phi)$ requires information from all points on the curve, that is, having information about *all* boundary conditions at the same time. Thus, bundle solving is an essential ingredient to solve our problem.

4.3 Solving the Einstein-Klein-Gordon equations

4.3.1 ODE setup

We now proceed to implement the EKG equations (2.3) in the NN. For this purpose, it is convenient to redefine the functions A , Σ , ϕ in such a way that the solutions are finite in the entire computational domain $u \in [0, 1]$. In view of the asymptotic behavior (2.4), this can be achieved through the following redefinitions:

$$\tilde{\Sigma} = u \Sigma, \quad \tilde{A} = u^2 A. \quad (4.1)$$

It is also convenient to transform the EKG equations into first-order ODEs by thinking of the first-derivatives of $\tilde{A}$, $\tilde{\Sigma}$ and ϕ as independent variables. Together with the constraint (2.3d), this results in 6 functions $(\nu_\Sigma, \nu_A, \nu_\phi, \tilde{\Sigma}, \tilde{A}, \phi)$ subject to the following seven, coupled, first-order ODEs

$$E_\alpha = 0, \quad \alpha = 1, \dots, 7, \quad (4.2)$$

where

$$E_1 = \nu_\Sigma - \tilde{\Sigma}', \quad (4.3a)$$

$$E_2 = \nu_A - \tilde{A}', \quad (4.3b)$$

$$E_3 = \nu_\phi - \phi', \quad (4.3c)$$

$$E_4 = \nu'_\Sigma + \frac{2}{3} \tilde{\Sigma} \nu_\phi^2, \quad (4.3d)$$

$$E_5 = u^2 \tilde{\Sigma} \nu'_A + \frac{8}{3} V(\phi) \tilde{\Sigma} + \nu_A \left(3u^2 \nu_\Sigma - 5u \tilde{\Sigma} \right) + \tilde{A} \left(8\tilde{\Sigma} - 6u \nu_\Sigma \right), \quad (4.3e)$$

$$E_6 = u^2 \tilde{\Sigma} \tilde{A} \nu'_\phi - \tilde{\Sigma} \frac{dV}{d\phi} + \nu_\phi \left(-3u \tilde{A} \tilde{\Sigma} + u^2 \tilde{\Sigma} \nu_A + 3u^2 \nu_\Sigma \tilde{A} \right), \quad (4.3f)$$

$$E_7 = \left(u \nu_\Sigma - \tilde{\Sigma} \right) \left(u^2 \tilde{\Sigma} \nu_A + 2u^2 \tilde{A} \nu_\Sigma - 4u \tilde{A} \tilde{\Sigma} \right) - \frac{2}{3} u \tilde{\Sigma}^2 \left(u^2 \tilde{A} \nu_\phi^2 - 2V(\phi) \right). \quad (4.3g)$$

In terms of these variables, the boundary conditions discussed in section 2 translate into

$$\tilde{A}|_{u=0} = 1, \quad (4.4a)$$

$$\tilde{\Sigma}|_{u=0} = 1, \quad (4.4b)$$

$$\phi|_{u=0} = 0, \quad (4.4c)$$

$$\nu_\phi|_{u=0} = 1, \quad (4.4d)$$

$$\tilde{A}|_{u=1} = 0, \quad (4.4e)$$

$$\nu_A|_{u=1} = -4\pi T, \quad (4.4f)$$

$$\tilde{\Sigma}_{u=1} = (S/\pi)^{1/3}. \quad (4.4g)$$

The conditions (4.4a)–(4.4c) correspond to imposing that the solution approaches AdS at $u = 0$, in agreement with (2.4). Condition (4.4d) corresponds to imposing the fact that the dual operator has dimension (2.5) and the convention (2.8) for the scale of the boundary theory. The condition (4.4e) corresponds to imposing that the horizon is at $u = 1$, consistently with (2.6a). Finally, (4.4f) and (4.4g) correspond to conditions (2.6b), (2.6c) and (2.7).

In the direct problem, we choose a value of the scalar field, we obtain the numerical solution of the black brane, and we read off the entropy and the temperature. In contrast, in the inverse problem we start from the equation of state $S(T)$ and use it as a boundary condition: S and T in (4.4f)–(4.4g) are related by the equation of state. In the inverse problem we do not impose a condition for the scalar field at the horizon. The reason is that this information is not directly contained in the equation of state but can only be determined in combination with the potential itself.

4.3.2 NN setup and procedure

We use a dual NN strategy to solve the problem — see figure 3. Although this architecture is computationally more demanding than having only a single NN, it is better suited for our problem for several reasons. The first and most important one is that we want the reconstructed potential to have information about the entire set of boundary conditions. Using two separate NNs, we allow NN-V to find the function that better fits all the given boundary conditions along the $S(T)$ curve. Another interesting feature of using a dual network architecture is that the response of the potential NN can be much more non-linear. This is an important property when the reconstructed function depends on the boundary conditions in a highly non-linear way. Our method is organized into focused steps, making it manageable and easier to interpret.

1. The neural network NN-Solver with parameters W_D takes as an input a discrete set of pairs (T_i, S_i) , with T_i and S_i related by the equation of state, along with a discrete set of values u_n for the u -coordinate, with $u_n \in [0, 1]$. This network then produces estimates for the dependent variables, namely for the functions $\tilde{A}$, $\tilde{\Sigma}$, ϕ , ν_A , ν_Σ and ν_ϕ . By construction, these functions obey the boundary conditions (4.4) exactly. This is accomplished as described in [47]. Each point in the equation of state, namely each pair (T_i, S_i) , corresponds to an entire geometry on the gravitational side described by the functions above. Thus, these functions should be considered to depend on one coordinate, u , and one parameter, the pair (T, S) . Therefore, we will write expressions such as $\phi(u, (T, S))$, and similarly for the other functions. In some cases, specifying T and S separately is redundant, but not in others. Therefore, we will keep track of both T and S . The derivatives with respect to the u -coordinate, ν'_A , ν'_Σ , and ν'_ϕ , are computed analytically using automatic differentiation.
2. The set of predicted values of the scalar field, $\phi_{i,n} = \phi(u_n, (T_i, S_i))$, are then fed into a second neural network, NN-V, with parameters W_V . This network predicts $V(\phi_{i,n})$ based on these values. Automatic differentiation is also used to find the derivative of $V(\phi)$ with respect to ϕ .
3. These predicted values are substituted back into the original set of ODEs (4.3). We then compute and minimize a loss function L over all the different weights. This minimization is achieved by fine-tuning the parameters W_D and W_V through stochastic gradient descent. The loss function is defined as

$$L \equiv \sum_{\alpha} \sum_n \sum_i E_{\alpha}(u_n, (T_i, S_i))^2 + \lambda(C_0^2 + C_1^2 + C_2^2). \quad (4.5)$$




Figure 3. Structure of the implemented NN-system of two NNs with Gaussian localization. The red and blue arrows represent the backpropagation process, in which the parameters W_j ($j = D, V$) are modified such that the NN’s state in parameter space flows in the direction of the gradient of the loss function, $\nabla_{W_j} L$.

