

Received October 29, 2020, accepted November 12, 2020, date of publication November 16, 2020, date of current version December 1, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3038333

Adaptive Quadrature Schemes for Bayesian Inference via Active Learning

FERNANDO LLORENTE FERNÁNDEZ¹, LUCA MARTINO², VICTOR ELVIRA³, DAVID DELGADO¹, AND JAVIER LÓPEZ-SANTIAGO⁴

¹Department of Statistics, Universidad Carlos III de Madrid, 28911 Leganés, Spain

²Department of Signal Processing, Universidad Rey Juan Carlos, 28933 Fuenlabrada, Spain

³School of Mathematics, The University of Edinburgh, Edinburgh EH9 3FD, U.K.

⁴Department of Signal Processing, Universidad Carlos III de Madrid, 28911 Leganés, Spain

Corresponding author: Fernando Lorente Fernández (felloren@est-econ.uc3m.es)

ABSTRACT We propose novel adaptive quadrature schemes based on an active learning procedure. We consider an interpolative approach for building a surrogate posterior density, combining it with Monte Carlo sampling methods and other quadrature rules. The nodes of the quadrature are sequentially chosen by maximizing a suitable acquisition function, which takes into account the current approximation of the posterior and the positions of the nodes. This maximization does not require additional evaluations of the true posterior. We introduce two specific schemes based on Gaussian and Nearest Neighbors bases. For the Gaussian case, we also provide a novel procedure for fitting the bandwidth parameter, in order to build a suitable emulator of a density function. With both techniques, we always obtain a positive estimation of the marginal likelihood (a.k.a., Bayesian evidence). An equivalent importance sampling interpretation is also described, which allows the design of extended schemes. Several theoretical results are provided and discussed. Numerical results show the advantage of the proposed approach, including a challenging inference problem in an astronomic dynamical model, with the goal of revealing the number of planets orbiting a star.

INDEX TERMS Active learning, Bayesian quadrature, emulation, experimental design, Monte Carlo methods, numerical integration.

I. INTRODUCTION AND BRIEF OVERVIEW

In this work, we consider the approximation of intractable integrals of type

$$I = \int_{\mathcal{X}} f(\mathbf{x})\bar{\pi}(\mathbf{x})d\mathbf{x},$$

where $f(\mathbf{x})$ is a generic integrable function and $\bar{\pi}(\mathbf{x})$ is a probability density function (pdf). These integrals usually appear in Bayesian inference problems where $\bar{\pi}(\mathbf{x})$ represents the posterior distribution of the variable of interest given the observed data. In the next subsections, we briefly review several approaches presented in the literature, which are related to the methodology presented this work.

A. MAIN FAMILIES OF QUADRATURE METHODS

With the term *numerical integration*, we refer to a broad family of algorithms for calculating definite integrals, and by extension, the term is also used to describe the numerical

The associate editor coordinating the review of this manuscript and approving it for publication was Wentao Fan.

solution of differential equations. Although exact analytical solutions to integrals are always desirable, such “unicorns” are rarely available, specially in real-world systems. Indeed, many applications in signal processing, statistics, and machine learning inevitably require the approximation of intractable integrals [1]–[3]. In particular, Bayesian methods need the computation of posterior expectations which, generally, are analytically intractable [2], [4]. The term numerical quadrature (or simply quadrature) is employed as a synonym for numerical integration [1]. More specifically, a quadrature formula is often stated as a weighted sum of integrand evaluations at specified points (a.k.a., nodes or knots) within the domain of integration.

1) DETERMINISTIC QUADRATURES

A first family of numerical integration methods are the deterministic quadrature rules. A subclass within this family is the Newton-Cotes quadrature rules [1]. The Newton-Cotes formulas are based on evaluating the integrand at equally spaced nodes and are obtained by substituting the

integrand function with a corresponding polynomial interpolation. Smaller approximation errors can often be obtained by using the Gaussian quadratures, where the nodes are optimally placed [1], [5], [6]. However, their applicability is restricted to certain particular cases.

2) MONTE CARLO (MC) METHODS

A second family is formed by stochastic quadrature rules based on MC sampling methods [2], [4], such as Markov chain Monte Carlo (MCMC) and importance sampling algorithms. In this framework, the nodes of the quadrature rules are randomly chosen. However, the resulting estimators often have a high variance, specially when the dimension of the problem grows.

3) VARIANCE REDUCTION

A third family, formed by the variance reduction techniques [2], [7], combines elements of the first two classes. In order to reduce the variance of the corresponding Monte Carlo estimators, deterministic procedures are included within the sampling algorithms, e.g., conditioning, stratification, anti-thetic sampling, and control variates [7]. Other interesting examples are the Riemann-based approximations which are combinations of a Riemann quadrature and random sampling [2, Chapter 4.3]. The Quasi-Monte Carlo (QMC) algorithms can be also included in this family. In QMC, deterministic sequences of points are generated (based on the concept of low-discrepancy) and then used as nodes of the corresponding quadrature [3]. Several other combinations of the previous classes above, mixing determinism with random sampling schemes, can be found in the literature [8]–[10].

4) BAYESIAN QUADRATURE (BQ)

The BQ framework represents a fourth approach which employs Gaussian Process (GP) regression algorithms for approximating the integrand function (and, as a consequence, the resulting integral as well) [11]–[13]. In the last years, this approach has raised the interest of several authors. One problem with this approach is that, in some cases, a negative estimation of the marginal likelihood can be obtained. Some possible solutions have been proposed, although they are quite complex based on successive approximations [14], [15]. In this work, we provide two novel and much simpler alternatives for solving this issue. Moreover, unlike this work, most contributions in BQ literature focus on the GP approximation of the function $f(\mathbf{x})$ [14]–[16], although other papers on BQ describe quite general frameworks where $f(\mathbf{x})$ can contain the likelihood or $\pi(\mathbf{x})$ [11]–[13]. A connection between classical quadratures and BQ can be found in [17]. Finally, theoretical guarantees for adaptive BQ schemes can be found in the insightful work of [18].

B. EMULATION OF COMPLEX MODELS

Many Bayesian inference problems involve the evaluation of computationally intensive models, because of **(i)** the use of particularly complex systems or **(ii)** a large number of

available data (or both). To overcome this issue, one possible strategy consists in replacing the true model by a surrogate model (a.k.a. an *emulator*), that could be also adaptively improved [19]–[21]. Then, Bayesian inference is carried out on this approximate, cheaper model.

1) USE OF THE EMULATOR

The emulator can be applied mainly in three different ways. **(a)** One possibility is to apply MC sampling methods considering the surrogate model as the target pdf [22], [23]. This is used to speed up the MC algorithms. **(b)** In order to improve the efficiency of MC estimators, a second option is to use the emulator as a proposal density within an MC technique, as we discuss in Section I-C [24]–[26]. **(c)** A third possibility is to replace the true posterior with the emulator in the integrals of interest, and computing them [11]–[13]. Here, we mainly focus on the last approach, also combining it with MC methods (and other quadrature rules).

2) CONSTRUCTION OF THE EMULATOR

In the literature, the surrogate model is often built by using a regression algorithm, like a GP model or similar techniques [27], [28]. This probabilistic approach provides also uncertainty quantification that is used for estimating the approximation error and adapting the emulator [29]. Sometimes, the approximation regards only some part of the model or is applied in a different domain (as the log-domain) [30]–[33]. Other authors employ density estimation techniques for building the surrogate model, and then using it as a proposal density within MC algorithms [34]–[36] or for replacing the true posterior (again within MC methods) [37].

C. INTERPOLATIVE PROPOSAL DENSITIES WITHIN MONTE CARLO SCHEMES

The first use of an interpolative procedure for building a proposal density is ascribable to the adaptive rejection sampling schemes [24], [38]–[40]. The proposal is formed by polynomial pieces (constant, linear, etc.). Several works have proposed the use of interpolative proposal densities within MCMC algorithms [25], [41]–[43]. For more details, see also [4, Chapters 4 and 7]. Their use within an importance sampling scheme is considered in [44]. The adaptation is carried out considering different statistical tests, by measuring the discrepancy between the emulator and the posterior [26].

The conditions needed for applying an emulator as an proposal density are discussed in [26]. For this purpose, we need to be able to: **(a)** update the construction of the emulator, **(b)** evaluate the emulator, **(c)** normalize the function defined by the emulator, and **(d)** draw samples from the emulator. It is not straightforward to find an interpolative construction which satisfies all those conditions jointly, for an arbitrary dimension of the problem. However, the resulting algorithms (when they can be applied) provide good performance, confirming that the interpolative approach deserves more attention.

D. CONTRIBUTIONS

In this work, we leverage the advances in different fields of numerical integration and emulation, in order to design algorithms which build (a) better emulators and (b) more efficient quadrature rules. The novel algorithms are adaptive schemes which automatically select the nodes of the quadrature and of the resulting emulator. Namely, the set of nodes used by the emulator is sequentially updated by maximizing a suitable acquisition function. Below, we list the main contributions of the work.

- We propose a novel design of a suitable acquisition function defined as product of the posterior and a diversity term, taking into account the current positions of the nodes. Note that, unlike several works in the literature, e.g., [19], [20], [45], [46], we consider jointly both: the information regarding the posterior and the distances among the current nodes. For the selection of the nodes, some authors also consider the use of MCMC runs [22] or more sophisticated procedures combining sampling and deterministic quadrature schemes for selecting the nodes [47]. Unlike [22], [47], our adaptive approach is based on an active learning procedure. We also provide cheap versions of the acquisition function. The cheap acquisition functions do not require the evaluation of the posterior but only the evaluation of the emulator. The overall schemes are then *parsimonious* techniques which require the evaluation of the posterior density only at the nodes, sequentially selected by optimizing a cheap acquisition function. The proposed active learning strategy is also connected to the idea of obtaining a finite set of weighted *representative points* which can summarize, in some sense, a distribution. This topic has gained attention in the last years [48]–[51].

- We consider an interpolative approximation of the posterior density $\bar{\pi}(\mathbf{x})$, where the interpolant is expressed as a linear combination of generic kernel-basis functions. Unlike several BQ techniques in [14]–[16], we approximate $\bar{\pi}(\mathbf{x})$ instead of the function $f(\mathbf{x})$ in the integral I . For this purpose, we also propose the combination of the interpolant approach with MC and other quadrature schemes.

- With respect to other schemes in the literature [12], [13], our assumptions regarding the kernel-basis functions are less restrictive, e.g., they do not need to be symmetric. We could also employ different type of bases jointly, e.g., one different basis for each node. For instance, our framework allows the use of nearest neighbors (NN) basis functions, which presents several advantages: it does not require any matrix inversion and the coefficients of the linear combination (which defines the interpolator) are always positive [52], obtaining always a positive estimation of the marginal likelihood. These benefits are very appealing as shown in [14], [15], [52], [53].

- Section V presents an importance sampling (IS) interpretation of the proposed schemes, where the weights involve the interpolant instead of the true posterior density. This again shows that we can improve the Monte Carlo approximations without requiring additional evaluations of $\bar{\pi}(\mathbf{x})$. Moreover, the alternative IS interpretation allows to design different

techniques. One possible example is given in the final part of Section V.

- We also introduce a novel procedure for fitting the bandwidth parameter of the Gaussian kernel in order to build an *emulator of a density function*. In this scenario, the proposed strategy performs better than the standard maximization of the marginal likelihood of the corresponding GP. Using this tuning procedure, we always obtain positive estimation of the marginal likelihood, even with Gaussian kernels (this is an important point; see [14], [15]).

We provide the theoretical support for the proposed methods in Section VII. Most of the convergence results are mainly known in the scattered data approximation literature [54]–[56]. The efficiency of the proposed schemes is also confirmed by several numerical experiments (in Section VIII) with different target pdfs and dimensions of the problem. One of them is also a challenging astronomical application, where the goal is to detect the number of exoplanets orbiting a star, and infer their orbital parameters.

II. INTERPOLATIVE QUADRATURES FOR BAYESIAN INFERENCE

In many signal processing applications, the goal is to infer a variable of interest given a set of observations or measurements. Let us denote the variable of interest by $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$, and let $\mathbf{y} \in \mathbb{R}^{d_y}$ be the observed data. The posterior pdf is then

$$\bar{\pi}(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}) = \frac{\ell(\mathbf{y}|\mathbf{x})g(\mathbf{x})}{Z(\mathbf{y})},$$

where $\ell(\mathbf{y}|\mathbf{x})$ is the likelihood function, $g(\mathbf{x})$ is the prior pdf, and $Z(\mathbf{y})$ is the model evidence (a.k.a. marginal likelihood). Generally, $Z(\mathbf{y})$ is unknown, so we are able to evaluate the unnormalized target function,

$$\pi(\mathbf{x}) = \ell(\mathbf{y}|\mathbf{x})g(\mathbf{x}).$$

Usually, the analytical computation of the posterior density $\bar{\pi}(\mathbf{x}) \propto \pi(\mathbf{x})$ is unfeasible, hence numerical approximations are required. Our goal is to approximate integrals of the form

$$I = \int_{\mathcal{X}} f(\mathbf{x})\bar{\pi}(\mathbf{x})d\mathbf{x} = \frac{1}{Z} \int_{\mathcal{X}} f(\mathbf{x})\pi(\mathbf{x})d\mathbf{x}, \tag{1}$$

where $f(\mathbf{x})$ is some integrable function, and

$$Z = \int_{\mathcal{X}} \pi(\mathbf{x})d\mathbf{x}. \tag{2}$$

In the literature, random sampling or deterministic quadratures are often used [2], [4], [57]. In this work, we consider alternative quadrature rules based on an adaptive interpolative procedure. The adaptation is obtained by applying an active learning scheme.

A. INTERPOLATIVE APPROACH

Let us consider a set of distinct nodes $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathcal{X}$ and some non-negative kernel or basis function, $k(\mathbf{x}, \mathbf{x}') : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+ \cup \{0\}$ (i.e., $k(\mathbf{x}, \mathbf{x}') \geq 0$). From now on, we use the

terms basis or kernel as synonyms. The interpolant of $\pi(\mathbf{x})$ is as follows

$$\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i), \quad (3)$$

where the coefficients β_i must be such that $\hat{\pi}(\mathbf{x})$ interpolates the points $\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)$, that is, $\hat{\pi}(\mathbf{x}_i) = \pi(\mathbf{x}_i)$ for $i = 1, \dots, N$. Hence, the β_i are the solutions to the following linear system

$$\begin{cases} \beta_1 k(\mathbf{x}_1, \mathbf{x}_1) + \dots + \beta_N k(\mathbf{x}_1, \mathbf{x}_N) = \pi(\mathbf{x}_1), \\ \beta_1 k(\mathbf{x}_2, \mathbf{x}_1) + \dots + \beta_N k(\mathbf{x}_2, \mathbf{x}_N) = \pi(\mathbf{x}_2), \\ \vdots \\ \beta_1 k(\mathbf{x}_N, \mathbf{x}_1) + \dots + \beta_N k(\mathbf{x}_N, \mathbf{x}_N) = \pi(\mathbf{x}_N). \end{cases} \quad (4)$$

Denoting $(\mathbf{K})_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$ ($1 \leq i, j \leq N$), $\boldsymbol{\beta} = [\beta_1, \dots, \beta_N]^T$ and $\mathbf{d} = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)]^T$, Eq. (4) can be written in matrix form as $\mathbf{K}\boldsymbol{\beta} = \mathbf{d}$. Thus, the coefficients are given by

$$\boldsymbol{\beta} = \mathbf{K}^{-1}\mathbf{d}. \quad (5)$$

Note that, depending on the choice of kernel and its parameters, these coefficients can be negative.

Remark 1: The only requirement regarding the functions $k(\mathbf{x}, \mathbf{x}')$ is that the interpolation matrix \mathbf{K} must be non-singular (i.e., invertible) for any set of distinct nodes. The symmetry of $k(\mathbf{x}, \mathbf{x}')$ is not required. Different type of bases can be employed, for instance, one for each node \mathbf{x}_i , i.e., $k_i(\mathbf{x}, \mathbf{x}_i)$.

Remark 2: For simplicity, in this first part of the paper, we consider a fixed number of nodes N . However, a key point of the work is the adaptation procedure in Section VI, where new nodes are sequentially added.

A detailed theoretical analysis is provided in Section VII.

B. INTERPOLATIVE QUADRATURE SCHEMES

We can approximate both Z and I by substituting the true $\pi(\mathbf{x})$ with its interpolant $\hat{\pi}(\mathbf{x})$.

1) APPROXIMATION OF Z

Let $\int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = C_i > 0$ be the measure of the i -th kernel. An approximation of Z can be obtained, by substituting Eq. (3) in (2),

$$\hat{Z} = \int_{\mathcal{X}} \hat{\pi}(\mathbf{x}) d\mathbf{x} = \sum_{i=1}^N \beta_i \int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = \sum_{i=1}^N \beta_i C_i. \quad (6)$$

If the kernels are normalized, i.e., $C_i = 1$, note that $\hat{Z} = \sum_{i=1}^N \beta_i$.

Remark 3: Although $Z > 0$, \hat{Z} can take negative values, since the coefficients β_i can be negative. However, in this work, we suggest two schemes (with Gaussian bases and a suitable tuning procedure, and with NN bases) which ensure a positive estimation of Z .

