



Teaching to extract spectral densities from lattice correlators to a broad audience of learning-machines

Michele Buzziotti^a, Alessandro De Santis^b , Nazario Tantalo^c

University and INFN of Roma Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy

Received: 11 July 2023 / Accepted: 5 January 2024 / Published online: 12 January 2024
© The Author(s) 2024

Abstract We present a new supervised deep-learning approach to the problem of the extraction of smeared spectral densities from Euclidean lattice correlators. A distinctive feature of our method is a model-independent training strategy that we implement by parametrizing the training sets over a functional space spanned by Chebyshev polynomials. The other distinctive feature is a reliable estimate of the systematic uncertainties that we achieve by introducing several ensembles of machines, the broad audience of the title. By training an ensemble of machines with the same number of neurons over training sets of fixed dimensions and complexity, we manage to provide a reliable estimate of the systematic errors by studying numerically the asymptotic limits of infinitely large networks and training sets. The method has been validated on a very large set of random mock data and also in the case of lattice QCD data. We extracted the strange-strange connected contribution to the smeared R -ratio from a lattice QCD correlator produced by the ETM Collaboration and compared the results of the new method with the ones previously obtained with the HLT method by finding a remarkably good agreement between the two totally unrelated approaches.

1 Introduction

The problem of the extraction of hadronic spectral densities from Euclidean correlators, computed from numerical lattice QCD simulations, has attracted a lot of attention since many years (see Refs. [1–30], the works on the subject of which we are aware of, and Refs. [31, 32] for recent reviews). At zero temperature, the theoretical and phenomenological importance of hadronic spectral densities, strongly emphasized in

the context of lattice field theory in Refs. [1, 11, 13, 14, 17–19], is associated with the fact that from their knowledge it is possible to extract all the information needed to study the scattering of hadrons and, more generally, their interactions.

From the mathematical perspective, the problem of the extraction of spectral densities from lattice correlators is equivalent to that of an inverse Laplace-transform operation, to be performed numerically by starting from a discrete and finite set of noisy input data. This is a notoriously ill-posed numerical problem that, in the case of lattice field theory correlators, gets even more complicated because lattice simulations have necessarily to be performed on finite volumes where the spectral densities are badly-behaving distributions.

In Ref. [14], together with M. Hansen and A. Lupo, one of the authors of the present paper proposed a method to cope with the problem of the extraction of spectral densities from lattice correlators that allows to take into account the fact that distributions have to be smeared with sufficiently well-behaved test functions. Once smeared, finite volume spectral densities become numerically manageable and the problem of taking their infinite volume limit is mathematically well defined. The method of Ref. [14] (HLT method in short) has been further refined in Ref. [21] where it has been validated by performing very stringent tests within the two-dimensional $O(3)$ non-linear σ -model.

In this paper we present a new method for the extraction of smeared spectral densities from lattice correlators that is based on a supervised deep-learning approach.

The idea of using machine-learning techniques to address the problem of the extraction of spectral densities from lattice correlators is certainly not original (see e.g. Refs. [15, 16, 22–29]). The great potential of properly-trained deep neural networks in addressing this problem is pretty evident from the previous works on the subject. These findings strongly motivated us to develop an approach that can be used to obtain

^a e-mail: michele.buzziotti@roma2.infn.it

^b e-mail: alessandro.desantis@roma2.infn.it (corresponding author)

^c e-mail: nazario.tantalo@roma2.infn.it




Fig. 1 By introducing a discrete functional-basis, with elements $B_n(E)$, that is dense in the space of square-integrable functions $f(E)$ in the interval $[E_0, \infty)$ with $E_0 > 0$, any such function can exactly be represented as $f(E) = \sum_{n=0}^{\infty} c_n B_n(E)$. With an infinite number of basis functions ($N_b = \infty$) and by randomly selecting an infinite number ($N_\rho = \infty$) of coefficient vectors $\mathbf{c} = (c_0, \dots, c_{N_b})$, one can get any possible spectral density. This is the situation represented by the filled blue disk. If the number of basis functions N_b and the number of randomly extracted spectral densities N_ρ are both finite one has a training set that is finite and that also depends on N_b . This is the situation represented in the first disk on the left. The other two disks schematically represent the situations in which either N_b or N_ρ is infinite

trustworthy theoretical predictions. To this end we had to address the following two pivotal questions

1. is it possible to devise a *model independent* training strategy?
2. if such a strategy is found, is it then possible to quantify reliably, together with the statistical errors, also the unavoidable *systematic uncertainties*?

The importance of the first question can hardly be underestimated. Under the working assumption, supported by the so-called universal reconstruction theorems (see Refs. [33–35]), that a sufficiently large neural network can perform any task, limiting either the size of the network or the information to which it is exposed during the training process means, in fact, limiting its ability to solve the problem in full generality. Addressing the second question makes the difference between providing a possibly efficient but qualitative solution to the problem and providing a scientific numerical tool to be used in order to derive theoretical predictions for phenomenological analyses.

In order to address these two questions, the method that we propose in this paper has been built on two pillars

1. the introduction of a *functional-basis* to parametrize the correlators and the smeared spectral densities of the training sets in a model independent way;
2. the introduction of the *ensemble of machines*, the broad audience mentioned in the title, to estimate the systematic errors.¹

¹ The idea of using ensembles of machines has already been explored within the machine-learning literature in other contexts, see e.g. Ref. [36] where the so-called “ensembling” technique to estimate the errors is discussed from a Bayesian perspective.

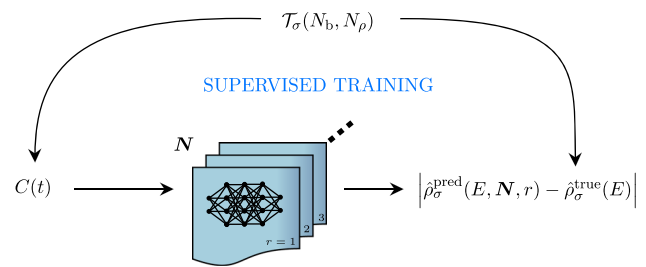


Fig. 2 We generate several training sets $\mathcal{T}_\sigma(N_b, N_\rho)$ built by considering randomly chosen spectral densities. These are obtained by choosing N_ρ random coefficient vectors with N_b entries, see Fig. 1. For each spectral density $\rho(E)$ we build the associated correlator $C(t)$ and smeared spectral density $\hat{\rho}_\sigma(E)$, where σ is the smearing parameter. We then distort the correlator $C(t)$, by using the information provided by the statistical variance of the lattice correlator (the one we are going to analyse at the end of the trainings), and obtain the input–output pair $(C_{\text{noisy}}(t), \hat{\rho}_\sigma(E))$ that we add to $\mathcal{T}_\sigma(N_b, N_\rho)$. We then implement different neural networks with N_n neurons and at fixed $\mathbf{N} = (N_n, N_b, N_\rho)$ we introduce an *ensemble of machines* with N_r replicas. Each machine $r = 1, \dots, N_r$ belonging to the ensemble has the same architecture and, before the training, differs from the other replicas for the initialization parameters. All the replicas are then trained over $\mathcal{T}_\sigma(N_b, N_\rho)$ and, at the end of the training process, each replica will give a different answer depending upon \mathbf{N}

A bird-eye view of the proposed strategy is given in Figs. 1, 2 and in Fig. 3. The method is validated by using both mock and true lattice QCD data. In the case of mock data the exact result is known and the validation test is quite stringent. In the case of true lattice QCD data the results obtained with the new method are validated by comparison with the results obtained with the HLT method.

The plan of the paper is as follows. In Sect. 2 we introduce and discuss the main aspects of the problem. In Sect. 3 we illustrate the numerical setup used to obtain the results presented in the following sections. In Sect. 4 we describe the construction of the training sets and the proposed model-independent training strategy. In Sect. 5 we illustrate the procedure that we use to extract predictions from our ensembles of trained machines and present the results of a large number of validation tests performed with random mock data. Further validation tests, performed by using mock data generated by starting from physically inspired models, are presented in Sect. 6. In Sect. 7 we present our results in the case of lattice QCD data and the comparison with the HLT method. We draw our conclusions in Sect. 8 and discuss some technical details in the two appendices.

2 Theoretical framework

The problem that we want to solve is the extraction of the infinite-volume spectral density $\rho(E)$ from the Euclidean

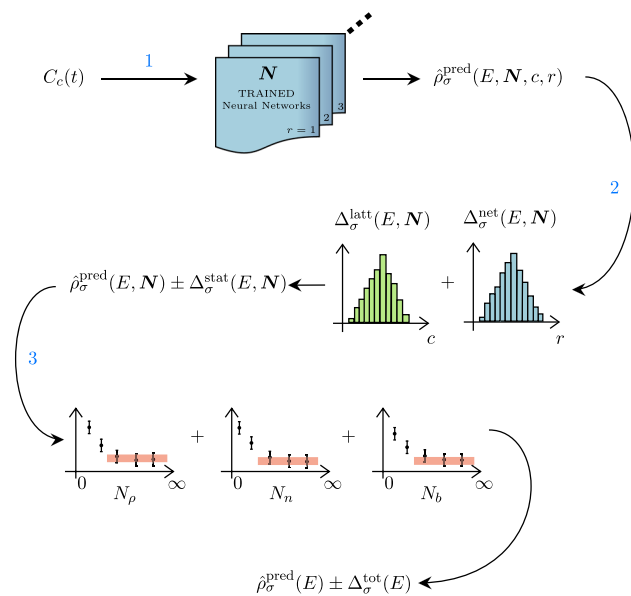


Fig. 3 Flowchart illustrating the three-step procedure that, after the training, we use to extract the final result. Here $C(t)$ represents the input lattice correlator that, coming from a Monte Carlo simulation, is affected by statistical noise. We call $C_c(t)$ the c -th bootstrap sample (or jackknife bin) of the lattice correlator with $c = 1, \dots, N_c$. In the first step, $C_c(t)$ is fed to all the trained neural networks belonging to the ensemble at fixed \mathbf{N} and the corresponding answers $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, c, r)$ are collected. In the second step, by combining in quadrature the widths of the distributions of the answers as a function of the index c ($\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$) and of the index r ($\Delta_\sigma^{\text{net}}(E, \mathbf{N})$) we estimate the error $\Delta_\sigma^{\text{stat}}(E, \mathbf{N})$ at fixed \mathbf{N} . At the end of this step we are left with a collection of results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}) \pm \Delta_\sigma^{\text{stat}}(E, \mathbf{N})$. In the third and last step, the limits $\mathbf{N} \mapsto \infty$ are studied numerically and an unbiased estimate of $\hat{\rho}_\sigma^{\text{pred}}(E)$ and of its error $\Delta_\sigma^{\text{tot}}(E)$, with the latter also taking into account the unavoidable systematics associated with these limits, is finally obtained

correlators

$$C_{LT}(t) = \int_{E_0}^{\infty} dE b_T(t, E) \rho_{LT}(E), \tag{1}$$

computed numerically in lattice simulations. These are performed by introducing a finite spatial volume L^3 , a finite Euclidean temporal direction T and by discretizing space-time,

$$x = a(\tau, \mathbf{n}), \quad 0 \leq \tau < N_T = \frac{T}{a}, \quad 0 \leq n_i < N_L = \frac{L}{a}, \tag{2}$$

where a is the so-called lattice spacing and (τ, \mathbf{n}) are the integer space-time coordinates in units of a . In order to obtain the infinite-volume spectral density one has to perform different lattice simulations, by progressively increasing L and/or T , and then study numerically the $L \mapsto \infty$ and $T \mapsto \infty$ limits. In the $T \mapsto \infty$ limit the basis functions $b_T(t, E)$ become

decaying exponentials and the correlator is given by

$$C_L(t) = \lim_{T \mapsto \infty} C_{TL}(t) = \int_{E_0}^{\infty} dE e^{-tE} \rho_L(E), \tag{3}$$

where $\rho_L(E) = 0$ for $E < E_0$ and the problem is that of performing a numerical inverse Laplace-transform operation, by starting from a discrete and finite set of noisy input data. This is a classical numerical problem, arising in many research fields, and has been thoroughly studied (see e.g. Refs. [37,38] for textbooks discussing the subject from a machine-learning perspective). The problem, as we are now going to discuss, is particularly challenging from the numerical point of view and becomes even more challenging and delicate in our Quantum Field Theory (QFT) context because, according to Wightman’s axioms, QFT spectral densities live in the space of tempered distributions and, therefore, cannot in general be considered smooth and well behaved functions.

Before discussing though the aspects of the problem that are peculiar to our QFT context, it is instructive to first review the general aspects of the numerical inverse Laplace-transform problem that, in fact, is ill-posed in the sense of Hadamard. To this end, we start by considering the correlator in the infinite L and T limits,

$$C(a\tau) = \lim_{L, T \mapsto \infty} C_{TL}(a\tau) = \int_{E_0}^{\infty} dE e^{-\tau a E} \rho(E), \tag{4}$$

and we assume that our knowledge of $C(t)$ is limited to the discrete and finite set of times $t = a\tau$. Moreover we assume that the input data are affected by numerical and/or statistical errors that we call $\Delta(a\tau)$.

The main point to be noticed is that, in general, the spectral density $\rho(E)$ contains an infinite amount of information that cannot be extracted from the limited and noisy information contained into the input data $C(a\tau)$. As a consequence, in any numerical approach to the extraction of $\rho(E)$ a discretization of Eq. (4) has to be introduced. Once a strategy to discretize the problem has been implemented, the resulting spectral density has then to be interpreted as a “filtered” or (as is more natural to call it in our context) *smearred* version of the exact spectral density,

$$\hat{\rho}(E) = \int_{E_0}^{\infty} d\omega K(E, \omega) \rho(\omega), \tag{5}$$

where $K(E, \omega)$ is the so-called smearing kernel, explicitly or implicitly introduced within the chosen numerical approach.

There are two main strategies (and many variants of them) to discretize Eq. (4). The one that we will adopt in this paper has been introduced and pioneered by Backus and Gilbert [39] and is based on the introduction of a smearing kernel in the first place. We will call this the “smearing” discretization approach. The other approach, that is more frequently

used in the literature and that we will therefore call the “standard” one, is built on the assumption that spectral densities are smooth and well-behaved functions. Before discussing the smearing approach we briefly review the standard one, by putting the emphasis on the fact that also in this case a smearing kernel is implicitly introduced.

In the standard discretization approach the infinite-volume correlator is approximated as a Riemann sum,

$$\hat{C}(a\tau) = \sigma \sum_{m=0}^{N_E-1} e^{-\tau a E_m} \hat{\rho}(E_m), \quad E_m = E_0 + m\sigma, \quad (6)$$

under the assumption that the infinite-volume spectral density is sufficiently regular to have

$$|C(a\tau) - \hat{C}(a\tau)| \ll \Delta(a\tau) \quad (7)$$

for N_E sufficiently large and σ sufficiently small. By introducing the “veilbein matrix”

$$\hat{\mathcal{E}}_{\tau m} \equiv e^{-\tau a E_m}, \quad (8)$$

and the associated “metric matrix” in energy space,

$$\begin{aligned} \hat{G}_{nm} &\equiv \left[\hat{\mathcal{E}}^T \hat{\mathcal{E}} \right]_{nm} \\ &= \sum_{\tau=1}^{N_T-1} e^{-\tau a(E_n+E_m)} = \frac{e^{-a(E_n+E_m)} - e^{-T(E_n+E_m)}}{1 - e^{-a(E_n+E_m)}}, \end{aligned} \quad (9)$$

Eq. (6) is then solved,

$$\begin{aligned} \hat{\rho}(E_n) &= \sum_{\tau=1}^{N_T-1} g_{\tau}(E_n) \hat{C}(a\tau), \\ g_{\tau}(E_n) &= \frac{1}{\sigma} \sum_{m=0}^{N_E-1} \hat{G}_{nm}^{-1} \hat{\mathcal{E}}_{\tau m}. \end{aligned} \quad (10)$$

By using the previous expressions we can now explain why the problem is particularly challenging and in which sense it is numerically ill-posed.

