

Improved Calibration of Numerical Integration Error in Sigma-Point Filters

Jakub Průher , Toni Karvonen , Chris J. Oates , Ondřej Straka , and Simo Särkkä 

Abstract—The sigma-point filters, such as the unscented Kalman filter, are popular alternatives to the ubiquitous extended Kalman filter. The classical quadrature rules used in the sigma-point filters are motivated via polynomial approximation of the integrand; however, in the applied context, these assumptions cannot always be justified. As a result, a quadrature error can introduce bias into estimated moments, for which there is no compensatory mechanism in the classical sigma-point filters. This can lead in turn to estimates and predictions that are poorly calibrated. In this article, we investigate the Bayes–Sard quadrature method in the context of sigma-point filters, which enables uncertainty due to quadrature error to be formalized within a probabilistic model. Our first contribution is to derive the well-known classical quadratures as special cases of the Bayes–Sard quadrature method. Based on this, a general-purpose moment transform is developed and utilized in the design of a novel sigma-point filter, which explicitly accounts for the additional uncertainty due to quadrature error.

Index Terms—Bayesian quadrature (BQ), Gaussian processes (GPs), Kalman filters, quantification of uncertainty, sigma points.

I. INTRODUCTION

This article is concerned with quantification of uncertainty associated with sigma-point approximations, which are widely employed in nonlinear filtering algorithms, such as the unscented Kalman filter (UKF). The goal of filtering algorithms is to estimate the state of a dynamical stochastic system based on all measurements obtained until the present. The applications of filters are manifold, ranging from global positioning [1], object tracking [2], [3], and simultaneous localization and mapping [4] to weather forecasting [5] and finance [6].

Manuscript received March 25, 2020; accepted April 17, 2020. Date of publication May 6, 2020; date of current version February 26, 2021. The work of Jakub Průher and Ondřej Straka was supported by the Czech Ministry of Education, Youth and Sports through Project LO1506. The work of Toni Karvonen was supported by the Aalto ELEC Doctoral School. The work of Toni Karvonen and Chris J. Oates was supported by the Lloyd's Register Foundation Program on Data-Centric Engineering at the Alan Turing Institute, U.K. The work of Simo Särkkä was supported by the Academy of Finland. Recommended by Associate Editor S. S. Saab. (*Corresponding author: Jakub Průher.*)

Jakub Průher and Ondřej Straka are with the Department of Cybernetics, University of West Bohemia, 30100 Pilsen, Czech Republic Department of Electrical Engineering and Automation, Aalto University, 02150 Espoo, Finland.

Toni Karvonen is with the Alan Turing Institute, London NW1 2DB, U.K. (e-mail: tkarvonen@turing.ac.uk).

Chris J. Oates is with the School of Mathematics, Statistics and Physics, Newcastle University, Newcastle Upon Tyne NE1 7RU, U.K., and also with the Alan Turing Institute, London NW1 2DB, U.K. (e-mail: chris.oates@ncl.ac.uk).

Simo Särkkä is with the Department of Electrical Engineering and Automation, Aalto University, 02150 Espoo, Finland (e-mail: simo.sarkka@aalto.fi).

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2020.2991698

Instead of keeping track of the whole state posterior, the sigma-point filters only work with mean and covariance of the state and the measurement. For nonlinear systems and/or measurements, the moments are defined by intractable integrals that have to be approximated using numerical quadratures, also known as the sigma-point rules (which is where the filters get their name). The classical quadrature rules, such as the Gauss–Hermite (GH) rule, are designed with the assumption that the nonlinear integrand is well approximated with a polynomial of a given maximal degree. Since these assumptions are almost never met in practice, there will always be a quadrature error involved. Standard sigma-point filters do not attempt to compensate for this source of error, and in practice, this can lead to estimates and predictions that are biased and overconfident [7], [8].

In recent years, the *Bayesian quadrature* (BQ) has received much attention in the probabilistic numerics community [9]–[12]. The BQ approach posits that the integrand can be modeled by a stochastic process defined on the domain of integration. This model is subsequently refined by conditioning on pointwise evaluations of the integrand, which induces a posterior distribution over the value of the integral. The posterior mean of this distribution is point estimate of the value of the integral, while the posterior variance quantifies the integration error.

Applications of the BQ in nonlinear filtering have appeared previously in [13] and [14] with encouraging results. These BQ-based filters do not generally coincide with any classical sigma-point filter, such as the UKF or Gauss–Hermite Kalman filter (GHKF), and tend to be rather sensitive to specification of the stochastic process model for the integrand. It has been shown that classical sigma-point rules can be cast as *degenerate* BQ rules [9], [13]. This is to say that the variance associated with the integral vanishes, being thus of no use in modeling integration error.

In this article, we utilize the recently proposed Bayes–Sard quadrature (BSQ) [15] for the design of novel sigma-point filters, which can be viewed as probabilistic versions of the well-known sigma-point filters. Namely, under certain conditions, the BSQ allows us to recover the classical sigma-point rules and, at the same time, endow the sigma-point rule with nondegenerate probabilistic output. We, thus, obtain versions of standard sigma-point filters that are, to some extent, capable of accounting for numerical integration error in filtering by inflating the error covariance. In some cases, such covariance inflation is known to improve stability of nonlinear Kalman filters (see, for instance, [16, Remark 1], [17, Sec. 3.3], and [18, Sec. V.C]).

The rest of this article is structured as follows. In II, we formally outline the nonlinear filtering problem and the nature of sigma-point approximations. III identifies the moment transformation problem as the central issue in sigma-point filtering and describes the structure of sigma-point moment transforms. The BSQ is formalized in IV, which is later used in V to design the Bayes–Sard quadrature moment transform (BSQMT). Finally, VII concludes this article.

II. SIGMA-POINT FILTERING

This section is devoted to the sigma-point filters, which are a subset of nonlinear filtering algorithms characterized by their reliance on a

Gaussian approximation together with a numerical quadrature method. Let the stochastic dynamical system and the process by which its state is observed be described by the state-space model

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}) + \mathbf{q}_{k-1} \quad (1)$$

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{r}_k \quad (2)$$

where the function $\mathbf{f} : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ is the system dynamics, $\mathbf{h} : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_z}$ is the measurement model, $\mathbf{x}_k \in \mathbb{R}^{d_x}$ is the latent state vector, and $\mathbf{z}_k \in \mathbb{R}^{d_z}$ is the measurement vector. Both the process noise $\mathbf{q}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ and the measurement noise $\mathbf{r}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ are zero-mean white Gaussian sequences, independent of each other and independent of the system initial condition $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{m}_0^x, \mathbf{P}_0^x)$.

The Bayesian formulation of the filtering problem can be summarized by the following two general relations. The state posterior is $p(\mathbf{x}_k | \mathbf{z}_{1:k}) \propto p(\mathbf{z}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{z}_{1:k-1})$, where the likelihood $p(\mathbf{z}_k | \mathbf{x}_k)$ is obtained from the measurement model (2) and $\mathbf{z}_{1:k} \triangleq \{\mathbf{z}_1, \dots, \mathbf{z}_k\}$. The predictive density is given by the Chapman–Kolmogorov equation $p(\mathbf{x}_k | \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}) d\mathbf{x}_{k-1}$, where the transition density $p(\mathbf{x}_k | \mathbf{x}_{k-1})$ is obtained from the system dynamics (1).