2) APPROXIMATION OF I

By substituting (3) and (6) in (1), we obtain an approximation of I as

$$I \approx \hat{I} = \frac{1}{\hat{Z}} \int_{\mathcal{X}} f(\mathbf{x}) \hat{\pi}(\mathbf{x}) d\mathbf{x}. \quad (7)$$

Note that, given $\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i)$, the approximation of I in (7) can be expressed as

$$\begin{aligned} \hat{I} &= \frac{1}{\hat{Z}} \sum_{i=1}^N \beta_i \int_{\mathcal{X}} f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = \frac{1}{\hat{Z}} \sum_{i=1}^N \beta_i J_i, \\ &= \frac{1}{\hat{Z}} \sum_{i=1}^N v_i \pi(\mathbf{x}_i), \end{aligned} \quad (8)$$

where $J_i = \int_{\mathcal{X}} f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x}$, $\mathbf{v} = [v_1, \dots, v_N]^T = \mathbf{K}^{-1}\boldsymbol{\zeta}$ with $\boldsymbol{\zeta} = [J_1, \dots, J_N]^T$ being the vector of integrals. Clearly, the performance of \hat{I} depends on the discrepancy between $\hat{\pi}(\mathbf{x})$ and $\pi(\mathbf{x})$, as shown by Theorem 1. This discrepancy is reduced by properly adding new nodes, as suggested in Section VI.

C. MONTE-CARLO BASED INTERPOLATIVE QUADRATURE SCHEMES

In this work, we assume that the evaluation of the target function $\pi(\mathbf{x})$ is the main computational bottleneck [19], [21]. We consider that other operations, such as sampling and evaluating different proposal densities, are negligible with respect to the target evaluation. The techniques, presented in this section, do not require additional target evaluations with respect to Eq. (8). In some specific cases, we can compute the integrals J_i and C_i analytically (e.g., see next section). Otherwise, we need to approximate J_i , and in some cases, also C_i . Some general ideas are described below.

1) NORMALIZED KERNELS ($C_i = 1$)

If the values $C_i = 1$ are known,¹ we can compute $\hat{Z} = \frac{1}{N} \sum_{i=1}^N \beta_i$. Moreover, if we are able to draw samples from each $k(\mathbf{x}, \mathbf{x}_i)$, we have

$$J_i = \int_{\mathcal{X}} f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} \approx \hat{J}_i = \frac{1}{M} \sum_{m=1}^M f(\mathbf{z}_{i,m}), \quad (9)$$

with $\mathbf{z}_{i,m} \sim k(\mathbf{x}, \mathbf{x}_i)$, hence

$$\hat{I} \approx \frac{1}{\hat{Z}M} \sum_{i=1}^N \beta_i \sum_{m=1}^M f(\mathbf{z}_{i,m}). \quad (10)$$

If we know C_i , another possible scenario is when we are not able to draw from $k(\mathbf{x}, \mathbf{x}_i)$. In this case, we can employ the importance sampling (IS) procedure described below to approximate the integrals J_i .

¹For the sake of simplicity and without loss of generality, we assume $C_i = 1$.

2) KERNELS WITH UNKNOWN C_i

In this case, we also have to approximate $\int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = C_i$. For this purpose, we can employ IS with proposal densities $q_i(\mathbf{x})$, with $i = 1, \dots, N$, obtaining

$$C_i \approx \widehat{C}_i = \frac{1}{M} \sum_{m=1}^M w_{i,m}, \tag{11}$$

where the weights are $w_{i,m} = \frac{k(\mathbf{z}_{i,m}, \mathbf{x}_i)}{q_i(\mathbf{z}_{i,m})}$ and $\mathbf{z}_{i,m} \sim q_i(\mathbf{x})$. Moreover, we also obtain

$$J_i \approx \widehat{J}_i = \frac{1}{M} \sum_{m=1}^M w_{i,m} f(\mathbf{z}_{i,m}). \tag{12}$$

Replacing (11)-(12) into (8), the final estimator is given by

$$\widehat{I} \approx \frac{1}{\sum_{i=1}^N \beta_i \sum_{m=1}^M w_{i,m}} \sum_{i=1}^N \beta_i \sum_{m=1}^M w_{i,m} f(\mathbf{z}_{i,m}), \tag{13}$$

$$= \sum_{m=1}^M \sum_{i=1}^N \bar{\rho}_{i,m} f(\mathbf{z}_{i,m}), \tag{14}$$

where $\bar{\rho}_{i,m} = \frac{\beta_i w_{i,m}}{\sum_{j=1}^N \sum_{k=1}^M \beta_j w_{j,k}}$.

Remark 4: Note that, in any of the scenarios above, we do not need to evaluate the target $\pi(\mathbf{x})$ at the samples $\mathbf{z}_{i,m}$. Namely, we do not require additional target evaluations with respect to Section II-B. Moreover, as $M \rightarrow \infty$, the estimators in Eqs. (10)-(14) converge to the expression (8), under standard MC arguments [2].

For further details, see the theoretical results in Section VII-B2 and Theorems 6 and 7. So far we have considered Monte Carlo approaches to estimate J_i and C_i . Other particular and more efficient approaches (such as deterministic quadratures) are possible if we consider specific kernel functions. In the next sections, we analyze two specific cases (with Gaussian and NN kernels).

III. INTERPOLATION WITH GAUSSIAN KERNELS

Let us consider the case of Gaussian kernels (with an unbounded support $\mathcal{X} = \mathbb{R}^{d_x}$),

$$k_G(\mathbf{x}, \mathbf{x}_i) = \frac{1}{(2\pi)^{\frac{d_x}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}_i)^\top \Sigma^{-1}(\mathbf{x} - \mathbf{x}_i)\right), \tag{15}$$

where Σ is a positive definite matrix. We take $\Sigma = h^2 \mathbf{I}$ where $h > 0$ is the bandwidth hyperparameter that needs to be tuned (see Section III-A). Alternatively, note that we can also use unnormalized Gaussian kernels $k_G(\mathbf{x}, \mathbf{x}_i) = A \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}_i)^\top \Sigma^{-1}(\mathbf{x} - \mathbf{x}_i)\right)$, where A is another parameter to possibly tune, and then consider $C_i = A(2\pi)^{\frac{d_x}{2}} |\Sigma|^{\frac{1}{2}}$.

Polynomial functions $f(\mathbf{x})$. The use of Gaussian kernel functions $k_G(\mathbf{x}, \mathbf{x}_i)$ with $f(\mathbf{x})$ being polynomial, ensures that the integrals in (8) are available in closed-form.

Let $\mathbf{f}(\mathbf{x}) = \mathbf{x}^r = [x_1^r, \dots, x_{d_x}^r]^\top$ be componentwise powers of $\mathbf{x} \in \mathbb{R}^{d_x}$ ($r = 1, 2, \dots$). Then,

$$J_i = \int_{\mathbb{R}^{d_x}} \mathbf{f}(\mathbf{x}) k_G(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = \int_{\mathbb{R}^{d_x}} \mathbf{x}^r k_G(\mathbf{x}, \mathbf{x}_i) d\mathbf{x},$$

corresponds to the r -th marginal moments of a multivariate Gaussian centered at \mathbf{x}_i . Note that the marginal moments of a Gaussian density are well-known. Some instances are

$$\int_{\mathbb{R}^{d_x}} \mathbf{x} k_G(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = \mathbf{x}_i \quad (r = 1),$$

$$\int_{\mathbb{R}^{d_x}} \mathbf{x}^2 k_G(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} = \mathbf{x}_i^2 + \text{diag}(\Sigma), \quad (r = 2),$$

where the power \mathbf{x}_i^2 is considered a componentwise operation. Then, in this case, we can directly replace the values of J_i in Eq. (8).

Generic functions $f(\mathbf{x})$. Each of the N integrals on the right hand of (8) may be also approximated efficiently with a Gauss-Hermite quadrature (GH) [5], [6], i.e.,

$$\int_{\mathbb{R}^{d_x}} f(\mathbf{x}) k_G(\mathbf{x}, \mathbf{x}_i) d\mathbf{x} \approx \widehat{J}_i = \sum_{m=1}^M \bar{w}_m^{\text{GH}} f(\mathbf{z}_{i,m}),$$

where \bar{w}_m^{GH} and $\mathbf{z}_{i,m}$ are the weights and nodes of the GH quadrature used for i -th integral. Note the quadrature weights are independent of i and are normalized, i.e., $\sum_{m=1}^M \bar{w}_m^{\text{GH}} = 1$. Moreover, we have $\mathbf{z}_{i,m} = \tilde{\mathbf{z}}_m + \mathbf{x}_i$, that is, the only difference is a translation of a single set of GH nodes $\tilde{\mathbf{z}}_m$ [6] (see also the Suppl. Material). Again, we do not need extra evaluations of the target $\pi(\mathbf{x})$. Note that, with enough number of points $\mathbf{z}_{i,m}$, Gauss-Hermite quadrature is also exact when $f(\mathbf{x})$ are polynomial functions [58]. Theoretical results, valid for positive definite radial basis functions, can be found in Section VII-B.

A. PROBABILISTIC INTERPRETATION

If $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}', \mathbf{x})$ (i.e., it is symmetric) and $k(\mathbf{x}, \mathbf{x}')$ is semi positive definite, as in the Gaussian case, we can interpret the construction of the interpolant $\widehat{\pi}(\mathbf{x})$ as a Gaussian process (GP) [59]. In our setting, $\mathbf{d} = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)]^\top$ represents the observed vector. The process starts by placing a GP prior on $\pi(\mathbf{x})$, $\pi(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k(\mathbf{x}, \mathbf{x}'))$, where the GP mean is $\mathbf{0}$ and $k(\mathbf{x}, \mathbf{x}')$ is the covariance function. Conditioning on \mathbf{d} , it can be shown that the posterior of $\pi(\mathbf{x})$ is given by

$$\pi(\mathbf{x}) | \mathbf{d} \sim \mathcal{GP}(\widehat{\pi}(\mathbf{x}), C(\mathbf{x}, \mathbf{x}')),$$

where the mean function is the interpolant $\widehat{\pi}(\mathbf{x})$ given in (3), and the posterior covariance function is $C(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') - \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}')$, with

$$\mathbf{k}(\mathbf{x}) = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_N)]^\top,$$

and $(\mathbf{K})_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. The variance at \mathbf{x} is

$$V(\mathbf{x}) = C(\mathbf{x}, \mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^\top \mathbf{K}^{-1} \mathbf{k}(\mathbf{x}). \tag{16}$$

Observe that $V(\mathbf{x}_i) = 0$ for all $i = 1, \dots, N$. If we assume that the vector of evaluations \mathbf{d} is noisy, we can relax the exact

fit requirement by introducing a regularization term, replacing \mathbf{K} with the matrix $\mathbf{K} + \sigma^2 \mathbf{I}$, where \mathbf{I} is an $N \times N$ identity matrix. The noise term σ^2 also provides numerical stability. The probabilistic interpretation of the integrals involving π is given in Appendix B.

B. TUNING OF HYPERPARAMETERS

Let us denote as θ the vector as hyperparameters of the kernel functions $k(\mathbf{x}, \mathbf{x}')$. A standard way of fitting the hyperparameters θ is to maximize the marginal likelihood of the GP [59]. In this case, the evaluations of $\pi(\mathbf{x})$ play the role of data. Given the evaluations $\mathbf{d} = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)]^\top$, the marginal likelihood is given by $p(\mathbf{d}|\theta) = \mathcal{N}(\mathbf{d}|\theta, \mathbf{K})$, and its log-version is

$$\log p(\mathbf{d}|\theta) = -\frac{1}{2} \mathbf{d}^\top \mathbf{K}^{-1} \mathbf{d} - \frac{1}{2} \log |\mathbf{K}| + c,$$

where c is a constant. Note that \mathbf{K} depends on θ . However, for fitting the bandwidth parameter h of the Gaussian kernels, we propose an alternative procedure described in Appendix A, specifically designed for building an emulator of a *density function*. In this context, the proposed procedure performs better than the maximization of $p(\mathbf{d}|\theta)$.

Remark 5: Using the novel tuning procedure in Appendix A, the corresponding estimator \hat{Z} takes always positive values.

IV. CONSTANT KERNELS BASED ON NEAREST NEIGHBORS

Given the set of nodes $\{\mathbf{x}_i\}_{i=1}^N$ in a bounded domain \mathcal{X} , consider now the use of constant kernels with finite support

$$k(\mathbf{x}, \mathbf{x}_i) = \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}), \quad (17)$$

where $\mathbb{I}_{\mathcal{R}_i}(\mathbf{x})$ is the indicator function in \mathcal{R}_i , i.e., $\mathbb{I}_{\mathcal{R}_i}(\mathbf{x}) = 1$ for all $x \in \mathcal{R}_i$ and zero otherwise. Each \mathcal{R}_i consists of the points $\mathbf{x} \in \mathcal{X}$ that are closest to \mathbf{x}_i , i.e.,

$$\mathcal{R}_i = \{\mathbf{x} \in \mathcal{X} : \|\mathbf{x} - \mathbf{x}_i\|_p \leq \min_{j \neq i} \|\mathbf{x} - \mathbf{x}_j\|_p\},$$

where $\|\cdot\|_p$ denotes the p -norm. That is, $\mathcal{X} = \cup_{i=1}^N \mathcal{R}_i$ is the Voronoi partition of \mathcal{X} using $\{\mathbf{x}_i\}_{i=1}^N$ as support points. In this case, solving (5) for the coefficients β is straightforward since the matrix \mathbf{K} is the identity matrix, and thus

$$\beta_i = \pi(\mathbf{x}_i) \quad \text{for } i = 1, \dots, N.$$

Note that all $\beta_i \geq 0$ with this kernel. Hence the interpolant is given by

$$\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \pi(\mathbf{x}_i) \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}). \quad (18)$$

Note that to evaluate $\hat{\pi}(\mathbf{x})$ at any \mathbf{x} we need to find just the closest node. We do not need to know the borders of regions $\{\mathcal{R}_i\}_{i=1}^N$ for this purpose. This choice of kernels has three clear advantages:

- (i) no need to solve the linear system in (5) since $\mathbf{K} = \mathbf{I}$ and hence $\beta = \mathbf{d}$,

- (ii) the coefficients $\beta = \mathbf{d}$ are always non-negative (this ensures that $\hat{Z} \geq 0$),
- (iii) no need of tuning the bandwidth hyperparameter.

The difficulty, however, is determining the Voronoi partition, as well as the measures $C_i = \int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x}$. We show how to address these issues in Section IV-A. In this case,

$$C_i = \int_{\mathcal{X}} \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}) d\mathbf{x} = |\mathcal{R}_i|,$$

where $|\mathcal{R}_i|$ denotes the measure of the i -th Voronoi region. The approximation of Z is given by

$$\hat{Z} = \sum_{i=1}^N \pi(\mathbf{x}_i) C_i, \quad (19)$$

and Eq. (8) is expressed as

$$\begin{aligned} \hat{I} &= \frac{1}{\hat{Z}} \sum_{i=1}^N \pi(\mathbf{x}_i) \int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x}, \\ &= \frac{1}{\sum_{k=1}^N \pi(\mathbf{x}_k) C_k} \sum_{i=1}^N \pi(\mathbf{x}_i) \int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (20)$$

The convergence of this scheme is guaranteed as N grows, as shown by Theorems 8 and 9. Further theoretical analysis are provided in Section VII-C. Note that we need to estimate the measures C_i , as well as the integrals $\int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x}$ to compute \hat{Z} and \hat{I} . The next section is devoted to this purpose.

A. APPROXIMATING VORONOI REGIONS AND RESULTING ESTIMATORS

In order to approximate C_i , we can generate M uniform vectors $\{\mathbf{z}_m\}_{m=1}^M$ in \mathcal{X} via Monte Carlo sampling or Quasi-Monte Carlo sequences (e.g. a Sobol sequence) [57]. Define the set \mathcal{U}_i as

$$\begin{aligned} \mathcal{U}_i &= \{\mathbf{z}_m : \|\mathbf{z}_m - \mathbf{x}_i\|_p \leq \min_{j \neq i} \|\mathbf{z}_m - \mathbf{x}_j\|_p\} \\ &= \{\mathbf{z}_{\ell_i}\}_{\ell_i=1}^{|\mathcal{U}_i|}, \end{aligned}$$

i.e., the $|\mathcal{U}_i|$ vectors closest to \mathbf{x}_i in p -norm, which form a discrete approximation of \mathcal{R}_i . Note that $\sum_{i=1}^N |\mathcal{U}_i| = M$. Hence, the measure C_i can be approximated by noting that $\frac{C_i}{|\mathcal{X}|} \approx \frac{|\mathcal{U}_i|}{M}$, hence

$$C_i \approx \frac{|\mathcal{U}_i|}{M} |\mathcal{X}|, \quad (21)$$

where $|\mathcal{X}|$ is the measure of \mathcal{X} . Thus, the estimator in Eq. (19) can be rewritten as

$$\hat{Z} \approx \frac{|\mathcal{X}|}{M} \sum_{i=1}^N \pi(\mathbf{x}_i) |\mathcal{U}_i|. \quad (22)$$

We can also obtain an approximation of the integral $J_i = \int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x}$ by leveraging a QMC or MC approximation of the Voronoi regions. Specifically, the uniform vectors \mathbf{z}_{ℓ_i} in \mathcal{U}_i can be used to approximate the integral in (20) as follows

$$J_i = \int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x} \approx \frac{C_i}{|\mathcal{U}_i|} \sum_{\ell_i=1}^{|\mathcal{U}_i|} f(\mathbf{z}_{\ell_i}) \approx \frac{|\mathcal{X}|}{M} \sum_{\ell_i=1}^{|\mathcal{U}_i|} f(\mathbf{z}_{\ell_i}), \quad (23)$$

where we used (21) again in (23). The procedure above can be seen as an accept-reject method, and the estimators are also unbiased [4, Chapter 3 and Section 6.6]. Note that a simpler possible approximation with one point is $J_i = \int_{\mathcal{R}_i} f(\mathbf{x})d\mathbf{x} \approx f(\mathbf{x}_i)C_i$. Thus, replacing the expressions (22)-(23) in (20), the final estimator becomes

$$\hat{I} \approx \frac{1}{\sum_{k=1}^N \pi(\mathbf{x}_k)|\mathcal{U}_k|} \sum_{i=1}^N \pi(\mathbf{x}_i) \sum_{\ell_i=1}^{|\mathcal{U}_i|} f(\mathbf{z}_{\ell_i}). \quad (24)$$

Connection with Section II-C. The estimators above can be interpreted as the application of an importance sampling (IS) scheme as described in Section II-C, for kernel functions with unknown C_i . However, unlike in Section II-C, here we consider a unique and uniform proposal density

$$q_i(\mathbf{x}) = q(\mathbf{x}) = \frac{1}{|\mathcal{X}|} \mathbb{1}_{\mathcal{X}}(\mathbf{x}), \quad \forall i = 1, \dots, N.$$

Then, we can also remove the subindex i in the sample $\mathbf{z}_{i,m} \sim q(\mathbf{x})$, i.e., we have only M samples $\mathbf{z}_m \sim q(\mathbf{x})$. Hence, following Eqs. (11)-(12), we have

$$C_i \approx \frac{1}{M} \sum_{m=1}^M w_{i,m}, \quad (25)$$

$$J_i = \int_{\mathcal{R}_i} f(\mathbf{x})d\mathbf{x} \approx \frac{1}{M} \sum_{m=1}^M w_{i,m}f(\mathbf{z}_m), \quad (26)$$

where $\mathbf{z}_m \sim q(\mathbf{x}) = \frac{1}{|\mathcal{X}|} \mathbb{1}_{\mathcal{X}}(\mathbf{x})$, and the weights are

$$w_{i,m} = \frac{k(\mathbf{z}_m, \mathbf{x}_i)}{q(\mathbf{z}_m)} = \begin{cases} |\mathcal{X}| & \text{if } \mathbf{z}_m \in \mathcal{R}_i, \\ 0 & \text{if } \mathbf{z}_m \notin \mathcal{R}_i. \end{cases} \quad (27)$$

Replacing the expression of the weights $w_{i,m}$ into the formulas above, we recover the estimators in (22) and (24).