On the numerical side, the metric matrix \hat{G} is very badly conditioned in the limit of large N_E and small σ . Consequently, the coefficients $g_{\tau}(E_n)$ become huge in magnitude and oscillating in sign in this limit and even a tiny distortion of the input data gets enormously amplified,

$$\sum_{\tau=1}^{N_T-1} g_{\tau}(E_n) \Delta(a\tau) \xrightarrow{\sigma \rightarrow 0} \infty. \quad (11)$$

In this sense the numerical solution becomes unstable and the problem ill-posed.

Another important observation concerning Eqs. (10), usually left implicit, concerns the interpretation of $\hat{\rho}(E_n)$ as a smeared spectral density. By introducing the smearing kernel

$$K(E_n, \omega) = \sum_{\tau=1}^{N_T-1} g_{\tau}(E_n) e^{-a\tau\omega}, \quad (12)$$

and by noticing that, as a matter of fact, the spectral density has to be obtained by using the correlator $C(a\tau)$ (and not its approximation $\hat{C}(a\tau)$, to be considered just as a theoretical device introduced in order to formalize the problem), we have

$$\hat{\rho}(E_n) = \int_{E_0}^{\infty} d\omega K(E_n, \omega) \rho(\omega). \quad (13)$$

Consistency would require that $K(E_n, \omega) = \delta(E_n - \omega)$ but this cannot happen at finite T and/or N_E . In fact $K(E_n, \omega)$ can be considered a numerical approximation of $\delta(E_n - \omega)$ that, as can easily be understood by noticing that

$$K(E_n, E_m) = \frac{\delta_{nm}}{\sigma}, \quad (14)$$

has an intrinsic energy-resolution proportional to the discretization interval σ of the Riemann’s sum. A numerical study of $K(E_n, \omega)$ at fixed E_n reveals that for $\omega \geq E_n$ the kernel behaves smoothly while it oscillates wildly for $\omega < E_n$ and small values of σ . A numerical example is provided in Fig. 4.

Once the fact that a smearing operation is unavoidable in any numerical approach to the inverse Laplace-transform problem has been recognized, the smearing discretization approach that we are now going to discuss appears more natural. Indeed, the starting point of the smearing approach is precisely the introduction of a smearing kernel and this allows to cope with the problem also in the case, relevant for our QFT applications, in which spectral densities are distributions.

By reversing the logic of the standard approach, that led us to Eq. (12), in the smearing approach the problem is discretized by representing the possible smearing kernels as functionals of the coefficients g_{τ} according to

$$K(\mathbf{g}, \omega) = \sum_{\tau=1}^{N_T-1} g_{\tau} b_T(a\tau, \omega). \quad (15)$$

Notice that we are now considering the correlator $C_{LT}(a\tau)$ at finite T and L and this allows to analyse the results of a single lattice simulation. The main observation is that in the $T \mapsto \infty$ limit any target kernel $K_{\text{target}}(E, \omega)$ such that

$$\|K_{\text{target}}(E)\|^2 = \int_{E_0}^{\infty} d\omega |K_{\text{target}}(E, \omega)|^2 < \infty, \quad (16)$$

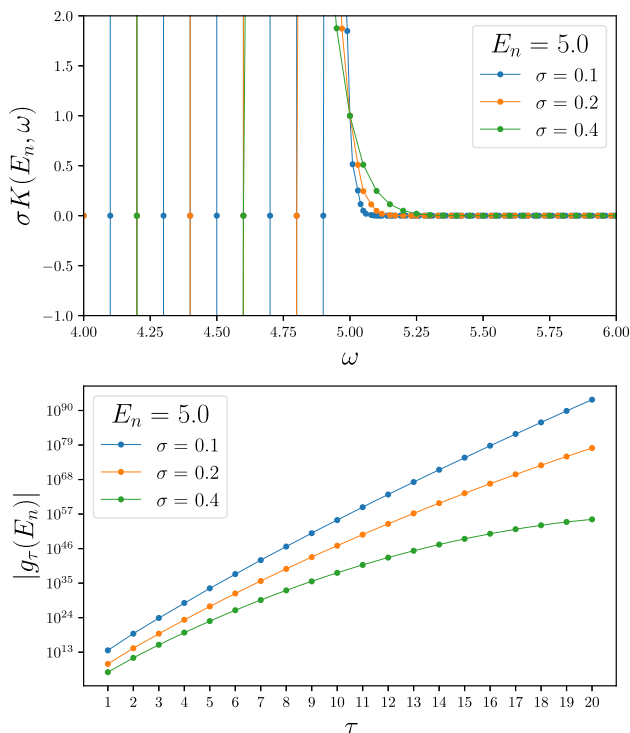


Fig. 4 Top panel: Smearing kernel of Eq. (12) at $E_n = 5$ for three values of σ . The reconstruction is performed by setting $E_0 = 1$ and $E_{\max} = 10$ as UV cutoff in Eq. (5) and by working in lattice units with $a = 1$. The kernel is smooth for $\omega \geq E_n$ while it presents huge oscillations for $\omega < E_n$. In the case $\sigma = 0.1$ these oscillations range from -10^{126} to $+10^{121}$. Bottom panel: Corresponding coefficients $g_\tau(E_n)$ (see Eq. 10) in absolute value for the first 20 discrete times. Extended-precision arithmetic is mandatory to invert the matrix \hat{G}_{nm} . This test, in which $\det(\hat{G}) = \mathcal{O}(10^{-15700})$ for $\sigma = 0.1$, has been performed by using 600-digits arithmetic

can exactly be represented as a linear combination of the $b_\infty(a\tau, \omega) = \exp(-a\omega\tau)$ basis functions (that are indeed dense in the functional-space $L_2[E_0, \infty]$ of square-integrable functions). With a finite number of lattice times, the smearing kernel is defined as the best possible approximation of $K_{\text{target}}(E, \omega)$, i.e. the coefficients $g_\tau(E)$ are determined by the condition

$$\left. \frac{\partial \|K(\mathbf{g}) - K_{\text{target}}(E)\|^2}{\partial g_\tau} \right|_{g_\tau = g_\tau(E)} = 0. \tag{17}$$

Once the coefficients are given, the smeared spectral density is readily computed by relying on the linearity of the problem,

$$\begin{aligned} \hat{\rho}_{LT}(E) &= \sum_{\tau=1}^{N_T-1} g_\tau(E) C_{LT}(a\tau) \\ &= \int_{E_0}^\infty d\omega K(E, \omega) \rho_{LT}(\omega), \end{aligned} \tag{18}$$

where now

$$K(E, \omega) \equiv K(\mathbf{g}(E), \omega). \tag{19}$$

In the smearing approach one has

$$\lim_{T \rightarrow \infty} \hat{\rho}_{LT}(E) = \int_{E_0}^\infty d\omega K_{\text{target}}(E, \omega) \rho_L(E). \tag{20}$$

Moreover, provided that $\rho(E)$ exists, it can be obtained by choosing as the target kernel a smooth approximation to $\delta(E - \omega)$ depending upon a resolution parameter σ , e.g.

$$\begin{aligned} K_{\text{target}}(E, \omega) &= \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(E-\omega)^2}{2\sigma^2}}, \\ \lim_{\sigma \rightarrow 0} K_{\text{target}}(E, \omega) &= \delta(E - \omega), \end{aligned} \tag{21}$$

and by considering the limits

$$\lim_{\sigma \rightarrow 0} \lim_{L, T \rightarrow \infty} \hat{\rho}_{LT}(E) = \rho(E), \tag{22}$$

in the specified order. When σ is very small, the coefficients $g_\tau(E)$ determined by using Eq. (17) become gigantic in magnitude and oscillating in sign, as in the case of the coefficients obtained by using Eq. (10). In fact the numerical problem is ill-posed independently of the approach used to discretize it.

We now come to the aspects of the problem that are peculiar to our QFT context. Hadronic spectral densities

$$\begin{aligned} \rho(p_1, \dots, p_{n-1}) &= \langle 0 | \hat{O}_1(2\pi)^3 \delta^4(\hat{P} - p_1) \hat{O}_2 \\ &\dots (2\pi)^3 \delta^4(\hat{P} - p_{n-1}) \hat{O}_n | 0 \rangle, \end{aligned} \tag{23}$$

are the Fourier transforms of Wightman’s functions in Minkowski space,

$$W(x_1, \dots, x_{n-1}) = \langle 0 | \hat{O}_1 e^{-i\hat{P}\cdot x_1} \hat{O}_2 \dots e^{-i\hat{P}\cdot x_{n-1}} \hat{O}_n | 0 \rangle, \tag{24}$$

where $\hat{P} = (\hat{H}, \hat{\mathbf{P}})$ is the QCD four-momentum operator and the \hat{O}_i ’s are generic hadronic operators. According to Wightman’s axioms, vacuum expectation values of field operators are tempered distributions. This implies that also their Fourier transforms, the spectral densities, are tempered distributions in energy–momentum space. It is therefore impossible, in general, to compute numerically a spectral density. On the other hand, it is always possible to compute *smeared* spectral densities,

$$\rho[K_1, \dots, K_{n-1}] = \int \left\{ \prod_i d^4 p_i K_i(p_i) \right\} \rho(p_1, \dots, p_{n-1}), \tag{25}$$

where the kernels $K_i(p)$ are Schwartz functions.

In this paper, to simplify the discussion and the notation, we focus on the dependence upon a single space-time variable,

$$W(x) = \langle 0 | \hat{O}_F e^{-i\hat{P}\cdot x} \hat{O}_I | 0 \rangle, \quad (26)$$

where the operators \hat{O}_I and \hat{O}_F might also depend upon other coordinates. The spectral density associated with $W(x)$ is given by

$$\begin{aligned} \rho(E) &= \frac{1}{2\pi} \int d^4x e^{ip\cdot x} W(x) \\ &= \langle 0 | \hat{O}_F \delta(\hat{H} - E) (2\pi)^3 \delta^3(\hat{\mathbf{P}} - \mathbf{p}) \hat{O}_I | 0 \rangle, \end{aligned} \quad (27)$$

and, to further simplify the notation, we don't show explicitly the dependence of $\rho(E)$ w.r.t. the fixed spatial momentum \mathbf{p} . The very same spectral density appears in the two-point Euclidean correlator

$$\begin{aligned} C(t) &= \int d^3x e^{-i\mathbf{p}\cdot\mathbf{x}} \langle 0 | \hat{O}_F e^{-t\hat{H} + i\hat{\mathbf{P}}\cdot\mathbf{x}} \hat{O}_I | 0 \rangle \\ &= \int_{E_0}^{\infty} dE e^{-tE} \rho(E) \end{aligned} \quad (28)$$

for positive Euclidean time t . In the last line of the previous equation we have used the fact that, because of the presence of the energy Dirac-delta function appearing in Eq. (27), the spectral density vanishes for $E < E_0$ where E_0 is the smallest energy that a state propagating between the two hadronic operators \hat{O}_I and \hat{O}_F can have.

On a finite volume the spectrum of the Hamiltonian is discrete and, consequently, the finite-volume spectral density $\rho_L(E)$ becomes a particularly wild distribution,

$$\rho_L(E) = \sum_n w_n(L) \delta(E_n(L) - E), \quad (29)$$

the sum of isolated Dirac-delta peaks in correspondence of the eigenvalues $E_n(L)$ of the finite volume Hamiltonian. By using the previous expression one has that

$$C_L(t) = \sum_n w_n(L) e^{-E_n(L)t}, \quad (30)$$

and this explains why we have discussed the standard approach to the discretization of the inverse Laplace-transform problem by first taking the infinite-volume limit. Indeed if the Riemann's discretization of Eq. (6) is applied to $C_L(a\tau)$ and all the E_m 's are different from all the $E_n(L)$'s one gets $\hat{C}_L(a\tau) = 0!$

On the one hand, also in infinite volume, spectral densities have to be handled with the care required to cope with tempered distributions and this *excludes* the option of using the

standard approach to discretize the problem in the first place. On the other hand, in some specific cases it might be conceivable to assume that the infinite volume spectral density is sufficiently smooth to attempt using the standard discretization approach. This requires that the infinite-volume limit of the correlator has been numerically taken and that the associated systematic uncertainty has been properly quantified (a non-trivial numerical task at small Euclidean times where the errors on the lattice correlators are tiny).

The smearing discretization approach can be used either in finite- and infinite-volume and the associated systematic errors can reliably be quantified by studying numerically the limits of Eq. (22). For this reason, as in the case of the HLT method of Ref. [14], the target of the machine-learning method that we are now going to discuss in details is the extraction of *smear*ed spectral densities.

3 Numerical setup

In order to implement the strategy sketched in Figs. 1, 2 and 3 and discussed in details in the following sections, the numerical setup needs to be specified. In this section we describe the data layout that we used to represent the input correlators and the output smeared spectral densities and then provide the details of the architectures of the neural networks that we used in our study.

3.1 Data layout

In this work we have considered both mock and real lattice QCD data and have chosen the algorithmic parameters in order to be able to extract the hadronic spectral density from the lattice QCD correlator described in Sect. 7. Since in the case of the lattice QCD correlator the basis function $b_T(t, E)$ (see Eq. (1)) is given by

$$b_T(t, \omega) = \frac{\omega^2}{12\pi^2} \left[e^{-t\omega} + e^{-(T-t)\omega} \right], \quad (31)$$

with $T = 64a$, we considered the same setup also in the case of mock data. These are built by starting from a model unsmeared spectral density $\rho(E)$ and by computing the associated correlator according to

$$C(t) = \int_{E_0}^{\infty} d\omega \frac{\omega^2}{12\pi^2} \left[e^{-t\omega} + e^{-(T-t)\omega} \right] \rho(\omega), \quad (32)$$

and the associated smeared spectral density according to

$$\hat{\rho}_\sigma(E) = \int_{E_0}^{\infty} d\omega K_\sigma(E, \omega) \rho(\omega). \quad (33)$$

Since in the case of the lattice QCD spectral density, in order to compare the results obtained with the proposed new method with the ones previously obtained in Ref. [40], we considered a Gaussian smearing kernel, also in the case of mock data we made the choice

$$K_\sigma(E, \omega) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(E-\omega)^2}{2\sigma^2}}. \quad (34)$$

We stress that there is no algorithmic reason behind the choice of the Gaussian as the smearing kernel and that any other kernel can easily be implemented within the proposed strategy.