A vast majority of well-known filters, such as extended Kalman filter (EKF), UKF and GHKF, can be recovered from the Bayesian formulation under a Gaussian approximation of the joint density of the state and measurement. That is, when the density $p(\mathbf{x}_k, \mathbf{z}_k | \mathbf{z}_{1:k-1}) = p(\mathbf{z}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{z}_{1:k-1})$ is approximated by a Gaussian density of the form

$$\mathcal{N} \left(\begin{bmatrix} \mathbf{x}_k \\ \mathbf{z}_k \end{bmatrix} \middle| \begin{bmatrix} \mathbf{m}_{k|k-1}^x \\ \mathbf{m}_{k|k-1}^z \end{bmatrix}, \begin{bmatrix} \mathbf{P}_{k|k-1}^x & \mathbf{P}_{k|k-1}^{xz} \\ \mathbf{P}_{k|k-1}^{zx} & \mathbf{P}_{k|k-1}^z \end{bmatrix} \right) \quad (3)$$

then the mean and covariance of the state posterior have analytical form, given by¹

$$\mathbf{m}_{k|k}^x = \mathbf{m}_{k|k-1}^x + \mathbf{G}_k \left(\mathbf{z}_k - \mathbf{m}_{k|k-1}^z \right) \quad (4)$$

$$\mathbf{P}_{k|k}^x = \mathbf{P}_{k|k-1}^x - \mathbf{G}_k \mathbf{P}_{k|k-1}^z \mathbf{G}_k^\top \quad (5)$$

where $\mathbf{G}_k = \mathbf{P}_{k|k-1}^{xz} (\mathbf{P}_{k|k-1}^z)^{-1}$ is the Kalman gain. The predictive moments of the state, $\mathbf{m}_{k|k-1}^x$ and $\mathbf{P}_{k|k-1}^x$, and the moments of measurements, $\mathbf{m}_{k|k-1}^z$, $\mathbf{P}_{k|k-1}^z$, and $\mathbf{P}_{k|k-1}^{xz}$, are defined as integrals of the form

$$\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})] \triangleq \int \mathbf{g}(\mathbf{x}) \mathcal{N}(\mathbf{x} | \mathbf{m}, \mathbf{P}) d\mathbf{x}. \quad (6)$$

Table I shows which quantities have to be substituted for $\mathbf{x}(\mathbf{x})$, \mathbf{x} , \mathbf{m} , and \mathbf{P} to obtain any of the above predictive moments. Since the function \mathbf{g} being integrated is nonlinear in each case, these integrals cannot be typically computed analytically, and some type of approximation needs to be employed. Each nonlinear filter is distinguished solely by the type of integral approximation it uses. For example, the EKF employs the first-order Taylor expansion to linearize \mathbf{g} in the vicinity of \mathbf{m} , which, in turn, facilitates analytic tractability of the moment integrals. On the other hand, the sigma-point filters, such as the UKF and the GHKF, leverage numerical quadrature for approximation of the integral. Since quadratures are typically designed to be used with standard Gaussian, the integrals of the form (6) need to be converted by employing a stochastic decoupling substitution $\mathbf{x}^{(n)} = \mathbf{m} + \mathbf{L}\xi^{(n)}$, which leads to

¹Note that $\mathbf{m}_{k|k}^x \triangleq \mathbb{E}_{\mathbf{x}}[\mathbf{x}_k | \mathbf{z}_{1:k}]$ and $\mathbf{P}_{k|k}^x \triangleq \mathbb{E}_{\mathbf{x}}[(\mathbf{x}_k - \mathbf{m}_{k|k}^x)(\mathbf{x}_k - \mathbf{m}_{k|k}^x)^\top | \mathbf{z}_{1:k}]$.

TABLE I
QUANTITIES THAT NEED TO BE SUBSTITUTED INTO THE GAUSSIAN INTEGRAL (6) IN ORDER TO OBTAIN EVERY PREDICTIVE MOMENT NECESSARY TO COMPUTE THE MOMENTS OF THE STATE POSTERIOR

Moment	$\mathbf{g}(\mathbf{x})$	\mathbf{x}	\mathbf{m}	\mathbf{P}
$\mathbf{m}_{k k-1}^x$	$\mathbf{f}(\mathbf{x}_{k-1})$	\mathbf{x}_{k-1}	$\mathbf{m}_{k-1 k-1}^x$	$\mathbf{P}_{k-1 k-1}^x$
$\mathbf{P}_{k k-1}^x$	$\Delta \mathbf{f} \Delta \mathbf{f}^\top$	\mathbf{x}_{k-1}	$\mathbf{m}_{k-1 k-1}^x$	$\mathbf{P}_{k-1 k-1}^x$
$\mathbf{m}_{k k-1}^z$	$\mathbf{h}(\mathbf{x}_k)$	\mathbf{x}_k	$\mathbf{m}_{k k-1}^x$	$\mathbf{P}_{k k-1}^x$
$\mathbf{P}_{k k-1}^z$	$\Delta \mathbf{h} \Delta \mathbf{h}^\top$	\mathbf{x}_k	$\mathbf{m}_{k k-1}^x$	$\mathbf{P}_{k k-1}^x$
$\mathbf{P}_{k k-1}^{xz}$	$\Delta \mathbf{x} \Delta \mathbf{h}^\top$	\mathbf{x}_k	$\mathbf{m}_{k k-1}^x$	$\mathbf{P}_{k k-1}^x$

The following shorthand notation is used: $\Delta \mathbf{f} = \mathbf{f}(\mathbf{x}_{k-1}) - \mathbf{m}_{k-1|k-1}^x$, $\Delta \mathbf{h} = \mathbf{h}(\mathbf{x}_k) - \mathbf{m}_{k|k-1}^z$, and $\Delta \mathbf{x} = \mathbf{x}_k - \mathbf{m}_{k|k-1}^x$.

an approximation

$$\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})] \approx \sum_{n=1}^N w_n \mathbf{g}(\mathbf{m} + \mathbf{L}\xi^{(n)}) = \sum_{n=1}^N w_n \tilde{\mathbf{g}}(\xi^{(n)}) \quad (7)$$

where $\xi^{(n)}$ denotes the n th unit sigma point, $w_n \in \mathbb{R}$ is the n th weight, N is the total number of sigma points, \mathbf{L} is a matrix factor such that $\mathbf{P} = \mathbf{L}\mathbf{L}^\top$, and $\tilde{\mathbf{g}}(\xi) \triangleq \mathbf{g}(\mathbf{m} + \mathbf{L}\xi)$. Note that various quadrature rules are distinguished by different weights and sigma points they prescribe to satisfy various optimality criteria.

III. SIGMA-POINT MOMENT TRANSFORMS

From the above exposition, it is apparent that the central issue in filtering is the design of the so-called moment transformations, which generate approximations to the moments of a random variable under a nonlinear transformation.

Let $\mathbf{x} \in \mathbb{R}^D$ be an input Gaussian random variable and $\mathbf{y} \in \mathbb{R}^E$ an output random variable defined by

$$\mathbf{y} = \mathbf{g}(\mathbf{x}), \quad \mathbf{x} \sim \mathcal{N}(\mathbf{m}, \mathbf{P}). \quad (8)$$

If the transformation \mathbf{g} is nonlinear, the joint density $p(\mathbf{x}, \mathbf{y})$ will be non-Gaussian in general. However, there are many applied situations, where \mathbf{g} is approximately linear in the region where probability mass is concentrated. In such situations, the principal error term in the moment transform is numerical quadrature error. This error is the focus this article, and therefore, in what follows, we proceed under the assumption that the Gaussian approximation

$$\mathcal{N} \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \middle| \begin{bmatrix} \mathbf{m} \\ \boldsymbol{\mu} \end{bmatrix}, \begin{bmatrix} \mathbf{P} & \mathbf{C} \\ \mathbf{C}^\top & \boldsymbol{\Pi} \end{bmatrix} \right) \quad (9)$$

of $p(\mathbf{x}, \mathbf{y})$ can be justified. In this setting, the moment transformation then reduces to computing the output mean $\boldsymbol{\mu}$, covariance $\boldsymbol{\Pi}$, and cross-covariance \mathbf{C} as accurately as possible, when supplied with the input moments, \mathbf{m} and \mathbf{P} . This is a specific instance of *uncertainty propagation* [19].