V. AN ALTERNATIVE IS INTERPRETATION

In this section, we discuss a special case of the IS scheme given in Section II-C, when a unique proposal $q_i(\mathbf{x}) = q(\mathbf{x})$ is employed and only M samples $\mathbf{z}_m \sim q(\mathbf{x})$ are drawn (as already considered in the previous section). In this scenario, the IS procedure in Section II-C has another relevant interpretation, which allows us to design other different schemes. Considering a generic kernel $k(\mathbf{x}, \mathbf{x}_i)$ and Eq. (25), we can rearrange \hat{Z} as

$$\begin{aligned} \hat{Z} &= \sum_{i=1}^N \beta_i C_i \approx \sum_{i=1}^N \beta_i \frac{1}{M} \sum_{m=1}^M w_{i,m} \\ &= \sum_{i=1}^N \beta_i \frac{1}{M} \sum_{m=1}^M \frac{k(\mathbf{z}_m, \mathbf{x}_i)}{q(\mathbf{z}_m)} \\ &= \frac{1}{M} \sum_{m=1}^M \frac{\sum_{i=1}^N \beta_i k(\mathbf{z}_m, \mathbf{x}_i)}{q(\mathbf{z}_m)}. \end{aligned}$$

Then, recalling that $\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i)$ and replacing this expression above, we finally obtain

$$\hat{Z} \approx \frac{1}{M} \sum_{m=1}^M \frac{\hat{\pi}(\mathbf{z}_m)}{q(\mathbf{z}_m)} = \frac{1}{M} \sum_{m=1}^M \gamma_m, \quad (28)$$

where $\gamma_m = \gamma(\mathbf{z}_m) = \frac{\hat{\pi}(\mathbf{z}_m)}{q(\mathbf{z}_m)}$ for $m = 1, \dots, M$. Moreover, with similar steps, we can obtain

$$\hat{I} \approx \frac{1}{M\hat{Z}} \sum_{m=1}^M \gamma_m f(\mathbf{z}_m), \quad (29)$$

Remark 6: The weights γ_m have the form of the standard IS weights with the target function $\hat{\pi}$ in the numerator, and the proposal density q in the denominator. Hence, the entire sampling procedure can be interpreted as a standard IS scheme where the target function is $\hat{\pi}$ instead of π . This shows again that we do not need extra target evaluations and, hence, we can employ an arbitrary large value of M .

Remark 7: Note that this result is valid for any kernel $k(\mathbf{x}, \mathbf{x}_i)$, and we use a unique proposal $q(\mathbf{x})$ in the procedure described in Section II-C.

Below, we consider the NN case with a uniform proposal $q(\mathbf{x})$, deriving the same formulas in Section IV-A.

A. UNIFORM PROPOSAL DENSITY AND NN INTERPOLATOR

Let us consider $q(\mathbf{x}) = \frac{1}{|\mathcal{X}|} \mathbb{1}_{\mathcal{X}}(\mathbf{x})$, i.e., a uniform density in \mathcal{X} , and the NN kernel function. For each sample \mathbf{z}_m , the corresponding weight γ_m is

$$\gamma_m = \gamma(\mathbf{z}_m) = \frac{\hat{\pi}(\mathbf{z}_m)}{\frac{1}{|\mathcal{X}|}} = \frac{\pi(\mathbf{x}_{k_m})}{\frac{1}{|\mathcal{X}|}} = |\mathcal{X}| \pi(\mathbf{x}_{k_m}),$$

where \mathbf{x}_{k_m} is the closest node to sample \mathbf{z}_m , i.e., $\mathbf{x}_{k_m} = \arg \min_j \|\mathbf{z}_m - \mathbf{x}_j\|_p$. Then, the IS approximation of \hat{Z} is

$$\hat{Z} \approx \frac{1}{M} \sum_{m=1}^M \gamma_m = \frac{|\mathcal{X}|}{M} \sum_{m=1}^M \pi(\mathbf{x}_{k_m}) = \frac{|\mathcal{X}|}{M} \sum_{k=1}^N \pi(\mathbf{x}_k) |\mathcal{U}_k|,$$

where $|\mathcal{U}_k|$ counts the number of \mathbf{z}_m whose closest node is \mathbf{x}_k ($k = 1, \dots, N$). Note that this expression is the same as in (22). Similarly, the IS estimate of \hat{I} is given by

$$\begin{aligned} \hat{I} &\approx \frac{1}{M\hat{Z}} \sum_{m=1}^M \gamma_m f(\mathbf{z}_m) = \frac{|\mathcal{X}|}{M\hat{Z}} \sum_{m=1}^M \pi(\mathbf{x}_{k_m}) f(\mathbf{z}_m) \\ &= \frac{|\mathcal{X}|}{M\hat{Z}} \sum_{k=1}^N \pi(\mathbf{x}_k) \sum_{\ell_k=1}^{|\mathcal{U}_k|} f(\mathbf{z}_{\ell_k}), \end{aligned}$$

which is the same expression as in (24). However, this alternative IS interpretation allows us to design different schemes using a different proposal density, as shown below.




FIGURE 1. Example of application of NN-AQ. The cross-marks represent the starting nodes, while the points added adaptively by NN-AQ are shown with dots. (a) The banana-shaped target and the starting nodes. (b)-(c)-(d) The NN-AQ emulator with $E = 50, 250, 10^3$ number of target evaluations.

B. GAUSSIAN MIXTURE PROPOSAL

We consider now an alternative to the uniform proposal in \mathcal{X} . More specifically, we propose drawing $\{\mathbf{z}_\ell\}_{\ell=1}^M$ from a Gaussian mixture proposal pdf built considering the set of nodes $\{\mathbf{x}_i\}_{i=1}^N$, i.e.,

$$\mathbf{z}_m \sim q(\mathbf{x}) = \sum_{i=1}^N \xi_i \mathcal{N}(\mathbf{x}|\mathbf{x}_i, \mathbf{C}_i),$$

where the mixture weights ξ_i are

$$\xi_i = \frac{\pi(\mathbf{x}_i)}{\sum_{n=1}^N \pi(\mathbf{x}_n)}, \quad i = 1, \dots, N,$$

and the covariances \mathbf{C}_i can be determined by the minimum distance of \mathbf{x}_i to its closest node. In this case, the IS weights are given by

$$\gamma(\mathbf{z}_m) = \frac{\widehat{\pi}(\mathbf{z}_m)}{\sum_{i=1}^N \xi_i \mathcal{N}(\mathbf{z}_m|\mathbf{x}_i, \mathbf{C}_i)} = \frac{\pi(\mathbf{x}_{k_m})}{\sum_{i=1}^N \xi_i \mathcal{N}(\mathbf{z}_m|\mathbf{x}_i, \mathbf{C}_i)},$$

where \mathbf{x}_{k_m} is the closest node to \mathbf{z}_m , with $m = 1, \dots, M$.

VI. ADAPTIVE PROCEDURE

In this section, we present an adaptive mechanism to add new nodes to the interpolant. Our algorithm adds nodes sequentially with the aim to discover high-valued regions of $\pi(\mathbf{x})$ while fostering the exploration of the state space. We employ an active learning procedure where a new point is obtained by maximizing a suitable acquisition function. The resulting adaptive algorithm is shown in Table 1. Note that the final number of nodes is $N_T = T + N_0$. The adaptive quadrature scheme based on the Gaussian kernels is denoted as GK-AQ, whereas the other scheme based on the Nearest Neighbors (NN) kernels is denoted as NN-AQ. Figure 1 depicts an example of application of the NN-AQ.

A. BUILDING SUITABLE ACQUISITION FUNCTIONS

Let us denote as $t \in \mathbb{N}$ the t th iteration of the algorithm. In the update stage, we decide to add a new node where the acquisition function, $A_t : \mathcal{X} \rightarrow \{0\} \cup \mathbb{R}^+$, is maximum. The acquisition function takes into account the shape of $\pi(\mathbf{x})$ and the spatial distribution of the current nodes. More specifically,

TABLE 1. Adaptive Quadrature Algorithm.

<p>Initialization: Set N_0 initial nodes and set $\mathbf{X}_0 = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_0}\}$, $\mathbf{d}_0 = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_{N_0})]^\top$.</p> <p>For $t = 0, \dots, T$:</p> <ol style="list-style-type: none"> 1) <i>Build the interpolator.</i> Use the set $\mathbf{X}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_t}\}$ and corresponding evaluations $\mathbf{d}_t = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_{N_t})]^\top$ to build $\widehat{\pi}_t(\mathbf{x})$ using Gaussian kernels (see Section III) or constant kernels (see Section IV). 2) <i>Build the acquisition function.</i> Use $\widehat{\pi}_t(\mathbf{x})$ and the set of current nodes \mathbf{X}_t to build the acquisition function $A_t(\mathbf{x})$, e.g., Eqs. (33)-(34). 3) <i>Update stage.</i> Obtain new node $\mathbf{x}_{N_{t+1}}$ by $\mathbf{x}_{N_{t+1}} = \arg \max_{\mathbf{x} \in \mathcal{X}} A_t(\mathbf{x}), \quad (31)$ append $\mathbf{X}_{t+1} = \{\mathbf{X}_t, \mathbf{x}_{N_{t+1}}\}$ and $\mathbf{d}_{t+1} = [\mathbf{d}_t, \pi(\mathbf{x}_{N_{t+1}})]^\top$. <p>Outputs: Build the final interpolant $\widehat{\pi}_T(\mathbf{x})$ and obtain the approximations \widehat{I} and \widehat{Z}.</p>
--

it must fulfill

$$A_t(\mathbf{x}_i) = 0 \quad \text{for all } t \text{ and } i = 1, \dots, N_t,$$

and grow as we move apart from the nodes. We consider acquisition functions $A_t(\mathbf{x})$ of the form

$$A_t(\mathbf{x}) = \pi(\mathbf{x})D_t(\mathbf{x}), \quad (31)$$

where $D_t(\mathbf{x})$ is a diversity term that penalizes the proximity to the current nodes. Note that the information of $f(\mathbf{x})$ could be also included as $A_t(\mathbf{x}) = f(\mathbf{x})\pi(\mathbf{x})D_t(\mathbf{x})$. In some settings, the function $A_t(\mathbf{x})$ above could be directly used after choosing a diversity term $D_t(\mathbf{x})$. However, in this work, we consider that evaluating $\pi(\mathbf{x})$ is costly, so we propose cheaper versions of (31).

B. CHEAP ACQUISITION FUNCTIONS

We recall that the most costly step is the evaluation of the target function $\pi(\mathbf{x})$. This is often due to the use of complex models and/or large amounts of data. For that reason, we propose a cheap type of $A_t(\mathbf{x})$,

$$A_t(\mathbf{x}) = \widehat{\pi}_t(\mathbf{x})D_t(\mathbf{x}), \quad (32)$$

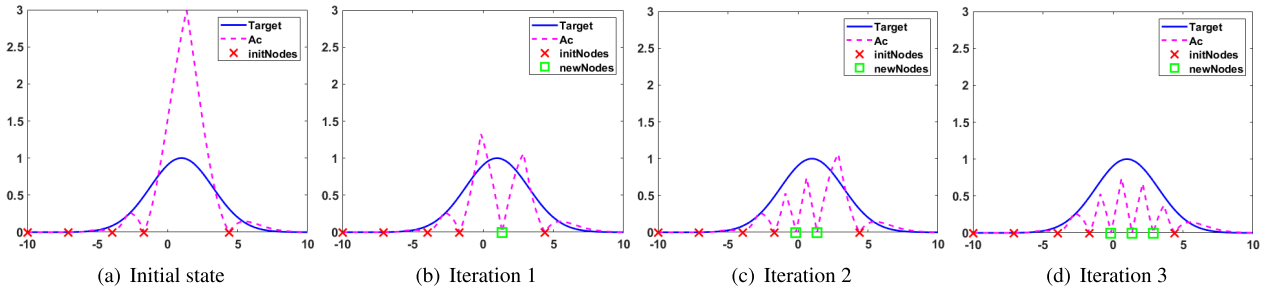


FIGURE 2. 1D example of application of $A_t(\mathbf{x}) = \pi(\mathbf{x})D_t(\mathbf{x})$ with the diversity term $D_t(\mathbf{x}) = \min_{i=1, \dots, N_t} |\mathbf{x} - \mathbf{x}_i|$. At each iteration, the new node, shown with a green square, is added where $A_t(\mathbf{x})$ is maximum.

so that no evaluations of the true $\pi(\mathbf{x})$ are required. In this case, in terms of posterior evaluations E , the cost of the overall algorithm in Table 1, is $E = N_0 + T$.

Remark 8: The particular case $A_t(\mathbf{x}) = D_t(\mathbf{x})$ corresponds to the space-filling experimental designs (e.g., see [46], [56], [60] and Theorem 4). In the other particular case with $A_t(\mathbf{x}) = \hat{\pi}_t(\mathbf{x})$, the resulting schemes are similar to other approaches in literature which combine sampling and optimization (e.g., see [9]).

In the Gaussian kernel scenario, we may use the variance in (16) as diversity term

$$A_t(\mathbf{x}) = \hat{\pi}_t(\mathbf{x})V_t(\mathbf{x}), \quad (33)$$

where we have set $D_t(\mathbf{x}) = V_t(\mathbf{x})$, that fulfills $V_t(\mathbf{x}_i) = 0$ for $i = 1, \dots, N_t$. This choice is motivated by the fact that the approximation error is bounded by the maximum value of $V_t(\mathbf{x})$ (e.g., see Theorem 3). Since the function $V_t(\mathbf{x})$ is unfeasible with constant NN kernels, we suggest a diversity term of the form

$$A_t(\mathbf{x}) = \hat{\pi}_t(\mathbf{x}) \min_{i=1, \dots, N_t} \|\mathbf{x} - \mathbf{x}_i\|_p. \quad (34)$$

Note that the term $D_t(\mathbf{x}) = \min_{i=1, \dots, N_t} \|\mathbf{x} - \mathbf{x}_i\|_p$ is zero when evaluated at any current node: for each $\mathbf{x}_j \in \mathbf{X}_t$ the minimum distance is w.r.t. itself, which is zero. This choice is motivated by Theorem 4, since the approximation error is also bounded by the maximum value of $D_t(\mathbf{x})$. Figure 2 depicts an example with this choice of $D_t(\mathbf{x})$. Note that the choice $D_t(\mathbf{x}) = \min_{i=1, \dots, N_t} \|\mathbf{x} - \mathbf{x}_i\|_p$ can be also employed in the Gaussian kernel scenario.

Another alternative is to consider tempering versions of the acquisition function,

$$A_t(\mathbf{x}) = [\hat{\pi}_t(\mathbf{x})]^\alpha [D_t(\mathbf{x})]^\beta, \quad (35)$$

where $\alpha \geq 0$ can be used to prioritize moving towards high-valued zones of $\hat{\pi}_t(\mathbf{x})$, while $\beta \geq 0$ to encourage exploration. The values α and β can also vary with the iteration t . The maximization of $A_t(\mathbf{x})$ can be performed by simulated annealing or other optimization techniques. The performance of different acquisition functions have been compared in Figure 4 (see Section VIII-A). One can observe that maximizing the

proposed acquisition functions provides much better results than adding uniformly random nodes.