Among the many computational paradigms available within the machine-learning framework, we opted for the most direct one and represented both the correlator and the smeared spectral density as finite dimensional vectors, that we used respectively as input and output of our neural networks. More precisely, the dimension of the input correlator vector has been fixed to $N_T = 64$, coinciding with the available number of Euclidean lattice times in the case of the lattice QCD correlator. The inputs of the neural networks are thus the 64-components vectors $\mathbf{C} = \{C(a), C(2a), \dots, C(64a)\}$. The output vectors are instead given by $\hat{\rho}_\sigma = \{\hat{\rho}_\sigma(E_{\min}), \dots, \hat{\rho}_\sigma(E_{\max})\}$. As in Ref. [40], we have chosen to measure energies in units of the muon mass, $m_\mu = 0.10566$ GeV, and set $E_{\min} = m_\mu$ and $E_{\max} = 24m_\mu$. The interval $[E_{\min}, E_{\max}]$ has been discretized in steps of size $m_\mu/2$. With these choices our output smeared spectral densities are the vectors $\hat{\rho}_\sigma$ with $N_E = 47$ elements corresponding to energies ranging from about 100 MeV to 2.5 GeV.

The noise-to-signal ratio in (generic) lattice QCD correlators increases exponentially at large Euclidean times. For this reason it might be numerically convenient to choose $N_T < T/a$ and discard the correlator at large times where the noise-to-signal ratio is bigger than one. According to our experience with the HLT method, that inherits the original Backus–Gilbert regularization mechanism, it is numerically inconvenient to discard part of the information available on the correlator provided that the information contained in the noise is used to conveniently regularize the numerical problem. Also in this new method, as we are going to explain in Sect. 4.2, we use the information contained in the noise of the lattice correlator during the training process and this, as shown in Appendix B, allows us to conveniently use all the available information on the lattice correlator, i.e. to set $N_T = T/a$, in order to extract the smeared spectral density.

We treat σ , the width of the smearing Gaussian, as a fixed parameter by including in the corresponding training sets only spectral functions that are smeared with the chosen value of σ . This is a rather demanding numerical strategy because in order to change σ the neural networks have to be trained anew, by replacing the smeared spectral densities in

the training sets with those corresponding to the new value of σ . Architectures that give the possibility to take into account a variable input parameter, and a corresponding variable output at fixed input vector, have been extensively studied in the machine learning literature and we leave a numerical investigation of this option to future work on the subject. In this work we considered two different values, $\sigma = 0.44$ GeV and $\sigma = 0.63$ GeV, that correspond respectively to the smallest and largest values used in Ref. [40].

3.2 Architectures

By reading the discussion on the data layout presented in the previous subsection from the machine-learning point of view, we are in fact implementing neural networks to solve a $\mathbb{R}^{N_T} \mapsto \mathbb{R}^{N_E}$ regression problem with $N_T = 64$ and $N_E = 47$ which, from now on, fix the dimension of the input and output layers of the neural networks.

There are no general rules to prefer a given network architecture among the different possibilities that have been considered within the machine-learning literature and it is common practice to make the choice by taking into account the details of the problem at hand. For our analysis, after having performed a comparative study at (almost) fixed number of parameters of the so-called Multilayer perceptron and convolutional neural networks, we used feed-forward convolutional neural networks based on the LeNet architecture introduced in Ref. [41].

We studied in details the dependence of the output of the neural networks upon their size N_n and, to this end, we implemented three architectures that we called *arcS*, *arcM* and *arcL*. These architectures, described in full details in Tables 1, 2 and 3, differ only for the number of maps in the convolutional layers. The number of maps are chosen so that the number of parameters of *arcS:arcM:arcL* are approximately in the proportion 1 : 2 : 3. For the implementation and the training we employed both Keras [42] and TensorFlow [43].

4 Model independent training

In the supervised deep-learning framework a neural network is trained over a training set which is representative of the problem that has to be solved. In our case the inputs to each neural network are the correlators \mathbf{C} and the target outputs are the associated smeared spectral densities $\hat{\rho}_\sigma$. As discussed in the Introduction, our main goal is to devise a model-independent training strategy. To this end, the challenge is that of building a training set which contains enough variability so that, once trained, the network is able to provide the correct answer, within the quoted errors, for any possible input correlator.

Table 1 *arcS*: the smallest neural network architecture used in this work. The architecture is of the type feed-forward and the structure can be read from top to bottom of the table. It consist of three 1D convolutional layers with an increasing number of maps followed by two fully connected layers. The two blocks are intermediated by one flatten layer. The column denoted by “Size” reports the shape of the signal produced by the corresponding layer. The stride of the filters is set to 2 in such a way that the dimension of the signal is halved at each 1D convolutional layer thus favouring the neural network to learn a more abstract, and possibly more effective, representation of the input data. As activation functions we use the LeakyReLU with negative slope coefficient set to -0.2 . The neurons with activation functions are also provided with biases. The output is devoid of activation function in order not to limit the output range. The bottom line reports the total number of trainable parameters

Type	Maps	Size	Kernel size	Stride	Activation
Input		64			
Conv1D	2	32x2	3	2	LeakyReLU
Conv1D	4	16x4	3	2	LeakyReLU
Conv1D	8	8x8	3	2	LeakyReLU
Flatten		384			
Fully conn		256			LeakyReLU
Fully conn		256			LeakyReLU
Output		47			
Parameters	94651				

As a matter of fact, the situation in which the neural network can exactly reconstruct any possible function is merely ideal. That would be possible only in absence of errors on the input data and with a neural network with an infinite number of neurons, trained on an infinitely large and complex training set. This is obviously impossible and in fact our goal is the realistic task of getting an output for the smeared spectral density as close as possible to the exact one by trading the unavoidable limited abilities of the neural network with a reliable estimate of the systematic error. In order to face this challenge we used the algorithmic strategy described in Figs. 1, 2 and in Fig. 3. In our strategy,

- the fact that the network cannot be infinitely large is parametrized by the fact that N_n (the number of neurons) is finite;
- the fact that during the training a network cannot be exposed to any possible spectral density is parametrized by the fact that N_b (the number of basis functions) and N_ρ (the number of spectral densities to which a network is exposed during the training) are finite (see Fig. 1);
- the fact that at fixed

$$\mathbf{N} = (N_n, N_b, N_\rho) \tag{35}$$

the answer of a network cannot be exact, and therefore has to be associated with an error, is taken into account by

Table 2 *arcM*: the medium-size architecture used in this work. See Table 1 for the description

Type	Maps	Size	Kernel size	Stride	Activation
Input		64			
Conv1D	12	32 × 12	3	2	LeakyReLU
Conv1D	24	16 × 24	3	2	LeakyReLU
Conv1D	48	8 × 48	3	2	LeakyReLU
Flatten		384			
Fully conn		256			LeakyReLU
Fully conn		256			LeakyReLU
Output		47			
Parameters	180,871				

Table 3 *arcL*: the largest architecture used in this work. See Table 1 for the description

Type	Maps	Size	Kernel size	Stride	Activation
Input		64			
Conv1D	32	32 × 32	3	2	LeakyReLU
Conv1D	64	16 × 64	3	2	LeakyReLU
Conv1D	128	8 × 128	3	2	LeakyReLU
Flatten		1024			
Fully conn		256			LeakyReLU
Fully conn		256			LeakyReLU
Output		47			
Parameters	371,311				

introducing an ensemble of machines, with N_r replicas, and by estimating this error by studying the distribution of the different N_r answers in the $N_r \mapsto \infty$ limit (see Fig. 2);

- once the network (and statistical) errors at fixed \mathbf{N} are given, we can study numerically the $\mathbf{N} \mapsto \infty$ limits and also quantify, reliably, the additional systematic errors associated with these unavoidable extrapolations (see Fig. 3).

We are now going to provide the details concerning the choice of the functional basis that we used to parametrize the spectral densities and to build our training sets.

4.1 The functional-basis

In our strategy we envisage studying numerically the limit $N_b \mapsto \infty$ and, therefore, provided that the systematic errors associated with this extrapolation are properly taken into account, there is no reason to prefer a particular basis w.r.t. any other. For our numerical study we used the Chebyshev polynomials of the first kind as basis functions (see for example Ref. [44]).

The Chebyshev polynomials $T_n(x)$ are defined for $x \in [-1, 1]$ and satisfy the orthogonality relations

$$\int_{-1}^{+1} dx \frac{T_n(x)T_m(x)}{\sqrt{1-x^2}} = \begin{cases} 0 & n \neq m, \\ \frac{\pi}{2} & n = m \neq 0, \\ \pi & n = m = 0 \end{cases} . \tag{36}$$

In order to use them as a basis for the spectral densities that live in the energy domain $E \in [E_0, \infty)$, we introduced the exponential map

$$x(E) = 1 - 2e^{-E} \tag{37}$$

and set

$$B_n(E) = T_n(x(E)) - T_n(x(E_0)) . \tag{38}$$

Notice that the subtraction of the constant term $T_n(x(E_0))$ has been introduced in order to be able to cope with the fact that hadronic spectral densities vanish below a threshold energy $E_0 \geq 0$ that we consider an unknown of the problem. With this choice, the unsmeared spectral densities that we use to build our training sets are written as

$$\rho(E; N_b) = \theta(E - E_0) \sum_{n=0}^{N_b} c_n [T_n(x(E)) - T_n(x(E_0))] , \tag{39}$$




Fig. 5 Examples of unsmeared spectral densities generated according to Eq. (39) using the Chebyshev polynomial as basis functions for different N_b . For all the examples we set $E_0 = 0.8$ GeV. As it can be seen, by generating the c_n coefficients according to Eq. (40), the larger the number of terms of the Chebyshev series the richer is the behaviour of the spectral density in terms of local structures

and vanish identically for $E \leq E_0$. Once E_0 and the coefficients c_n that define $\rho(E; N_b)$ are given, the correlator and the smeared spectral density associated with $\rho(E; N_b)$ can be calculated by using Eqs. (32) and (33).²

Each $\rho(E; N_b)$ that we used in order to populate our training sets has been obtained by choosing E_0 randomly in the interval $[0.2, 1.3]$ GeV with a uniform distribution and by inserting in Eq. (39) the coefficients

$$c_0 = r_0 ; \quad c_n = \frac{r_n}{n^{1+\varepsilon}} , \quad n > 0 , \tag{40}$$

where the r_n 's are N_b uniformly distributed random numbers in the interval $[-1, 1]$ and ε is a non-negative parameter that we set to 10^{-7} . Notice that with this choice of the coefficients c_n the Chebyshev series of Eq. (39) is convergent in the $N_b \mapsto \infty$ limit and that the resulting spectral densities can be negative and in general change sign several times in the interval $E \in [0, \infty)$. Notice also that up to a normalization constant, that it will turn to be irrelevant in our training strategy in which the input data are standardized as explained in Sect. 4.3, the choice of the interval $[-1, 1]$ for the r_n random numbers is general.

A few examples of unsmeared spectral densities generated according to Eq. (39) are shown in Fig. 5. As it can be seen, with the choice of the coefficients c_n of Eq. (40), the larger is the number of terms in the Chebyshev series of Eq. (39) the richer is the behaviour of the resulting unsmeared spectral densities in terms of local structures.

This is an important observation and a desired feature. Indeed, a natural concern about the choice of Chebyshev

² By representing also the smearing kernels and the basis functions on a Chebyshev basis, the orthogonality relations of Eq. (36) can conveniently be exploited to speedup this step of the numerical calculations.




Fig. 6 The black dashed curve is a generic continuous unsmeared spectral density $\rho(E)$ while the vertical dashed lines indicate the energy sampling defining the support of $\rho_\delta(\omega)$ according to Eq. (42). The blue and red curves are the smeared version of respectively $\rho(\omega)$ and $\rho_\delta(\omega)$. The green curve represents $\Sigma_{\sigma, \Delta E}(E) = \hat{\rho}_\sigma(E) - \hat{\rho}_{\delta, \sigma}(E)$. The spacing is set to $\Delta E = 0.1$ GeV. The smearing kernel is a normalized Gaussian function of width σ and as resolution width we refer to the full width at half maximum $2\sqrt{2 \ln 2} \sigma$. Top-panel: the resolution width is equal to ΔE and it is such that the single peaks are resolved. Bottom-panel: in this case $2\sqrt{2 \ln 2} \sigma \gg \Delta E$ and the two smeared functions are almost undistinguishable

polynomials as basis functions is the fact that we are thus sampling the space of regular functions while, as we pointed out in Sect. 2, on a finite volume the lattice QCD spectral density is expected to be a discrete sum of Dirac-delta peaks (see Eq. (29)). Here we are relying on the fact that the inputs of our networks will be Euclidean correlators, that in fact are smeared spectral densities, and will also be asked to provide as output the smeared spectral densities $\hat{\rho}_\sigma(E)$. This allows us to assume that, provided that the energy smearing parameter σ and N_b are sufficiently large, the networks will be exposed to sufficiently general training sets to be able to extract the correct smeared spectral densities within the quoted errors. To illustrate this point we consider in Fig. 6 a continuous spectral density $\rho(E)$ and approximate its smeared version through a Riemann's sum according to

$$\hat{\rho}_\sigma(E) = \int_{E_0}^{\infty} d\omega K_\sigma(E, \omega) \rho(\omega)$$

$$\begin{aligned} &= \Delta E \sum_{n=0}^{\infty} K_\sigma(E, \omega_n) \rho(\omega_n) + \Sigma_{\sigma, \Delta E}(E) \\ &= \int_{E_0}^{\infty} d\omega K_\sigma(E, \omega) \rho_\delta(\omega) + \Sigma_{\sigma, \Delta E}(E), \end{aligned} \quad (41)$$

where

$$\rho_\delta(\omega) \equiv \Delta E \sum_{n=0}^{\infty} \rho(\omega) \delta(\omega - \omega_n), \quad \omega_n = E_0 + n \Delta E. \quad (42)$$

The previous few lines of algebra show that the smearing of a continuous function $\rho(\omega)$ can be written as the smearing of the distribution defined in Eq. (42), a prototype of the finite-volume distributions of Eq. (29), plus the approximation error $\Sigma_{\sigma, \Delta E}(E) = \hat{\rho}_\sigma(E) - \hat{\rho}_{\delta, \sigma}(E)$. The quantity $\Sigma_{\sigma, \Delta E}(E)$, depending upon the spacing ΔE of the Dirac-delta peaks and the smearing parameter σ , is expected to be sizeable when $\sigma \leq \Delta E$ and to become irrelevant in the limit $\sigma \gg \Delta E$. A quantitative example is provided in Fig. 6 where $\Sigma_{\sigma, \Delta E}(E)$ is shown at fixed ΔE for two values of σ . In light of this observation, corroborated by the extensive numerical analysis that we have performed at the end of the training sessions to validate our method (see Sect. 5.2 and, in particular, Fig. 16), we consider justified the choice of Chebyshev polynomials as basis functions provided that, as is the case in this work, the energy smearing parameter σ is chosen sufficiently large to not be able to resolve the discrete structure of the finite-volume spectrum.

4.2 Building the training sets

Having provided all the details concerning the functional basis that we use to parametrize the unsmeared spectral densities, we can now explain in details the procedure that we used to build our training sets.

A training set $\mathcal{T}_\sigma(N_b, N_\rho)$ contains N_ρ input–output pairs. Each pair is obtained by starting from a random unsmeared spectral density, parametrized at fixed N_b according to Eq. (39) and generated as explained in the previous subsection. Given the unsmeared spectral density $\rho(E; N_b)$, we then compute the corresponding correlator vector \mathbf{C} (by using Eq. (32)) and, for the two values of σ that we used in this study, the smeared spectral densities vectors $\hat{\rho}_\sigma$ (by using Eq. (33)). From each pair $(\mathbf{C}, \hat{\rho}_\sigma)$ we then generate an element $(\mathbf{C}_{\text{noisy}}, \hat{\rho}_\sigma)$ of the training set at fixed N_b and σ ,

$$\begin{aligned} \rho^i(E; N_b) &\mapsto (\mathbf{C}, \hat{\rho}_\sigma)^i \mapsto (\mathbf{C}_{\text{noisy}}, \hat{\rho}_\sigma)^i \in \mathcal{T}_\sigma(N_b, N_\rho), \\ i &= 1, \dots, N_\rho, \end{aligned} \quad (43)$$

that we obtain, as we are now going to explain, by distorting the correlator \mathbf{C} using the information provided by the noise of the lattice correlator \mathbf{C}_{latt} (see Sect. 7).