In this article, we focus on the sigma-point approximations, exemplified by (7), to the moment integrals in Table I. The well-known classical approximations, such as the GH, the spherical–radial, and the unscented transform (UT), are conventionally written in the form

$$\boldsymbol{\mu} \approx \hat{\boldsymbol{\mu}} = \sum_{n=1}^N w_n \tilde{\mathbf{g}}(\xi^{(n)}) \quad (10)$$

$$\mathbf{\Pi} \approx \hat{\mathbf{\Pi}} = \sum_{n=1}^N w_n (\tilde{\mathbf{g}}(\xi^{(n)}) - \hat{\boldsymbol{\mu}})(\tilde{\mathbf{g}}(\xi^{(n)}) - \hat{\boldsymbol{\mu}})^\top \quad (11)$$

$$\mathbf{C} \approx \hat{\mathbf{C}} = \mathbf{L} \sum_{n=1}^N w_n \xi^{(n)} (\tilde{\mathbf{g}}(\xi^{(n)}) - \hat{\boldsymbol{\mu}})^\top \quad (12)$$

which, under the assumption that $\sum w_n = 1$ and $\sum w_n \xi^{(n)} = 0$, we will prefer to write using the matrix notation as

$$\hat{\boldsymbol{\mu}} = \mathbf{Y}^\top \mathbf{w} \quad (13a)$$

$$\hat{\mathbf{\Pi}} = \mathbf{Y}^\top \mathbf{W} \mathbf{Y} - \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top \quad (13b)$$

$$\hat{\mathbf{C}} = \mathbf{L} \mathbf{\Xi} \mathbf{W}_c \mathbf{Y} \quad (13c)$$

where $\mathbf{\Xi} = [\xi^{(1)} \dots \xi^{(N)}]$ and the matrix of integrand evaluations is given by $[\mathbf{Y}]_{ne} \triangleq \tilde{g}_e(\xi^{(n)})$, where e indexes outputs of $\tilde{\mathbf{g}}$ and $[\cdot]_{ne}$ denotes the matrix element at position (n, e) . The vector \mathbf{w} contains the weights and $\mathbf{W} = \mathbf{W}_c = \text{diag}(\mathbf{w})$ for any classical sigma-point moment transform. Each moment transform uses a different set of sigma points and weights.

A. Unscented Transform

The UT of D -dimensional input uses $N = 2D + 1$ sigma points, which exploit symmetry of the Gaussian distribution, given, for $d = 1, \dots, D$, by

$$\xi^{(0)} = 0, \quad \xi^{(d)} = \sqrt{c} \mathbf{e}_d, \quad \xi^{(D+d)} = -\sqrt{c} \mathbf{e}_{D+d} \quad (14)$$

where \mathbf{e}_d is the standard unit vector and $c = D + \kappa$ for a scaling parameter κ . The weights are defined as $w_0 = \frac{\kappa}{c}$ and $w_d = w_{D+d} = \frac{1}{2c}$. This selection of sigma points and weights yields a quadrature rule that integrates exactly all polynomials of (total) degree at most three; the derivation is essentially contained in the proof of 1. The *spherical-radial rule*, which is used in the cubature Kalman filter [20], is equivalent to the UT with $\kappa = 0$; it, therefore, lacks the central sigma point.

B. GH Rule

From nonsingularity of the Vandermonde matrix $[\mathbf{V}]_{nm} = x_n^{m-1}$ for any distinct sigma points $x_1, \dots, x_p \in \mathbb{R}$, it follows that there are unique weights such that $\sum_{n=1}^p w_n x_n^m = \int x^m \mathbf{N}(x | 0, 1) dx$ for every $m \leq p - 1$ (i.e., the rule has a degree of exactness $p - 1$). However, degree of exactness $2p - 1$ can be achieved with p sigma points if these are selected to be the roots of the p th-degree Hermite polynomial H_p . The weights are then given by $w_n = \frac{p!}{p^2 H_{p-1}(\xi^{(n)})^2}$. This is the GH rule [21]–[24]. In multivariate versions, the sigma points are formed as Cartesian products of the aforementioned one-dimensional points, and the weights are products of w_n . The multivariate GH rule exactly integrates functions in the space

$$\mathbb{P}_{2p-1}^{\max} \triangleq \text{span} \left\{ \mathbf{x}^\alpha : \alpha \in \mathbb{N}_0^D, \max_{d=1, \dots, D} \alpha_d \leq 2p - 1 \right\} \quad (15)$$

where $\mathbf{x}^\alpha = \prod_{d=1}^D x_d^{\alpha_d}$ denotes multivariate monomial. Because of the Cartesian product design, the number of points, $N = p^D$, in the GH rule grows exponentially with dimension, which makes it practically unattractive for $D > 5$ [22]. The problem can be partially mitigated by using sparse grids [25].

IV. BAYESIAN QUADRATURE

This section reviews the underlying philosophy of the BQ as an alternative perspective on numerical integration and describes the BSQ as a necessary stepping stone on the way to building the Bayes–Sard

moment transform proposed in V. A general formulation of the BQ is presented for integrals

$$\mathbb{E}_{\mathbf{x}}[g^\dagger(\mathbf{x})] = \int g^\dagger(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \quad (16)$$

with arbitrary density function p . Vector-valued integrands are discussed in Section IV-B. The moment transform proposed in Section V then specializes to the case $p(\mathbf{x}) = \mathbf{N}(\mathbf{x} | 0, \mathbf{I})$. Throughout this section, the true integrand will be denoted by g^\dagger to distinguish it from the stochastic model of the integrand.

From (7), it is clear that the quadrature approximation of the integral (6) is based on limited knowledge about the behavior of the integrand, because it only relies on finitely many evaluations. The design of classical quadrature rules typically involves formation of polynomial interpolant passing through the observed function values, which is then integrated instead of the intractable integrand. The polynomial interpolation of the integrand consequently implies that the classical rules are only able to integrate polynomial integrands exactly. Another downside of the classical rules is that they are unable to account for the functional uncertainty (interpolation error), which occurs when the integrand is not a polynomial.

The Bayesian approach to quadrature [11], [26], [27] aims to address these limitations by treating the numerical approximation of intractable integrals as a problem of Bayesian statistical inference, where a prior for the integrand is specified by a stochastic process model $g(\mathbf{x})$ with the user-defined mean function $m(\mathbf{x}) = \mathbb{E}_g[g(\mathbf{x})]$ and the covariance (or kernel) function $k(\mathbf{x}, \mathbf{x}') = \mathbb{C}_g[g(\mathbf{x}), g(\mathbf{x}')]$, where \mathbf{x}' denotes the second argument (not a transpose). The dataset $\mathcal{D} = \{(\mathbf{x}^{(n)}, g^\dagger(\mathbf{x}^{(n)}))\}_{n=1}^N$ comprises evaluations of the integrand $g^\dagger(\mathbf{x}^{(n)})$ at predefined points $\mathbf{x}^{(n)}$. Conditioning on \mathcal{D} leads to a posterior stochastic process, with mean $m_{\mathcal{D}}(\mathbf{x}) = \mathbb{E}_{g|\mathcal{D}}[g(\mathbf{x})]$ and covariance $k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}') = \mathbb{C}_{g|\mathcal{D}}[g(\mathbf{x}), g(\mathbf{x}')]$, which, in turn, induces a posterior marginal distribution on the value of the integral $\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]$, with the first two moments given by [28]

$$\mathbb{E}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}_{g|\mathcal{D}}[g(\mathbf{x})]] \quad (17)$$

$$\mathbb{V}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}, \mathbf{x}'}[\mathbb{C}_{g|\mathcal{D}}[g(\mathbf{x}), g(\mathbf{x}')]]. \quad (18)$$

The mean is a convenient point estimate, while the full posterior serves as a probabilistic model of the integration error. The most common stochastic process model of the integrand is a Gaussian process (GP), which has been studied extensively [11], [29].