Observations: For the GK-AQ algorithm, the most costly step corresponds to the inversion of the $N_t \times N_t$ matrix \mathbf{K}_t , needed to be done in order to build the acquisition function in Eq. (33). Note that the inverse \mathbf{K}_t^{-1} is used for both evaluating the interpolant $\hat{\pi}_t(\mathbf{x})$ and computing the variance $V_t(\mathbf{x})$. We can alleviate the cost of this step by building \mathbf{K}_t^{-1} iteratively from \mathbf{K}_{t-1}^{-1} . The recursion formula is given in Appendix C. In the case of NN-AQ, evaluating the acquisition function in (34) requires only to calculate the distances with respect to each node. This computation can be used for both evaluating the interpolant and the diversity term $D_t(\mathbf{x}) = \min \|\mathbf{x} - \mathbf{x}_i\|_p$. Note that the cost of searching for the nearest neighbor has only a weak dependence on the dimension of the space.

VII. THEORETICAL SUPPORT

In this section, we provide some theoretical results supporting the proposed schemes. We consider $\tilde{\pi}(\mathbf{x}) = \frac{1}{Z}\pi(\mathbf{x})$ a bounded target pdf and a bounded domain $\mathcal{X} \subset \mathbb{R}^{d_x}$. Let also $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ be an integrable function. In this section, we consider $J = \int_{\mathcal{X}} f(\mathbf{x})\pi(\mathbf{x})d\mathbf{x}$ as the integral of interest. For a generic $f(\mathbf{x})$, J corresponds to the numerator of the integral I in Eq. (1). For $f(\mathbf{x}) = 1$, J becomes the normalizing constant of $\pi(\mathbf{x})$, i.e., $J = Z$, which is the denominator of I . Thus, working with J is equivalent to working with I . Let also $\tilde{J} = \int_{\mathcal{X}} f(\mathbf{x})\hat{\pi}(\mathbf{x})d\mathbf{x}$, be the approximation of J given by substituting the interpolant $\hat{\pi}(\mathbf{x})$. A first general result valid for any interpolation procedure is given below.

Theorem 1: The error incurred by substituting $\pi(\mathbf{x})$ with $\hat{\pi}(\mathbf{x})$ in J is bounded,

$$\begin{aligned} |J - \tilde{J}| &\leq \|f(\pi - \hat{\pi})\|_1 \\ &\leq \|f\|_2 \|\pi - \hat{\pi}\|_2 \\ &\leq |\mathcal{X}| \|f\|_\infty \|\pi - \hat{\pi}\|_\infty, \end{aligned}$$

where $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_\infty$ denote the L^1 , L^2 and L^∞ norms respectively.

Proof: See Appendix D-A. \square

Therefore, if are able to build an interpolant $\hat{\pi}$ in a way such $\|\pi - \hat{\pi}\|_\infty$ vanishes to zero, then the approximation \tilde{J}

will converge to J . Note that, in this section, we ensure the convergence of numerator J and denominator Z of $I = \frac{J}{Z}$, independently. A complete treatment (yet more complicated) should consider the convergence of the two quantities at the same time. For the rest of results, we need to distinguish between the case of Gaussian kernel and constant kernel interpolators. To establish convergence of both schemes we need to make some preliminary definitions and considerations.

A. SPACE-FILLING MEASURES AND RELATED RESULTS

We introduce two well-known measures of dispersion widely employed in the function approximation literature. In this section, we always consider a bounded support \mathcal{X} .

1) FILL DISTANCE

Given the set of nodes $\{\mathbf{x}_i\}_{i=1}^N \subset \mathcal{X}$, let us define the following quantity

$$r = \max_{\mathbf{x} \in \mathcal{X}} \min_{1 \leq i \leq N} \|\mathbf{x} - \mathbf{x}_i\|_2, \tag{36}$$

which is the fill distance.

2) SEPARATION DISTANCE

The separation distance is defined as

$$s = \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\|_2, \tag{37}$$

i.e., the minimal distance between two nodes. Note that $s \leq 2r$. Having a small s increases the numerical instability and can have a detrimental effect in the error bounds. The adaptive procedure described in Sect. VI produces a sequence of nodes that sequentially minimizes r .

Proposition 1: Consider the acquisition function given in Eq. (35) with $\alpha = 0$ and $\beta = 1$, and the choice $A_t(\mathbf{x}) = \min_{i=1, \dots, N_t} \|\mathbf{x} - \mathbf{x}_i\|_2$, where $\{\mathbf{x}_i\}_{i=1}^{N_t}$ are the current nodes of the interpolator. The maximum of this function is the fill distance r_t in Eq. (36), at iteration t . Adding the point \mathbf{x}_{N_t+1} corresponding to r_t to the set of current nodes ensures that

$$r_{t+1} = \max_{i=1, \dots, N_{t+1}} \min_{1 \leq i \leq N_{t+1}} \|\mathbf{x} - \mathbf{x}_i\|_2 \leq r_t,$$

and that $r_t \rightarrow 0$ when $t \rightarrow \infty$.

Proof: See Sect. 4.1 in [56] and [45]. This procedure is related to the ‘‘coffee house design’’ in [60]. \square

Proposition 2: For isotropic kernels, the variance function $V(\mathbf{x})$ given in Eq. (16) satisfies that $\max_{\mathbf{x} \in \mathcal{X}} [V(\mathbf{x})]^{1/2} \leq \Phi(r)$, where $\Phi(r)$ is an increasing function of r , depending on the kernel function. In the case of Gaussian kernels, $\Phi(r)$ is an exponential function.

Proof: See Sect. 2.1 in [56] and Sect. 2 in [45]. \square

Proposition 3: Consider the acquisition function given in Eq. (35) with $\alpha = 0$ and $\beta = 1$, i.e., and the choice $A_t(\mathbf{x}) = V_t(\mathbf{x})$. Let us set also $\varphi_t = \max_{\mathbf{x} \in \mathcal{X}} V_t(\mathbf{x})$. By adding new nodes according to the rule

$$\mathbf{x}_{N_t+1} = \arg \max A_t(\mathbf{x}),$$

we are minimizing φ_t over the iterations t , i.e., φ_t is a non-increasing function of t and $\varphi_t \rightarrow 0$ as $t \rightarrow \infty$.

Proof: This algorithm is known as p -greedy algorithm in [61]. See the behavior of the variance of a GP interpolant [59]. This acquisition function is commonly used in the kriging literature. For instance, see [62] and [46]. \square

Proposition 4: Consider the acquisition function given in Eq. (34) with $\alpha = 0$ and $\beta = 1$, and the choice $A_t(\mathbf{x}) = \min_{i=1, \dots, N_t} \|\mathbf{x} - \mathbf{x}_i\|_2$, where $\{\mathbf{x}_i\}_{i=1}^{N_t}$ are the current nodes of the interpolator. The sequence of nodes obtained as $\mathbf{x}_{N_t+1} = \arg \max A_t(\mathbf{x})$, for $t \in \mathbb{N}^+$, is a uniform low-discrepancy sequence in a bounded \mathcal{X} [63].

Proof: This procedure can be interpreted as deterministic and sequential version of the well-known latin hypercube sampling (LHS) [63]. \square

Remark 9: Note that the proposed schemes do not need that the space is covered uniformly. The only requirement, for decreasing the fill distance r , is to be able to reach any subset of the domain \mathcal{X} with a non-null probability (strictly positive).

B. RESULTS FOR INTERPOLATORS BASED ON RADIAL BASIS FUNCTIONS (RBFs)

In this section, we consider that $k(\mathbf{x}, \mathbf{x}')$ is the Gaussian kernel considered in Sect. III. More generally, the results from this section are valid for any $k(\mathbf{x}, \mathbf{x}')$ that is a (positive definite) radial basis function (RBF).

1) EXACT COMPUTATION OF J_j

Recall $\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i)$, where the weights are $\boldsymbol{\beta} = [\beta_1, \dots, \beta_N] = \mathbf{K}^{-1} \mathbf{d}$ using the interpolation matrix \mathbf{K} and the vector of target evaluations \mathbf{d} . The approximation \tilde{J} can be written as

$$\tilde{J} = \int_{\mathcal{X}} f(\mathbf{x}) \hat{\pi}(\mathbf{x}) d\mathbf{x} = \sum_{i=1}^N \beta_i J_i = \sum_{i=1}^N v_i \pi(\mathbf{x}_i),$$

where $J_i = \int_{\mathcal{X}} f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x}$, and the weights $\mathbf{v} = [v_1, \dots, v_N]^T$ are given by $\mathbf{v} = \mathbf{K}^{-1} \boldsymbol{\zeta}$ with $\boldsymbol{\zeta}$ being the vector of J_i 's. In this form, \tilde{J} is expressed as a combination of evaluations of $\pi(\mathbf{x})$, i.e., a quadrature. The following theorem establishes that the weights $\mathbf{v} = \mathbf{K}^{-1} \boldsymbol{\zeta}$ are optimal for a quadrature of this kind. Note that the Gaussian kernels are symmetric positive definite functions, and are special cases of radial basis functions (RBF).

Theorem 2: Let us consider a symmetric kernel function $k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_j, \mathbf{x}_i)$ which always defines a positive definite matrix \mathbf{K} . The native space related to $k(\mathbf{x}, \mathbf{x}')$ is a reproducing kernel Hilbert space (RKHS) [64], [65]. Given the points $\{\mathbf{x}_i\}_{i=1}^N$ and $\mathbf{v} = \mathbf{K}^{-1} \boldsymbol{\zeta}$, the quadrature $\tilde{J} = \sum_{i=1}^N v_i \pi(\mathbf{x}_i)$ is optimal in the sense of Golomb-Weinberg [66], i.e., the weights v_i minimizes the norm of the integration error functional in the dual space [64], [65].

Proof: A sketch of the proof is in App. D-B. See also [67] and [16] and references therein. \square

Theorem 3: Suppose that $\pi(\mathbf{x})$ belongs to the RKHS generated by the kernel function $k(\mathbf{x}, \mathbf{x}')$. The interpolant

$\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i)$ satisfies $|\pi(\mathbf{x}) - \hat{\pi}(\mathbf{x})| \leq \|\pi\|_{\mathcal{H}} [V(\mathbf{x})]^{\frac{1}{2}}$ for all $\mathbf{x} \in \mathcal{X}$ and hence $\|\pi - \hat{\pi}\|_{\infty} \leq \|\pi\|_{\mathcal{H}} \max_{\mathbf{x} \in \mathcal{X}} [V(\mathbf{x})]^{\frac{1}{2}}$, where $\|\cdot\|_{\mathcal{H}}$ denotes the norm in the RKHS, and $V(\mathbf{x})$ is the variance function given in Eq. (16). Hence, from Theorem 1, we have

$$|J - \tilde{J}| \leq |\mathcal{X}| \|f\|_{\infty} \|\pi\|_{\mathcal{H}} \max_{\mathbf{x} \in \mathcal{X}} [V(\mathbf{x})]^{\frac{1}{2}}.$$

Proof: See Sect. 2.1 in [56] and Sect. 2 in [45]. \square

The theorem above, jointly with Proposition 3, justify the choice of the diversity term $D_t(\mathbf{x}) = V_t(\mathbf{x})$ in Section VI-B. The next theorem, based on results from the literature on approximating functions with RBFs, establishes that the approximation error tends to zero when $r \rightarrow 0$, and that the rate of convergence can be exponentially fast in the case of infinitely smooth RBFs, such as the Gaussian kernels.

Theorem 4: The error of the quadrature \tilde{J} is

$$|J - \tilde{J}| \leq |\mathcal{X}| \|f\|_{\infty} \|\pi - \hat{\pi}\|_{\infty} = \mathcal{O}(\lambda(r)),$$

where $\lambda(r) \rightarrow 0$ as $r \rightarrow 0$, with r being the fill distance given in Eq. (36). The convergence rate depends on the regularity degree of $\pi(\mathbf{x})$. For $\pi(\mathbf{x})$ sufficiently regular (technically, belonging to the RKHS induced by the RBF kernel), and Gaussian RBF the bound $\lambda(r)$ decreases exponentially

$$\lambda(r) = e^{-c_h |\log r|/r},$$

with a certain constant $c_h > 0$, which generally depends on the bandwidth h .

Proof: See Sect. 11.3 and table in page 188 of [55]. \square

Recall that the diversity term in (34) produces a monotonically decreasing sequence of fill distances that converges to zero in the limit of $t \rightarrow \infty$, as stated in Proposition 1. The next theorem states that the approximation error tends to zero as $N \rightarrow \infty$, and provides a quite pessimistic upper bound.

Theorem 5: Given a sequence of nodes $\{\mathbf{x}_i\}_{i=1}^N$ generated as in Proposition 4, it can be shown that $r \leq C_{d_x, \mathcal{X}} N^{-1/d_x} \log N$, where $C_{d_x, \mathcal{X}}$ is a constant that probably depends on the dimension d_x and the measure of \mathcal{X} . Then, the following (pessimistic) upper bound can be provided

$$|J - \tilde{J}| = \mathcal{O} \left(e^{-c_1 \frac{1}{N^{-1/d_x} \log N}} - c_2 \frac{|\log(N^{-1/d_x} \log N)|}{N^{-1/d_x} \log N} \right),$$

where $c_1 > 0$ and $c_2 > 0$ are constants depending on h , d_x and the measure of \mathcal{X} .

Proof: See Sect. 2.5.1 in [56] and [63]. \square

2) NOISY COMPUTATION OF J_j

Theorem 4 above states that the convergence of \tilde{J} is achieved when the fill distance r goes to zero. Recall that in $\tilde{J} = \sum_{i=1}^N \beta_i J_i$ we consider the exact computation of $J_i = \int_{\mathcal{X}} f(\mathbf{x}) k(\mathbf{x}, \mathbf{x}_i) d\mathbf{x}$. In this section, we consider of approximating J_i by the estimator \hat{J}_i , so that we finally have a noisy version of \tilde{J} , i.e., $\hat{\tilde{J}} = \sum_{i=1}^N \beta_i \hat{J}_i$. Below, we show some results related to $\hat{\tilde{J}}$, but we need some previous definitions.

a: STABILITY

The numerical stability of the solution depends on the inversion of the interpolation matrix \mathbf{K} and it is connected to the separation distance s . Clearly, if two nodes are very close, then the corresponding two rows of the interpolation matrix are almost identical and the matrix becomes ill-conditioned [55], [68].

b: REPRODUCTION QUALITY

Roughly speaking, an interpolant built with more nodes (i.e., N grows) filling the space, generally yields a better approximation. This concept is connected to the fill distance r in Eq. (36). Recall that the fill distance is a measure of how well the data fills the space [55].

c: UNCERTAINTY PRINCIPLE

A typical problem when reconstructing functions is the trade-off between reproduction quality and numerical stability. Let us consider RBF kernels with a fixed bandwidth, as N grows. Generally, when one aims at a very good approximation of the function of interest, the numerical stability gets compromised, and conversely, if one aims to have good numerical stability, the approximation will be poor. This is known in the literature as uncertainty principle [68].

Let us denote as h the parameter which controls the bandwidth of the RBFs, as $\Sigma = h^2 \mathbf{I}$ in the Gaussian kernel. The next theorem illustrates the case where the numerical instability combined with the error in computing the vector of integrals $\xi = [J_1, \dots, J_N]^T$ deteriorates the error bound of Theorem 4 (for a fixed h). Let us denote the vector of approximated integrals by $\hat{\xi} = [\hat{J}_1, \dots, \hat{J}_N]^T$ and recall $\mathbf{d} = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)]^T$ is the vector of evaluations of π .

Theorem 6 (For a Fixed Bandwidth h): Let us consider a bounded support \mathcal{X} . If we take into account the error in the evaluation of the integrals $\xi = [J_1, \dots, J_N]^T$, denoted by $\hat{\xi} = [\hat{J}_1, \dots, \hat{J}_N]^T$, the corresponding approximation $\hat{\tilde{J}} = \sum_{i=1}^N \beta_i \hat{J}_i$ has an error of

$$|J - \hat{\tilde{J}}| \leq |\mathcal{X}| \|f\|_{\infty} \|\pi - \hat{\pi}\|_{\infty} + \|\mathbf{K}^{-1}\|_2 \|\mathbf{d}\|_2 \|\xi - \hat{\xi}\|_2 = \mathcal{O}(\lambda(r)) + \mathcal{O}(v(s, h)) \|\xi - \hat{\xi}\|_2,$$

where $\lambda(r) \rightarrow 0$ as $r \rightarrow 0$, $v(s, h) \rightarrow \infty$ as $s \rightarrow 0$, with r and s being, respectively, the fill distance and separation distance given in Eqs. (36) and (37). The parameter h , which determines the bandwidth of the radial kernel, is considered fixed. The function $v(s, h)$ is an upper bound for $\|\mathbf{K}^{-1}\|_2$, which is a measure of stability (note that $\|\mathbf{K}^{-1}\|_2$ corresponds to the inverse of the lowest eigenvalue of \mathbf{K}).

Proof: See Appendix D-C. For the bound $v(s, h)$ see Corollary 12.4 in [55]. \square

The bound in Theorem 6 expresses the uncertainty relation. Indeed, we see that making $s \rightarrow 0$ poses a problem if we use a fixed bandwidth h . Indeed, the interpolation matrix \mathbf{K} becomes ill-conditioned as two nodes are too close, and the error $\|\xi - \hat{\xi}\|_2$ is amplified. The growing rate of $v(s, h)$, as $\lambda(r)$, depends on the smoothness of the RBF. For Gaussian

kernels, the rates of $\nu(s, h)$ and $\lambda(r)$ are both exponential. However, with a Monte Carlo approximation, we can always improve the approximation $\hat{\zeta}$ by increasing the number of samples M , so that $\|\zeta - \hat{\zeta}\|_2 \rightarrow 0$. Recall that the increase of the number of Monte Carlo samples M does not require additional evaluations of the target π in the proposed schemes. Furthermore, even with a fixed M , we can control the value $\|\mathbf{K}^{-1}\|_2$ by decreasing the bandwidth h of the kernel function. The following results consider these two cases.