The first term is the sum of the residues of all equations of motion (4.2) (sum over α), evaluated on all the values of the u -coordinate (sum over n), and for all the different solutions (sum over i). The second term is the sum of the three boundary conditions (3.2) on the potential. The larger the coefficient of this term, the more efficiently these conditions are enforced by the NN. We choose an initial value of $\lambda = 0$ for the first million epochs. After that, we set $\lambda = 50$ to enforce the boundary conditions (3.2), and then train the model for an additional $(0.5\text{--}1) \cdot 10^6$ epochs.

We have chosen the **Adam** optimizer² for the optimization process, which uses stochastic gradient descent with information about higher-order momenta (see [48] for details).

4. The entire process is iterated for a fixed number of training cycles, known as epochs. During each epoch, we evaluate the error (loss) in our predicted field values for each equation to ensure convergence below a predefined error threshold.

The described NN system has been implemented in the **Python** language through the open source **neurodiffEq** library [47], built on **PyTorch**.

The hyper-parameters that we use include an initial learning rate of 0.001 decreasing 1.5% every five thousand epochs. All other parameters are kept at the **PyTorch** default. For the solver network, we use 6 different networks with three hidden layers of 32 nodes with

²In the context of NN-training, the optimizer is the method that implements some algorithm for the optimization of the loss function, generally some form of stochastic gradient descent.

$\tanh$ activation. For the V-network, we employ 4 hidden layers with 16 nodes and SiLU activation. This architecture was determined through trial and error, exploring configurations that have previously proven effective for similar problems.

4.4 Technical aspects of setup and training

In this section, we describe in detail the main technical aspects of our method. The setup is shown in figure 3. As already mentioned, the NN-solver takes as inputs both the sampling of points of the independent variable u , and pairs of points (T, S) along the thermodynamic curve, each of them corresponding to different boundary conditions on the functions $\tilde{\Sigma}$ and ν_A . Let us describe some aspects of each input separately.

4.4.1 Independent variable generator

The independent variable u runs from the AdS boundary ($u = 0$) to the location of the black brane horizon ($u = 1$). Since this is the independent variable of the differential equations and one of the inputs of our model, we first need to sample it in such a way that it covers the range of interest. To do so, we have decided to select 48 points sampled using Chebyshev of the second kind nodes,³ that is, by placing a higher density of points near the endpoints of the interval. This results in a better resolution of the value of the functions near both boundaries.

We have also explored the possibility of sampling a different number of points, and using a different sampling technique. The main advantage of reducing the number of points in the u -direction is that the training process is faster but, in general, we obtain less precise solutions to the differential equations and, thus, a less precise reconstruction of the unknown potential. We have also explored the possibility of using an equally spaced distribution of points randomly displaced a small quantity (noise) at each training epoch. However, this sampling technique does not allow us to resolve the boundary region as well as with Chebyshev sampling and it results in slightly worse results.

4.4.2 Points on the equation of state

An important aspect regarding the inputs of the network is the number of points sampled along the $S(T)$ curve. On the one hand, a large number of points will lead to long training times. On the other hand, a small number of points will not provide the network with enough information to reconstruct the potential correctly. We have found that the number of points that provides a good balance between training time and precise recovered potentials is in between 65–70, sampled for values of the temperatures ranging from $T/\Lambda = 0$ to $T/\Lambda \simeq 0.6$. This range of temperatures translates into a range of entropies from $S/\Lambda^3 = 0$ to $S/\Lambda^3 \simeq 30$. The maximum value of the temperature is chosen so that it is higher than the transition/crossover temperature, which is $T_c/\Lambda \simeq 0.4$, but not so high as to generate too large a hierarchy for the values of the entropy density. Moreover, this sampling is not uniformly done for all values of the temperature. Instead, a higher density of points is chosen in the phase transition or crossover regions, as well as in the IR region of the curve — see figure 4. The reason for this is that the UV region of the potential is easily constrained,

³For a mathematical definition of Chebyshev polynomials, see [49].

whereas more precision is needed in order to resolve correctly the parts of the potential that control the transition and the IR physics.

4.4.3 Architecture and activations

As illustrated in figure 3, the architectures of each of the two NNs are different. The first one, NN-Solver, is composed of 6 different nets, each one of them taking u and the boundary conditions (T, S) as inputs and giving the solutions to the DEs as outputs. These nets have 3 hidden layers with 32 neurons each. The second NN, NN-V, acts as the scalar potential function $V(\phi)$, taking the solution $\phi(u)$ coming from solving the DEs as input and giving the scalar potential as an output. This net is composed of 4 hidden layers with 16 neurons each. In summary, we have:

$$\begin{aligned} [\text{NN-Solver}]_i &: [32, 32, 32]_i, \quad \text{with } i = 1, \dots, 6; \\ \text{NN-V} &: [16, 16, 16, 16] \end{aligned}$$

We have chosen the standard hyperbolic tangent, $\tanh(x)$, for the activation functions of all neurons in the ODE nets NN-Solver. In contrast, we use the Sigmoid-weighter Linear Unit function (SiLU, a.k.a. Swish), $\text{SiLU}(x) = x(1 + e^{-x})^{-1}$, as the activation for the neurons in the hidden layers of the potential net NN-V. The fact that the SiLU function is not bounded from above, as opposed to $\tanh(x)$, allows for more general behavior of the output $V(\phi)$. This is a desirable property when predicting free-form functions that could a priori be extremely complicated, since they must reproduce the thermodynamics of a variety of different theories.

4.4.4 Gaussian localization in V-NN

We have developed a novel feature that we have dubbed ‘‘Gaussian localization’’ (GL), pictorially shown in figure 3, that has improved the reconstructed potential. Here we will describe the method and in section 6 we will discuss its effect. The NN-V network is composed of 4 layers whose response can be expressed as

$$\mathbf{h}_i = \text{SiLU} \left\{ \mathbf{h}_{i-1} \cdot \mathbf{W}_{\mathbf{V},i}^T \right\}, \quad (4.6)$$

where $\mathbf{h}_i$ is the output vector of the i -th layer. In the case of $\mathbf{h}_1$ (response of the first hidden layer of NN-V), $h_0 = \phi$, where ϕ is the scalar field that is one of the outputs of the solver network. $\mathbf{W}_{\mathbf{V},i}$ is the weight matrix⁴ (except $\mathbf{W}_{\mathbf{V},0}^T$ which is a vector).