A. Bayes–Sard GP Model

Let π be a linear function space spanned by $Q \leq N$ functions $\phi_1, \dots, \phi_Q : \mathbb{R}^D \rightarrow \mathbb{R}$. Modeling of the scalar integrand $g^\dagger : \mathbb{R}^D \rightarrow \mathbb{R}$ in the BSQ begins by considering a hierarchical GP prior given by

$$\gamma \sim \mathbf{N}(0, \Sigma_\pi), \quad \gamma \in \mathbb{R}^Q \quad (19)$$

$$m(\mathbf{x}) = \sum_{q=1}^Q \gamma_q \phi_q(\mathbf{x}) \quad (20)$$

$$g(\mathbf{x}) \sim \text{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'; \theta)) \quad (21)$$

where the prior mean function $m(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}$ is composed of basis functions $\phi_q(\mathbf{x})$ of Q -dimensional linear space π and the prior covariance function (kernel) $k(\mathbf{x}, \mathbf{x}'; \theta) : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ can be any symmetric positive-definite function parameterized by the vector θ (see Section IV-E for concrete example). The dependence on θ will be tacitly assumed and explicitly denoted only when required. Discussion about the particular choice of the kernel and its effects is postponed to IV-E.

The above model differs from the one often used in GP-based BQ, in that the prior mean function is nonzero and its coefficients are random.

The next phase in modeling is to consider a flat prior limit on the mean function coefficients, such that $\Sigma_\pi \rightarrow \infty$ [30, Ch. 4]. In order for the GP posterior to be well defined, the set $\chi = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$ of sigma points must meet the following condition of π -unisolvency, which is related to existence of interpolants formed out of linear combinations of ϕ_1, \dots, ϕ_q .

Definition 1 (π -unisolvency): Let π be a Q -dimensional linear space spanned by $\{\phi_1, \dots, \phi_q\}$. A point set χ is said to be π -unisolvant if and only if the $N \times Q$ alternant matrix $[\Phi]_{qn} \triangleq \phi_q(\mathbf{x}^{(n)})$ is of full rank.

We further restrict the model to the case when $N = Q$, which means that the alternant matrix Φ is square and, due to π -unisolvency of χ , invertible.

With all the assumptions laid out, the final step is to condition the GP on the set of sigma points χ and the corresponding integrand evaluations, to arrive at the posterior moments of the Bayes–Sard GP model given by [15] as

$$\mathbb{E}_{g|\mathcal{D}}[g(\mathbf{x})] = \phi(\mathbf{x})^\top \Phi^{-1} \mathbf{y} \quad (22)$$

$$\begin{aligned} \mathbb{C}_{g|\mathcal{D}}[g(\mathbf{x}), g(\mathbf{x}')] &= k(\mathbf{x}, \mathbf{x}') - 2k(\mathbf{x}, \mathbf{x}')^\top \Phi^{-1} \phi(\mathbf{x}') \\ &\quad + \phi(\mathbf{x})^\top [\Phi^\top \mathbf{K}^{-1} \Phi]^{-1} \phi(\mathbf{x}') \end{aligned} \quad (23)$$

where $[\mathbf{k}(\mathbf{x})]_n \triangleq k(\mathbf{x}, \mathbf{x}^{(n)})$, $[\phi(\mathbf{x})]_q \triangleq \phi_q(\mathbf{x})$, and $[\mathbf{y}]_n \triangleq g^\dagger(\mathbf{x}^{(n)})$, where $[\cdot]_n$ denotes the n th element of the given vector. Note that the posterior mean now only depends on the choice of the function space π , and the kernel affects only the posterior covariance. It is worth pointing out that all sigma-point sets in the established classical filters are π -unisolvant.

B. Vector-Valued Integrand

Until now, we have only considered scalar-valued integrands. The model specified by (22) and (23) can be straightforwardly extended to vector-valued integrands $\mathbf{g}^\dagger: \mathbb{R}^D \rightarrow \mathbb{R}^E$ that comply with the specification of the moment transformation problem in (8). Noting that we can decompose the integrand as $\mathbf{g}^\dagger(\mathbf{x}) = [g_1^\dagger(\mathbf{x}) \dots g_E^\dagger(\mathbf{x})]^\top$, the simplest solution is to use (22) and (23) to model each coordinate function independently, either using a common kernel parameter for all outputs, which is accomplished by

$$g_e(\mathbf{x}) | \mathcal{D} \sim \text{GP}(m_{\mathcal{D}}(\mathbf{x}), k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta)) \quad (24)$$

or using a different kernel parameter values for each output, so that

$$g_e(\mathbf{x}) | \mathcal{D} \sim \text{GP}(m_{\mathcal{D}}(\mathbf{x}), k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta_e)) \quad (25)$$

for all $e = 1, \dots, E$. In both cases, the GP posterior mean function is given as

$$\mathbf{m}_{\mathcal{D}}(\mathbf{x}) \triangleq \mathbb{E}_{g|\mathcal{D}}[g(\mathbf{x})] = \mathbf{Y}^\top \Phi^{-1} \phi(\mathbf{x}) \quad (26)$$

where $[\mathbf{Y}]_{ne} = g_e^\dagger(\mathbf{x}^{(n)})$. For the single-parameter model (24), the posterior covariance becomes

$$\mathbf{K}_{\mathcal{D}}(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{C}_{g|\mathcal{D}}[\mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x}')] = k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta) \mathbf{I}_E \quad (27)$$

and for the multiparameter model (25), we get

$$\mathbf{K}_{\mathcal{D}}(\mathbf{x}, \mathbf{x}') = \text{diag}([k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta_1) \dots k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta_E)]). \quad (28)$$

Both of these modeling choices assume that the outputs are conditionally independent given the inputs. Alternatively, the use of multioutput GPs [31], [32] would make it possible to model correlations between coordinate functions and use coordinate-dependent sigma points at the expense of increased computational cost.

C. Bayes–Sard Quadrature

The advantage of using a GP for modeling the integrand is that as it gets transformed by the integral, which is a linear operator, the resulting distribution over the value of the integral is also Gaussian. The BSQ [15], [27], [33] enables enforcing exactness conditions of the form

$$\mathbb{E}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] = \int \mathbf{g}^\dagger(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}$$

for all functions $\mathbf{g}^\dagger: \mathbb{R}^D \rightarrow \mathbb{R}^E$ such that $g_e^\dagger \in \pi$ for each $e = 1, \dots, E$. As shown in IV-D, the classical quadrature methods can be replicated by judicious choice of the function space π . The posterior integral mean and variance under the BSQ are straightforwardly derived by plugging the Bayes–Sard GP model moments from (26) and (28) into the general BQ expressions in (17) and (18).