Theorem 7 (For a Fixed Bandwidth h and $M \rightarrow \infty$): Given a bounded support \mathcal{X} , consider the application of a Monte Carlo method to approximate ζ , then $\|\zeta - \hat{\zeta}\|_2 \rightarrow 0$ as $M \rightarrow \infty$, where M is the number of samples. Hence, the approximation $\hat{J} = \sum_{i=1}^N \beta_i \hat{J}_i$ has an error

$$|J - \hat{J}| = \mathcal{O}(\lambda(r)),$$

where $\lambda(r) \rightarrow 0$ as the fill distance $r \rightarrow 0$ and $M \rightarrow \infty$.

Proof: The term $\|\zeta - \hat{\zeta}\|_2 \rightarrow 0$ as the number of Monte Carlo samples $M \rightarrow \infty$ [2]. \square

Conjecture 1 (For a Decreasing Bandwidth h and Fixed M) Given a bounded support \mathcal{X} , consider a noisy approximation $\hat{\zeta}$ of ζ . Assume that we decrease h as the number of nodes N grows (in order to control the instability term, i.e., the magnitude of $\|\mathbf{K}^{-1}\|_2$). Hence, the approximation $\hat{J} = \sum_{i=1}^N \beta_i \hat{J}_i$ has an error

$$|J - \hat{J}| = \mathcal{O}(\lambda(r)) + b,$$

where b is some constant bias, $\lambda(r) \rightarrow 0$ as $r \rightarrow 0$, and making $h \rightarrow 0$ when $N \rightarrow \infty$.

Note that, as h approaches 0, the interpolation matrix \mathbf{K} becomes a diagonal matrix, with the maximum values of the kernels in the diagonal. Thus, controlling the maximum values of the kernel functions, we can control the minimum value of the eigenvalues, such that the interpolation matrix \mathbf{K} be well-conditioned. Moreover, recall that we are using an interpolative approach and the probabilistic interpretation in Section III-A is not strictly required. Therefore, we have more flexibility in the choice and/or tuning of the kernel functions. Indeed, one could consider different bandwidths (one for each kernel function), bigger in regions with lower density of points, while smaller bandwidths in regions with a higher density of nodes. This would improve the numerical stability.

Remark 10: The interpolant based on NN kernels does not suffer the uncertainty problem, since they have compact non-overlapping supports. Namely, we can interpret that the bandwidths are automatically tuned.

C. RESULTS FOR LOCAL INTERPOLATORS

In a local interpolation method, the addition and/or a change of one node, only affects the solution in a subset of the support domain. This scenario corresponds to the use of the constant NN kernels. Recall that the interpolant based on

constant kernels,

$$\hat{\pi}(\mathbf{x}) = \sum_{i=1}^N \pi(\mathbf{x}_i) \mathbb{1}_{\mathcal{R}_i}(\mathbf{x}),$$

where \mathcal{R}_i denotes the Voronoi region associated with node \mathbf{x}_i . Let us first state a result for sufficiently smooth $\pi(\mathbf{x})$. If $\pi(\mathbf{x})$ is Lipschitz continuous, i.e., for all $\mathbf{x}, \mathbf{z} \in \mathcal{X}$ we have $|\pi(\mathbf{z}) - \pi(\mathbf{x})| \leq L_0 \|\mathbf{z} - \mathbf{x}\|$ for some constant L_0 , then we have the following result.

Theorem 8: Given the NN interpolant $\hat{\pi}(\mathbf{x})$, if $\pi(\mathbf{x})$ is Lipschitz continuous we have that $\|\pi - \hat{\pi}\|_\infty \leq L_0 r$, where L_0 is the Lipschitz constant and r is the fill distance introduced in Eq. (36). Then, from Theorem 1, we have

$$|J - \tilde{J}| \leq |\mathcal{X}| \|f\|_\infty L_0 r.$$

Moreover, given a sequence of nodes $\{\mathbf{x}_i\}_{i=1}^N$ generated as in Proposition 4, and since $r \leq C_{d_x, \mathcal{X}} N^{-1/d_x} \log N$, we have the following (pessimistic) bound

$$|J - \tilde{J}| = \mathcal{O}\left(N^{1/d_x} \log N\right).$$

Proof: See Appendix D-D. \square

Now, recall the approximation of \tilde{J} given by

$$\tilde{J} = \int_{\mathcal{X}} f(\mathbf{x}) \hat{\pi}(\mathbf{x}) d\mathbf{x} \approx S_N = \sum_{i=1}^N \pi(\mathbf{x}_i) f(\mathbf{x}_i) C_i,$$

where S_N is the Riemann approximation, which has been also discussed in Sect. IV-A, and $C_i = \int_{\mathcal{R}_i} d\mathbf{x}$, i.e., the measure of \mathcal{R}_i . Here, we used the approximation $\int_{\mathcal{R}_i} f(\mathbf{x}) d\mathbf{x} \approx f(\mathbf{x}_i) C_i$. We will show that S_N converges to $J = \int_{\mathcal{X}} f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}$ as we add more nodes according to one of the proposed acquisition functions, that is, as $t \rightarrow \infty$. As with Gaussian kernels, the convergence is related with how well the nodes fill space. Here, the role of fill distance is played by the maximum of the measures C_i . The theorem below states that, as we fill the space, the measures C_i converges to zero. Recall that the Voronoi partition $\{\mathcal{R}_i\}_{i=1}^N$ generated from the set of nodes $\{\mathbf{x}_i\}_{i=1}^N$ corresponds to the subdivision of \mathcal{X} in N non-overlapping pieces.

Proposition 5: Consider a sequence of points $\mathbf{x}_1, \dots, \mathbf{x}_N$ covering the space \mathcal{X} , then for the associated Voronoi regions \mathcal{R}_i , we have that $\max_i C_i \rightarrow 0$ as $N \rightarrow \infty$.

Proof: See the proofs of Theorems 1 and 4 in [69]. \square

Theorem 9: Let $\pi(\mathbf{x})$ be a continuous and bounded target pdf (up to a normalizing constant) defined on a bounded support $\mathcal{X} \subset \mathbb{R}^{d_x}$. Let $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ bounded on \mathcal{X} . Consider the integral $J = \int_{\mathcal{X}} f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}$. Let us consider a Voronoi partition of \mathcal{X} , generated by the nodes $\{\mathbf{x}_i\}_{i=1}^N$, defined as $\mathcal{R}_1, \dots, \mathcal{R}_N$ (recall that $C_i = |\mathcal{R}_i|$). Given the Riemann sum $S_N = \sum_{i=1}^N f(\mathbf{x}_i) \pi(\mathbf{x}_i) C_i$, the convergence of $S_N \rightarrow J$ is guaranteed as $\max_i C_i \rightarrow 0$ when $N \rightarrow \infty$.

Proof: See Sect 8.3 in [70]. \square

Above, we have assumed that C_i are known. However, we can have very accurate Monte Carlo estimates without requiring additional evaluations of the target $\pi(\mathbf{x})$ (but just




FIGURE 3. (a) Rel-MSE in log-scale for Z as function of number of target evaluations E . (b) Rel-MSE in log-scale for μ as function of number of target evaluations E . (c) Rel-MSE in log-scale for estimating $[\sigma_1^2, \sigma_2^2]$ as function of number of target evaluations E .

of the interpolant $\hat{\pi}(\mathbf{x})$), i.e., only with a slight increase in the overall computation cost.

VIII. NUMERICAL EXPERIMENTS

In this section, we provide several numerical tests in order to show the performance of the proposed adaptive quadrature schemes and compare them with benchmark approaches in the literature. The first example corresponds to a nonlinear banana-shaped density in dimension $d_x = 2, 3, 4$ and 5 . The second test is a multimodal scenario with dimension $d_x = 10$. Finally, we test our schemes in a challenging astronomic inference problem of detecting the number of exoplanets orbiting a star.

A. BANANA TARGET

As a first example, we consider a banana-shaped target pdf,

$$\bar{\pi}(\mathbf{x}) \propto \exp \left\{ -\frac{(\eta_1 - Bx_1 - x_2^2)^2}{2\eta_0^2} - \sum_{i=1}^{d_x} \frac{x_i^2}{2\eta_i^2} \right\}, \quad (38)$$

with $\mathbf{x} \in \mathcal{X} = [-10, 10]^{d_x}$, $B = 4$, $\eta_0 = 4$ and $\eta_i = 3.5$ for $i = 1, \dots, d_x$. We consider $d_x = \{2, 3, 4, 5\}$ (i.e., different dimensions) and compute in advance the true moments of the target (i.e., the groundtruth) by using a costly grid, in order to check the performance of the different techniques.

1) EXPERIMENT 1

We set $d_x = 2$ and test the different algorithms in order to compute the vector mean $\mu = [-0.4, 0]$ and the diagonal of the covariance matrix $[\sigma_1^2, \sigma_2^2] = [1.3813, 8.9081]$. Moreover, our schemes are also able to estimate Z , whose ground-truth is $Z = 7.9979$, thus we also measure the error in this estimation. We compare the performance in terms of Relative Mean Square Error (Rel-MSE), averaged over 500 independent runs, using different methodologies: (a) NN-AQ starting with $N_0 = 10$ nodes randomly chosen in $[-10, 10] \times [-10, 10]$ and $M = 10^5$; (b) an independent MH algorithm (I-MH) with random initialization in $[-10, 10] \times [-10, 10]$; (c) random-walk MH algorithms (RW-MH) with

different proposal variance, and random initialization in $[-10, 10] \times [-10, 10]$; (d) an IS algorithm. The proposal density for both I-MH and IS is a uniform in $[-10, 10] \times [-10, 10]$, whereas for the RW-MHs is a Gaussian density centered at the current state of the chain with covariance matrix $v^2\mathbf{I}$ where $v \in \{1, 2, 5\}$ (so we consider 3 different RW-MHs).

For a fair comparison, we need that all methods have the same number E of target evaluations (fixing $E = 70$). Since NN-AQ, I-MH and RW-MH require one new target evaluation per iteration, we run $T = 70$ iterations for I-MH and RW-MH ($E = T$), and $T - N_0 = 60$ iterations for NN-AQ. In this regard, the IS algorithm use 70 samples drawn from the uniform proposal. Hence, all methods need $T = 70$ target evaluations. The results are given in Figures 3(a)-(b). Note that the estimation of Z via MCMC techniques is not straightforward (e.g., see [71]).

Discussion 1: We can observe that NN-AQ outperforms the other methods in terms of Rel-MSE in estimation. Moreover, in Fig. 3(a)-(b) we can see that the decrease is much greater, as E grows, than the other methods. Namely, NN-AQ has more benefits with new evaluations of $\pi(\mathbf{x})$.

2) EXPERIMENT 2

In this case, we fix the number of target evaluations E , and vary $d_x = \{2, 3, 4, 5\}$. The Rel-MSE in the estimation of Z is given in Table 2 (with $E \in \{100, 1000\}$).

TABLE 2. Relative MSE of Z With $E \in \{100, 1000\}$ for Different d_x .

methods	E	$d_x = 2$	$d_x = 3$	$d_x = 4$	$d_x = 5$
NN-AQ	100	0.0027	0.1127	0.3798	1.9730
	1000	$4 \cdot 10^{-4}$	0.0023	0.0140	0.0374
IS	100	0.2645	0.4427	0.7627	1.1115
	1000	0.0226	0.0378	0.0641	0.1094

Discussion 2: In this experiment, E is fixed along different dimensions. The results given in Table 2, with fixed E , does not show all the potential of NN-AQ. However, NN-AQ outperforms IS in all the dimensions d_x considered when $E = 1000$.




FIGURE 4. (a) Rel-MSE in log-scale for Z as function of number of target evaluations E . (b) Rel-MSE in log-scale for μ as function of number of target evaluations E . (c) Rel-MSE in log-scale for estimating $[\sigma_1^2, \sigma_2^2]$ as function of number of target evaluations E .




FIGURE 5. (a) Example of application of GK-AQ with 10 starting points (red cross-marks) and $T=60$ iterations (red dots), i.e., $E = 70$ target evaluations. (b) Rel-MSE in log-scale for Z as function of number of target evaluations E . (c) Rel-MSE in log-scale for μ as function of number of target evaluations E . (d) Rel-MSE in log-scale for estimating $[\sigma_1^2, \sigma_2^2]$ as function of number of target evaluations E .

3) EXPERIMENT 3

For $d_x = 2$, we compare now IS, NN-AQ, and three variants of NN-AQ: (i) NN-U, where the optimization step in (31) is substituted with sampling uniformly the new node in $[-10, 10] \times [-10, 10]$ (i.e., without using an acquisition function), (ii) NN-AQ *only diversity*, which uses the acquisition in (35) with $\alpha = 0, \beta = 1$, i.e., with only the diversity term $D_t(\mathbf{x})$, and (iii) NN-AQ *tempered*, which uses the acquisition in (35) with $\alpha = 0, \beta_t = \frac{200}{t}$, i.e., $A_t(\mathbf{x}) = [D_t(\mathbf{x})]^{\beta_t}$. Note that the adaptation in NN-AQ *only diversity* can be viewed as filling the space in a deterministic way. Note also that the adaptation in NN-AQ *tempered* will encourage more exploration than NN-AQ in the early iterations. Again, we compare the error in estimating Z, μ and $[\sigma_1^2, \sigma_2^2]$ as a function of target evaluations E (up to $E = 70$). The results are given in Figures 4(a)-(b).

Discussion 3: We can observe that NN-AQ and NN-AQ *tempered* outperform the others in terms of Rel-MSE in estimation. Moreover, in Fig. 4(a)-(b) we can see that for NN-AQ and NN-AQ *tempered*, the RMSE decreases at a faster rate as E grows, than the NN-U and NN-AQ *only diversity*, highlighting the importance of taking into account the current interpolant to locate the new nodes. It can be seen that NN-AQ *only diversity* works much better than NN-U in the early iterations. We explain these results by the fact that NN-AQ *only diversity* tends to cover the space more efficiently in

these early iterations since it avoids placing new nodes near the existing ones. However, as E grows, the performance of NN-U and NN-AQ *only diversity* is similar since both end up filling uniformly the space. Interestingly, NN-U performs better than IS as E increases, which demonstrate the power of the interpolative approach even when the new nodes are randomly chosen.

4) EXPERIMENT 4

For $d_x = 2$, we investigate the performance of GK-AQ in the estimation of Z, μ and $[\sigma_1^2, \sigma_2^2]$ as function of E . NN-GK employs the acquisition in (33). The kernel bandwidth h is fitted using the procedure in Appendix A. As commented in Sect. III-B, we consider a small noise of $\sigma = 10^{-2}$ for numerical stability. We will compare the performance against NN-AQ. The results are given in Figures 5(a)-(d), along with an example of GK-AQ interpolant, with $E = 70$, obtained in a specific run.

Discussion 4: The results are shown in Figures 5(b)-(d). GK-AQ outperforms NN-AQ in this particular experiment. However, it is important to remark that the results of GK-AQ may worsen considerably if h is not selected adequately (we have used the procedure in App. A), in contrast to NN-AQ which is free of hyperparameter tuning and hence more robust.

B. MULTIMODAL TARGET

In this experiment, we consider a multimodal Gaussian target in $d_x = 10$,

$$\bar{\pi}(\mathbf{x}) = \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) + \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) + \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_3, \boldsymbol{\Sigma}_3),$$

with $\boldsymbol{\mu}_1 = [5, 0, \dots, 0]$, $\boldsymbol{\mu}_2 = [-7, 0, \dots, 0]$, $\boldsymbol{\mu}_3 = [1, \dots, 1]$ and $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}_3 = 4^2\mathbf{I}_{10}$. We want to test the performance of the different methods in estimating the normalizing constant $Z = 1$. We consider an application of GK-AQ with $N_0 = 500$ initial nodes, random in $[-15, 15]^{10}$, and $T = 1000 - N_0$, hence fixing the number of evaluations to $E = 1000$. We compare it against three sophisticated AIS schemes, namely PMC, LAIS and AMIS [72]. For PMC, we choose Gaussian proposal pdfs and test different number of proposals $L \in \{10, 100, 200, 500\}$, whose means are also initialized at random in $[-15, 15]^{10}$. At each iteration one sample is drawn from each proposal, hence the algorithm is run for $T_{\text{PMC}} = \frac{1000}{L}$ iterations for a fair comparison. As a second alternative, we consider the deterministic mixture weighting approach for PMC, which is shown to have better overall performance, denoted DM-PMC. For LAIS, we also consider different number of proposals $L \in \{10, 100, 200, 500\}$. More specifically, we consider two versions of LAIS: the *one-chain* version and an *ideal* version. In ideal LAIS, the means of the L Gaussian proposals are drawn exactly from $\bar{\pi}(\mathbf{x})$. The one-chain application of LAIS (OC-LAIS) requires to run a MCMC algorithm targeting $\bar{\pi}(\mathbf{x})$ to obtain the L proposal means, hence it requires L evaluations of the target. At each iteration one sample is drawn from the mixture of the L Gaussian proposals, hence we run the algorithm for $T_{\text{LAIS}} = 1000 - L$ iterations for a fair comparison. We used a Gaussian random walk Metropolis to obtain the L means in the one-chain scenario. Finally, we consider AMIS with several combinations of number of iterations T_{AMIS} and number of samples per iteration R . At each iteration, R samples are drawn from a single Gaussian proposal, hence the total number of evaluations is $E = RT_{\text{AMIS}}$. In this case, we test $E \in \{1000, 2000, 3000, 5000\}$, so the comparison is not fair except for $E = 1000$. For PMC, LAIS and AMIS, as well as for the random walk proposal within the Metropolis algorithm, the covariance of the Gaussian proposals was fixed to $h^2\mathbf{I}_{10}$ (for $h = 1, \dots, 6$), where h is the initial bandwidth parameter used in GK-AQ.² All the methods are compared through the mean absolute error (MAE) in estimating Z , and the results are averaged over 500 independent simulations. The results are shown in Table 3 and Table 4. For each method, the best and worst MAE are boldfaced.