In order to cope with the presence of noise in the input data that have to be processed at the end of the training, it is extremely useful (if not necessary) to teach to the networks during the training to distinguish the physical content of the input data from noisy fluctuations. This is particularly important when dealing with lattice QCD correlators for which, as discussed in Sect. 3.1, the noise-to-signal ratio grows exponentially for increasing Euclidean times (see Fig. 7). A strategy to cope with this problem, commonly employed in the neural network literature, is to add Gaussian noise to the input data used in the training. There are several examples in the literature where neural networks are shown to be able to learn rather efficiently by employing this strategy of data corruption (see the already cited textbooks Refs. [37,38]). According to our experience, it is crucially important that the structure of the noise used to distort the training input data resembles as much as possible that of the true input data. In fact, it is rather difficult to model the noise structure generated in Monte Carlo simulations and, in particular, the off-diagonal elements of the covariance matrices of lattice correlators. In the light of these observations we decided to use the covariance matrix $\hat{\Sigma}_{\text{latt}}$ of the lattice correlator \mathbf{C}_{latt} that we are going to analyse in Sect. 7 to obtain $\mathbf{C}_{\text{noisy}}$ from \mathbf{C} . More precisely, given a correlator \mathbf{C} , we generate $\mathbf{C}_{\text{noisy}}$ by extracting a random correlator vector from the multivariate Gaussian distribution

$$\mathbb{G} \left[\mathbf{C}, \left(\frac{C(a)}{C_{\text{latt}}(a)} \right)^2 \hat{\Sigma}_{\text{latt}} \right] \tag{44}$$

having \mathbf{C} as mean vector and as covariance the matrix $\hat{\Sigma}_{\text{latt}}$ normalized by the factor $\left(\frac{C(a)}{C_{\text{latt}}(a)} \right)^2$.

In order to be able to perform a numerical study of the limits $\mathbf{N} \mapsto \infty$, we have generated with the procedure just described several training sets, corresponding to $N_b = \{16, 32, 128, 512\}$ and $N_\rho = \{50, 100, 200, 400, 800\} \times 10^3$. At fixed N_b , each training set includes the smaller ones, i.e. the training set $\mathcal{T}_\sigma(N_b, 100 \times 10^3)$ includes the training set $\mathcal{T}_\sigma(N_b, 50 \times 10^3)$ enlarged with 50×10^3 new samples.

4.3 Data pre-processing

A major impact in the machine-learning performance is played by the way the data are presented to the neural network. The learning turns to be more difficult when the input variables have different scales and this is exactly the case of Euclidean lattice correlators since the components of the input vectors decrease exponentially fast. The risk in this case is that the components small in magnitude are given less




Fig. 7 Noise-to-signal ratio of the lattice correlator $C_{\text{latt}}(t)$, discussed in Sect. 7, whose covariance matrix is used to inject the statistical noise in the training sets

importance during the training or that some of the weights of the network become very large in magnitude, thus generating instabilities in the training process. We found that a considerable improvement in the validation loss minimization is obtained by implementing the standardization of the input correlators (see next subsection and the central panel of Fig. 8) and therefore used this procedure.

The standardization procedure of the input consists in rescaling the data in such a way that all the components of the input correlator vectors in the training set are distributed with average 0 and variance 1. For a given training set $\mathcal{T}_\sigma(N_b, N_\rho)$ we calculate the N_T -component vectors $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ whose components are

$$\mu(\tau) = \frac{1}{N_\rho} \sum_{i=1}^{N_\rho} C_i(a\tau) \tag{45}$$

and

$$\gamma(\tau) = \sqrt{\frac{\sum_{i=1}^{N_\rho} (C_i(a\tau) - \mu(\tau))^2}{N_\rho}}. \tag{46}$$

Each correlator in the training is then replaced by

$$C_{\text{noisy}}(a\tau) \mapsto C'_{\text{noisy}}(a\tau) = \frac{C_{\text{noisy}}(a\tau) - \mu(\tau)}{\gamma(\tau)}. \tag{47}$$

where $\mathbf{C}_{\text{noisy}}$ is the distorted version of \mathbf{C} discussed in the previous subsection. Notice that the vectors $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ are determined from the training set before including the noise. Although the pre-processed correlators look quite different from the original ones, since the components are no longer exponentially decaying, the statistical correlation in time is preserved by the standardization procedure.

At the end of the training, the correlators fed into the neural network for validation or prediction have also to be pre-processed by using the same vectors $\boldsymbol{\mu}$ and $\boldsymbol{\gamma}$ used in the training.

4.4 Training an ensemble of machines

Given a machine with N_n neurons we train it over the training set $\mathcal{T}_\sigma(N_b, N_\rho)$ by using as loss function the Mean Absolute Error (MAE)

$$\ell(\mathbf{w}) = \frac{1}{N_\rho} \sum_{i=1}^{N_\rho} \left| \hat{\rho}_\sigma^{\text{pred},i}(\mathbf{w}) - \hat{\rho}_\sigma^i \right|, \tag{48}$$

where $\hat{\rho}_\sigma^{\text{pred},i}(\mathbf{w})$ is the output of the neural network in correspondence of the input correlator $\mathbf{C}_{\text{noisy}}^i$. In Eq. (48) we used the norm $|\hat{\rho}| = \sum_{j=1}^{N_E=47} |\hat{\rho}_j|$ and we have explicitly shown the dependence of the predicted spectral density $\hat{\rho}_\sigma^{\text{pred},i}(\mathbf{w})$ upon the weights \mathbf{w} of the network.

At the beginning of each training session each weight w_n (with $n = 1, \dots, N_n$) is extracted from a Gaussian distribution with zero mean value and variance 0.05. To end the training procedure we rely on the early stopping criterion: the training set is split into two subsets containing respectively the 80% and 20% of the entries of the training set $\mathcal{T}_\sigma(N_b, N_\rho)$. The larger subset is used to update the weights of the neural network with the gradient descent algorithm. The smaller subset is the so-called *validation set* and we use it to monitor the training process. At the end of each epoch we calculate the loss function of Eq. (48) for the validation set, the so-called validation loss, and stop the training when the drop in the validation loss is less than 10^{-5} for 15 consecutive epochs. In the trainings performed in our analysis this occurs typically between epoch 150 and 200. The early stopping criterion provides an automatic way to prevent the neural network from overfitting the input data.

We implement a Mini-Batch Gradient Descent algorithm, with Batch Size (BS) set to 32, by using the Adam optimizer [45] combined with a learning rate decaying according to

$$\eta(e) = \frac{\theta(e - 25)\eta(e - 1)}{1 + e \cdot 4 \times 10^{-4}}, \tag{49}$$

where $e \in \mathbb{N}$ is the epoch and $\eta(0) = 2 \times 10^{-4}$ is the starting value. The step-function is included so that the learning rate is unchanged during the first 25 epochs. Although a learning rate scheduler is not strictly mandatory, since the Adam optimizer already includes adaptive learning rates, we found that it provides an improvement in the convergence with less noisy fluctuations in the validation loss (see the top panel of Fig. 8). Concerning the BS, we tested the neural network performance by starting from BS = 512 and by halving it up

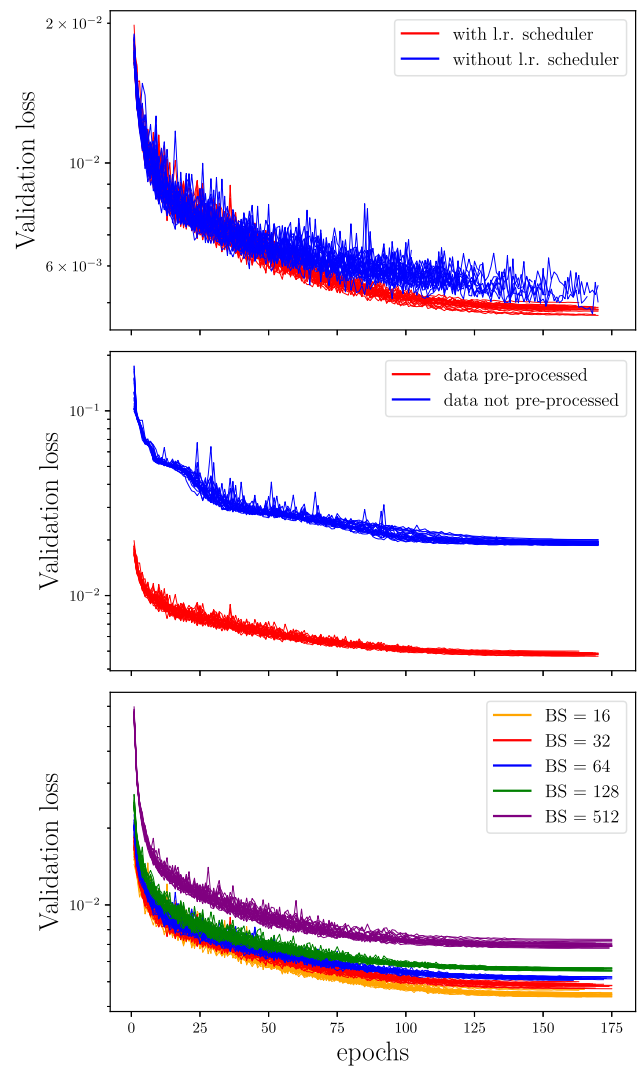


Fig. 8 Numerical tests to optimize the training performances. The validation loss functions of the trainings of the $N_f = 20$ replica machines belonging to the same ensemble are plotted with the same color. All the trainings refer to $arcL$, $N_\rho = 50 \times 10^3$, $N_b = 512$ and $\sigma = 0.44$ GeV. If not otherwise specified the data are pre-processed, BS=32 and the learning rate scheduler is implemented. Top panel: Comparison of the validation losses with and without the learning scheduler rate defined in Eq. (49). Central panel: Comparison of the validation losses by implementing or not the input data pre-processing procedure described in Sect. 4.3. Bottom panel: Comparison of the validations losses at different values of the BS parameter

to BS = 16 (see bottom panel of Fig. 9). Although the performance improves as BS decreases we set BS = 32 in order to cope with the unavoidable slowing down of the training for smaller values of BS.

As we already stressed several times, at fixed \mathbf{N} the answer of a neural network cannot be exact. In order to be able to study numerically the $\mathbf{N} \mapsto \infty$ limits, the error associated with the limited abilities of the networks at finite \mathbf{N} has to be quantified. To do this we introduce the *ensemble of machines*




Fig. 9 Validation losses at the end of the trainings as functions of the training set dimension N_ρ . The top-plot corresponds to $\sigma = 0.44$ GeV and the bottom-plot to $\sigma = 0.63$ GeV. The different colors correspond to the different architectures and different markers to the different number of basis functions N_b considered in this study. The error on each point correspond to the standard deviation of the distribution of the validation loss in an ensemble of $N_r = 20$ machines. As expected, the validation loss decreases by increasing N_ρ (more general training) and/or N_n (larger and therefore smarter networks). Moreover, at fixed N_n and N_ρ the neural network performs better at smaller values of N_b as consequence of the fact that the training set exhibits less complex features and learning them is easier

by considering at fixed \mathbf{N}

$$N_r = 20 \tag{50}$$

machines with the same architecture and trained by using the same strategy. More precisely, each machine of the ensemble is trained by using a training set $\mathcal{T}_\sigma(N_b, N_\rho)$ obtained by starting from the same unsmeared spectral densities, and therefore from the same pairs $(\mathbf{C}, \hat{\rho}_\sigma)$ (see Eq. (43)), but with different noisy input correlator vectors $\mathbf{C}_{\text{noisy}}$.

In Fig. 9 we show the validation loss as a function of N_ρ for the different values of N_b , the three different architectures and the two values of σ considered in this study. Each point in the figure has been obtained by studying the distribution of the validation loss at the end of the training within the corresponding ensemble of machines and by using respectively the mean and the standard deviation of each distribution as central value and error. The figure provides numerical evidences concerning the facts that

- networks with a finite number N_n of neurons, initialized with different weights and exposed to training sets containing a finite amount of information, provide different answers and this is a source of errors that have to be quantified;
- the validation loss decreases for larger N_n because larger networks are able to assimilate a larger amount of information;
- the validation loss decreases for larger N_ρ since, if N_n is sufficiently large, the networks learn more from larger and more general training sets;
- at fixed N_n and N_ρ the networks perform better at smaller values of N_b because the training sets exhibit less complex features and learning them is easier.

In order to populate the plots in Fig. 9 we considered several values of N_ρ and N_b and this is rather demanding from the computational point of view. The training sets that we found to be strictly necessary to quote our final results are listed in Table 4.

5 Estimating the total error and validation tests

Having explained in the previous section the procedure that we used to train our ensembles of machines, we can now explain in details the procedure that we use to obtain our results. We start by discussing the procedure that we use to estimate the total error, taking into account both statistical and systematics uncertainties, and then illustrate the validation tests that we have performed in order to assess the reliability of this procedure.