This general framework is modified as follows by our GL: we multiply the output given by (4.6) of the first hidden layer by a Gaussian function, resulting in

$$\tilde{\mathbf{h}}_1 = \mathbf{h}_1 \times \exp \left[- \left(\frac{\mathbf{h}_1 - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \right)^2 \right], \quad (4.7)$$

where $\tilde{\mathbf{h}}_1$ the output vector of neurons of the first layer after the localization. The difference is thus that the Gaussian multiplies the output of the neurons from the 1st layer, which is (4.6) after passing it through the activation SiLU. The parameters $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are vectors of the means

⁴It is customary to absorb the bias in the weight matrix by simply adding a column of ones.




Figure 4. Sampling of the $S(T)$ curve for a crossover (top), a second-order transition (middle) and a first-order transition (bottom), in the range of temperatures 0 to 0.7. The points along the curves are sampled more densely around the phase transition or crossover ($T/\Lambda \approx 0.4$), as well as in the IR ($T/\Lambda \approx 0$).

and standard deviations of the Gaussians corresponding to the localization for each neuron. For the Gaussian widths we have chosen $\sigma = 0.1$, and we have let their centers, μ , be a learnable parameter. The result of this learning is that these centers are approximately evenly distributed throughout the range of values of ϕ . As a consequence, each neuron in the first layer, and the subsequent paths that follow from its output, specialize in recovering a certain region of $V(\phi)$. To understand the enhanced performance of this approach, consider how the network approximates the function $V(\phi)$ across the entire input domain. Given that we employ a fully connected network, a modification in the weights of the first hidden layer that may reduce the loss at a specific input value ϕ could potentially have a contrasting impact at another input value ϕ' . It is imperative that the gradients of the loss with respect to the weights of the first layer exhibit a low correlation for different values of ϕ . Essentially, the quantity $(\nabla_W L|_{\phi}) \cdot (\nabla_W L|_{\phi'})$ should be minimized when ϕ and ϕ' are significantly distant. While there exist various methods to achieve this objective, a straightforward and efficient approach involves scaling the outputs of each node in the first layer by a Gaussian function, as previously outlined. From a physics perspective, an alternative interpretation lies in the “multi-scale” and “multi-entangled” characteristics inherent in the gravitational aspect of the problem. The multi-scale aspect refers to the fact that the five-dimensional geometries can develop a large hierarchy of scales between the near-boundary and the near-horizon regions. We will come back to this point in section 6. Multi-entanglement refers to the fact that the value of the potential at a certain value $\phi = \phi_0$ affects the thermodynamic properties of *all* the black brane geometries with horizons such that $\phi_H \geq \phi_0$. Our results suggest that a NN in which each neuron specialises in a certain region or scale is better suited for solving a multi-scale and multi-entangled problem.

We have also examined the effect of GL on the network NN-Solver that we have used to solve the DEs. Note that since this model has three inputs, u , T and S , three independent localizations are needed, one for each of these variables. Localizing in a multi-dimensional space presents significant challenges due to the extensive parameterization. Currently, the methods that we have employed for localizing in dimensions such as u , T , and S have not yielded the desired improvements in training. In fact, they have hindered progress rather than facilitating it. As we continue our efforts, we are exploring alternative approaches and refining our strategies to enhance the localization process and improve overall training outcomes. For now, we have decided not to include it in this work.

We have also explored the possibility of letting the network learn the values of the Gaussian widths σ_j in NN-V. However, this leads to only a couple neurons in the first layer controlling the whole output of the layer, resulting in an augmentation of free parameters that does not improve the fit.

4.5 Numerical tests

Here we perform two numerical checks of the model. First, we verify that the NN is solving the DEs correctly. For this purpose, rather than reconstructing the potential, we fix it to be one of the potentials in figure 2. Then we ask the NN-Solver to solve the direct problem with this potential, namely, to solve the system of ODEs (4.2) for three different points along the $S(T)$ -curve. We then compare these solutions with their “theoretical” values obtained with a traditional method [36]. For each function, we define the error as the square of the

difference between the function produced by the NN and the theoretical function. Since these functions take values of order unity in most of their domain (see figures 10 and 11), the square root of this error is a good measure of the relative error between the theoretical and the NN-predicted functions.

The three different pairs (T, S) that we have given to the NN correspond to different boundary conditions of the functions $\tilde{\Sigma}(u)$, $\tilde{A}(u)$ and $\phi(u)$ at the horizon. We have chosen these points to probe the quality of the solutions in three different regions of the $S(T)$ -curve. The first one (high temperature) is chosen to be around $T/\Lambda \simeq 0.5$ and tests the solution near the AdS boundary. The second one is placed around $T/\Lambda \simeq 0.4$ (intermediate temperature) and it probes the most “non-conformal” region of the boundary theory. The third point is chosen to be near $T/\Lambda \simeq 0.25$ (low temperature) and it probes the solutions near the IR fixed point. We have performed this test for two different theories characterized by $\phi_M = 5$, which possesses a crossover, and $\phi_M = 1$, which corresponds to a first-order phase transition. For the first case and for $1 \cdot 10^6$ epochs, the highest error is found for the low-temperature solution and it has a value of approximately $1.4 \cdot 10^{-4}$. This maximum error decreases substantially when increasing the number of epochs to $5 \cdot 10^6$, giving a final value of $7 \cdot 10^{-6}$. For the case $\phi_M = 1.0$ we have also found the maximum error for the low-temperature solution, with a value of $1.0 \cdot 10^{-3}$ when the model is trained for $1 \cdot 10^6$ epochs. As before, if we increase the number of epochs up to $5 \cdot 10^6$ then this error decreases down to $2.5 \cdot 10^{-4}$. The main conclusion of this test is that the NN-Solver is able to correctly solve the system of ODEs with sufficiently high precision when the function $V(\phi)$ is specified. The fact that the largest errors occur for the low-temperature solutions is consistent with the multiscale, multientangled nature of the problem discussed in 4.4.4, since these solutions depend on the value of the potential over a longer range of ϕ values.

We have performed several tests regarding the choice of the sampling technique and the number of points of the independent variable. As we have mentioned, for the sampling techniques we have used Chebyshev nodes and equally-spaced points with noise. For each of these methods we have tested the model’s performance for 32 and 48 points, for the $S(T)$ -curve corresponding to the $\phi_M = 5$ case. In general, we have observed that the Chebyshev sampling yields better results, in the sense that we have obtained losses of the order of 10^{-6} using equally-spaced noisy points, and losses of 10^{-7} using Chebyshev nodes. Moreover, we have observed that this last method is much more controlled and consistent. Indeed, in both cases the learning rate needs to be adjusted manually at some point, but the random component of the noisy equally spaced points makes it very difficult to estimate *when* this should be done. In contrast, for the Chebyshev technique it is very clear that this must be done when the loss starts flattening (see figure 5).