Discussion: We can observe that GK-AQ obtains the best range of MAE values [0.078, 0.4782] and the best results for $h = 1$. For $h > 1$, we can see in Tables 3-4 that the lowest MAE values are obtained by ideal LAIS with $L = 500$ and $h = 3$. We stress that ideal LAIS is not available in

²Recall that, for GK-AQ, the final bandwidth is tuned as described in App. A.

TABLE 3. MAE of Z with $E = 1000$ (Best and Worst MAE of Each Method are Boldfaced).

Methods		h = 1	h = 2	h = 3	h = 4	h = 5	h = 6
GK-AQ		0.4782	0.1741	0.0780	0.1362	0.1497	0.2322
PMC	$N = 10$	0.9993	0.9526	0.8603	0.6743	0.6024	0.6155
	$N = 100$	0.9998	0.9896	0.8853	0.6761	0.5192	0.4544
	$N = 200$	1.0002	0.9893	0.8816	0.7099	0.6389	0.5384
	$N = 500$	0.9995	0.9916	0.9741	0.8700	0.7421	0.6544
DM-PMC	$N = 10$	0.9991	0.9478	0.8505	0.6009	0.5352	0.5814
	$N = 100$	0.9997	0.8719	0.4490	0.2425	0.1901	0.2193
	$N = 200$	0.9999	0.9321	0.5708	0.3257	0.2374	0.2524
	$N = 500$	1.0000	0.9888	0.7969	0.5009	0.3684	0.3800
Ideal LAIS	$N = 10$	0.9992	0.8114	0.2579	0.0863	0.0819	0.1091
	$N = 100$	0.9918	0.3638	0.0547	0.0407	0.0598	0.1053
	$N = 200$	0.9846	0.2486	0.0352	0.0411	0.0680	0.1093
	$N = 500$	0.9687	0.1852	0.0335	0.0473	0.0891	0.1353
OC-LAIS	$N = 10$	1.0000	1.0000	0.9992	0.9883	0.9468	0.9079
	$N = 100$	0.9999	0.8731	0.4434	0.2785	0.2392	0.2870
	$N = 200$	0.9982	0.7028	0.2418	0.1243	0.1406	0.2070
	$N = 500$	0.9937	0.4949	0.1221	0.0857	0.1195	0.1786

TABLE 4. MAE of Z of AMIS With $E \in \{1000, 2000, 3000, 5000\}$.

Methods		h = 1	h = 2	h = 3	h = 4	h = 5	h = 6
GK-AQ ($E=1000$)		0.4782	0.1741	0.0780	0.1362	0.1497	0.2322
AMIS $E = 1000$	$M = 10$	0.9998	0.9997	0.9997	0.9996	0.9996	0.9995
	$M = 100$	1.0000	1.0000	1.0000	0.9999	0.9997	0.9990
	$M = 200$	1.0000	1.0000	1.0000	1.0000	0.9998	0.9994
	$M = 500$	1.0000	1.0000	1.0000	1.0000	0.9998	0.9989
AMIS $E = 2000$	$M = 10$	0.9155	0.9117	0.8981	0.8987	0.8891	0.8878
	$M = 100$	0.9998	0.9986	0.9934	0.9784	0.9559	0.9072
	$M = 200$	1.0000	1.0000	0.9998	0.9981	0.9888	0.9712
	$M = 500$	1.0000	1.0000	1.0000	0.9998	0.9984	0.9953
AMIS $E = 3000$	$M = 10$	0.3293	0.3402	0.3051	0.3381	0.3540	0.3443
	$M = 100$	0.9725	0.9040	0.7963	0.6384	0.4964	0.3816
	$M = 200$	0.9998	0.9977	0.9884	0.9527	0.8308	0.7119
	$M = 500$	1.0000	1.0000	0.9998	0.9988	0.9859	0.9566
AMIS $E = 5000$	$M = 10$	0.0766	0.0768	0.0695	0.0722	0.0699	0.0725
	$M = 100$	0.1626	0.1176	0.0957	0.0810	0.0737	0.0656
	$M = 200$	0.8771	0.6040	0.2824	0.1473	0.1163	0.0899
	$M = 500$	1.0000	0.9982	0.9904	0.9449	0.7944	0.4532

practice, since we usually cannot sample directly from $\bar{\pi}(\mathbf{x})$. Regardless of the ideal LAIS scheme (not applicable in practice), GK-AQ provides the best results. Moreover, we see that GK-AQ with $h = 3$ is the best performing method in this experiment, since it achieves a lower MAE than PMC, DM-PMC and OC-LAIS for every combination of L and h . Table 4 shows that AMIS performs worse than GK-AQ for $E = 1000$ (fair comparison), but even with much more AMIS evaluations $E \in \{2000, 3000\}$ (unfair comparison in favor of AMIS). AMIS needs to reach a big enough value of E ($E = 5000$), to beat GK-AQ in terms of MAE.

C. APPLICATIONS TO EXOPLANET DETECTION

In recent years, the problem of revealing objects orbiting other stars has acquired large attention. Different techniques have been proposed to discover exo-objects but, nowadays, the radial velocity technique is still the most used [73]–[76]. The problem consists in fitting a dynamical model to data acquired at different moments spanning during long time periods (up to years). The model is highly non-linear and, for certain sets of parameters, its evaluation is quite costly in terms of computation time. This is due to the fact that its evaluation involves numerically integrating a differential equation, or using an iterative procedure for solving a non-linear

TABLE 5. Description of Parameters in Eq. (39).

Parameter	Description	Units
For each planet		
K_i	amplitude of the curve	m s^{-1}
$u_{i,t}$	true anomaly	rad
ω_i	longitude of periastron	rad
e_i	orbit's eccentricity	...
P_i	orbital period	s
τ_i	time of periastron passage	s
Below: not depending on the number of objects/satellite		
V_0	mean radial velocity	m s^{-1}

equation (until a certain condition is satisfied). This loop can be very long for some sets of parameters.

1) LIKELIHOOD FUNCTION

When analyzing radial velocity data of an exoplanetary system, it is commonly accepted that the *wobbling* of the star around the center of mass is caused by the sum of the gravitational force of each planet independently and that they do not interact with each other. Each planet follows a Keplerian orbit and the radial velocity of the host star is given by

$$y_t = V_0 + \sum_{i=1}^S K_i [\cos(u_{i,t} + \omega_i) + e_i \cos(\omega_i)] + \xi_t, \quad (39)$$

with $t = 1, \dots, T$.³ The number of objects in the system is S . Both $y_t, u_{i,t}$ depend on time t , and ξ_t is a Gaussian noise perturbation with variance σ_e^2 . We consider the noise variance σ_e^2 an unknown parameter as well. The meaning of each parameter in Eq. (39) is given in Table 5. The likelihood function is jointly defined by (39) and some indicator variables described below. The angle $u_{i,t}$ is the true anomaly of the planet i and it can be determined from

$$\frac{du_{i,t}}{dt} = \frac{2\pi}{P_i} \frac{(1 + e_i \cos u_{i,t})^2}{(1 - e_i)^{\frac{3}{2}}}$$

This equation has an analytical solution. As a result, the true anomaly $u_{i,t}$ can be determined from the mean anomaly $M_{i,t}$. However, the analytical solution contains a non-linear term that needs to be determined by iterating. First, we define the mean anomaly $M_{i,t}$ as

$$M_{i,t} = \frac{2\pi}{P_i} (t - \tau_i),$$

where τ_i is the time of periastron passage of the planet i and P_i is the period of its orbit (see Table 5). Then, through the Kepler's equation,

$$M_{i,t} = E_{i,t} - e_i \sin E_{i,t}, \quad (40)$$

where $E_{i,t}$ is the eccentric anomaly. Equation (40) has no analytic solution and it must be solved by an iterative procedure. A Newton-Raphson method is typically used to find the roots

³More generally, we can have y_{t_j} with $j = 1, \dots, T$.

of this equation [77]. For certain sets of parameters, this iterative procedure can be particularly slow and the computation of the likelihood becomes quite costly. We also have

$$\tan \frac{u_{i,t}}{2} = \sqrt{\frac{1 + e_i}{1 - e_i}} \tan \frac{E_{i,t}}{2}, \quad (41)$$

Therefore, the variable of interest \mathbf{x} is the vector of dimension $d_X = 1 + 5S$ (where S is the number of planets),

$$\mathbf{x} = [V_0, K_1, \omega_1, e_1, P_1, \tau_1, \dots, K_S, \omega_S, e_S, P_S, \tau_S],$$

For a single object (e.g., a planet or a natural satellite), the dimension of \mathbf{x} is $d_X = 5 + 1 = 6$, with two objects the dimension of \mathbf{x} is $d_X = 11$, etc. All the Eqs. from (39) to (41) induce a likelihood function $\ell(\mathbf{y}|\mathbf{x}, \sigma_e) = \prod_{t=1}^T \ell(y_t|\mathbf{x}, \sigma_e)$, where $\mathbf{y} = \{y_1, \dots, y_T\}$.

2) PRIOR AND POSTERIOR DENSITIES

The prior $g(\mathbf{x})$ is defined as multiplication of indicator variables $V_0 \in [-20, 20], K_i \in [0, \max y_{i,t} - \min y_{i,t}], e_i \in [0, 1], P_i \in [0, 365], \omega_{i,t} \in [0, 2\pi], \tau_i \in [0, 30]$, (i.e., the prior is zero outside these intervals), for all $i = 1, \dots, S$. This means that the prior density is zero when the particles fall out of these intervals. Note that the interval of τ_i is conditioned to the value P_i . This parameter is the time of periastron passage, i.e. the time passed since the object crossed the closest point in its orbit. It has the same units of P_i and can take values from 0 to P_i . The complete posterior is

$$p(\mathbf{x}|\mathbf{y}, \sigma_e) = \frac{1}{p(\mathbf{y}|\sigma_e)} \ell(\mathbf{y}|\mathbf{x}, \sigma_e) g(\mathbf{x}).$$

We are interested in inferring the parameters \mathbf{x} and, more specifically, computing the marginal likelihood

$$Z = p(\mathbf{y}|\sigma_e) = \int_{\mathcal{X}} \ell(\mathbf{y}|\mathbf{x}, \sigma_e) g(\mathbf{x}) d\mathbf{x},$$

obtained integrating out \mathbf{x} , in order to infer the number of planets. The noise variance σ_e^2 is also inferred after the sampling, by maximizing $Z = p(\mathbf{y}|\sigma_e)$, i.e., $\hat{\sigma}_e^2 = \arg \max_{\sigma_e} p(\mathbf{y}|\sigma_e)$.

3) EXPERIMENTS

Given a set of data \mathbf{y} generated according to the model (see the initial parameter values below), our goal is to infer the number S of planets in the system. For this purpose, we have to approximate the model evidence $Z = p(\mathbf{y}|\sigma_e)$ of each model. In all experiments, we consider 60 total number of observations. We consider three different experiments: **(E1)** $S = 0$, i.e., no object, **(E2)** $S = 1$ (one object) and **(E3)** the case of two objects $S = 2$. We set $V = 2$, in all cases. For the first object in **E1** and **E2**, we set $K_1 = 25, \omega_1 = 0.61, e_1 = 0.1, P_1 = 15, \tau_1 = 3$. For **E2**, we also consider a second object with $K_2 = 5, \omega_2 = 0.17, e_2 = 0.3, P_2 = 115, \tau_2 = 25$ (in that case $S = 2$). All the data are generated with $\sigma_e^2 = 2$. The rest of trajectories are generated according to the transition model (and the corresponding measurements y_t according to the observation model).




FIGURE 6. Plot of marginal likelihood estimates of Model 1 (one-planet) and Model 2 (two-planets) versus σ_e for the three data sets. The straight lines represent the known marginal likelihoods of Model 0 (zero planets) for each data set. (a) data set with zero planets, (b) data set with one planet, (c) data set with two planets.

4) METHODS

For each experiment, three models (i.e. three different target pdfs) are considered: a model with $S = 0$ (Zero-Planets), a model with $S = 1$ (One-Planet) and a model with $S = 2$ (Two-Planets). The goal is to estimate the marginal likelihood of these models and then correctly detect the number of planets, i.e., $S = 0$ for (E1), $S = 1$ for (E2) and $S = 2$ for (E3). The marginal likelihoods corresponding to the Zero-planets models are available in closed form and need not be estimated (the model is simply Gaussian in that case). For this purpose, we apply NN-AQ (with $M = 10^7$) and an IS procedure. We allocate a budget of $4 \cdot 10^6$ evaluations of the target. In IS, this budget is used to draw $4 \cdot 10^6$ samples from the priors. While NN-AQ uses first $4 \cdot 10^6 - 5000$ of these samples to look for a good initialization, more specifically, the sample with the highest target evaluation is kept, along with 9 more samples taken at random, to use them as initial nodes. Then, NN-AQ is run for 5000 iterations. Both One-Planet and Two-Planets models are estimated for different values of $\sigma_e = 1, 2, \dots, 15$. Note that we do not need to evaluate the target again when considering different σ_e , i.e., a single target evaluation can be reused for all values of σ_e . The results are shown in Figure 6.

5) RESULTS

For each experiment (E1)–(E3), Figure 6(a)–(c) depicts the estimations of Z of the different models provided by NN-AQ and IS, versus σ_e . The horizontal lines correspond to the known marginal likelihoods of Zero-Planets models. Overall, NN-AQ outperforms IS and predicts correctly the number of planets as well as the true value of σ_e (indeed, the curves corresponding to NN-AQ reach a maximum at $\sigma_e = 2$). Figure 6(a) shows that the estimations provided by NN-AQ and IS correctly rank the Zero-planets model ($S = 0$) as the most probable one. Figure 6(b) shows both NN-AQ and IS predict correctly the One-Planet model ($S = 1$) to be the correct one. However, for $\sigma_e = 2$, IS barely differentiates between the Zero-Planet and One-Planets models. Further, for $\sigma_e = 1$, it wrongly predicts Zero-Planets as the best one. Conversely, NN-AQ is able to predict the correct model

for every value of σ_e , and besides, also predicts the true value $\sigma_e = 2$. In Figure 6(c), the difference in performance of NN-AQ and IS is more acute. While NN-AQ is able to correctly predict the Two-Planets model ($S = 2$) as the most probable for all values of σ_e , IS is unable to detect that second planet and, therefore, considers the One-Planet model more probable. As in the previous case, IS fails at detecting any planet for small values of σ_e . Again, NN-AQ predict the correct value of σ_e .

IX. CONCLUSION

In this work, we have described a general framework for adaptive interpolative quadrature schemes, leveraging an in-depth study of different fields and related techniques in the literature, such as Bayesian quadrature algorithms, scattered data approximations, emulation, experimental design and active learning schemes. The nodes of the quadrature are adaptively chosen by maximizing a suitable acquisition function, which depends on the current interpolant and the positions of the nodes. This maximization does not require extra evaluations of the true posterior. The proposed methods supply also a surrogate model (emulator) which approximates the true posterior density, that can be also employed in further statistical analyses. Two specific schemes, based on Gaussian and NN bases, have been described. In both cases, a non-negative estimation \widehat{Z} of the marginal likelihood Z is ensured.

In the proposed framework, we also relax the assumptions regarding the kernel-basis functions with respect to other approaches in the literature, e.g., the bases could be non-symmetric. For instance, the NN bases are non-symmetric functions and their use has different important benefits: (a) they ensure obtaining non-negative interpolation coefficients and estimators \widehat{Z} , (b) the linear system is directly solved without the need of inverting any matrix (the interpolation matrix is always diagonal), and (c) the bandwidth of the bases are automatically selected. Our scheme also allows selecting different kernel functions for each node point. Therefore, the quadrature rules in Bayesian quadrature are a special case of our proposed scheme. Indeed, Bayesian quadrature considers a single symmetric and semi




FIGURE 7. (a) GK based interpolant with $N_T = 70$ nodes and $h = 2.7$, fitted by maximizing the marginal likelihood. (b) Plot of \widehat{Z} as function of h . The value of h at which \widehat{Z} attains the local maximum is used to build the interpolant in our procedure. (d) GK based interpolant with $N_T = 70$ nodes and $h = 0.51$ fitted with the heuristic.

positive definite kernel function. An importance sampling interpretation has been also provided. It is important to remark that the true posterior is only evaluated at the nodes selected sequentially by the algorithm, and the rest of other computations does not query the true model. The convergence of the proposed quadrature rules has been discussed, jointly with other theoretical results. The new algorithms are powerful techniques as also shown by several numerical experiments.