5.1 Procedure to estimate the total error

The procedure that we use to quote our results for smeared spectral densities by using the trained ensembles of machines, illustrated in Fig. 3, is the following

- given an ensemble of N_r machines trained at fixed \mathbf{N} , we feed in each machine r belonging to the ensemble all the different bootstrap samples (or jackknife bins) $C_c(a\tau)$ of the input correlator ($c = 1, \dots, N_c$) and obtain a collection of results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, c, r)$ for the smeared spectral density that depend upon \mathbf{N} , c and r ;
- at fixed \mathbf{N} and r we compute the bootstrap (or jackknife) central values $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, r)$ and statistical errors $\Delta_\sigma^{\text{latt}}(E, \mathbf{N}, r)$ and average them over the ensemble of machines,

$$\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}) = \frac{1}{N_r} \sum_{r=1}^{N_r} \hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, r),$$

$$\Delta_\sigma^{\text{latt}}(E, \mathbf{N}) = \frac{1}{N_r} \sum_{r=1}^{N_r} \Delta_\sigma^{\text{latt}}(E, \mathbf{N}, r); \tag{51}$$

- by computing the standard deviation over the ensemble of machines of $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, r)$, we obtain an estimate of the error, that we call $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$, associated with the fact that at fixed \mathbf{N} the answer of a network cannot be exact;
- both $\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$ and $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$ have a statistical origin. The former, $\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$, comes from the limited statistics of the lattice Monte Carlo simulation while the latter, $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$, comes from the statistical procedure that we used to populate our training sets $\mathcal{T}_\sigma(N_b, N_\rho)$ and to train our ensembles of machines at fixed N_b . We sum them in quadrature and obtain an estimate of the statistical error at fixed \mathbf{N} ,

$$\Delta_\sigma^{\text{stat}}(E, \mathbf{N}) = \sqrt{[\Delta_\sigma^{\text{latt}}(E, \mathbf{N})]^2 + [\Delta_\sigma^{\text{net}}(E, \mathbf{N})]^2}; \tag{52}$$

- we then study numerically the $\mathbf{N} \mapsto \infty$ limits by using a data-driven procedure. We quote as central value and statistical error of our final result

$$\begin{aligned} \hat{\rho}_\sigma^{\text{pred}}(E) &\equiv \rho_\sigma^{\text{pred}}(E, \mathbf{N}^{\text{max}}), \\ \Delta_\sigma^{\text{stat}}(E) &\equiv \Delta_\sigma(E, \mathbf{N}^{\text{max}}), \end{aligned} \tag{53}$$

where \mathbf{N}^{max} , given in Table 4, is the vector with the largest components among the vectors \mathbf{N} considered in this study;

- in order to check the numerical convergence of the $\mathbf{N} \mapsto \infty$ limits and to estimate the associated systematic uncertainties $\Delta_\sigma^X(E)$, where

$$X = \{\rho, n, b\}, \tag{54}$$

we define the reference vectors $\mathbf{N}_X^{\text{ref}}$ listed in Table 4. We then define the pull variables

$$P_\sigma^X(E) = \frac{|\hat{\rho}_\sigma^{\text{pred}}(E) - \hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}_X^{\text{ref}})|}{\sqrt{[\Delta_\sigma^{\text{stat}}(E)]^2 + [\Delta_\sigma^{\text{stat}}(E, \mathbf{N}_X^{\text{ref}})]^2}}, \tag{55}$$

and then we weight the absolute value of the difference $|\hat{\rho}_\sigma^{\text{pred}}(E) - \hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}_X^{\text{ref}})|$ with the Gaussian probability that this is not due to a statistical fluctuation,

$$\Delta_\sigma^X(E) = \left| \hat{\rho}_\sigma^{\text{pred}}(E) - \hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}_X^{\text{ref}}) \right| \operatorname{erf} \left(\frac{|P_\sigma^X(E)|}{\sqrt{2}} \right); \tag{56}$$

Table 4 List of the vectors \mathbf{N} used to estimate the systematic errors associated with the $\mathbf{N} \mapsto \infty$ limits. With these choices the errors $\Delta_\sigma^X(E)$ defined in Eq. (56) are maximised on our setup and, consequently, Eq. (57) provides a conservative estimate of the total error

Label	N_n	N_b	N_ρ
\mathbf{N}^{max}	arcL	512	800×10^3
$\mathbf{N}_b^{\text{ref}}$	arcL	16	800×10^3
$\mathbf{N}_n^{\text{ref}}$	arcS	512	800×10^3
$\mathbf{N}_\rho^{\text{ref}}$	arcL	512	50×10^3

- the total error that we associate to our final result $\hat{\rho}_\sigma^{\text{pred}}(E)$ is finally obtained according to

$$\begin{aligned} \Delta_\sigma^{\text{tot}}(E) \\ = \sqrt{[\Delta_\sigma^{\text{stat}}(E)]^2 + [\Delta_\sigma^\rho(E)]^2 + [\Delta_\sigma^b(E)]^2 + [\Delta_\sigma^n(E)]^2}. \end{aligned} \tag{57}$$

Some remarks are in order here. We don't have a theoretical understanding neither of the dependence of the statistical errors $\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$ and $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$ upon E and \mathbf{N} (see Appendix C for a numerical investigation) nor of the rates of convergence of the $\mathbf{N} \mapsto \infty$ limits. Gaining this theoretical understanding is a task that goes far beyond the scope of this paper. The procedure that we have devised to quote our results, that at first sight might appear too complicated, has therefore to be viewed as just one of the possible ways to perform conservative plateau-analyses of the $\mathbf{N} \mapsto \infty$ limits. This explains our choice of the points $\mathbf{N}_X^{\text{ref}}$ given in Table 4 that, in our numerical setup, provide the most conservative estimates of the systematic errors. The stringent validation tests that we have performed with mock data, and that we are going to discuss below, provide a quantitative evidence that the results obtained by implementing this procedure can in fact be trustworthy used in phenomenological applications.

5.2 Validation

In order to illustrate how the method described in the previous subsection can be applied in practice, we now consider an unsmeared spectral density $\rho(E)$ that has not been used in any of the trainings and that has been obtained as described in Sect. 4.1, i.e. by starting from Eq. (39), but this time with $N_b = 1024$, i.e. a number of basis functions which is twice as large as the largest N_b employed in the training sessions.

From this $\rho(E)$ we then calculate the associated correlator \mathbf{C} and the smeared spectral density corresponding to $\sigma = 0.44$ GeV. We refer to the true smeared spectral density as $\hat{\rho}_\sigma^{\text{true}}(E)$, while we call $\hat{\rho}_\sigma^{\text{pred}}(E)$ the final predicted result. Both the unsmeared spectral density $\rho(E)$ (grey dotted curve, partially visible) and the expected exact result $\hat{\rho}_\sigma^{\text{true}}(E)$ (solid black curve) are shown in the top panel of Fig. 12.

Starting from the exact correlator \mathbf{C} corresponding to $\rho(E)$ we have then generated $N_c = 800$ bootstrap samples




Fig. 10 Each row corresponds to a different energy, as written on the top of each plot, and in all cases we set $\sigma = 0.44$ GeV and $\mathbf{N} = \mathbf{N}^{\text{max}} = (\text{arL}, 512, 800 \times 10^3)$. Left panels: the x -axes corresponds to the replica index $r = 1, \dots, N_r = 20$ running over the entries of the ensemble of machines. The plotted points correspond to the results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, r) \pm \Delta_{\text{att}}^{\text{latt}}(E, \mathbf{N}, r)$ obtained by computing the bootstrap averages and errors of the $c = 1, \dots, N_c = 800$ results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, c, r)$. Right panels: distributions of the central values $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, r)$. The means of these distributions correspond to the results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N})$ that are shown in the left panels as solid blue lines. The widths correspond to the network errors $\Delta_{\text{net}}^{\text{net}}(E, \mathbf{N})$ that are represented in the left panels with two dashed blue lines. The blue bands in the left panels correspond to $\Delta_{\text{stat}}^{\text{stat}}(E, \mathbf{N})$, defined in Eq. (52)

from the distribution of Eq. (44), thus simulating the outcome of a lattice Monte Carlo simulation. The N_c samples have been fed into each trained neural network and we collected the answers $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}, c, r)$ that, as shown in Fig. 10 for a selection of the considered values of E , we have then analysed to obtain $\Delta_{\text{att}}^{\text{latt}}(E, \mathbf{N})$ and $\Delta_{\text{net}}^{\text{net}}(E, \mathbf{N})$.

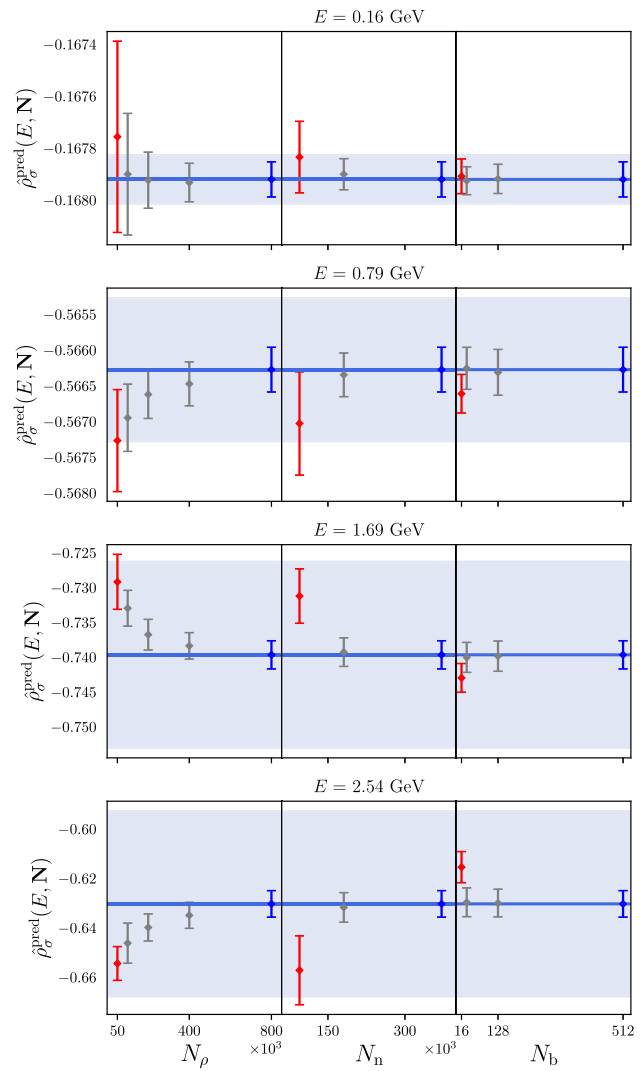


Fig. 11 Numerical study of the $\mathbf{N} \mapsto \infty$ limits. Each row corresponds to a different energy, as written on the top of each plot, and all cases we set $\sigma = 0.44$ GeV. Left panels: Study of the limit $N_\rho \mapsto \infty$ at fixed $N_b = 512$ and $N_n = \text{arL}$. Central panels: Study of the limit $N_n \mapsto \infty$ at fixed $N_b = 512$ and $N_\rho = 800 \times 10^3$. Right panels: Study of the limit $N_b \mapsto \infty$ at fixed $N_n = \text{arL}$ and $N_\rho = 800 \times 10^3$. All panels: The horizontal blue line is the final central value, corresponding to $\mathbf{N} = \mathbf{N}^{\text{max}} = (\text{arL}, 512, 800 \times 10^3)$ (blue points) and it is common to all the panels within the same row. The blue band represents instead $\Delta_{\text{tot}}^{\text{tot}}(E)$ calculated by the combination in quadrature of all the errors (see Eq. 57). The red points correspond to $\mathbf{N}_X^{\text{ref}}$ and the grey ones to the other trainings, see Table 4

The next step is now the numerical study of the limits $\mathbf{N} \mapsto \infty$ that we illustrate in Fig. 11 for the same values of E considered in Fig. 10. The limit $N_\rho \mapsto \infty$ (left panels) is done at fixed $N_b = 512$ and $N_n = \text{arL}$. The limit $N_n \mapsto \infty$ (central panels) is done at fixed $N_b = 512$ and $N_\rho = 800 \times 10^3$. The limit $N_b \mapsto \infty$ (right panels) is done at fixed $N_n = \text{arL}$ and $N_\rho = 800 \times 10^3$. As it can be seen, the blue points, corresponding to our results $\hat{\rho}_\sigma^{\text{pred}}(E) \pm \Delta_{\text{stat}}^{\text{stat}}(E)$, are always statistically compatible with the first grey point on the

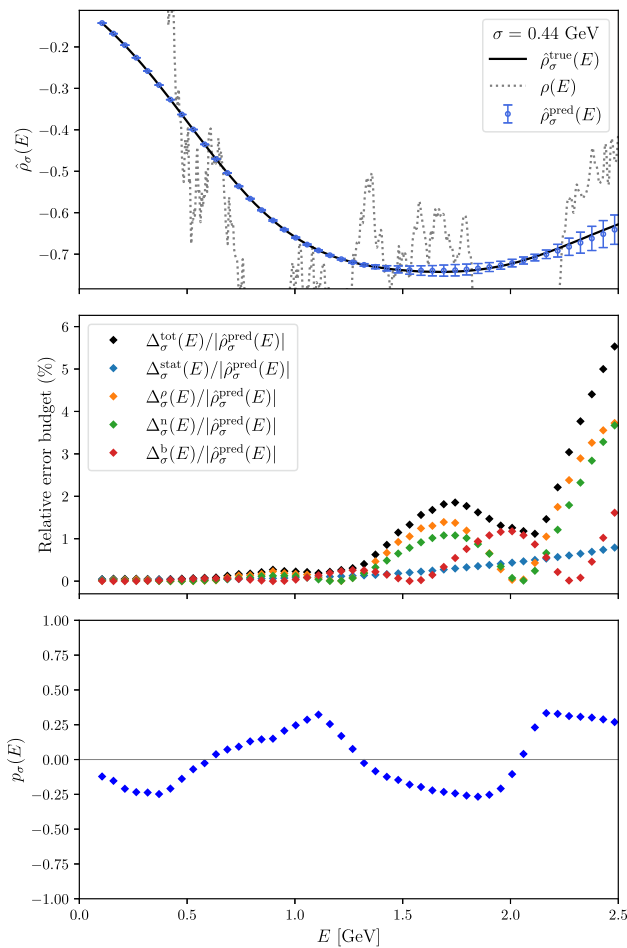


Fig. 12 Final results for $\sigma = 0.44$ GeV. Top panel: The final results predicted by the neural networks (blue points) is compared with the true result (solid black). The dotted grey curve in the background represents the unsmeared spectral density. Central panel: Relative error budget as a function of the energy. Bottom panel: Deviation of the predicted result from the true one in units of the standard deviation, see Eq. (58)

left of the blue one in each plot. The red points correspond to the results $\hat{\rho}_\sigma^{\text{pred}}(E, \mathbf{N}_X^{\text{ref}}) \pm \Delta_\sigma^{\text{stat}}(E, \mathbf{N}_X^{\text{ref}})$ that we use to estimate the systematic uncertainties $\Delta_\sigma^X(E)$. The blue band correspond to our estimate of the total error $\Delta_\sigma^{\text{tot}}(E)$.

In the top panel of Fig. 12 we show the comparison of our final results $\hat{\rho}_\sigma^{\text{pred}}(E) \pm \Delta_\sigma^{\text{tot}}(E)$ (blue points) with the true smeared spectral density $\hat{\rho}_\sigma^{\text{true}}(E)$ (black curve) that in this case is exactly known. The central panel in Fig. 12 shows the relative error budget as a function of the energy. As it can be seen, the systematics errors represent a sizeable and important fraction of the total error, particularly at large energies. The bottom panel of Fig. 12 shows the pull variable

$$p_\sigma(E) = \frac{\hat{\rho}_\sigma^{\text{pred}}(E) - \hat{\rho}_\sigma^{\text{true}}(E)}{\Delta_\sigma^{\text{tot}}(E)} \tag{58}$$

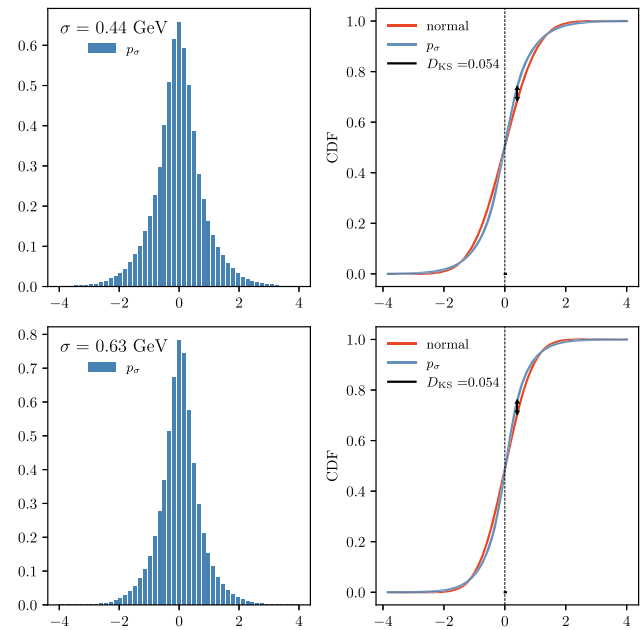


Fig. 13 Left panel: Normalized distribution of the significance defined in Eq. (58). The distribution is calculated over 2000 validation samples generated on the Chebyshev functional space. Right panel: comparison of the normalized cumulative distribution functions (CDF) of the observed distribution of p_σ (blue) and of the normal one obtained from the mean and variance of p_σ (red). The black arrow represents the Kolmogorov–Smirnov statistics (D_{KS}), i.e. the magnitude and the position of the maximum deviation between the two CDFs in absolute value

and, as it can be seen, by using the proposed procedure to estimate the final results and their errors, no significant deviations from the true result have been observed in this particular case.