Regarding the number of points, we have observed that the overall loss can be lower for 32 points than for 48. For example, the lowest loss we found was $8.18 \cdot 10^{-8}$, which occurred for the $\phi_M = 5$ model with Chebyshev nodes and 32 points in the u -direction. In contrast, with 48 points we obtained a loss of $1.16 \cdot 10^{-7}$. Nevertheless, we have also observed that the recovered potential captures the UV behaviour near $\phi = 0$ more precisely with 48 points. For this reason, we decided to use 48 points in the u -direction in the final algorithm.

The training process of the NN is subject to a certain stochasticity due both to the nature of the stochastic gradient descent and to the random distribution of the initial weights

and biases. In order to estimate the magnitude of this uncertainty, we have trained the same model multiple times for a given number of epochs. The small dispersion of the results gives an idea of the robustness of our method. In addition, we have found that this is a more efficient way of obtaining a good solution, as compared to training the model only once for much larger number of epochs. In summary, we performed 10 different runs of order $\sim 10^6$ epochs for each model and, in some cases, we continued the training for the best run.

4.6 Computational considerations

The total number of free parameters of our model (weights, biases and Gaussian parameters) is 14519, of which 881 belong to W_V and 13638 belong to W_D . This gives us approximately $1.5 \cdot 10^4$ adjustable parameters to train our model to solve the differential equations and to reconstruct the potential.

Training and discovery of a solution with sub-per cent precision in the potential and per cent precision in the recovered equation of state requires about 8–16 hours and a couple of million epochs in a dedicated NVidia A40 GPU. As explained above, we performed ten runs for each (T, S) -curve. The solutions discovered could be used for other problems using transfer learning and will be added to the [neurodiffHUB](#) repository. These solutions, together with our code, will be made publicly available in the near future.

5 Results

We test our NN algorithm using boundary conditions associated to the $S(T)$ -curves in figure 1. Different values of ϕ_M correspond to different theories with a crossover, a second- or a first-order phase transition. These differences are encoded in the different shapes of the thermodynamic curves $S(T)$, as can be seen in figure 1. As explained in sections 2 and 4, we find solutions for the dependent variables $\tilde{\Sigma}(u)$, $\tilde{A}(u)$ and $\phi(u)$, one for each boundary condition, i.e. one for each point in $S(T)$. These are solutions to the system of ODEs (4.2), which depends on the unknown potential $V(\phi)$. This potential is the same for, and is informed about, all the boundary conditions, and we find it (necessarily) at the same time as the rest of the functions. We obtain the following results.

5.1 Loss function

Our method is based on minimizing the loss function (4.5). Figure 5 shows the typical form of this function, in this case for the best run for the $\phi_M = 1$ model, which possesses a first-order phase transition case. We see how L greatly decreases in the first half a million epochs and keeps doing so at an increasingly slower rate for the next 2.5 million epochs. It is interesting to note that, for our choice of hyperparameters, the 1.5% decrease in the learning rate every 5000 epochs (see section 4.3.2) makes the training more efficient only up to a certain point, after which the learning rate becomes too small compared to the value of the loss, and the loss flattens. At this point, we say that the NN is “trapped” in some local minimum. To continue with efficient training, one must increase the learning rate by hand. This can be seen at 2 and 3 million epochs in figure 5. Automatization of this procedure is part of our future work that could incorporate techniques such as annealing.




Figure 5. Loss function (squared residual from our ODEs) in log scale as a function of the epoch, for the best run of the first order transition case ($\phi_M = 1$). This NN has been trained for 3.5 million epochs. The minimum loss achieved is $\approx 7.4 \cdot 10^{-7}$.

5.2 Recovered potential and equation of state

In the top panels of figures 6, 7 and 8 we show the potentials produced by our method when supplied with the $S(T)$ curves corresponding to theories with a crossover, a second- and a first-order phase transition, respectively. For each case, we run 10 realizations, and we select the best run as the one with the lowest value of the loss function. The resulting potentials for these best runs are shown as solid red curves in the figures. The remaining 9 runs are shown as dash-dotted green curves. The small dispersion between these curves illustrates the robustness of our method. Finally, the theoretical potentials (3.3) are shown as dashed blue curves. As is clear from the curves, the recovered potential is in general in excellent agreement with the exact one. For reasons discussed in section 6, the precision of the recovery degrades as we move from a crossover to a second- to a first-order phase transition. We quantify this with the relative error

$$|\delta V(\phi)| = \left| \frac{V_{\text{theory}}(\phi) - V_{\text{PINN}}(\phi)}{V_{\text{theory}}(\phi)} \right|. \quad (5.1)$$

We plot this error in figure 9(top), where we see that it is below $\sim 1\%$ for all values of ϕ for all theories. In this figure we also list the relative root mean square⁵ (rms), which is below 0.5%.

As explained in section 2, the test for accuracy shown in figure 9(top) can only be done if one *knows* the theoretical potential $V_{\text{theory}}(\phi)$. This implies having previously solved

⁵This is defined as $\text{rms} = \sqrt{\frac{1}{n} \sum_{i=1}^n x_i^2}$ with x_i is the relative error defined in equations (5.1) and (5.2).




Figure 6. (Top) Best scalar field potential predicted by NN-V (solid red) compared to the theoretical potential (dashed blue) for a theory with a crossover. The plot extends up to the position of the predicted minimum of the potential, while the vertical line (dashed blue) indicates the position of the theoretical minimum. (Bottom) Recovered equation of state (solid red), obtained by solving the direct problem with the recovered potential, compared to the theoretical one (dashed blue).




Figure 7. (Top) Best scalar field potential predicted by NN-V (solid red) compared to the theoretical potential (dashed blue) for a theory with a second-order phase transition. The plot extends up to the position of the predicted minimum of the potential, while the vertical line (dashed blue) indicates the position of the theoretical minimum. (Bottom) Recovered equation of state (solid red), obtained by solving the direct problem with the recovered potential, compared to the theoretical one (dashed blue).




Figure 8. (Top) Best scalar field potential predicted by NN-V (solid red) compared to the theoretical potential (dashed blue) for a theory with a first-order phase transition. The plot extends up to the position of the predicted minimum of the potential, while the vertical line (dashed blue) indicates the position of the theoretical minimum. (Bottom) Recovered equation of state (solid red), obtained by solving the direct problem with the recovered potential, compared to the theoretical one (dashed blue).