For the mean of the posterior distribution of the integral, we have

$$\mathbb{E}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}}[m_{\mathcal{D}}(\mathbf{x})], \mathbf{Y}^\top \Phi^{-1} \bar{\phi} \quad (29)$$

where $[\bar{\phi}]_q = \mathbb{E}_{\mathbf{x}}[\phi_q(\mathbf{x})] = \int \phi_q(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}$. Recognizing that the vector of quadrature weights is $\mathbf{w} = \Phi^{-1} \bar{\phi}$, we see that the posterior mean of the integral

$$\mathbb{E}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] = \mathbf{Y}^\top \mathbf{w} = \sum_{n=1}^N w_n \mathbf{g}^\dagger(\mathbf{x}^{(n)}) \quad (30)$$

takes on the form of weighted sum from (7). The integral covariance becomes

$$\mathbb{V}_{g|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}, \mathbf{x}'}[\mathbf{K}_{\mathcal{D}}(\mathbf{x}, \mathbf{x}')] = \text{diag}([\bar{k}_D^1 \dots \bar{k}_D^E]) \quad (31)$$

where $\bar{k}_D^e \triangleq \mathbb{E}_{\mathbf{x}, \mathbf{x}'}[k_{\mathcal{D}}(\mathbf{x}, \mathbf{x}'; \theta_e)]$ and

$$\bar{k}_D^e = \bar{k} - 2\bar{\mathbf{K}}^\top \Phi^{-1} \bar{\phi} + \bar{\phi}^\top [\Phi^\top \mathbf{K}^{-1} \Phi]^{-1} \bar{\phi}. \quad (32)$$

Since the single-parameter model in (27) is a special case of (28), the posterior integral variance under this model would be a trivial modification of (31).

D. Relationship to Classical Sigma-Point Rules

As stated in the previous section, careful selection of π (via the basis functions ϕ_q) allows for recovery of many well-known classical quadrature rules used in nonlinear filtering. In the following, we show that the UT and the GH rule are special cases of the BSQ whenever the space π is spanned by suitably selected polynomial basis. Similar results can be proved for many other sigma-point rules. Note that the BSQ reports a nonzero integral variance even for \mathbf{g}^\dagger whose coordinate functions are in π (and hence integrated exactly). This behavior is desirable because, given only a finite set of function values, one can never tell with certainty the true nature of the integrand.

Theorem 1: Consider the standard Gaussian distribution $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{I})$. Select the $2D + 1$ dimensional function space

$$\pi = \text{span}\{1, x_1, \dots, x_D, x_1^2, \dots, x_D^2\} \quad (33)$$

and the $N = 2D + 1$ UT points (14). Then, the Bayes–Sard weights $\mathbf{w} = \Phi^{-1} \bar{\phi}$ that determine the posterior mean (29) coincide with the UT weights.

Proof: Because $\dim(\pi) = N$, the Bayes–Sard weights \mathbf{w} solve the linear system $\Phi \mathbf{w} = \bar{\phi}$. That is, they are the unique weights such that

$$\sum_{n=0}^{2D} w_n v(\boldsymbol{\xi}^{(n)}) = \int v(\mathbf{x}) \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{I}) d\mathbf{x} \quad (34)$$

for every polynomial $v \in \pi$. In the following, let $d = 1, \dots, D$. We have $\int N(\mathbf{x}|0, \mathbf{I})d\mathbf{x} = 1$ and $\int x_d^2 N(\mathbf{x}|0, \mathbf{I})d\mathbf{x} = 1$. Consequently, (34) is equivalent to $\sum_{n=0}^{2D} w_n = 1$, $\sum_{n=0}^{2D} w_n x_{n,d} = 0$, and $\sum_{n=0}^{2D} w_n x_{n,d}^2 = 1$. Because $\xi_d = -\xi_{D+d}$, the second of these equations implies that $w_d = w_{D+d}$, while the third one yields $w_d = w_{D+d} = \frac{1}{2c}$. Furthermore, $w_0 = \frac{c}{c}$ due to the weights summing up to 1. We have thus solved the BSQ weights $\mathbf{w} = \Phi^{-1}\bar{\phi}$ and see that they are precisely UT weights in Section III-A. ■

Theorem 2: Consider the standard Gaussian distribution $p(\mathbf{x}) = N(\mathbf{x}|\mathbf{0}, \mathbf{I})$, and let $p \geq 1$. Select the p^D dimensional function space $\pi = \Pi_{p-1}^{\max} \triangleq \text{span}\{\mathbf{x}^\alpha : \alpha \in \mathbb{N}_0^D, \max_{d=1, \dots, D} \alpha_d \leq p-1\}$, and the points that constitute the Cartesian product of the roots of the p th-degree Hermite polynomial. Then, the Bayes–Sard weights $\mathbf{w} = \Phi^{-1}\bar{\phi}$ that determine the posterior mean (29) coincide with the classical GH weights from III-B.

Proof: Since the Bayes–Sard weights yield, by their definition, a quadrature rule exact for functions in π and it is known that, given the GH points, the GH weights are the unique weights that determine a quadrature rule exact for this very same function space (see III-B), the result follows. ■

E. Choice of Kernel

As already noted, the posterior mean for the integral produced by the BSQ depends only on π and the kernel controls the posterior variance of the integral. The reasonableness of the BSQ output depends on the reasonableness of the assumption that g^\dagger is “well modeled” by the GP specified by the kernel k . Consequently, selection of the kernel is important in order to ensure that the integral variance is meaningful in modeling the integration error. At the same time, the functional form of the kernel is constrained by the requirement in the BSQ to analytically compute the integral of the kernel. To facilitate analytic tractability of the Bayes–Sard moment transform, introduced next, we use the radial basis function (RBF) kernel

$$k(\mathbf{x}, \mathbf{x}') = \alpha^2 \prod_{d=1}^D \exp\left(-\frac{(x_d - x'_d)^2}{2\ell_d^2}\right) \quad (35)$$

throughout the remainder. The parameters θ of this kernel consist of the scale parameter $\alpha > 0$ and dimensionwise lengthscale parameters $\ell_1, \dots, \ell_D > 0$. A particular modeling assumption associated with this kernel is that the integrand is infinitely differentiable. If this is not the case (i.e., there is model misspecification), the proposed method still works, but the uncertainty quantification for the integral may be rendered less meaningful. For certain classes of kernels, it has been shown that convergence rates to the true integral as $N \rightarrow \infty$ are not much affected by model misspecification [34].

V. BAYES–SARD MOMENT TRANSFORM

The simplest way to design a moment transform is to use the BSQ directly for approximation of the moment integrals in (13a)–(13c). However, this design does not reflect integral uncertainty, which is the key advantage of BQ, not to mention the fact that we would only obtain the classical rules as a result. To resolve this issue, we employ the same general conceptual framework used in the design of the GPQ moment transform in [14], which can account for the variance of the mean integral (13a).

A. Incorporating Integration Error

First, it is important to realize that the output variable \mathbf{y} is now subject to an additional source of uncertainty in \mathbf{g} introduced by the

model. The key idea is to account for all sources of uncertainty in the computed moments, which can be achieved with the following:

$$\mu = \mathbb{E}_{\mathbf{x}}[\mathbf{g}^\dagger(\mathbf{x})] \approx \hat{\mu} = \mathbb{E}_{\mathbf{x}, \mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x})] \quad (36)$$

$$\mathbf{\Pi} = \mathbb{C}_{\mathbf{x}}[\mathbf{g}^\dagger(\mathbf{x}), \mathbf{g}^\dagger(\mathbf{x})] \approx \hat{\mathbf{\Pi}} = \mathbb{C}_{\mathbf{x}, \mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x})] \quad (37)$$

$$\mathbf{C} = \mathbb{C}_{\mathbf{x}}[\mathbf{x}, \mathbf{g}^\dagger(\mathbf{x})] \approx \hat{\mathbf{C}} = \mathbb{C}_{\mathbf{x}, \mathbf{g}|\mathcal{D}}[\mathbf{x}, \mathbf{g}(\mathbf{x})]. \quad (38)$$