In order to validate our method we repeated the test just described 2000 times. We have generated random spectral densities, not used in the trainings, by starting again from Eq. (39) and by selecting random values for E_0 and N_b respectively in the intervals $[0.3, 1.3]$ GeV and $[8, 1024]$.

The values of $p_\sigma(E)$ for all the energies and for all the samples is collected in one set, p_σ for short, whose normalized distribution is shown in the left panels of Fig. 13 (top panel for $\sigma = 0.44$ GeV, bottom panel for $\sigma = 0.63$ GeV). The distributions are nicely bell-shaped around 0 and the right panels in Fig. 13 show the comparison between the normalized cumulative distribution functions of p_σ with the normal distribution obtained from the mean and variance of p_σ . The p-value calculated from the Kolmogorov–Smirnov test is much less than 0.05 in both the cases and therefore the distribution of p_σ cannot be considered as a normal one. On the other hand, given the procedure that we use to quote the total error, we observe deviations $\hat{\rho}_\sigma^{\text{pred}}(E) - \hat{\rho}_\sigma^{\text{true}}(E)$ smaller than 2 standard deviations in 95% of the cases and smaller than 3 standard deviations in 99% of the cases. The

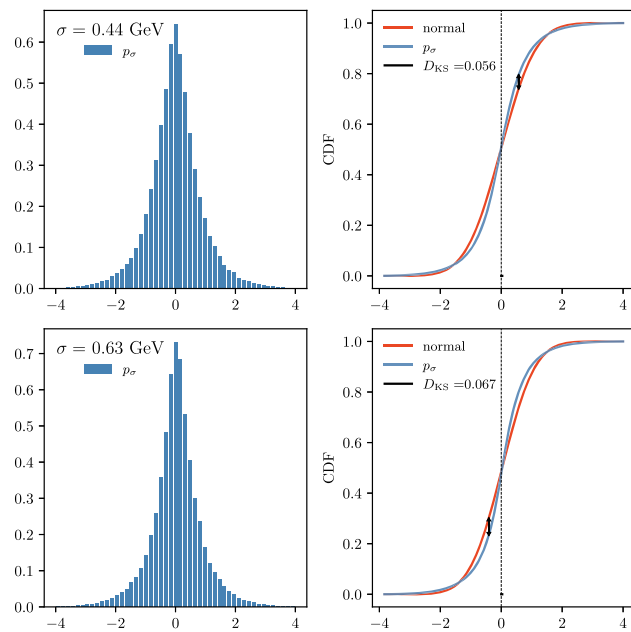


Fig. 14 The same as Fig. 13 but for 2000 samples of unsmeared spectral densities generated according to Eq. (59) and thus totally unrelated to the data used during the training

fact that deviations smaller than 1 standard deviation occur in $\sim 80\%$ of the cases for $\sigma = 0.44$ GeV and $\sim 85\%$ of the cases for $\sigma = 0.63$ GeV (to be compared with the expected 68% in the case of a normal distribution) is an indication of the fact that our estimate of the total error is indeed conservative. Moreover, from Fig. 13 it is also evident that our ensembles of neural networks are able to generalize very efficiently outside the training set.

In order to perform another stringent validation test of the proposed method, we also considered unsmeared spectral densities that cannot be written on the Chebyshev basis. We repeated the analysis described in the previous paragraph for a set of spectral densities generated according to

$$\rho(E) = \sum_{n=1}^{N_{\text{peaks}}} c_n \delta(E - E_n), \tag{59}$$

where $0 < E_0 \leq E_1 \leq E_2 \leq \dots \leq E_{\text{peaks}}$. By plugging Eq. (59) into Eqs. (32) and (33) we have calculated the correlator and smeared spectral density associated to each unsmeared $\rho(E)$. We have generated 2000 unsmeared spectral densities with $N_{\text{peaks}} = 5000$. The position E_n of each peak has been set by drawing independent random numbers uniformly distributed in the interval $[E_0, 15 \text{ GeV}]$ while E_0 has been randomly chosen in the interval $[0.3, 1.3]$ GeV. The coefficients c_n have also been generated randomly in the interval $[-0.01, 0.01]$. The 2000 trains of isolated Dirac-delta peaks are representative of unsmeared spectral densities that might arise in the study of finite volume lattice correla-

Table 5 Additional parameters used to generate the mock unsmeared spectral densities based on a Gaussian mixture model, Eq. (60). The corresponding functions are plotted in Fig. 15 (dashed grey)

Model	C	M (GeV)	Γ (GeV)	δ_2 (GeV)
$\rho_1^{\text{GM}}(E)$	10	1	0.1	1.5
$\rho_2^{\text{GM}}(E)$	5	1	0.1	1.5
$\rho_3^{\text{GM}}(E)$	5	0.8	0.05	1.2

tors. These are rather wild and irregular objects that our neural networks have not seen during the trainings. Figure 14 shows the plots equivalent to those of Fig. 13 for this new validation set. As it can be seen, the distributions of $p_\sigma(E)$ are basically unchanged, an additional reassuring evidence of the ability of our ensembles of neural networks to generalize very efficiently outside the training set and, more importantly, on the robustness of the procedure that we use to estimate the errors.

6 Results for mock data inspired by physical models

In the light of the results of the previous section, providing a solid quantitative evidence of the robustness and reliability of the proposed method, we investigate in this section the performances of our trained ensembles of machines in the case of mock spectral densities coming from physical models. A few more validation tests, unrelated to physical models, are discussed in Appendix A.

For each test discussed in the following subsections we define a model spectral density $\rho(E)$ and then use Eqs. (32) and (33) to calculate the associated exact correlator and true smeared spectral density. We then generate $N_c = 800$ bootstrap samples $C_c(t)$, by sampling the distribution of Eq. (44), and quote our final results by using the procedure described in details in the previous section.

6.1 A resonance and a multi-particle plateau

The first class of physically motivated models that we investigate is that of unsmeared spectral densities exhibiting the structures corresponding to an isolated resonance and a multi-particle plateau. To build mock spectral densities belonging to this class we used the same Gaussian mixture model used in Ref. [24] and given by

$$\rho^{\text{GM}}(E) = \hat{\theta}(E, \delta_1, \zeta_1) \left[C e^{-\left(\frac{E-M}{\Gamma}\right)^2} \left(1 - \hat{\theta}(E, \delta_2, \zeta_2)\right) \right] + C_0 \hat{\theta}(E, \delta_1, \zeta_1) \hat{\theta}(E, \delta_2, \zeta_2), \tag{60}$$




Fig. 15 Results for three different unsmoothed spectral densities $\rho(E)$ (dashed grey lines) generated from the Gaussian mixture model of Eq. (60). The solid lines correspond to $\rho_\sigma^{true}(E)$. The points, with error bars, refer to the predicted spectral densities $\hat{\rho}_\sigma^{pred}(E)$

where

$$\hat{\theta}(E, \delta_i, \zeta_i) = \frac{1}{1 + \exp\left(-\frac{E - \delta_i}{\zeta_i}\right)}. \tag{61}$$

The sigmoid function $\hat{\theta}(E, \delta_1, \zeta_1)$ with $\delta_1 = 0.1$ GeV and $\zeta_1 = 0.01$ GeV dumps $\rho^{GM}(E)$ at low energies while the other sigmoid $\hat{\theta}(E, \delta_2, \zeta_2)$, with $\zeta_2 = 0.1$ GeV, connects the resonance to the continuum part whose threshold is δ_2 . The parameter $C_0 = 1$ regulates the height of the continuum plateaux which also coincides with the asymptotic behaviour of the spectral density, i.e. $\rho^{GM}(E) \mapsto C_0$ for $E \mapsto \infty$. We have generated three different spectral densities with this model that, in the following, we call $\rho_1^{GM}(E)$, $\rho_2^{GM}(E)$ and $\rho_3^{GM}(E)$ and whose parameters are given in Table 5.




Fig. 16 Reconstructed smeared spectral densities (points with errors) compared to the true ones (solid lines) for two unsmoothed spectral densities simulating a finite volume distribution, see Eq. (62). The unsmoothed $\rho(E)$ is represented by vertical lines located in correspondence of $E = E_n$ with heights proportional to c_n

The predicted smeared spectral densities for smearing widths $\sigma = 0.44$ GeV and $\sigma = 0.63$ GeV are compared to the true ones in Fig. 15. In all cases the predicted result agrees with the true one, within the quoted errors, for all the explored values of E . The quality of the reconstruction of the smeared spectral densities is excellent for $E < 1.5$ GeV while, at higher energies, the quoted error is sufficiently large to account for the deviation of $\hat{\rho}_\sigma^{pred}(E)$ from $\hat{\rho}_\sigma^{true}(E)$. This is particularly evident in the case of $\rho_1^{GM}(E)$ shown in the top panel.

The same class of physical models can also be investigated by starting from unsmoothed spectral densities that might arise in finite volume calculations by considering

$$\rho^{peak}(E) = C_{peak} \delta(E - E_{peak}) + \sum_{n=1}^{N_{peaks}} c_n \delta(E - E_n). \tag{62}$$

The parameter E_{peak} parametrizes the position of an isolated peak while the multi-particle part is introduced with the other peaks located at $E_{peak} < E_1 \leq \dots \leq E_{N_{peaks}}$. We have generated two unsmoothed spectral densities, $\rho_1^{peak}(E)$ and $\rho_2^{peak}(E)$ by using this model. In both cases we set $N_{peaks} = 10,000$, selected random values for the E_n 's up to 50 GeV and random values for the c_n 's in the range $[0, 0.01]$. In the case of

$\rho_1^{\text{peak}}(E)$ we set $E_{\text{peak}} = 0.8 \text{ GeV}$, $C_{\text{peak}} = 1$ and $E_1 = 1.2 \text{ GeV}$. In the case of $\rho_2^{\text{peak}}(E)$ we set instead $E_{\text{peak}} = 0.7 \text{ GeV}$, $C_{\text{peak}} = 1$ and $E_1 = 1.5 \text{ GeV}$.

The predicted smeared spectral densities are compared with the true ones in Fig. 16. In both cases $\hat{\rho}_\sigma^{\text{pred}}(E)$ is in excellent agreement with $\hat{\rho}_\sigma^{\text{true}}(E)$.

It is worth emphasizing once again that the model in Eq. (62) cannot be represented by using the Chebyshev basis of Eq. (39) and, therefore, it is totally unrelated to the smooth unsmeared spectral densities that we used to populate the training sets.

6.2 $O(3)$ non-linear σ -model

In this subsection we consider a model for the unsmeared spectral density that has already been investigated by using the HLT method in Ref. [21]. More precisely, we consider the two-particles contribution to the vector-vector spectral density in the the 1 + 1 dimensional $O(3)$ non-linear σ -model (see Ref. [21] for more details) given by

$$\begin{aligned} \rho^{O(3)}(E) &= \theta(E - E_{\text{th}}) \frac{3\pi^3}{4\theta^2} \frac{\theta^2 + \pi^2}{\theta^2 + 4\pi^2} \tanh^3 \frac{\theta}{2} \Big|_{\theta=2 \cosh^{-1} \frac{E}{E_{\text{th}}}} \end{aligned} \tag{63}$$

where E_{th} is the two-particle threshold. We considered three mock unsmeared spectral densities, that we call $\rho_1^{O(3)}(E)$, $\rho_2^{O(3)}(E)$ and $\rho_3^{O(3)}(E)$, differing for the position of the multi-particle threshold, which has been set respectively to $E_{\text{th}} = 0.5 \text{ GeV}$, 1 GeV and 1.5 GeV . The reconstructed smeared spectral densities for $\sigma = 0.44 \text{ GeV}$ and $\sigma = 0.63 \text{ GeV}$ are compared with the exact ones in Fig. 17.

The predicted smeared spectral densities are in remarkably good agreement with the true ones in the full energy range and in all cases. This result can be read as an indication of the fact that the smoothness of the underlying unsmeared spectral density plays a crucial rôle in the precision that one can get on $\hat{\rho}_\sigma^{\text{pred}}(E)$, also if the problem is approached by using a neural network approach. Indeed, this fact had already been observed and exploited in Ref. [21], where the authors managed to perform the $\sigma \mapsto 0$ limit of $\hat{\rho}_\sigma^{O(3)}(E)$ with controlled systematic errors by using the HLT method.

6.3 The R -ratio with mock data

The last case that we consider is that of a model spectral density coming from a parametrization of the experimental data of the R -ratio. The R -ratio, denoted by $R(E)$, is defined as the ratio between the inclusive cross section of $e^+e^- \rightarrow$ hadrons and $e^+e^- \rightarrow \mu^+\mu^-$ and plays a crucial rôle in particle physics phenomenology (see e.g. Refs. [40,47]).




Fig. 17 Predicted smeared spectral densities for the two-particles contribution to the vector-vector correlator in the 1 + 1 dimensional $O(3)$ non-linear σ model for different values of the two-particle threshold: 0.5 GeV (top panel), 1 GeV (central panel) and 1.5 GeV (third panel). The solid lines and the dashed curve correspond respectively to $\hat{\rho}_\sigma^{\text{pred}}(E)$ and $\rho^{O(3)}(E)$

In the next section we will present results for a contribution to the smeared R -ratio $R_\sigma(E)$ that we obtained by feeding into our ensembles of trained neural networks a lattice QCD correlator that has been already used in Ref. [40] to calculate the same quantity with the HLT method.

Before doing that, however, we wanted also to perform a test with mock data generated by starting from the parametrization of $R(E)$ given in Ref. [46]. This parametrization, that we use as model unsmeared spectral density by setting $\rho(E) \equiv R(E)$, is meant to reproduce the experimental measurements of $R(E)$ for energies $E < 1.1 \text{ GeV}$, i.e. in the low-energy region where there are two dominant structures associated with the mixed ρ and ω resonances, a rather broad peak at $E \simeq 0.7 \text{ GeV}$, and the narrow resonance $\phi(1020)$.




Fig. 18 Results for the R -ratio using the parametrization of Ref. [46]. The gray dashed line is $R(E)$ while the solid lines correspond to $\hat{R}_\sigma^{\text{true}}(E)$

Table 6 ETMC gauge ensemble used in this work. See Refs. [40,48] for more details

ID	$L^3 \times T$	a fm	aL fm	m_π GeV
B64	$64^3 \times 128$	0.07957(13)	5.09	0.1352(2)

Given this rich structure, shown as the dashed grey curve in Fig. 18, this is a much more challenging validation test w.r.t. the one of the $O(3)$ model discussed in the previous subsection.

In Fig. 18, $\hat{R}_\sigma^{\text{pred}}(E)$ is compared with $\hat{R}_\sigma^{\text{true}}(E)$ for both $\sigma = 0.44$ GeV (orange points and solid curve) and $\sigma = 0.63$ GeV (blue points and solid curve). In both cases the difference $\hat{R}_\sigma^{\text{pred}}(E) - \hat{R}_\sigma^{\text{true}}(E)$ doesn't exceed the quoted error for all the considered values of E .

7 The R -ratio with lattice QCD data

In this section we use our trained ensembles of neural networks to extract the smeared spectral density from a real lattice QCD correlator.