Figure 9. Relative error for the potential (top) and for the recovered equation of state (bottom) for the best run. These quantities have been computed using the expressions (5.1) and (5.2). In the legend, we show the rms for each value of ϕ_M . Note that $|\delta V|$ (top) extends over a longer range of ϕ for the cases with lower ϕ_M . This is because the position of the minimum of the scalar potential $V(\phi)$ is at a larger ϕ for cases with smaller ϕ_M (see figure 2). Also, note that $|\delta T|$ (bottom) tends to be a constant for $S/\Lambda^3 \rightarrow 0$. The small value of this constant indicates that the PINN reconstructs the IR CFT with precision.

the “direct problem” to obtain $S_{\text{theory}}(T)$ from $V_{\text{theory}}(\phi)$. However, if the input curve $S_{\text{input}}(T)$ has been obtained in a different manner than solving the direct problem (from the lattice, from observations, etc.), then one would not have the theoretical potential to compare $V_{\text{PINN}}(\phi)$ with.

We thus turn to a more stringent test for accuracy. We take the best-recovered potential $V_{\text{PINN}}(\phi)$ and use it to solve the direct problem. This produces a curve $S_{\text{PINN}}(T)$ that we then compare with $S_{\text{input}}(T)$. The results are shown in the bottom panels of figures 6, 7 and 8, respectively. We plot S/T^3 instead of S because this exhibits the UV and IR behaviors more clearly. As in the case of the potential, we quantify the relative error as follows.

$$|\delta T(S)| = \left| \frac{T_{\text{input}}(S) - T_{\text{PINN}}(S)}{T_{\text{input}}(S)} \right|. \quad (5.2)$$

We use $T(S)$ instead of $S(T)$ because the latter is multi-valued, whereas the former is not. This error is shown in figure 9(bottom). The phase transition is replicated fairly accurately. For example, the largest relative error in the recovered $S(T)$ occurs in the case of a second-order phase transition illustrated in figure 7, and it is $\sim 5\%$. We also note that not only the PINN “discovers” the presence of an IR fixed point, but it predicts the number of IR degrees of freedom with remarkable precision. This is measured by the small error, of order $\sim 0.2\%$, in $\delta T(S = 0)$, as illustrated in figure 9(bottom). In addition, the rms, which is an “integrated” measure of the overall error, is below $\sim 1\%$ for theories with a crossover, and around $\sim 1\%$ for theories with a phase transition. We thus conclude that the entire RG flow is correctly reconstructed.

5.3 Solutions to the ODEs

The correct reconstruction of the potential requires solving the EKG accurately. Here we show the results for these solutions, namely, for the functions $\tilde{\Sigma}(u)$, $\tilde{A}(u)$ and $\phi(u)$. There is one solution for each boundary condition, i.e. for each point in $S(T)$. We show these solutions for the theories with $\phi_M = 1$ (figure 10) and $\phi_M = 5$ (figure 11). For each theory, we show the solutions for 3 different points on the $S(T)$ curve that correspond to high, intermediate and low temperature in the language of section 4.5. We can see an excellent agreement for both theories between the “theoretical” solutions (dash-dotted curves) and the NN-predicted ones (solid curves), with a mean squared error⁶ (MSE) of order 10^{-6} – 10^{-5} . Note that, since the functions $\tilde{\Sigma}(u)$, $\tilde{A}(u)$ and $\phi(u)$ take values of order unity in most of their domain, the square root of the MSE is a good measure of the relative error between the theoretical and the NN-reconstructed functions.

6 Discussion

Holography maps the quantum properties of a gauge theory in four dimensions to the classical properties of a gravitational theory in five dimensions. If the gravitational theory is known, the equation of state of the dual gauge theory can be found by solving a direct problem, i.e. by finding all the black hole solutions of the gravitational theory. The inverse problem,

⁶This is defined as $\text{MSE} = \frac{1}{n} \sum_{i=1}^n \left(x_i^{(\text{th})} - x_i^{(\text{NN})} \right)^2$.




Figure 10. (Left) Comparison between the theoretical (dashed curves) and the PINN (solid curves) solutions for the theory with $\phi_M = 1$. (Right) Squared differences between the theoretical and the PINN solutions.




Figure 11. (Left) Comparison between the theoretical (dashed curves) and the PINN (solid curves) solutions for the theory with $\phi_M = 5$. (Right) Squared differences between the theoretical and the PINN solutions.

namely to determine a gravitational theory that gives rise to a prescribed equation of state, is much more challenging. We have shown that this problem can be solved using Physics Informed Neural Networks. The resulting algorithm reconstructs not just a specific black hole solution but the gravitational theory itself. We have illustrated the method in a simple setup in which the gravitational theory is completely specified by one function, the potential $V(\phi)$ for a scalar field ϕ . However, we expect that the method can be generalized to gravitational theories with a general field content, as well as to inverse problems outside the holographic context involving highly non-linear partial differential equations.

Figures 6, 7 and 8 illustrate the reconstructed potentials, as well as the equations of state that they yield upon solving the direct problem with them. The quality of the results is quantified in figure 9, where we see that the relative error is at the sub-per cent level for the potential and of a few per cent for the equation of state. To obtain these results we provided the PINN with the boundary conditions (3.2) for the potential. These are encoded in the high-temperature behaviour (3.1) of the equation of state and correspond to the fact that the gauge theory is a CFT deformed by a relevant operator of dimension (2.5).

We have also trained the PINN without providing it with the boundary conditions (3.2) for the potential. In figure 12 we compare the results to those obtained by providing the PINN with (3.2) for the theory with $\phi_M = 1$. We see that, when the PINN is not provided with (3.2), it is still able to recover the entire potential with high precision. In particular, it is able to identify the presence of a maximum in $V(\phi)$ at $\phi = 0$, and to estimate $V(0)$ and $V''(0)$ with good precision. This means that not only does the PINN “discover” by itself the presence of an UV fixed point deformed by a relevant operator, but it also estimates with good precision the number of degrees of freedom at the fixed point, N_{UV} , and the dimension of the relevant operator, Δ . These two quantities can be extracted from $V(0)$ and $V''(0)$. Based on these, we find that the relative errors in the values reconstructed by the PINN are in the following ranges for theories with $1 \leq \phi_M \leq 5$:

$$10^{-6} \lesssim \frac{\delta N_{UV}}{N_{UV}} \lesssim 10^{-5}, \quad 10^{-3} \lesssim \frac{\delta \Delta}{\Delta} \lesssim 10^{-2}. \quad (6.1)$$

In the case of $V''(0)$, this error is illustrated in figure 12(middle). Figure 12(bottom) compares the reconstructed equations of state when the PINN is or is not provided with the boundary conditions (3.2). We see that the result is more precise in the first case. This illustrates the high sensitivity of the $S(T)$ curve to the value of $V''(0)$ or, equivalently, to the dimension of the relevant operator that triggers the RG flow.