Using the law of total expectation and covariance, the approximate moments of the output can be written as

$$\hat{\mu} = \mathbb{E}_{\mathbf{g}|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}_{\mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x})]] \quad (39)$$

$$\hat{\mathbf{\Pi}} = \mathbb{C}_{\mathbf{g}|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]] + \mathbb{E}_{\mathbf{g}|\mathcal{D}}[\mathbb{C}_{\mathbf{x}}[\mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x})]] \quad (40)$$

$$= \mathbb{C}_{\mathbf{x}}[\mathbb{E}_{\mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x})]] + \mathbb{E}_{\mathbf{x}}[\mathbb{C}_{\mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x})]] \quad (41)$$

$$\hat{\mathbf{C}} = \mathbb{E}_{\mathbf{x}}[\mathbf{x}\mathbb{E}_{\mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x})]] - \mathbb{E}_{\mathbf{x}}[\mathbf{x}]\mathbb{E}_{\mathbf{g}|\mathcal{D}, \mathbf{x}}[\mathbf{g}(\mathbf{x})]. \quad (42)$$

The first equality exposes the fact that integral mean is obtained by integrating the mean function of the integrand model. The way the integral uncertainty is incorporated into the output covariance is revealed by (40). Note that since the model of the integrand has conditionally independent outputs, the covariance of the integral, $\mathbb{C}_{\mathbf{g}|\mathcal{D}}[\mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})]]$, and the model covariance, $\mathbb{C}_{\mathbf{g}|\mathcal{D}}[\mathbf{g}(\mathbf{x}), \mathbf{g}(\mathbf{x})]$, are diagonal matrices. When either of the covariances approaches zero, (39)–(42) approach their true values. From now on, we will work with the output covariance in the form (41) because it is easier to analyze and implement.

B. Derivation of Transformed Moments

In the following derivations, explicit conditioning on \mathcal{D} in the expectations is omitted to reduce notational clutter. We also assume that the stochastic decoupling substitution has taken place in the integrals, so that $\tilde{\mathbf{g}}(\xi) = \mathbf{g}(\mathbf{m} + \mathbf{L}\xi)$.

Taking the expression for the mean function of the model in (26) and plugging it into (39), the output mean of the BSQMT becomes

$$\hat{\mu} = \mathbb{E}_{\xi}[\mathbb{E}_{\mathbf{g}}[\tilde{\mathbf{g}}(\xi)]] = \mathbf{Y}^\top \Phi^{-\top} \mathbb{E}_{\xi}[\phi(\xi)] = \mathbf{Y}^\top \mathbf{w} \quad (43)$$

where $\mathbf{w} = \Phi^{-\top} \mathbb{E}_{\xi}[\phi(\xi)]$ are the mean weights. The output covariance becomes

$$\begin{aligned} \hat{\mathbf{\Pi}} &= \mathbb{E}_{\xi}[\mathbb{E}_{\mathbf{g}}[\tilde{\mathbf{g}}(\xi)]\mathbb{E}_{\mathbf{g}}[\tilde{\mathbf{g}}(\xi)]^\top] - \hat{\mu}\hat{\mu}^\top + \mathbb{E}_{\xi}[\mathbb{C}_{\mathbf{g}}[\tilde{\mathbf{g}}(\xi), \tilde{\mathbf{g}}(\xi)]] \\ &= \mathbf{Y}^\top \mathbf{W} \mathbf{Y} - \hat{\mu}\hat{\mu}^\top + \bar{\sigma}^2 \mathbf{I}_E \end{aligned} \quad (44)$$

where the expected model variance is

$$\bar{\sigma}^2 = \mathbb{E}_{\xi}[k(\xi, \xi)] - \text{tr}[\mathbf{D}^\top \Phi^{-\top} + \mathbf{D} \Phi^{-1} - \mathbf{W} \mathbf{K}]. \quad (45)$$

Here, $\mathbf{D} = \mathbb{E}_{\xi}[\mathbf{k}(\xi)\phi(\xi)^\top]$ and the covariance weights are $\mathbf{W} = \Phi^{-\top} \mathbb{E}_{\xi}[\phi(\xi)\phi(\xi)^\top] \Phi^{-1}$. Finally, the covariance between the input and output becomes

$$\begin{aligned} \hat{\mathbf{C}} &= \mathbb{E}_{\xi}[(\mathbf{m} + \mathbf{L}\xi)\mathbb{E}_{\mathbf{g}}[\tilde{\mathbf{g}}(\xi)]] - \mathbb{E}_{\xi}[\mathbf{m} + \mathbf{L}\xi]\mathbb{E}_{\mathbf{g}, \xi}[\tilde{\mathbf{g}}(\xi)] \\ &= \mathbf{L}\mathbb{E}_{\xi}[\xi\phi(\xi)]\Phi^{-1}\mathbf{Y} = \mathbf{L}\mathbf{W}_c \mathbf{Y} \end{aligned} \quad (46)$$

where the cross-covariance weights are $\mathbf{W}_c = \mathbb{E}_{\xi}[\xi\phi(\xi)]\Phi^{-1}$.

It has now become evident that the output moments depend on the expectations of the basis functions. In IV, we have shown that the classical moment transforms can be recovered when the basis functions are multivariate polynomials. When this basis and the RBF kernel (35) are used, the expectations above are available in closed form. The complete algorithm of the Bayes–Sard moment transform is summarized in Algorithm 1.

Theorem 3: The BSQ output covariance $\hat{\mathbf{\Pi}}$ is positive semidefinite.

Algorithm 1: Bayes-Sard Quadrature Moment Transform.

Input: The mean \mathbf{m} and the covariance \mathbf{P} of the input variable \mathbf{x} , the integrand $\mathbf{g}(\mathbf{x})$, the matrix of unit sigma-points Ξ and the kernel parameters θ .

Output: Approximate mean $\hat{\boldsymbol{\mu}}$ and covariance $\hat{\boldsymbol{\Pi}}$ of the output variable $\mathbf{y} = \mathbf{g}(\mathbf{x})$, and approximate input-output covariance $\hat{\mathbf{C}}$.

Function BSQMT ($\mathbf{g}(\mathbf{x})$, \mathbf{m} , \mathbf{P} , Ξ , θ)

```

1 // form sigma-points
2  $\mathbf{L} \leftarrow \text{MatrixFactor}(\mathbf{P})$ 
3  $\mathbf{X} \leftarrow \mathbf{m} + \mathbf{L}\Xi$ 
4  $\bar{k} \leftarrow \mathbb{E}_{\xi}[k(\xi, \xi; \theta)]$ 
5  $\bar{\phi} \leftarrow \mathbb{E}_{\xi}[\phi(\xi)]$ 
6  $\mathbf{A} \leftarrow \mathbb{E}_{\xi}[\phi(\xi)\phi(\xi)^{\top}]$ 
7  $\mathbf{B} \leftarrow \mathbb{E}_{\xi}[\xi\phi(\xi)^{\top}]$ 
8 for  $n \leftarrow 1$  to  $N$  do
9    $[\mathbf{Y}]_{n*} \leftarrow \mathbf{g}(\mathbf{x}^{(n)})$ 
10   $[\Phi]_{n*} \leftarrow [\phi_1(\xi^{(n)}) \dots \phi_N(\xi^{(n)})]$ 
11   $[\mathbf{D}]_{n*} \leftarrow \mathbb{E}_{\xi}[k(\xi, \xi^{(n)}; \theta)\phi(\xi)^{\top}]$ 
12  for  $m \leftarrow 1$  to  $N$  do
13     $[\mathbf{K}]_{nm} \leftarrow k(\xi^{(n)}, \xi^{(m)}; \theta)$ 
14  end
15 end
16 // compute BSQ weights
17  $\mathbf{w} \leftarrow \Phi^{-\top} \bar{\phi}$ 
18  $\mathbf{W} \leftarrow \Phi^{-\top} \mathbf{A} \Phi^{-1}$ 
19  $\mathbf{W}_c \leftarrow \mathbf{B} \Phi^{-1}$ 
20 // compute transformed moments
21 // using eqs. (43) to (46)
22 return  $\hat{\boldsymbol{\mu}}$ ,  $\hat{\boldsymbol{\Pi}}$ ,  $\hat{\mathbf{C}}$ 
end