We have considered a lattice correlator, measured by the ETMC on the ensemble described in Table 6, that has already been used in Ref. [40] to extract the so-called strange-strange connected contribution to the smeared R -ratio by using the HLT method of Ref. [14]. The choice is motivated by the fact that in this case the exact answer for the smeared spectral density is not known and we are going to compare our results $\hat{R}_\sigma^{\text{pred}}(E)$ with the ones, that we call $\hat{R}_\sigma^{\text{HLT}}(E)$, obtained in Ref. [40].




Fig. 19 The results obtained with our ensembles of neural networks in the case of a real lattice QCD correlator are compared with those obtained with the HLT method (grey, Refs. [14,40]). The top-panel shows the two determinations of the strange-strange connected contribution to the smeared R -ratio for $\sigma = 0.44$ GeV while the case $\sigma = 0.63$ GeV is shown in the bottom-panel

In QCD the R -ratio can be extracted from the lattice correlator

$$\begin{aligned}
 C_{\text{latt}}(t) &= -\frac{1}{3} \sum_{i=1}^3 \int d^3x T\langle 0|J_i(x)J_i(0)|0\rangle \\
 &= \int_0^\infty d\omega \frac{\omega^2}{12\pi^2} e^{-t\omega} R(\omega) .
 \end{aligned}
 \tag{64}$$

where $J_\mu(x)$ is the hadronic electromagnetic current. The previous formula (in which we have neglected the term proportional to $e^{-(T-t)\omega}$ that vanishes in the $T \mapsto \infty$ limit) explains the choice that we made in Eq. (32) to represent all the correlators that we used as inputs to our neural networks. In Ref. [40] all the connected and disconnected contributions to $C_{\text{latt}}(t)$, coming from the fact that $J_\mu = \sum_f q_f \bar{\psi}_f \gamma_\mu \psi_f$ with $f = \{u, d, s, c, b, t\}$, $q_{u,c,t} = 2/3$ and $q_{d,s,b} = -1/3$, have been taken into account. Here we consider only the connected strange-strange contribution, i.e. we set $J_\mu = -\frac{1}{3} \bar{s} \gamma_\mu s$ and neglect the contribution associated with the fermionic-disconnected Wick contraction. From $C_{\text{latt}}(t)$ we have calculated the covariance matrix $\hat{\Sigma}_{\text{latt}}$ that we used to inject noise in all the training sets that we have built and used in this work (see Sect. 4.2).

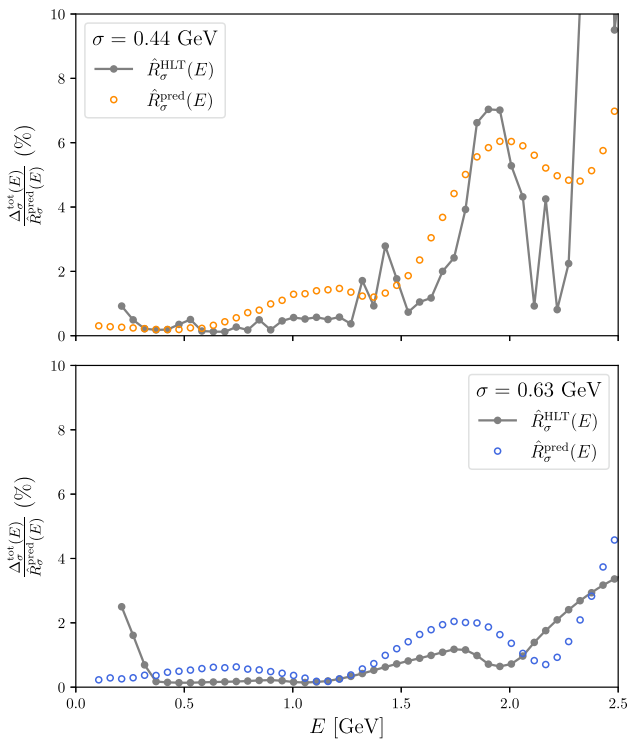


Fig. 20 Relative error on the smeared spectral density obtained from our ensembles of neural networks (colored points) and the HLT method (grey points) for $\sigma = 0.44$ GeV (top panel) and $\sigma = 0.63$ GeV (bottom panel)

Although, as already stressed, in this case we don't know the exact smeared spectral density, we do expect from phenomenology to see a sizeable contribution to the unsmeared spectral density coming from the $\phi(1020)$ resonance. Therefore, the shape of the smeared spectral density is expected to be similar to those shown in Fig. 16 for which the results obtained from our ensembles of neural networks turned out to be reliable. The comparison of $\hat{R}_\sigma^{\text{pred}}(E)$ and $\hat{R}_\sigma^{\text{HLT}}(E)$ is shown in Fig. 19 and, as it can be seen, the agreement between the two determinations is remarkably good.

Some remarks are in order here. The HLT method and the new method that we propose here are radically different and the fact that the results obtained with the two procedures are in such a good agreement, especially for $E < 1.5$ GeV where both methods provide very small errors, is at the same time reassuring and encouraging in view of future applications of spectral reconstruction techniques.

Concerning the errors, there is no significant evidence that one method has better performances than the other. This can be better appreciated in Fig. 20 where the relative error is shown in both cases as a function of the energy.

We want to stress that in the new proposed method, as in the HLT case, no prior information on the expected spectral density has to be provided and, therefore, the results for $\hat{R}_\sigma^{\text{pred}}(E)$ shown in Fig. 19 have to be considered as first-

principles model-independent determinations of the smeared strange-strange contribution to the R -ratio that we managed to obtain by using supervised deep-learning techniques.

We leave to future work on the subject the task of analysing all the contributions to $R(E)$, as done in Ref. [40], in order to obtain a machine-learning determination of $\hat{R}_\sigma(E)$ to be compared with experiments.

8 Conclusions

In this work we have proposed a new method to extract smeared hadronic spectral densities from lattice correlators. The method has been built by using supervised deep-learning techniques and is characterized by the distinctive features to implement a model-independent training strategy and to provide a reliable estimate of both the statistical and systematic uncertainties.

We managed to implement a model-independent training strategy by introducing a basis in the functional space from which we extract the spectral densities that we use to populate our training sets. In order to obtain a reliable estimate of the systematic errors, we introduced the ensembles of machines.

We have shown that, by studying the distribution of the answers of the different machines belonging to an ensemble, at fixed and finite number of neurons, dimension and complexity of the training set, it is possible to quantify reliably the systematic errors associated with the proposed method. To do that, we presented a large number of stringent validation tests, performed with mock data, providing quantitative evidence of the reliability of the procedure that we use to estimate the total error on our final results (see Figs. 13, 14).

In addition to mock data, we have applied the new proposed method also in the case of a lattice QCD correlator obtained from a simulation performed by the ETM Collaboration. We have extracted the strange-strange connected contribution to the smeared R -ratio and compared the predictions obtained by using our ensembles of trained machines with the ones previously obtained by using the HLT method [14,40]. We found a remarkably good agreement between the results obtained by using the two totally unrelated methods that provide total errors of the same order of magnitude.

The proposed method requires the training of many machines with different architectures and dimensions of the training sets. Admittedly, although this is a task that can be completed in a few days on a modern desktop computer, the procedure might end up to be computationally demanding. On the one hand, we don't exclude the possibility that the proposed method can be simplified in order to speed up the required computations. At the same time, on the basis of our experience, we are firmly convinced that a careful study of the different sources of systematic uncertainties is mandatory when dealing with machine-learning techniques

and when the aim is to compare theoretical predictions with experiments. In fact, the computational cost of the proposed method is the price that we had to pay for the reliability of the results.

As a final remark we want to stress that in this paper we have taught a lesson to a broad audience of learning-machines. The subject of the lesson, the extraction of smeared spectral densities from lattice correlators, is just a particular topic. The idea of teaching *systematically* to a broad audience of machines is much more general and can be used to estimate reliably the systematic errors in many other applications.

Acknowledgements We warmly thank our colleagues of the ETMC for very useful discussions and for providing us the data of the lattice QCD correlator used in this study. We also thank L. Del Debbio, R. Frezzotti, A. Lupo, M. Panero and M. Sbragaglia for enlightening discussions and B. Lucini for his extremely valuable comments on a preliminary version of the manuscript. MB acknowledges support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 882340).

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: The material used in this paper is available upon request.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

Appendix A: Results in exceptional cases

In this appendix we illustrate some additional validation tests that we performed by using mock data designed in order to challenge our ensembles of trained machines in extreme situations, i.e. in the case of unsmeared spectral densities that are very different from the ones used in the trainings. To this end, we considered the following unsmeared spectral densities

$$\begin{aligned}
 \rho_1(E) &= 0.03 \cdot \delta(E - 0.6) + 0.05 \cdot \delta(E - 1.2) , \\
 \rho_2(E) &= 0.03 \cdot \delta(E - 0.6) + 0.06 \cdot \delta(E - 1.8) , \\
 \rho_3(E) &= 0.01 \cdot \delta(E - 0.6) - 0.01 \cdot \delta(E - 1.3) \\
 &\quad + 0.02 \cdot \delta(E - 2.2) ,
 \end{aligned}
 \tag{A1}$$

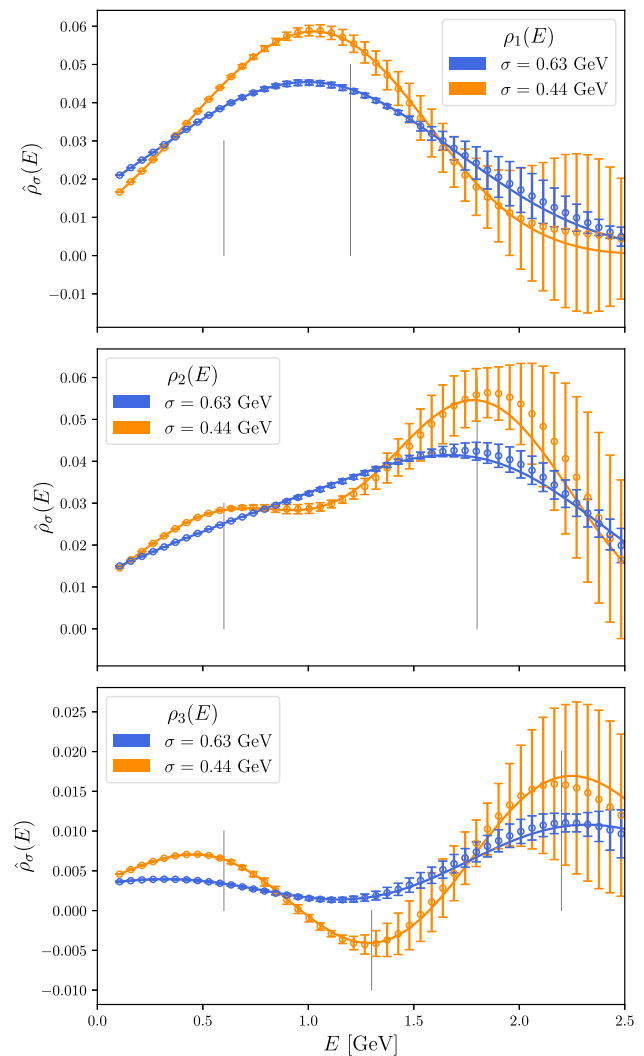


Fig. 21 Predicted smeared spectral densities for the three models of unsmeared spectral densities given in Eq. (A1). The solid lines correspond to $\rho_\sigma^{\text{true}}(E)$. The vertical grey lines are in correspondence of the isolated Dirac-delta peaks and the height is proportional to the associated weights

where the arguments of the Dirac-delta functions are in GeV units. The above spectral densities correspond to either two or three very well separated Dirac-delta peaks and one of the peaks of $\rho_3(E)$ has a negative coefficient. The final predicted results, compared with the exact ones, are shown in Fig. 21 for $\sigma = 0.44$ GeV and $\sigma = 0.63$ GeV. The agreement with $\hat{\rho}_\sigma^{\text{true}}(E)$ is excellent in all cases. As it can be seen, a difference w.r.t. the other models considered in Sects. 5 and 6 is the size of the errors that are much larger in these cases for $\sigma = 0.44$ GeV. This had to be expected and is a desired feature. Indeed, the inverse problem becomes harder when one tries to reconstruct sharp peaks with high resolution. The presence of statistical noise in the input data prevents any method to provide a very precise result, especially at high energies. Our neural network strategy, in these exceptional

situations, provides a large error and this is a further reassuring evidence of the robustness of the proposed procedure.

Appendix B: Dependence on the number of input time slices

In Sect. 3 we emphasized that the noise-to-signal ratio of generic lattice correlators grows exponentially with the Euclidean time (see Fig. 7) and that for this reason it might be convenient to reduce the number N_T of components of the input vector \mathbf{C} .

The approaches based on the Backus–Gilbert regularization method have a built-in mechanism to suppress noisy input data and by using these methods one can safely feed the whole information contained in the correlator into the algorithm. In general, this cannot be expected to be true when supervised deep-learning approaches are employed. Even though neural networks have been proven to be robust in handling noisy data, too much noise can have a bad impact on the performances since a big fraction of the effort in the minimization algorithm is put in the suppression of the noise and in learning the distinction between outliers and effective information rather than in learning general features of the problem.

In the following we investigate the dependence of the training performances, obtained by injecting the noise of the lattice correlator in the training sets as explained in Sect. 4.2, upon N_T . In Fig. 22 we show the validation loss as a function of N_T . As it can be seen, the performances of the networks improve as N_T increases but the validation loss reaches a saturation point around $N_T^{\text{sat}} = 40$. On the one hand, N_T^{sat} can be considered as the maximum number of time slices of the correlator from which meaningful physical information can

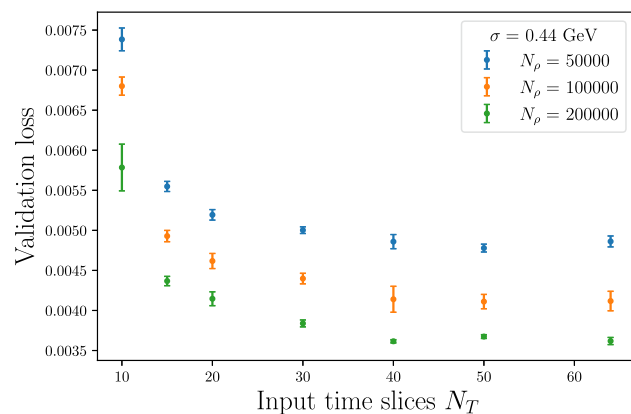


Fig. 22 Validation losses at the end of the training for different number of time slices in the input data and for different training set size N_ρ . This test is performed with $N_b = 512$ and $\sigma = 0.44$ GeV. Each point with the associated error bar comes from the average of $N_r = 20$ trainings (see Sect. 4)

be extracted. On the other hand, despite the corruption of the input data for $N_T > N_T^{\text{sat}}$, including more time slices does not compromise the quality of the trainings and we could safely set $N_T = T/a = 64$.