Overall, we regard the results summarized in the previous paragraphs as a remarkable success given the challenges implied by the multi-entangled and multi-scale nature of the problem on the gravity side. Multi-entanglement refers to the fact that the value of the potential at a certain point, $V(\phi = \phi_0)$, affects the thermodynamic properties of all the black brane geometries with horizons such that $\phi_H \geq \phi_0$. In other words, a single potential must be compatible with an infinite number of solutions with different boundary conditions. This feature makes two of the ingredients that we have implemented indispensable. One is the architecture based on two independent NNs (see figure 3), which allowed us to determine both the spacetime geometries and the potential at the same time. The other is the use




Figure 12. Comparison between the reconstructed potentials (top), their second derivatives (middle) and the corresponding equations of state (bottom) when the PINN is or is not provided with the boundary conditions (3.2) for the potential.

of solution bundles (see section 4.2), which allowed us to feed the PINN the information about all the boundary conditions simultaneously.

The multiscale aspect refers to the fact that the five-dimensional geometries can develop a large hierarchy of scales between the boundary and the horizon. In the gauge theory, this is due to the fact that the RG flow can generate a large hierarchy between the number of degrees of freedom in the UV and in the IR. This aspect is illustrated in figure 13, where we plot the ratio between the spacetime curvatures at the horizon ($\phi = \phi_H$) and at the boundary ($\phi = 0$), as measured by the corresponding Ricci scalars. The intermediate points on some curves indicate solutions with $T = T_c$. The horizontal lines mark the endpoints of each curve, which correspond to the zero-temperature, ground-state solutions. In order to reconstruct correctly the part of the potential that gives rise to the phase transition, the PINN must be able to resolve the hierarchy associated to the intermediate points. In order to reconstruct the deep IR, the PINN must resolve the even larger hierarchy associated to the endpoints of the curves. We see that these hierarchies increase as we move from a crossover to a second- to a first-order phase transition.

As with most numerical methods, we expect that the appearance of a separation of scales will make the problem more challenging. The new feature of “Gaussian localization” (GL) that we have developed (see section 4.4.4) helps address this challenge. Intuitively, it results in the specialisation of each neuron on a specific part of the potential. In order to illustrate its effect, in figure 14 we compare the potentials for a theory with $\phi_M = 0.8$ that were reconstructed by a NN with GL and by a Fully Connected Neural Network (FCNN) with no GL. The improvement in the results with GL can be seen with the naked eye. Nevertheless, theories like the $\phi_M = 0.8$ theory, in which the hierarchy of scales is sufficiently large, remain challenging for our method even with GL. This can already be seen by comparing figure 14 to figure 8(top). To illustrate it further, in figure 15(top) we show the additional 9 runs that led to the reconstructed potential with GL in figure 14. We see that the dispersion is much larger than in the $\phi_M = 1$ case of figure 8(top). Finally, in figure 15(bottom) we show the reconstructed equation of state. Although the presence of a first-order phase transition is recovered, the quality of the result is clearly inferior to that in figure 8(bottom). Motivated by these challenges, we are currently improving our algorithm by implementing transfer learning. This will allow the PINN to use an already-determined solution to be used as a seed to find the solution for a harder problem.

The above discussion provides the right perspective with which to frame our results. On the one hand, we regard them as an important proof of concept that the gravitational theory can be reconstructed from the equation of state of the dual gauge theory. On the other hand, the previous paragraph illustrates that much work is still necessary in order to perfect the method. Once this is achieved, we expect far-reaching implications for a wide range of physical systems including, but not limited to, those described by gauge theories. Indeed, although we have focused on the thermodynamics of the latter, the equation of state is one of the most fundamental properties of any physical system. Since it only requires knowledge of the equilibrium properties, it can often be determined with a variety of methods. In contrast, the out-of-equilibrium dynamics of many physical systems remains an outstanding challenge. In fact, the exact simulation of this dynamics may be beyond the reach of classical




Figure 13. The Ricci scalar at the horizon, normalized to the Ricci scalar at the boundary, varies with the temperature of the solution, i.e. with ϕ_H . For theories with a first-order phase transition, the three circles on the curve indicate the three different solutions at the critical temperature. These points coalesce into a single one for a second-order phase transition, and disappear for a crossover. The horizontal lines indicate the endpoints of each curve, which correspond to the zero-temperature, ground-state solutions of each theory. The hierarchy of curvatures increases as we move from a crossover to a second- to a first-order phase transition. The missing intermediate part of the $\phi_M = 0.58$ curve is due to the fact that some values of ϕ_H do not correspond to thermal states in the dual gauge theory [36].




Figure 14. Comparison between the reconstructed potentials for a theory with $\phi_M = 0.8$ with (GL) and without (FCNN) Gaussian localization.

calculations and require quantum computers [50, 51]. Given the scale of this challenge, the gravitational dual provided by holography may provide a valuable, if not the only, analytical approximation to the out-of-equilibrium dynamics of many systems.

Acknowledgments

We thank Jorge Casalderrey-Solana and Javier Gomez Subils for very useful discussions. This work was supported by the “Center of Excellence Maria de Maeztu 2020–2023” award to the ICCUB (CEX2019-000918-M) funded by MCIN/AEI/10.13039/501100011033. YB, DM and PeT acknowledge support from grants PID2019-105614GB-C22 and 2021-SGR-872. The work of YB is also funded by a Maria Zambrano postdoctoral fellowship from the University of Barcelona. PeT is supported by the project “Dark Energy and the Origin of the Universe” (PRE2022-102220), funded by MCIN/AEI/10.13039/501100011033. Funding for the work of RJ was partially provided by the projects PGC2018-098866-B-I00 and FEDER “Una manera de hacer Europa”. The work by PaT is supported by the project “Accurate Cosmology and the Laws of Nature” (PID2022-141125NB-I00) with grant PREP2022-000507, funded by MCIN/AEI/10.13039/501100011033.




Figure 15. Reconstructed potential (top) and equation of state (bottom) for a theory with a stronger first-order phase transition than that in figure 8.