**APPENDIX A
PROCEDURE FOR TUNING THE GAUSSIAN
KERNEL BANDWIDTH**

In this Appendix, we propose a procedure for fitting the bandwidth parameter h of the Gaussian kernel (GK),

$$k_G(\mathbf{x}, \mathbf{x}_i) = \frac{1}{(2\pi)^{\frac{d_x}{2}} h^{d_x}} \exp\left(-\frac{1}{2h^2}(\mathbf{x} - \mathbf{x}_i)^\top(\mathbf{x} - \mathbf{x}_i)\right), \tag{42}$$

when building the GK based interpolant of Sect. III for a given number of nodes. Assume we have run the GK-AQ algorithm (with some fixed h_0), so we have a total of N_T nodes. Now, for any h , we may solve the linear system (Eq. (5)), obtain the coefficients $\{\beta_i\}_{i=1}^{N_T}$ and calculate

$$\widehat{Z} = \sum_{i=1}^{N_T} \beta_i. \tag{43}$$

Note that, although not explicit, the β_i 's, and hence \widehat{Z} , depend on h . The proposed procedure consists of taking h as the value where \widehat{Z} attains its first local maximum. Starting from a small value h close to zero and increasing it, the estimation \widehat{Z} is growing reaching a maximum. Then, h is starting to become “too big”, producing too much overlapping among the kernel areas. The values of the elements out the diagonal of \mathbf{K} grow, and some of the coefficients β_i are negative, and the estimation \widehat{Z} decreases. As h becomes greater and greater, the matrix \mathbf{K} tends to become ill-conditioned, and the absolute values of β_i 's grows. Figure 7 compares the GK

based interpolant of the target from Sect VIII-A with two different choices of h and $N_T = 70$ nodes. Figure 7(a) plots the interpolant taking h as the value which minimizes the marginal likelihood (see Sect. III-B). Note that this value of h is too big given the dispersion of the nodes. While Figure 7(c) plots the interpolant taking h as the value where the curve of \widehat{Z} (Figure 7(b)) attains its local maximum. This choice of h seems to fit better the existing nodes. Note also that, for some values of h , \widehat{Z} may be negative.

**APPENDIX B
PROBABILISTIC INTERPRETATION OF J**

Let us consider $J = \int_{\mathcal{X}} f(\mathbf{x})\pi(\mathbf{x})d\mathbf{x}$, which is the numerator of (1), our integral of interest I . In section III-A, we have seen that, when $k(\mathbf{x}, \mathbf{x}_i) = k(\mathbf{x}_i, \mathbf{x})$ (i.e., a symmetric basis function), the interpolant $\widehat{\pi}(\mathbf{x}) = \sum_{i=1}^N \beta_i k(\mathbf{x}, \mathbf{x}_i)$ has the probabilistic interpretation of being the mean of the posterior distribution of (the “unknown”) $\pi(\mathbf{x})$ after observing $\mathbf{d} = [\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_N)]^\top$, i.e., $\mathbb{E}[\pi(\mathbf{x})|\mathbf{d}] = \widehat{\pi}(\mathbf{x})$. The distribution on $\pi(\mathbf{x})$ induces a posterior distribution on J , which is a Gaussian with mean

$$\mathbb{E}[J|\mathbf{d}] = \widetilde{J} = \int_{\mathcal{X}} f(\mathbf{x})\widehat{\pi}(\mathbf{x})d\mathbf{x}, \tag{44}$$

and variance given by

$$\text{var}[J|\mathbf{d}] = \int \int k(\mathbf{x}, \mathbf{x}')f(\mathbf{x})f(\mathbf{x}')d\mathbf{x}d\mathbf{x}' - \zeta^\top \mathbf{K}^{-1} \zeta, \tag{45}$$

where $\zeta = [J_1, \dots, J_N]$ and $J_i = \int_{\mathcal{X}} f(\mathbf{x})k(\mathbf{x}, \mathbf{x}_i)d\mathbf{x}$. This interpretation corresponds to the so-called Bayesian quadrature, which uses Eq. (44) as approximation of J . Note that Eq. (44) is the quadrature obtained by substituting the true $\pi(\mathbf{x})$ with its interpolant $\widehat{\pi}(\mathbf{x})$, which coincides with the numerator of \widehat{T} in Eq. (7).

**APPENDIX C
RECURSIVE INVERSION OF A BORDERED MATRIX**

The most costly step when calculating β in (5) consists in inverting the $N \times N$ matrix $(\mathbf{K})_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$ ($i, j \in \{1, \dots, N\}$). Moreover, every time a new node is added, the β_i

must be recomputed, so the step of computing the inverse has to be done again. This time the matrix is bigger due to adding a new node, that is, it has an additional row and column. We show that knowing \mathbf{K}^{-1} help us to compute the inverse of augmented matrices (called “bordered matrix”, i.e., adding a “border” of new row and column to an existing matrix).

Let us denote with \mathbf{K}_N the matrix built using N nodes, and let \mathbf{K}_{N+1} be the matrix with $N + 1$ nodes. Of course we have

$$\mathbf{K}_{N+1} = \begin{pmatrix} \mathbf{K}_N & \mathbf{k}_N \\ \mathbf{k}_N^T & k \end{pmatrix} \quad (46)$$

where $\mathbf{k}_N = (k(\mathbf{x}_1, \mathbf{x}_{N+1}), k(\mathbf{x}_2, \mathbf{x}_{N+1}), \dots, k(\mathbf{x}_N, \mathbf{x}_{N+1}))^T$ and $k = k(\mathbf{x}_{N+1}, \mathbf{x}_{N+1})$. The $(N + 1) \times (N + 1)$ inverse of \mathbf{K}_{N+1} can be expressed in terms of \mathbf{K}_N^{-1} as follows

$$\mathbf{K}_{N+1}^{-1} = \begin{pmatrix} \mathbf{A} & \mathbf{b} \\ \mathbf{c} & s \end{pmatrix}, \quad (47)$$

where

$$\begin{aligned} \mathbf{A} &= \mathbf{K}_N^{-1} + \mathbf{K}_N^{-1} \mathbf{k}_N \left(k - \mathbf{k}_N^T \mathbf{K}_N^{-1} \mathbf{k}_N \right)^{-1} \mathbf{k}_N^T \mathbf{K}_N^{-1} \in \mathbb{R}^{N \times N}, \\ \mathbf{b} &= -\mathbf{K}_N^{-1} \mathbf{k}_N \left(k - \mathbf{k}_N^T \mathbf{K}_N^{-1} \mathbf{k}_N \right)^{-1} \in \mathbb{R}^{N \times 1}, \\ \mathbf{c} &= -\left(k - \mathbf{k}_N^T \mathbf{K}_N^{-1} \mathbf{k}_N \right)^{-1} \mathbf{k}_N^T \mathbf{K}_N^{-1} \in \mathbb{R}^{1 \times N}, \\ s &= \left(k - \mathbf{k}_N^T \mathbf{K}_N^{-1} \mathbf{k}_N \right)^{-1} \in \mathbb{R}. \end{aligned}$$

Note that computing $s = \left(k - \mathbf{k}_N^T \mathbf{K}_N^{-1} \mathbf{k}_N \right)^{-1}$ is not costly since it is a scalar value.

APPENDIX D

PROOFS

A. PROOF TO THEOREM 1

We have that

$$\begin{aligned} |J - \widehat{J}| &= \left| \int_{\mathcal{X}} f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x} - \int_{\mathcal{X}} f(\mathbf{x}) \widehat{\pi}(\mathbf{x}) d\mathbf{x} \right| \\ &= \left| \int_{\mathcal{X}} f(\mathbf{x}) (\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})) d\mathbf{x} \right|. \end{aligned}$$

It is easy to see that, for any $g(\mathbf{x})$ we have $-|g(\mathbf{x})| \leq g(\mathbf{x}) \leq |g(\mathbf{x})|$ for all \mathbf{x} , and that $-\int |g(\mathbf{x})| d\mathbf{x} \leq \int g(\mathbf{x}) d\mathbf{x} \leq \int |g(\mathbf{x})| d\mathbf{x}$, so we have $|\int g(\mathbf{x}) d\mathbf{x}| \leq \int |g(\mathbf{x})| d\mathbf{x}$. Using this result we can state the first inequality

$$\begin{aligned} |J - \widehat{J}| &= \left| \int_{\mathcal{X}} f(\mathbf{x}) (\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})) d\mathbf{x} \right| \\ &\leq \int_{\mathcal{X}} |f(\mathbf{x})| |\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})| d\mathbf{x} \\ &= \|f(\pi - \widehat{\pi})\|_1. \end{aligned}$$

The second inequality of the theorem follows from Holder’s inequality

$$\|f(\pi - \widehat{\pi})\|_1 \leq \|f\|_2 \|\pi - \widehat{\pi}\|_2.$$

Finally, the last inequality of the theorem is obtained after manipulating the $\|f\|_2$ and $\|\pi - \widehat{\pi}\|_2$,

$$\begin{aligned} \|f\|_2 \|\pi - \widehat{\pi}\|_2 &= \left(\int_{\mathcal{X}} |f(\mathbf{x})|^2 d\mathbf{x} \right)^{\frac{1}{2}} \left(\int_{\mathcal{X}} |\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})|^2 d\mathbf{x} \right)^{\frac{1}{2}} \\ &\leq \left(|\mathcal{X}| \max |f(\mathbf{x})|^2 \right)^{\frac{1}{2}} \left(|\mathcal{X}| \max |\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})|^2 \right)^{\frac{1}{2}} \\ &= |\mathcal{X}| \max |f(\mathbf{x})| \max |\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})| \\ &= |\mathcal{X}| \|f\|_{\infty} \|\pi - \widehat{\pi}\|_{\infty}. \end{aligned}$$

B. PROOF TO THEOREM 2

We provide the main concepts and elements of the proof. For more details, see [16], [67]. Let $J = \int_{\mathcal{X}} f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}$ and $\widetilde{J} = \sum_{i=1}^N v_i \pi(\mathbf{x}_i)$ be the integral of interest and the quadrature using points $\{\mathbf{x}_i\}_{i=1}^N$, respectively. Recall that we also denote $\mathbf{v} = [v_1, \dots, v_N]^T$.

Consider that π is a function belonging to the reproducing kernel Hilbert space of functions \mathcal{H} originated from the symmetric and positive definite kernel function $k(\mathbf{x}, \mathbf{x}')$. Hence, J and \widetilde{J} are functionals over that RKHS

$$\begin{aligned} J[\pi] &= \int_{\mathcal{X}} f(\mathbf{x}) \pi(\mathbf{x}) d\mathbf{x}, \\ \widetilde{J}[\pi] &= \sum_{i=1}^N v_i \pi(\mathbf{x}_i), \quad \pi \in \mathcal{H}. \end{aligned}$$

where we write explicitly $J[\cdot]$ is the functional that integrates w.r.t. $f(\mathbf{x})$, while $\widetilde{J}[\cdot]$ is the functional that integrates w.r.t. the weighted sum $\sum_{i=1}^N v_i \delta_{\mathbf{x}_i}$, where $\delta_{\mathbf{x}_i}$ denotes the point evaluation in \mathbf{x}_i . The integration error associated with \widetilde{J} is characterized by the norm, in the dual space \mathcal{H}^* , of the error functional

$$\|J - \widetilde{J}\|_{\mathcal{H}^*} = \sup_{\|\pi\|_{\mathcal{H}} \leq 1} |\widetilde{J}[\pi] - J[\pi]|, \quad (48)$$

where $\|\cdot\|_{\mathcal{H}}$ and $\|\cdot\|_{\mathcal{H}^*}$ denote the norm in \mathcal{H} and \mathcal{H}^* respectively. Eq. (48) is also called worst-case error (WCE). Define the functions

$$k_f(\mathbf{x}) = \int_{\mathcal{X}} f(\mathbf{x}') k(\mathbf{x}, \mathbf{x}') d\mathbf{x}', \quad (49)$$

and

$$k_{\widetilde{J}}(\mathbf{x}) = \sum_{i=1}^N v_i k(\mathbf{x}, \mathbf{x}_i), \quad (50)$$

where $k_f, k_{\widetilde{J}} \in \mathcal{H}$. These functions exist as consequence of $\int_{\mathcal{X}} k(\mathbf{x}, \mathbf{x}) f(\mathbf{x}) d\mathbf{x} < \infty$. It can be shown that $\|J - \widetilde{J}\|_{\mathcal{H}^*} = \|k_f - k_{\widetilde{J}}\|_{\mathcal{H}}$, and

$$\begin{aligned} \|J - \widetilde{J}\|_{\mathcal{H}^*}^2 &= \mathbf{v}^T \mathbf{K} \mathbf{v} - 2\mathbf{v}^T \boldsymbol{\zeta} \\ &\quad + \int_{\mathcal{X}} \int_{\mathcal{X}} f(\mathbf{x}) f(\mathbf{x}') k(\mathbf{x}, \mathbf{x}') d\mathbf{x} d\mathbf{x}', \quad (51) \end{aligned}$$

for a vector of weights $\mathbf{v} \in \mathbb{R}^N$, the matrix $(\mathbf{K})_{1 \leq i, j \leq N} = k(\mathbf{x}_i, \mathbf{x}_j)$, and the vector of integrals $\boldsymbol{\zeta} = [k_f(\mathbf{x}_1), \dots, k_f(\mathbf{x}_N)]^T$.

Conditional on the fixed states $\{\mathbf{x}_i\}_{i=1}^N$, the weights \mathbf{v} that minimizes the above expression are given by $\mathbf{v} = \mathbf{K}^{-1}\boldsymbol{\zeta}$. These are the weights that arises if we build the interpolant $\widehat{\pi}$ of π at points $\{\mathbf{x}_i\}_{i=1}^N$, using $k(\mathbf{x}, \mathbf{x}')$ as the basis function, and substitute it in J to obtain the quadrature.

C. PROOF TO THEOREM 6

Let J be the integral of interest, and let $\tilde{J} = \sum_{i=1}^N \beta_i J_i$ and $\widehat{J} = \sum_{i=1}^N \beta_i \widehat{J}_i$ be the approximations using, respectively, the exact J_i and the noisy estimation \widehat{J}_i . Recall that the coefficients β_i are written in matrix form as $\boldsymbol{\beta} = \mathbf{K}^{-1}\mathbf{d}$ where \mathbf{K} is the interpolation matrix and \mathbf{d} is the vector of evaluations of π . Let us denote $\boldsymbol{\zeta} = [J_1, \dots, J_N]^\top$ and $\widehat{\boldsymbol{\zeta}} = [\widehat{J}_1, \dots, \widehat{J}_N]^\top$. Denoting the dot product in \mathbb{R}^N as $\langle \cdot, \cdot \rangle$, we can express $\tilde{J} = \langle \boldsymbol{\zeta}, \boldsymbol{\beta} \rangle$ and $\widehat{J} = \langle \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle$. Thus

$$\begin{aligned} |J - \widehat{J}| &= |J - \langle \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle| \\ &= |J - \langle \boldsymbol{\zeta} - \widehat{\boldsymbol{\zeta}} + \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle| \\ &= |J - \langle \boldsymbol{\zeta}, \boldsymbol{\beta} \rangle + \langle \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle - \langle \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle| \\ &\leq |J - \tilde{J}| + |\langle \widehat{\boldsymbol{\zeta}}, \boldsymbol{\beta} \rangle - \langle \boldsymbol{\zeta}, \boldsymbol{\beta} \rangle| \\ &= |J - \tilde{J}| + |\langle \mathbf{K}^{-1}(\widehat{\boldsymbol{\zeta}} - \boldsymbol{\zeta}), \mathbf{d} \rangle| \\ &\leq |f(\pi - \widehat{\pi})|_1 + \|\mathbf{K}^{-1}(\widehat{\boldsymbol{\zeta}} - \boldsymbol{\zeta})\|_2 \|\mathbf{d}\|_2 \\ &\leq |\mathcal{X}| \|f\|_\infty \|\pi - \widehat{\pi}\|_\infty + \|\mathbf{K}^{-1}\|_2 \|\widehat{\boldsymbol{\zeta}} - \boldsymbol{\zeta}\|_2 \|\mathbf{d}\|_2 \end{aligned}$$

where the norm $\|\mathbf{K}^{-1}\|_2$ represents the largest singular value of \mathbf{K}^{-1} . The bounds $\|\pi - \widehat{\pi}\|_\infty = \lambda(r)$ and $\|\mathbf{K}^{-1}\|_2 = \mathcal{O}(\nu(\sigma, h))$ for different RBF can be found respectively in Chapters 11.3 and 12.2 of [55]. For further details, see Proposition 1 in [67].

D. PROOF TO THEOREM 8

Let us consider the target $\pi(\mathbf{x})$ and the interpolant $\widehat{\pi}(\mathbf{x})$ based on NN constant kernels. Note that for all $\mathbf{x} \in \mathcal{X}$ we have $\widehat{\pi}(\mathbf{x}) = \pi(\mathbf{x}^*)$, where $\mathbf{x}^* = \arg \min_i \|\mathbf{x} - \mathbf{x}_i\|$, i.e., the node that is closest to \mathbf{x} . Lipschitz continuity implies that $|\pi(\mathbf{z}) - \pi(\mathbf{x})| \leq L_0 \|\mathbf{z} - \mathbf{x}\|$ for all $\mathbf{z}, \mathbf{x} \in \mathcal{X}$. Hence,

$$\begin{aligned} \|\pi - \widehat{\pi}\|_\infty &= \max_{\mathbf{x} \in \mathcal{X}} |\pi(\mathbf{x}) - \widehat{\pi}(\mathbf{x})| \\ &= \max_{\mathbf{x} \in \mathcal{X}} |\pi(\mathbf{x}) - \pi(\mathbf{x}^*)| \\ &\leq L_0 \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - \mathbf{x}^*\| \\ &= L_0 \max_{\mathbf{x} \in \mathcal{X}} \min_i \|\mathbf{x} - \mathbf{x}_i\| \\ &= L_0 r, \end{aligned}$$

where we used the definition of fill distance r , i.e.,

$$r = \max_{\mathbf{x} \in \mathcal{X}} \min_i \|\mathbf{x} - \mathbf{x}_i\|.$$

For further details, see [27], [56].