Appendix C: Analysis of the statistical uncertainties

In illustrating our method we have stressed that there are two sources of statistical errors that contribute to $\Delta_\sigma^{\text{stat}}(E, \mathbf{N})$ defined in Eq. (52): $\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$, coming from the Monte Carlo simulation, and $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$, coming from our ensemble of trained machines. In this appendix we study the depen-




Fig. 23 Behaviour of $\Delta_\sigma^{\text{latt}}(E, \mathbf{N})$ (blue) and $\Delta_\sigma^{\text{net}}(E, \mathbf{N})$ (red) with respect to $\mathbf{N} = \{N_\rho, N_n, N_b\}$ at four different values of the energy. The data refer to $\sigma = 0.44$ GeV and have been obtained by analysing the 2000 samples built on the Chebyshev basis already used in Sect. 5.2. The central values and the associated error bars come from the average over the 2000 samples

dence of $\Delta_\sigma^{\text{latt}}$ and $\Delta_\sigma^{\text{net}}$ upon the energy E and upon N . The investigation is carried out by collecting $\Delta_\sigma^{\text{latt}}$ and $\Delta_\sigma^{\text{net}}$ for the 2000 mock samples built on the Chebyshev basis and used in Sect. 5.2 to validate our procedure. For each combination of N the 2000 results are averaged and the standard deviation is calculated. The results are shown in Fig. 23 at $\sigma = 0.44$ GeV and at four different energies spanning the output energy range. We observe that both $\Delta_\sigma^{\text{latt}}$ and $\Delta_\sigma^{\text{net}}$ increase in magnitude as the energy increases. This phenomenology is a well-known feature of the inverse Laplace transform problem in which the contributions to the correlation function coming from high energies are strongly suppressed by the exponential basis (see Eq. (3)). Therefore, even small statistical fluctuations in the correlator turn into enhanced errors in the spectral density at high energy. At fixed energy the Monte Carlo error $\Delta_\sigma^{\text{latt}}$ does not show a significant dependence upon N . This is a reassuring evidence that $\Delta_\sigma^{\text{latt}}$ and $\Delta_\sigma^{\text{net}}$ are two decorrelated sources of uncertainties. Conversely, at fixed energy $\Delta_\sigma^{\text{net}}$ decreases by increasing N_ρ and/or N_n , thus confirming our working hypothesis that an infinitely large neural network, trained over an infinitely large dataset, is able to provide the exact answer independently of the initialization parameters. $\Delta_\sigma^{\text{net}}$ is instead stable with respect to changes in the dimension N_b of the functional-basis. This stability can be ascribed to the fact that we are targeting the extraction of smeared spectral densities at sufficiently large values of σ so that the local complexity introduced in the unsmeared spectral densities by increasing N_b is washed out. A trend in $\Delta_\sigma^{\text{net}}$ as a function of N_b may thus become visible when considering smaller values of σ .

References

- J.C.A. Barata, K. Fredenhagen, Particle scattering in Euclidean lattice field theories. *Commun. Math. Phys.* **138**, 507 (1991). <https://doi.org/10.1007/BF02102039>
- M. Jarrell, J.E. Gubernatis, Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data. *Phys. Rep.* **269**, 133 (1996). [https://doi.org/10.1016/0370-1573\(95\)00074-7](https://doi.org/10.1016/0370-1573(95)00074-7)
- Y. Nakahara, M. Asakawa, T. Hatsuda, Hadronic spectral functions in lattice QCD. *Phys. Rev. D* **60**, 091503 (1999). <https://doi.org/10.1103/PhysRevD.60.091503>. [arXiv:hep-lat/9905034](https://arxiv.org/abs/hep-lat/9905034)
- M. Asakawa, T. Hatsuda, Y. Nakahara, Maximum entropy analysis of the spectral functions in lattice QCD. *Prog. Part. Nucl. Phys.* **46**, 459 (2001). [https://doi.org/10.1016/S0146-6410\(01\)00150-8](https://doi.org/10.1016/S0146-6410(01)00150-8). [arXiv:hep-lat/0011040](https://arxiv.org/abs/hep-lat/0011040)
- G. Aarts, J.M. Martinez Resco, Transport coefficients, spectral functions and the lattice. *JHEP* **04**, 053. <https://doi.org/10.1088/1126-6708/2002/04/053>. [arXiv:hep-ph/0203177](https://arxiv.org/abs/hep-ph/0203177)
- G. Aarts, J.M. Martinez Resco, Continuum and lattice meson spectral functions at nonzero momentum and high temperature. *Nucl. Phys. B* **726**, 93 (2005). <https://doi.org/10.1016/j.nuclphysb.2005.08.012>. [arXiv:hep-lat/0507004](https://arxiv.org/abs/hep-lat/0507004)
- M. Laine, How to compute the thermal quarkonium spectral function from first principles? *Nucl. Phys. A* **820**, 25C (2009). <https://doi.org/10.1016/j.nuclphysa.2009.01.015>. [arXiv:0810.1112](https://arxiv.org/abs/0810.1112) [hep-ph]
- H.B. Meyer, Energy-momentum tensor correlators and spectral functions. *JHEP* **08**, 031. <https://doi.org/10.1088/1126-6708/2008/08/031>. [arXiv:0806.3914](https://arxiv.org/abs/0806.3914) [hep-lat]
- Y. Burnier, A. Rothkopf, Bayesian approach to spectral function reconstruction for Euclidean quantum field theories. *Phys. Rev. Lett.* **111**, 182003 (2013). <https://doi.org/10.1103/PhysRevLett.111.182003>. [arXiv:1307.6106](https://arxiv.org/abs/1307.6106) [hep-lat]
- H.B. Meyer, Transport properties of the quark-gluon plasma: a lattice QCD perspective. *Eur. Phys. J. A* **47**, 86 (2011). <https://doi.org/10.1140/epja/i2011-11086-3>. [arXiv:1104.3708](https://arxiv.org/abs/1104.3708) [hep-lat]
- M.T. Hansen, H.B. Meyer, D. Robaina, From deep inelastic scattering to heavy-flavor semileptonic decays: Total rates into multihadron final states from lattice QCD. *Phys. Rev. D* **96**, 094513 (2017). <https://doi.org/10.1103/PhysRevD.96.094513>. [arXiv:1704.08993](https://arxiv.org/abs/1704.08993) [hep-lat]
- R.-A. Tripolt, P. Gubler, M. Ulybyshev, L. Von Smekal, Numerical analytic continuation of Euclidean data. *Comput. Phys. Commun.* **237**, 129 (2019). <https://doi.org/10.1016/j.cpc.2018.11.012>. [arXiv:1801.10348](https://arxiv.org/abs/1801.10348) [hep-ph]
- J. Bulava, M.T. Hansen, Scattering amplitudes from finite-volume spectral functions. *Phys. Rev. D* **100**, 034521 (2019). <https://doi.org/10.1103/PhysRevD.100.034521>. [arXiv:1903.11735](https://arxiv.org/abs/1903.11735) [hep-lat]
- M. Hansen, A. Lupo, N. Tantalo, Extraction of spectral densities from lattice correlators. *Phys. Rev. D* **99**, 094508 (2019). <https://doi.org/10.1103/PhysRevD.99.094508>. [arXiv:1903.06476](https://arxiv.org/abs/1903.06476) [hep-lat]
- L. Kades, J.M. Pawłowski, A. Rothkopf, M. Scherzer, J.M. Urban, S.J. Wetzel, N. Wink, F.P.G. Ziegler, Spectral reconstruction with deep neural networks. *Phys. Rev. D* **102**, 096001 (2020). <https://doi.org/10.1103/PhysRevD.102.096001>. [arXiv:1905.04305](https://arxiv.org/abs/1905.04305) [physics.comp-ph]
- J. Karpie, K. Orginos, A. Rothkopf, S. Zafeiropoulos, Reconstructing parton distribution functions from Ioffe time data: from Bayesian methods to Neural Networks. *JHEP* **04**, 057. [https://doi.org/10.1007/JHEP04\(2019\)057](https://doi.org/10.1007/JHEP04(2019)057). [arXiv:1901.05408](https://arxiv.org/abs/1901.05408) [hep-lat]
- G. Bailas, S. Hashimoto, T. Ishikawa, Reconstruction of smeared spectral function from Euclidean correlation functions. *PTEP* **2020**, 043B07 (2020). <https://doi.org/10.1093/ptep/ptaa044>. [arXiv:2001.11779](https://arxiv.org/abs/2001.11779) [hep-lat]
- P. Gambino, S. Hashimoto, Inclusive semileptonic decays from lattice QCD. *Phys. Rev. Lett.* **125**, 032001 (2020). <https://doi.org/10.1103/PhysRevLett.125.032001>. [arXiv:2005.13730](https://arxiv.org/abs/2005.13730) [hep-lat]
- M. Bruno, M.T. Hansen, Variations on the Maiani-Testa approach and the inverse problem. *JHEP* **06**, 043. [https://doi.org/10.1007/JHEP06\(2021\)043](https://doi.org/10.1007/JHEP06(2021)043). [arXiv:2012.11488](https://arxiv.org/abs/2012.11488) [hep-lat]
- J. Horak, J.M. Pawłowski, J. Rodríguez-Quintero, J. Turnwald, J.M. Urban, N. Wink, S. Zafeiropoulos, Reconstructing QCD spectral functions with Gaussian processes. *Phys. Rev. D* **105**, 036014 (2022). <https://doi.org/10.1103/PhysRevD.105.036014>. [arXiv:2107.13464](https://arxiv.org/abs/2107.13464) [hep-ph]
- J. Bulava, M.T. Hansen, M.W. Hansen, A. Patella, N. Tantalo, Inclusive rates from smeared spectral densities in the two-dimensional O(3) non-linear σ -model. *JHEP* **07**, 034. [https://doi.org/10.1007/JHEP07\(2022\)034](https://doi.org/10.1007/JHEP07(2022)034). [arXiv:2111.12774](https://arxiv.org/abs/2111.12774) [hep-lat]
- L. Wang, S. Shi, K. Zhou, Automatic differentiation approach for reconstructing spectral functions with neural networks, in *35th Conference on Neural Information Processing Systems* (2021). [arXiv:2112.06206](https://arxiv.org/abs/2112.06206) [physics.comp-ph]
- S.-Y. Chen, H.-T. Ding, F.-Y. Liu, G. Papp, C.-B. Yang, Machine learning hadron spectral functions in lattice QCD. *PoS LATTICE2021*, 148 (2022). <https://doi.org/10.22323/1.396.0148>. [arXiv:2112.00460](https://arxiv.org/abs/2112.00460) [hep-lat]

24. S. Y. Chen, H. T. Ding, F. Y. Liu, G. Papp, C. B. Yang, Machine learning spectral functions in lattice QCD. (2021). [arXiv:2110.13521](https://arxiv.org/abs/2110.13521) [hep-lat]
25. L. Wang, S. Shi, K. Zhou, Reconstructing spectral functions via automatic differentiation. *Phys. Rev. D* **106**, L051502 (2022). <https://doi.org/10.1103/PhysRevD.106.L051502>. [arXiv:2111.14760](https://arxiv.org/abs/2111.14760) [hep-ph]
26. M. Zhou, F. Gao, J. Chao, Y.-X. Liu, H. Song, Application of radial basis functions neural networks in spectral functions. *Phys. Rev. D* **104**, 076011 (2021). <https://doi.org/10.1103/PhysRevD.104.076011>. [arXiv:2106.08168](https://arxiv.org/abs/2106.08168) [hep-ph]
27. T. Lechien, D. Dudal, Neural network approach to reconstructing spectral functions and complex poles of confined particles. *SciPost Phys.* **13**, 097 (2022). <https://doi.org/10.21468/SciPostPhys.13.4.097>. [arXiv:2203.03293](https://arxiv.org/abs/2203.03293) [hep-lat]
28. D. Boyda et al., Applications of machine learning to lattice quantum field theory, in *2022 Snowmass Summer Study* (2022). [arXiv:2202.05838](https://arxiv.org/abs/2202.05838) [hep-lat]
29. S. Shi, L. Wang, K. Zhou, Rethinking the ill-posedness of the spectral function reconstruction: why is it fundamentally hard and how artificial neural networks can help. *Comput. Phys. Commun.* **282**, 108547 (2023). <https://doi.org/10.1016/j.cpc.2022.108547>. [arXiv:2201.02564](https://arxiv.org/abs/2201.02564) [hep-ph]
30. T. Bergamaschi, W. I. Jay, P. R. Oare, Hadronic structure, conformal maps, and analytic continuation (2023). [arXiv:2305.16190](https://arxiv.org/abs/2305.16190) [hep-lat]
31. A. Rothkopf, Inverse problems, real-time dynamics and lattice simulations. *EPJ Web Conf.* **274**, 01004 (2022). <https://doi.org/10.1051/epjconf/202227401004>. [arXiv:2211.10680](https://arxiv.org/abs/2211.10680) [hep-lat]
32. J. Bulava, The spectral reconstruction of inclusive rates. *PoS LATTICE2022*, 231 (2023). <https://doi.org/10.22323/1.430.0231>. [arXiv:2301.04072](https://arxiv.org/abs/2301.04072) [hep-lat]
33. K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators. *Neural Netw.* **2**, 359 (1989). [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
34. I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning* (MIT Press, 2016). <http://www.deeplearningbook.org>
35. G. Cybenko, Approximation by superpositions of a sigmoidal function. *Math. Control Signals Syst.* **2**, 303 (1989)
36. T. Pearce, F. Leibfried, A. Brintrup, Uncertainty in neural networks: approximately bayesian ensembling, in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, series Proceedings of Machine Learning Research*, vol. 108, ed. by S. Chiappa, R. Calandra (PMLR, 2020), p. 234–244. <https://proceedings.mlr.press/v108/pearce20a.html>
37. C. E. Rasmussen, C. K. I. Williams, Gaussian processes for machine learning. *Adaptive Computation and Machine Learning* (MIT Press, 2006), p. I–XVIII, 1–248
38. K. P. Murphy, https://www.amazon.com/Machine-Learning-Probabilistic-Perspective-Computation/dp/0262018020/ref=sr_1_2?ie=UTF8&qid=1336857747&sr=8-2 *Machine learning : a probabilistic perspective* (MIT Press, Cambridge, 2013)
39. G. Backus, F. Gilbert, The resolving power of gross earth data. *Geophys. J. Int.* **16**, 169 (1968). <https://doi.org/10.1111/j.1365-246X.1968.tb00216.x>
40. C. Alexandrou et al. (collaboration Extended Twisted Mass Collaboration (ETMC)), Probing the Energy-Smeared R Ratio Using Lattice QCD. *Phys. Rev. Lett.* **130**, 241901 (2023). <https://doi.org/10.1103/PhysRevLett.130.241901>. [arXiv:2212.08467](https://arxiv.org/abs/2212.08467) [hep-lat]
41. Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278 (1998)
42. F. Chollet et al., Keras, howpublished <https://keras.io> (2015)
43. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, <https://www.tensorflow.org/> TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from tensorflow.org
44. J. P. Boyd, *Chebyshev and Fourier Spectral Methods* (Courier Corporation, 2001)
45. D. P. Kingma, J. Ba, Adam: a method for stochastic optimization (2017). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG]
46. David Bernecker, Harvey B. Meyer, Vector correlators in lattice QCD: methods and applications. *Eur. Phys. J. A* (2011). <https://doi.org/10.1140/epja/i2011-11148-6>
47. T. Aoyama et al., The anomalous magnetic moment of the muon in the Standard Model. *Phys. Rep.* **887**, 1–166 (2020). <https://doi.org/10.1016/j.physrep.2020.07.006>
48. C. Alexandrou et al., Lattice calculation of the short and intermediate time-distance hadronic vacuum polarization contributions to the muon magnetic moment using twisted-mass fermions (2022). [arXiv:2206.15084](https://arxiv.org/abs/2206.15084) [hep-lat]