A High-temperature behaviour

Consider a four-dimensional CFT deformed by a relevant operator $\mathcal{O}$ of conformal dimension Δ with source Λ . The action takes the form

$$S = S_{\text{CFT}} + \int d^4x \Lambda \mathcal{O}(x). \quad (\text{A.1})$$

For consistency, Λ must have dimension $4 - \Delta$. For homogeneous, thermal equilibrium states, the trace Ward identity takes the form

$$\mathcal{E} - 3\mathcal{P} = \Lambda \langle \mathcal{O} \rangle, \quad (\text{A.2})$$

with $\mathcal{E}$ and $\mathcal{P}$ the energy density and the pressure, respectively. The existence of a conformal, UV fixed point implies that the trace of the stress tensor must vanish at leading order at high temperatures. In this limit, all thermodynamic quantities scale with the temperature as dictated by dimensional analysis, so we have:

$$\mathcal{E} \simeq 3\mathcal{P} \sim T^4. \quad (\text{A.3})$$

Assuming that $\langle \mathcal{O} \rangle$ vanishes in the undeformed theory, its value at leading order in the deformed theory must be linear in the deformation parameter Λ , i.e.

$$\langle \mathcal{O} \rangle \propto \Lambda T^{2\Delta-4}, \quad (\text{A.4})$$

where the power of T is fixed by dimensional analysis. Through the Ward identity (A.2), this fixes the leading-order correction to the asymptotic behaviour (A.3) and, together with the thermodynamic identity $S = \partial \mathcal{P} / \partial T$, it results in the entropy density

$$S = c_1 T^3 + c_2 T^{2\Delta-5} + \dots, \quad (\text{A.5})$$

where c_1, c_2 are numerical coefficients. The main conclusion is that the dimension of the operator responsible for triggering the flow is encoded in the thermodynamic curve $S(T)$. In particular, Δ can be extracted from the subleading term in the high-temperature expansion of the equation of state.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] J.M. Maldacena, *The large N limit of superconformal field theories and supergravity*, *Adv. Theor. Math. Phys.* **2** (1998) 231 [[hep-th/9711200](#)] [[INSPIRE](#)].
- [2] E. Witten, *Anti-de Sitter space and holography*, *Adv. Theor. Math. Phys.* **2** (1998) 253 [[hep-th/9802150](#)] [[INSPIRE](#)].
- [3] S.S. Gubser, I.R. Klebanov and A.M. Polyakov, *Gauge theory correlators from noncritical string theory*, *Phys. Lett. B* **428** (1998) 105 [[hep-th/9802109](#)] [[INSPIRE](#)].

- [4] S.S. Gubser and A. Nellore, *Mimicking the QCD equation of state with a dual black hole*, *Phys. Rev. D* **78** (2008) 086007 [[arXiv:0804.0434](#)] [[INSPIRE](#)].
- [5] S.S. Gubser, A. Nellore, S.S. Pufu and F.D. Rocha, *Thermodynamics and bulk viscosity of approximate black hole duals to finite temperature quantum chromodynamics*, *Phys. Rev. Lett.* **101** (2008) 131601 [[arXiv:0804.1950](#)] [[INSPIRE](#)].
- [6] J. Alanen, K. Kajantie and V. Suur-Uski, *A gauge/gravity duality model for gauge theory thermodynamics*, *Phys. Rev. D* **80** (2009) 126008 [[arXiv:0911.2114](#)] [[INSPIRE](#)].
- [7] R. Rougemont et al., *Hot QCD phase diagram from holographic Einstein-Maxwell-Dilaton models*, *Prog. Part. Nucl. Phys.* **135** (2024) 104093 [[arXiv:2307.03885](#)] [[INSPIRE](#)].
- [8] K. Hashimoto, K. Ohashi and T. Sumimoto, *Deriving the dilaton potential in improved holographic QCD from the meson spectrum*, *Phys. Rev. D* **105** (2022) 106008 [[arXiv:2108.08091](#)] [[INSPIRE](#)].
- [9] K. Hashimoto, K. Ohashi and T. Sumimoto, *Deriving the dilaton potential in improved holographic QCD from the chiral condensate*, *PTEP* **2023** (2023) 033B01 [[arXiv:2209.04638](#)] [[INSPIRE](#)].
- [10] X. Chen and M. Huang, *Machine learning holographic black hole from lattice QCD equation of state*, *Phys. Rev. D* **109** (2024) L051902 [[arXiv:2401.06417](#)] [[INSPIRE](#)].
- [11] W.-C. Gan and F.-W. Shu, *Holography as deep learning*, *Int. J. Mod. Phys. D* **26** (2017) 1743020 [[arXiv:1705.05750](#)] [[INSPIRE](#)].
- [12] Y.-Z. You, Z. Yang and X.-L. Qi, *Machine Learning Spatial Geometry from Entanglement Features*, *Phys. Rev. B* **97** (2018) 045153 [[arXiv:1709.01223](#)] [[INSPIRE](#)].
- [13] K. Hashimoto, S. Sugishita, A. Tanaka and A. Tomiya, *Deep learning and the AdS/CFT correspondence*, *Phys. Rev. D* **98** (2018) 046019 [[arXiv:1802.08313](#)] [[INSPIRE](#)].
- [14] K. Hashimoto, S. Sugishita, A. Tanaka and A. Tomiya, *Deep Learning and Holographic QCD*, *Phys. Rev. D* **98** (2018) 106014 [[arXiv:1809.10536](#)] [[INSPIRE](#)].
- [15] H.-Y. Hu, S.-H. Li, L. Wang and Y.-Z. You, *Machine Learning Holographic Mapping by Neural Network Renormalization Group*, *Phys. Rev. Res.* **2** (2020) 023369 [[arXiv:1903.00804](#)] [[INSPIRE](#)].
- [16] K. Hashimoto, *AdS/CFT correspondence as a deep Boltzmann machine*, *Phys. Rev. D* **99** (2019) 106017 [[arXiv:1903.04951](#)] [[INSPIRE](#)].
- [17] X. Han and S.A. Hartnoll, *Deep Quantum Geometry of Matrices*, *Phys. Rev. X* **10** (2020) 011069 [[arXiv:1906.08781](#)] [[INSPIRE](#)].
- [18] J. Tan and C.-B. Chen, *Deep learning the holographic black hole with charge*, *Int. J. Mod. Phys. D* **28** (2019) 1950153 [[arXiv:1908.01470](#)] [[INSPIRE](#)].
- [19] Y.-K. Yan, S.-F. Wu, X.-H. Ge and Y. Tian, *Deep learning black hole metrics from shear viscosity*, *Phys. Rev. D* **102** (2020) 101902 [[arXiv:2004.12112](#)] [[INSPIRE](#)].
- [20] T. Akutagawa, K. Hashimoto and T. Sumimoto, *Deep Learning and AdS/QCD*, *Phys. Rev. D* **102** (2020) 026020 [[arXiv:2005.02636](#)] [[INSPIRE](#)].
- [21] K. Hashimoto, H.-Y. Hu and Y.-Z. You, *Neural ordinary differential equation and holographic quantum chromodynamics*, *Mach. Learn. Sci. Tech.* **2** (2021) 035011 [[arXiv:2006.00712](#)] [[INSPIRE](#)].
- [22] N. Jokela and A. Pönni, *Towards precision holography*, *Phys. Rev. D* **103** (2021) 026010 [[arXiv:2007.00010](#)] [[INSPIRE](#)].