```

Proof: Using the expression for the BSQ mean weights from Algorithm 1, we can write the output covariance as $\hat{\boldsymbol{\Pi}} = \mathbf{Y}^{\top} \Phi^{-\top} (\mathbf{A} - \bar{\phi} \bar{\phi}^{\top}) \Phi^{-1} \mathbf{Y} + \bar{\sigma}^2 \mathbf{I}_E$. Define $\mathbf{Z} = \Phi^{-1} \mathbf{Y}$ and $\tilde{\mathbf{A}} = \mathbf{A} - \bar{\phi} \bar{\phi}^{\top}$; then, $\boldsymbol{\Pi} = \mathbf{Z}^{\top} \tilde{\mathbf{A}} \mathbf{Z} + \bar{\sigma}^2 \mathbf{I}_E$. We recognize that $\tilde{\mathbf{A}} = \mathbb{V}[\phi(\xi)] = \mathbb{E}[\phi(\xi)\phi(\xi)^{\top}] - \mathbb{E}[\phi(\xi)]\mathbb{E}[\phi(\xi)]^{\top} \succeq 0$, which follows from the properties of covariance matrices. This implies that $\mathbf{Z}^{\top} \tilde{\mathbf{A}} \mathbf{Z} \succeq 0$ for any matrix \mathbf{Z} . Because $\bar{\sigma}^2 \geq 0$, we have that $\hat{\boldsymbol{\Pi}} \succeq 0$.

C. Relationship to the Gaussian Process Quadrature Moment Transform

The recently proposed Gaussian process quadrature moment transform (GPQMT) [14] and the BSQMT are both instances of the general BQ framework. The GPQMT uses a zero-mean GP prior model of the integrand as opposed to the more sophisticated hierarchical prior in (19) and (21). As a result, the GPQMT weights are affected by the choice of kernel and its parameter values, which is not the case in the BSQMT, where the kernel only affects the last term of the transformed covariance and the weights depend only on the sigma points and the choice of the function space π . Consequently, this makes BSQMT much less sensitive to misspecification of the kernel parameters, which is a notorious problem plaguing GPQMT. Discussion of the choice of kernel parameters can be found in the original publication [14].

Compared to the zero-mean GP employed in GPQMT, the Bayes-Sard GP is a stronger prior, which means it can provide better fit to the

Algorithm 2: Bayes-Sard Quadrature Kalman Filter.

Input: Sequence of measurements $\{\mathbf{z}_k\}_{k=1}^K$, initial conditions $\mathbf{m}_{0|0}^x$, $\mathbf{P}_{0|0}^x$, kernel parameters θ_f and θ_h , unit sigma-points Ξ

Output: Sequence of state estimates and covariances $\{\mathbf{m}_{k|k}^x, \mathbf{P}_{k|k}^x\}_{k=1}^K$

for $k \leftarrow 1$ **to** K **do**

// predictive state moments

$\mathbf{m}_{k|k-1}^x, \mathbf{P}_{k|k-1}^x \leftarrow$

BSQMT ($\mathbf{f}(\mathbf{x}_{k-1})$, $\mathbf{m}_{k-1|k-1}^x$, $\mathbf{P}_{k-1|k-1}^x$, Ξ , θ_f)

$\mathbf{P}_{k|k-1}^x \leftarrow \mathbf{P}_{k|k-1}^x + \mathbf{Q}$

// predictive measurement moments

$\mathbf{m}_{k|k-1}^z, \mathbf{P}_{k|k-1}^z, \mathbf{P}_{k|k-1}^{xz} \leftarrow$

BSQMT ($\mathbf{h}(\mathbf{x}_k)$, $\mathbf{m}_{k|k-1}^x$, $\mathbf{P}_{k|k-1}^x$, Ξ , θ_h)

$\mathbf{P}_{k|k-1}^z \leftarrow \mathbf{P}_{k|k-1}^z + \mathbf{R}$

// measurement update (filtering)

// using eqs. (4) and (5)

end

integrand when conditioned on smaller datasets, such as the UT sigma points, which are especially attractive in nonlinear filtering applications.

D. BSQ Moment Transform in Sigma-Point Filtering

As outlined in II, the filtering algorithms use the moment transformations for computing the predictive moments of the system state and measurement. Algorithm 2 summarizes the Bayes-Sard quadrature Kalman filter (BSQKF), which employs the proposed BSQ moment transform for computing the predictive moments from I. The BSQKF takes two different kernel parameter values, θ_f and θ_h , because there are two different functions that need to be integrated [see (1) and (2)].

VI. NUMERICAL EXPERIMENT

In order to test the sigma-point filters based on the BSQ moment transform, we consider the univariate nonstationary growth model (UNGM), which is often used to benchmark particle filters [35]. The system dynamics and the observation model are given by

$$x_k = \frac{1}{2}x_{k-1} + \frac{25x_{k-1}}{1+x_{k-1}^2} + 8\cos(1.2k) + q_{k-1} \quad (47)$$

$$z_k = \frac{1}{20}x_{k-1}^2 + r_k \quad (48)$$

with the state noise $q_{k-1} \sim \mathcal{N}(0, 10)$, measurement noise $r_k \sim \mathcal{N}(0, 1)$, and initial conditions $x_0 = x_{0|0} \sim \mathcal{N}(0, 5)$. Kernel scaling used in the BSQ with for the UT ($\kappa = 2$) and the seventh-order GH points was set to $\alpha = 3$ and the lengthscales to $\ell = 0.3$ and $\ell = 0.4$, respectively. For the fifth-order GH points, the kernel parameters were set to $\alpha = 5$ and $\ell = 0.6$.

The root-mean-square error (RMSE) was used to measure the tracking performance. The inclination indication (INC) [36], given by

$$\text{INC} = \frac{10}{K} \sum_{k=1}^K \log_{10} \frac{(\mathbf{x}_k - \mathbf{m}_{k|k}^x)^{\top} (\mathbf{P}_{k|k}^x)^{-1} (\mathbf{x}_k - \mathbf{m}_{k|k}^x)}{(\mathbf{x}_k - \mathbf{m}_{k|k}^x)^{\top} \boldsymbol{\Sigma}_k^{-1} (\mathbf{x}_k - \mathbf{m}_{k|k}^x)} \quad (49)$$

where $\boldsymbol{\Sigma}_k$ is the mean-squared error matrix of the state, was used to measure the credibility of the estimates. A perfectly balanced estimate has $\text{INC} = 0$. For $\text{INC} > 0$, the estimate is said to be optimistic, which

TABLE II
FILTER RMSE FOR THE UNGM EXAMPLE

	Classical	GPQ	BSQ
UT	10.81 (0.14)	10.37 (0.08)	9.70 (0.12)
GH-5	10.03 (0.14)	9.01 (0.11)	8.82 (0.08)
GH-7	9.74 (0.13)	8.80 (0.10)	8.61 (0.09)

TABLE III
FILTER INC FOR THE UNGM EXAMPLE

	Classical	GPQ	BSQ
UT	12.17 (0.06)	4.87 (0.01)	4.57 (0.03)
GH-5	10.33 (0.07)	5.26 (0.03)	1.85 (0.02)
GH-7	9.27 (0.07)	4.95 (0.03)	2.52 (0.03)

is to say the covariance is smaller than it should be, while negative values indicate pessimism. We refer to [36] for other related credibility measures. We simulated the model for $K = 500$ time steps and averaged the performance scores over 100 simulations. The variance of the average scores was estimated by bootstrapping. The parentheses in contain the uncertainty as the least significant digits of two standard deviations.