REFERENCES

- [1] R. L. Burden and J. D. Faires, *Numerical Analysis*. Boston, MA, USA: Brooks Cole, 2000.
- [2] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York, NY, USA: Springer, 2004.
- [3] H. Niederreiter, *Random Number Generation and Quasi-Monte Carlo Methods*. Philadelphia, PA, USA: Society for Industrial Mathematics, 1992.
- [4] L. Martino, D. Luengo, and J. Míguez, *Independent Random Sampling Methods*. New York, NY, USA: Springer, 2018.
- [5] Q. Liu and D. A. Pierce, "A note on Gauss–Hermite quadrature," *Biometrika*, vol. 81, no. 3, pp. 624–629, 1994.
- [6] P. Jäckel, *A Note on Multivariate Gauss–Hermite Quadrature*. London, U.K.: ABN-Amro. Re, 2005.
- [7] A. Owen. (2013). *Monte Carlo Theory, Methods and Examples*. [Online]. Available: <http://statweb.stanford.edu/~owen/mc/>
- [8] V. Elvira, L. Martino, and P. Closas, "Importance Gaussian quadrature," 2020, pp. 1–13, *arXiv:2001.03090*. [Online]. Available: <http://arxiv.org/abs/2001.03090>
- [9] Ö. D. Akyildiz and J. Míguez, "Nudging the particle filter," *Statist. Comput.*, vol. 30, no. 2, pp. 305–330, Mar. 2020.
- [10] Q. Liu and D. J. Lee, "Black-box importance sampling," in *Proc. Artif. Intell. Statist. (AISTATS)*, 2017, pp. 952–961.
- [11] A. O'Hagan, "Bayes–Hermite quadrature," *J. Stat. Planning Inference*, vol. 29, no. 3, pp. 245–260, Nov. 1991.
- [12] M. Kennedy, "Bayesian quadrature with non-normal approximating functions," *Statist. Comput.*, vol. 8, no. 4, pp. 365–375, 1998.
- [13] C. E. Rasmussen and Z. Ghahramani, "Bayesian Monte Carlo," in *Proc. Adv. Neural Inf. Process. Syst.*, 2003, pp. 505–512.
- [14] M. Osborne, R. Garnett, Z. Ghahramani, D. K. Duvenaud, S. J. Roberts, and C. E. Rasmussen, "Active learning of model evidence using Bayesian quadrature," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 46–54.
- [15] T. Gunter, M. A. Osborne, R. Garnett, P. Hennig, and S. J. Roberts, "Sampling for inference in probabilistic models with fast Bayesian quadrature," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2789–2797.
- [16] F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic, "Probabilistic integration: A role in statistical computation?" *Stat. Sci.*, vol. 34, no. 1, pp. 1–22, 2019.
- [17] T. Karvonen and S. Särkkä, "Classical quadrature rules via Gaussian processes," in *Proc. IEEE 27th Int. Workshop Mach. Learn. Signal Process. (MLSP)*, Sep. 2017, pp. 1–6.
- [18] M. Kanagawa and P. Hennig, "Convergence guarantees for adaptive Bayesian quadrature methods," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 6234–6245.
- [19] D. Busby, "Hierarchical adaptive experimental design for Gaussian process emulators," *Rel. Eng. Syst. Saf.*, vol. 94, no. 7, pp. 1183–1193, Jul. 2009.
- [20] M. T. Pratola, C. D. Lin, and P. F. Craigmile, "Optimal design emulators: A point process approach," 2018, *arXiv:1804.02089*. [Online]. Available: <http://arxiv.org/abs/1804.02089>
- [21] D. H. Svendsen, L. Martino, and G. Camps-Valls, "Active emulation of computer codes with Gaussian processes—Application to remote sensing," *Pattern Recognit.*, vol. 100, Apr. 2020, Art. no. 107103.
- [22] H. Ying, K. Mao, and K. Mosegaard, "Moving target Monte Carlo," 2020, *arXiv:2003.04873*. [Online]. Available: <http://arxiv.org/abs/2003.04873>
- [23] C. E. Rasmussen, J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West, "Gaussian processes to speed up hybrid Monte Carlo for expensive Bayesian integrals," in *Bayesian Statist*, vol. 7. Oxford, U.K.: Oxford Univ. Press, 2003, pp. 651–659.
- [24] W. R. Gilks and P. Wild, "Adaptive rejection sampling for Gibbs sampling," *J. Roy. Stat. Soc. C, Appl. Statist.*, vol. 41, no. 2, pp. 337–348, Jun. 1992.
- [25] W. R. Gilks, N. G. Best, and K. K. C. Tan, "Adaptive rejection metropolis sampling within Gibbs sampling," *Appl. Statist.*, vol. 44, no. 4, pp. 455–472, 1995.
- [26] L. Martino, R. Casarin, F. Leisen, and D. Luengo, "Adaptive independent sticky MCMC algorithms," *EURASIP J. Adv. Signal Process.*, vol. 2018, no. 1, p. 5, Dec. 2018.
- [27] T. Butler, L. Graham, S. Mattis, and S. Walsh, "A measure-theoretic interpretation of sample based numerical integration with applications to inverse and prediction problems under uncertainty," *SIAM J. Sci. Comput.*, vol. 39, no. 5, pp. A2072–A2098, Jan. 2017.
- [28] H. Wang and J. Li, "Adaptive Gaussian process approximation for Bayesian inference with expensive likelihood functions," *Neural Comput.*, vol. 30, no. 11, pp. 3072–3094, Nov. 2018.
- [29] A. M. Stuart and A. L. Teckentrup, "Posterior consistency for Gaussian process approximations of Bayesian posterior distributions," *Math. Comput.*, vol. 87, no. 310, pp. 721–753, Aug. 2017.

- [30] N. Bliznyuk, D. Ruppert, C. Shoemaker, R. Regis, S. Wild, and P. Mugunthan, "Bayesian calibration and uncertainty analysis for computationally expensive models using optimization and radial basis function approximation," *J. Comput. Graph. Statist.*, vol. 17, no. 2, pp. 270–294, Jun. 2008.
- [31] P. R. Conrad, Y. M. Marzouk, N. S. Pillai, and A. Smith, "Accelerating asymptotically exact MCMC for computationally intensive models via local approximations," *J. Amer. Stat. Assoc.*, vol. 111, no. 516, pp. 1591–1607, Oct. 2016.
- [32] K. Kandasamy, J. Schneider, and B. Póczos, "Query efficient posterior estimation in scientific experiments via Bayesian active learning," *Artif. Intell.*, vol. 243, pp. 45–56, Feb. 2017.
- [33] M. Järvenpää, M. U. Gutmann, A. Vehtari, and P. Marttinen, "Parallel Gaussian process surrogate Bayesian inference with noisy likelihood evaluations," *Bayesian Anal.*, Mar. 2020. [Online]. Available: <https://projecteuclid.org/euclid.ba/1584583229>, doi: 10.1214/20-BA1200.
- [34] D. Chauveau and P. Vandekerckhove, "Improving convergence of the Hastings-metropolis algorithm with an adaptive proposal," *Scandin. J. Statist.*, vol. 29, no. 1, pp. 13–29, Mar. 2002.
- [35] T. E. Hanson, J. V. D. Monteiro, and A. Jara, "The polya tree sampler: Toward efficient and automatic independent Metropolis–Hastings proposals," *J. Comput. Graph. Statist.*, vol. 20, no. 1, pp. 41–62, Jan. 2011.
- [36] L. Martino, V. Elvira, D. Luengo, and J. Corander, "Layered adaptive importance sampling," *Statist. Comput.*, vol. 27, no. 3, pp. 599–623, May 2017.
- [37] B. Delyon and F. Portier, "Integral approximation by kernel smoothing," *Bernoulli*, vol. 22, no. 4, pp. 2177–2208, Nov. 2016.
- [38] W. Hörmann, "A rejection technique for sampling from t-concave distributions," *ACM Trans. Math. Softw.*, vol. 21, no. 2, pp. 182–193, Jun. 1995.
- [39] D. Görür and Y. W. Teh, "Concave-convex adaptive rejection sampling," *J. Comput. Graph. Statist.*, vol. 20, no. 3, pp. 670–691, Jan. 2011.
- [40] L. Martino and J. Míguez, "A generalization of the adaptive rejection sampling algorithm," *Statist. Comput.*, vol. 21, no. 4, pp. 633–647, Oct. 2011.
- [41] R. Meyer, B. Cai, and F. Perron, "Adaptive rejection metropolis sampling using Lagrange interpolation polynomials of degree 2," *Comput. Statist. Data Anal.*, vol. 52, no. 7, pp. 3408–3423, Mar. 2008.
- [42] L. Martino, J. Read, and D. Luengo, "Independent doubly adaptive rejection metropolis sampling within gibbs sampling," *IEEE Trans. Signal Process.*, vol. 63, no. 12, pp. 3123–3138, Jun. 2015.
- [43] L. Martino, H. Yang, D. Luengo, J. Kannianen, and J. Corander, "A fast universal self-tuned sampler within gibbs sampling," *Digit. Signal Process.*, vol. 47, pp. 68–83, Dec. 2015.
- [44] J. Felip, N. Ahuja, and O. Tickoo, "Tree pyramidal adaptive importance sampling," 2019, *arXiv:1912.08434*. [Online]. Available: <http://arxiv.org/abs/1912.08434>
- [45] Y. Auffray, P. Barbillon, and J.-M. Marin, "Maximin design on non hypercube domains and kernel interpolation," *Statist. Comput.*, vol. 22, no. 3, pp. 703–712, May 2012.
- [46] L. Pronzato and W. G. Müller, "Design of computer experiments: Space filling and beyond," *Statist. Comput.*, vol. 22, no. 3, pp. 681–701, May 2012.
- [47] L. M. M. van den Bos, B. Sanderse, and W. A. A. M. Bierbooms, "Adaptive sampling-based quadrature rules for efficient Bayesian prediction," *J. Comput. Phys.*, vol. 417, Sep. 2020, Art. no. 109537.
- [48] Y. Chen, M. Welling, and A. Smola, "Super-samples from kernel herding," in *Proc. 26th Conf. Uncertainty Artif. Intell.*, 2010, pp. 1–8.
- [49] S. Lacoste-Julien, F. Lindsten, and F. Bach, "Sequential kernel herding: Frank-Wolfe optimization for particle filtering," in *Proc. 18th Int. Conf. Artif. Intell. Statist.*, 2015, pp. 544–552.
- [50] S. Mak and V. Roshan Joseph, "Support points," 2016, pp. 1–55, *arXiv:1609.01811*. [Online]. Available: <http://arxiv.org/abs/1609.01811>
- [51] L. Martino and V. Elvira, "Compressed Monte Carlo for distributed Bayesian inference," 2018, *viXra:1811.0505*. [Online]. Available: <https://vrxiv.org/abs/1811.0505>
- [52] T. Karvonen, M. Kanagawa, and S. Särkkä, "On the positivity and magnitudes of Bayesian quadrature weights," *Statist. Comput.*, vol. 29, no. 6, pp. 1317–1333, Nov. 2019.
- [53] T. Karvonen and S. Särkkä, "Fully symmetric kernel quadrature," *SIAM J. Sci. Comput.*, vol. 40, no. 2, pp. A697–A720, Jan. 2018.
- [54] R. Schaback, "Reconstruction of multivariate functions from scattered data," Ph.D. dissertation, Inst. Numer. Appl. Math., Univ. Göttingen, Göttingen, Germany, 1997.
- [55] H. Wendland, *Scattered Data Approximation*, vol. 17. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [56] L. Pronzato, "Minimax and maximin space-filling designs: Some properties and methods for construction," *J. de la Societe Franaise de Statistique*, vol. 158, no. 1, pp. 7–36, 2017.
- [57] R. E. Caflisch, "Monte Carlo and quasi-Monte Carlo methods," *Acta Numerica*, vol. 7, pp. 1–49, Jan. 1998.
- [58] V. Elvira, L. Martino, and P. C. Cloas, "Importance Gaussian quadrature," 2020, *arXiv:2001.03090*. [Online]. Available: <http://arxiv.org/abs/2001.03090>
- [59] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [60] W. G. Müller, "Coffee-house designs," in *Optimum Design*. Boston, MA, USA: Springer, 2001, pp. 241–248.
- [61] G. Santin and B. Haasdonk, "Convergence rate of the data-independent P-greedy algorithm in kernel-based approximation," 2017, *arXiv:1612.02672*. [Online]. Available: <https://arxiv.org/abs/1612.02672>
- [62] D. J. C. MacKay, "Information-based objective functions for active data selection," *Neural Comput.*, vol. 4, no. 4, pp. 590–604, Jul. 1992.
- [63] H. Niederreiter, *Random Number Generation and Quasi-Monte Carlo Methods*, vol. 63. Philadelphia, PA, USA: SIAM, 1992.
- [64] N. Aronszajn, "Theory of reproducing kernels," *Trans. Amer. Math. Soc.*, vol. 68, no. 3, pp. 337–404, 1950.
- [65] R. Schaback, "Native Hilbert spaces for radial basis functions I," in *New Developments in Approximation Theory*. New York, NY, USA: Springer, 1999, pp. 255–282.
- [66] M. Golomb and H. F. Weinberger, "Optimal approximations and error bounds," Wisconsin Univ.-Madison Math. Res. Center, Madison, WI, USA, Tech. Rep. 1, 1958.
- [67] A. Sommariva and M. Vianello, "Numerical cubature on scattered data by radial basis functions," *Computing*, vol. 76, nos. 3–4, p. 295, 2006.
- [68] R. Schaback, "Error estimates and condition numbers for radial basis function interpolation," *Adv. Comput. Math.*, vol. 3, no. 3, pp. 251–264, Apr. 1995.
- [69] L. Devroye, L. Györfi, G. Lugosi, and H. Walk, "On the measure of Voronoi cells," *J. Appl. Probab.*, vol. 54, no. 2, pp. 394–408, Jun. 2017.
- [70] M. H. Protter and C. B. Morrey, Jr., *A First Course in Real Analysis*. New York, NY, USA: Springer, 2012.
- [71] F. Llorente, L. Martino, D. Delgado, and J. Lopez-Santiago, "Marginal likelihood computation for model selection and hypothesis testing: An extensive review," 2019, *viXra:2001.0052*. [Online]. Available: <https://vrxiv.org/abs/2001.0052>
- [72] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Míguez, and P. M. Djuric, "Adaptive importance sampling: The past, the present, and the future," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 60–79, Jul. 2017.
- [73] P. C. Gregory, "Bayesian re-analysis of the gliese 581 exoplanet system," *Monthly Notices Roy. Astronomical Soc.*, vol. 415, no. 3, pp. 2523–2545, Aug. 2011.
- [74] S. C. C. Barros et al., "WASP-113b and WASP-114b, two inflated hot jupiters with contrasting densities," *Astron. Astrophys.*, vol. 593, p. A113, Sep. 2016.
- [75] L. Affer et al., "HADES RV program with HARPS-N at TNG. IX. A super-earth around the m dwarf Gl686," Feb. 2019, *arXiv:1901.05338*. [Online]. Available: <http://arxiv.org/abs/1901.05338>
- [76] T. Trifonov, S. Stock, T. Henning, S. Reffert, M. Kürster, M. H. Lee, B. Bitsch, R. P. Butler, and S. S. Vogt, "Two Jovian planets around the giant star HD 202696: A growing population of packed massive planetary pairs around massive stars?" *Astronomical J.*, vol. 157, no. 3, p. 93, Mar. 2019.
- [77] W. H. W. Press, A. S. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*. Cambridge, U.K.: Springer, 2002.



FERNANDO LLORENTE FERNÁNDEZ received the B.Sc. degree in physics from the Universidad Autónoma de Madrid, in 2016, and the M.Sc. degree in mathematical engineering from the Universidad Carlos III de Madrid, in 2018, where he is currently pursuing the Ph.D. degree with the Department of Statistics. His current research interests include Bayesian inference and statistical signal processing.



LUCA MARTINO was born in Palermo, Italy, in 1980. He received the M.Sc. degree in electronic engineering from the Politecnico di Milano, Italy, and the Ph.D. degree in statistical signal processing from the Universidad Carlos III de Madrid, Spain, in 2011. He spent two years with the Department of Statistics, University of Helsinki, Finland. He carried out a Postdoctoral Research at the Sao Paulo Research Foundation FAPESP and also at the Universitat de Valencia. From

September 2017 to December 2018, he was an Assistant Professor with the Universidad Carlos III de Madrid. He is currently a Professor with Universidad Rey Juan Carlos.



VICTOR ELVIRA received the M.Sc. and Ph.D. degrees in electrical engineering from the Universidad de Cantabria, Spain, in 2007, 2008, and 2011, respectively. Since 2012, he has been a Postdoctoral Researcher with the Universidad Carlos III de Madrid, Spain. He has been a Visiting Scholar with the IHP Institute, Frankfurt an der Oder, Germany, also with the University of Helsinki, Finland, and also with Stony Brook University, NY, USA. He is currently a Professor with

The University of Edinburgh. His research interests include statistical signal processing and biomedical signal processing.



DAVID DELGADO received the M.Sc. degree in mathematics from the Autonomous University of Madrid, Madrid, Spain, in 1999, and the Ph.D. degree in applied mathematics from the Technical University of Denmark, Lyngby, in 2005. He combines research and teaching as an Associate Professor with the Universidad Carlos III de Madrid. He has received several teaching and research awards. He has published 39 JCR articles. His research interests include applied multivariate

statistical analysis, machine learning, and pattern recognition.



JAVIER LÓPEZ-SANTIAGO received the Ph.D. degree in physics from the Universidad Complutense de Madrid, in 2010. He is currently an Assistant Professor with the Department of Signal Theory and Communications, Universidad Carlos III de Madrid. He has published more than 60 JCR articles. His research interests include statistical signal processing, Monte Carlo methods, and the simulation of digital communication systems.

...