- [23] M. Song, M.S.H. Oh, Y. Ahn and K.-Y. Kima, *AdS/Deep-Learning made easy: simple examples*, *Chin. Phys. C* **45** (2021) 073111 [[arXiv:2011.13726](#)] [[INSPIRE](#)].
- [24] K. Hashimoto and R. Watanabe, *Bulk reconstruction of metrics inside black holes by complexity*, *JHEP* **09** (2021) 165 [[arXiv:2103.13186](#)] [[INSPIRE](#)].
- [25] K. Li, Y. Ling, P. Liu and M.-H. Wu, *Learning the black hole metric from holographic conductivity*, *Phys. Rev. D* **107** (2023) 066021 [[arXiv:2209.05203](#)] [[INSPIRE](#)].
- [26] N. Jokela et al., *Holographic spacetime from lattice Yang-Mills theory*, *EPJ Web Conf.* **274** (2022) 08012 [[arXiv:2211.09563](#)] [[INSPIRE](#)].
- [27] N. Jokela et al., *Progress in the lattice evaluation of entanglement entropy of three-dimensional Yang-Mills theories and holographic bulk reconstruction*, *JHEP* **12** (2023) 137 [[arXiv:2304.08949](#)] [[INSPIRE](#)].
- [28] W.-B. Xu and S.-F. Wu, *Reconstructing black hole exteriors and interiors using entanglement and complexity*, *JHEP* **07** (2023) 083 [[arXiv:2305.01330](#)] [[INSPIRE](#)].
- [29] C. Park, S. Kim and J.H. Lee, *Holography Transformer*, [arXiv:2311.01724](#) [[INSPIRE](#)].
- [30] B. Ahn, H.-S. Jeong, K.-Y. Kim and K. Yun, *Deep learning bulk spacetime from boundary optical conductivity*, *JHEP* **03** (2024) 141 [[arXiv:2401.00939](#)] [[INSPIRE](#)].
- [31] Z.-F. Gu, Y.-K. Yan and S.-F. Wu, *Neural ODEs for holographic transport models without translation symmetry*, [arXiv:2401.09946](#) [[INSPIRE](#)].
- [32] J. Erlich, E. Katz, D.T. Son and M.A. Stephanov, *QCD and a holographic model of hadrons*, *Phys. Rev. Lett.* **95** (2005) 261602 [[hep-ph/0501128](#)] [[INSPIRE](#)].
- [33] L. Da Rold and A. Pomarol, *Chiral symmetry breaking from five dimensional spaces*, *Nucl. Phys. B* **721** (2005) 79 [[hep-ph/0501218](#)] [[INSPIRE](#)].
- [34] E. Witten, *Anti-de Sitter space, thermal phase transition, and confinement in gauge theories*, *Adv. Theor. Math. Phys.* **2** (1998) 505 [[hep-th/9803131](#)] [[INSPIRE](#)].
- [35] A. Karch, E. Katz, D.T. Son and M.A. Stephanov, *Linear confinement and AdS/QCD*, *Phys. Rev. D* **74** (2006) 015005 [[hep-ph/0602229](#)] [[INSPIRE](#)].
- [36] Y. Bea and D. Mateos, *Heating up Exotic RG Flows with Holography*, *JHEP* **08** (2018) 034 [[arXiv:1805.01806](#)] [[INSPIRE](#)].
- [37] M. Dissanayake and N. Phan-Thien, *Neural-network-based approximations for solving partial differential equations*, *Communi. Numer. Meth. Engng.* **10** (1994) 195.
- [38] I.E. Lagaris, A. Likas and D.I. Fotiadis, *Artificial neural networks for solving ordinary and partial differential equations*, *IEEE Trans. Neural Networks* **9** (1998) 987.
- [39] C. Bernardi and Y. Maday, *Spectral methods*, in *Handbook of Numerical Analysis*, vol. 5, Elsevier (1997), p. 209–485 [[DOI:10.1016/s1570-8659\(97\)80003-8](#)].
- [40] I.E. Lagaris, A. Likas and D.I. Fotiadis, *Artificial neural network methods in quantum mechanics*, *Comput. Phys. Commun.* **104** (1997) 1 [[quant-ph/9705029](#)] [[INSPIRE](#)].
- [41] I.E. Lagaris, A.C. Likas and D.G. Papageorgiou, *Neural-network methods for boundary value problems with irregular boundaries*, *IEEE Trans. Neural Networks* **11** (2000) 1041.
- [42] K.S. McFall and J.R. Mahan, *Artificial Neural Network Method for Solution of Boundary Value Problems With Exact Satisfaction of Arbitrary Boundary Conditions*, *IEEE Trans. Neural Networks* **20** (2009) 1221.

- [43] M. Baymani, A. Kerayechian and S. Effati, *Artificial Neural Networks Approach for Solving Stokes Problem*, *Applied Mathematics* **01** (2010) 288.
- [44] S. Desai et al., *One-Shot Transfer Learning of Physics-Informed Neural Networks*, [arXiv:2110.11286](#).
- [45] L. Lu, X. Meng, Z. Mao and G.E. Karniadakis, *DeepXDE: A deep learning library for solving differential equations*, *SIAM Rev.* **63** (2021) 208 [[arXiv:1907.04502](#)].
- [46] A. Koryagin, R. Khudorozkov and S. Tsimfer, *PyDEns: a Python Framework for Solving Differential Equations with Neural Networks*, [arXiv:1909.11544](#).
- [47] F. Chen et al., *NeuroDiffEq: A Python package for solving differential equations with neural networks*, *J. Open Source Softw.* **5** (2020) 1931.
- [48] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [49] J. Mason and D. Handscomb, *Chebyshev Polynomials*, CRC Press LLC (2000), N.W. Corporate Blvd., Boca Raton, Florida (2003).
- [50] C.W. Bauer et al., *Quantum Simulation for High-Energy Physics*, *PRX Quantum* **4** (2023) 027001 [[arXiv:2204.03381](#)] [[INSPIRE](#)].
- [51] C.W. Bauer, Z. Davoudi, N. Klco and M.J. Savage, *Quantum simulation of fundamental particles and forces*, *Nature Rev. Phys.* **5** (2023) 420 [[arXiv:2404.06298](#)] [[INSPIRE](#)].