The BSQ filters with classical points were tested against the well-known sigma-point filters as well as the GPQ filters from [14]. As seen in II, the filters based on the BSQ outperform the classical sigma-point filters in terms of RMSE. Assuming that the GH points are used, BSQKF's can outperform the GPQ filters as well. In comparison with the classical filters, the proposed BSQ filters also provide much more balanced estimates as evidenced by the values of the INC in III.

VII. CONCLUSION AND DISCUSSION

In this article, we designed a general-purpose moment transformation based on BSQ, which allows for explicit modeling of numerical integration error through the use of a stochastic process model. The hierarchical GP prior was shown to be key in developing probabilistic models, which lead to the classical quadrature rules used in the sigma-point filters and whose variance is statistically meaningful. We designed the BSQ Kalman filter by leveraging the proposed BSQ moment transform for computation of the predictive moments. Overall, the BSQ-based filters report more balanced estimates and tend to err on the side of caution (the reported estimates are more likely to be pessimistic).

REFERENCES

- [1] M. S. Grewal, L. R. Weill, and A. P. Andrews, *Global Positioning Systems, Inertial Navigation, and Integration*. Hoboken, NJ, USA: Wiley, 2007.
- [2] A. J. Haug, *Bayesian Estimation and Tracking: A Practical Guide*. Hoboken, NJ, USA: Wiley-Blackwell, 2012.
- [3] L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell, *Bayesian Multiple Target Tracking*. Norwood, MA, USA: Artech House, 2014.
- [4] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part I," *IEEE Robot. Autom. Mag.*, vol. 13, no. 2, pp. 99–110, Jun. 2006.
- [5] S. Gillijns, O. Mendoza, J. Chandrasekar, B. De Moor, D. Bernstein, and A. Ridley, "What is the ensemble Kalman filter and how well does it work?" in *Proc. Amer. Control Conf.*, Jun. 2006, p. 6.
- [6] R. Bhar, *Stochastic Filtering with Applications in Finance*. Singapore: World Scientific, 2010.
- [7] S. Holmes, G. Klein, and D. W. Murray, "A square root unscented Kalman filter for visual monoSLAM," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2008, pp. 3710–3716.
- [8] K. Lowrey, J. Dao, and E. Todorov, "Real-time state estimation with whole-body multi-contact dynamics: A modified UKF approach," in *Proc. IEEE-RAS 16th Int. Conf. Humanoid Robots*, Nov. 2016, pp. 1225–1232.
- [9] T. Karvonen and S. Särkkä, "Classical quadrature rules via Gaussian processes," in *Proc. IEEE 27th Int. Workshop Mach. Learn. Signal Process.*, 2017, pp. 1–6.
- [10] C. J. Oates, S. Niederer, A. Lee, F.-X. Briol, and M. Girolami, "Probabilistic models for integration error in the assessment of functional cardiac models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 110–118.
- [11] F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic, "Probabilistic integration: A role in statistical computation? (with discussion and rejoinder)," *Statist. Sci.*, vol. 34, no. 1, pp. 1–22, 2019, doi: 10.1214/18-STS660.
- [12] T. Karvonen, S. Särkkä, and C. Oates, "Symmetry exploits for Bayesian cubature methods," *Statist. Comput.*, vol. 29, pp. 1231–1248, 2019.
- [13] S. Särkkä, J. Hartikainen, L. Svensson, and F. Sandblom, "On the relation between Gaussian process quadratures and sigma-point methods," *J. Adv. Inf. Fusion*, vol. 11, pp. 31–46, Jun. 2016.
- [14] J. Prüher and O. Straka, "Gaussian process quadrature moment transform," *IEEE Trans. Autom. Control*, vol. 63, no. 9, pp. 2844–2854, Sep. 2018.
- [15] T. Karvonen, C. J. Oates, and S. Särkkä, "A Bayes-Sard cubature method," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 5882–5893.
- [16] K. Xiong, H. Zhang, and C. W. Chan, "Performance evaluation of UKF-based nonlinear filtering," *Automatica*, vol. 42, no. 2, pp. 261–270, 2006.
- [17] K. Xiong, L. D. Liu, and H. Y. Zhang, "Modified unscented Kalman filtering and its application in autonomous satellite navigation," *Aerosp. Sci. Technol.*, vol. 13, no. 4, pp. 238–246, 2009.
- [18] T. Karvonen, S. Bonnabel, E. Moulines, and S. Särkkä, "On stability of a class of filters for non-linear stochastic systems," Sep. 2018, arXiv:1809.05667v1.
- [19] A. A. Clifford, *Multivariate Error Analysis: A Handbook of Error Propagation and Calculation in Many-Parameter Systems*. New York, NY, USA: Wiley, 1973.
- [20] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Trans. Autom. Control*, vol. 54, no. 6, pp. 1254–1269, Jun. 2009.
- [21] I. Arasaratnam, S. Haykin, and R. Elliott, "Discrete-time nonlinear filtering algorithms using Gauss-Hermite quadrature," *Proc. IEEE*, vol. 95, no. 5, pp. 953–977, May 2007.
- [22] Y. Wu, D. Hu, M. Wu, and X. Hu, "A numerical-integration perspective on Gaussian filters," *IEEE Trans. Signal Process.*, vol. 54, no. 8, pp. 2910–2921, Aug. 2006.
- [23] K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE Trans. Autom. Control*, vol. 45, no. 5, pp. 910–927, May 2000.
- [24] G. H. Golub and J. H. Welsch, "Calculation of Gauss quadrature rules," *Math. Comput.*, vol. 23, no. 106, pp. 221–221, May 1969.
- [25] B. Jia, M. Xin, and Y. Cheng, "Sparse-grid quadrature nonlinear filtering," *Automatica*, vol. 48, no. 2, pp. 327–341, 2012.
- [26] F. Larkin, "Gaussian measure in Hilbert space and applications in numerical analysis," *Rocky Mountain J. Math.*, vol. 2, pp. 379–421, 1972.
- [27] A. O'Hagan, "Bayes-Hermite quadrature," *J. Statist. Planning Inference*, vol. 29, no. 3, pp. 245–260, 1991.
- [28] C. E. Rasmussen and Z. Ghahramani, "Bayesian monte carlo," in *Proc. 15th Int. Conf. Neural Inf. Process.*, 2003, pp. 505–512.
- [29] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.
- [30] T. J. Santner, B. J. Williams, and W. I. Notz, *The Design and Analysis of Computer Experiments* (ser. Springer Series in Statistics). New York, NY, USA: Springer, 2003.
- [31] M. Álvarez, L. Rosasco, and N. Lawrence, "Kernels for vector-valued functions: A review," *Found. Trends Mach. Learn.*, vol. 4, no. 3, pp. 195–266, 2012.
- [32] X. Xi, F.-X. Briol, and M. Girolami, "Bayesian quadrature for multiple related integrals," in *Proc. 35th Int. Conf. Mach. Learn.*, vol. 80, Jul. 2018, pp. 5373–5382.
- [33] F. M. Larkin, "Probabilistic error estimates in spline interpolation and quadrature," in *Information Processing 74 (Proceedings of IFIP Congress, Stockholm, 1974)*, vol. 74. Amsterdam, The Netherlands: North Holland, 1974, pp. 605–609.
- [34] M. Kanagawa, B. K. Sriperumbudur, and K. Fukumizu, "Convergence guarantees for kernel-based quadrature rules in misspecified settings," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 3288–3296.
- [35] G. Kitagawa, "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models," *J. Comput. Graph. Statist.*, vol. 5, no. 1, pp. 1–25, 1996.
- [36] X. R. Li and Z. Zhao, "Measuring estimator's credibility: Noncredibility index," in *Proc. 9th Int. Conf. Inf. Fusion.*, 2006, pp. 1–8.