This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Hyvönen, Nuutti; Mustonen, Lauri

Thermal tomography with unknown boundary

Published in:
SIAM Journal on Scientific Computing

DOI:
[10.1137/16M1104573](https://doi.org/10.1137/16M1104573)

Published: 01/01/2018

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Hyvönen, N., & Mustonen, L. (2018). Thermal tomography with unknown boundary. *SIAM Journal on Scientific Computing*, 40(3), B662-B683. <https://doi.org/10.1137/16M1104573>

THERMAL TOMOGRAPHY WITH UNKNOWN BOUNDARY*

NUUTTI HYVÖNEN[†] AND LAURI MUSTONEN[‡]

Abstract. Thermal tomography is an imaging technique for deducing information about the internal structure of a physical body from temperature measurements on its boundary. This work considers time-dependent thermal tomography modeled by a parabolic initial/boundary value problem without accurate information on the exterior shape of the examined object. The adaptive sparse pseudospectral approximation method is used to form a polynomial surrogate for the dependence of the temperature measurements on the thermal conductivity, the heat capacity, the boundary heat transfer coefficient, and the body shape. These quantities can then be efficiently reconstructed via nonlinear, regularized least squares minimization employing the surrogate and its derivatives. The functionality of the resulting reconstruction algorithm is demonstrated by numerical experiments based on simulated data in two spatial dimensions.

Key words. thermal tomography, inverse boundary value problems, sparse pseudospectral approximation, inaccurate measurement model

AMS subject classifications. 35R30, 41A10, 65M32

DOI. 10.1137/16M1104573

1. Introduction. The aim of thermal tomography is to gather information about the interior properties of a physical object based on temperature measurements on its boundary. The potential applications include nondestructive testing of materials; see, e.g., [1, 19] for further introduction to the topic as well as for a list of previous works. The model used in this article is time dependent: the boundary temperatures are measured several times while heating the object. Thermal conductivity and heat capacity are both considered as unknown quantities as in [14]. In addition, the boundary heat transfer coefficient between the heater elements is reconstructed. Such a setting was considered with simulated data in [19] and subsequently with experimental data in [20]. As a novelty of this work, the exterior shape of the imaged object is also reconstructed.

In this article, we approximate the dependence of the temperature measurements on the four unknown quantities by multivariate polynomials. The unknown functions (i.e., the thermal conductivity, heat capacity, boundary heat transfer coefficient, and boundary shape) are first expressed in terms of a finite number of parameters. The polynomial forward surrogate, which takes these parameters as input, is then constructed by the (adaptive) *sparse pseudospectral approximation method* (SPAM) [4, 3, 21, 22] that requires the solutions to a large number of parabolic initial/boundary value problems describing the heat flow in the examined object. When a measurement vector is given, the inverse problem of determining the unknowns is treated as a nonlinear, regularized least squares minimization involving the forward surrogate.

*Submitted to the journal's Computational Methods in Science and Engineering section November 21, 2016; accepted for publication (in revised form) December 28, 2017; published electronically May 1, 2018.

<http://www.siam.org/journals/sisc/40-3/M110457.html>

Funding: This work was supported by the Academy of Finland (decision 267789) and the Finnish Foundation for Technology Promotion TES.

[†]Aalto University, Department of Mathematics and Systems Analysis, P.O. Box 11100, FI-00076 Aalto, Finland (nuutti.hyvonen@aalto.fi, <https://users.aalto.fi/nhyvonen/>).

[‡]Department of Mathematics and Computer Science, Emory University, 400 Dowman Dr., Atlanta, GA 30322 (lauri.mustonen@emory.edu).

The advantage of this approach is that only polynomial evaluation and differentiation are needed when reconstructing the quantities of interest from the measurement data, which leads to a fast reconstruction process. More precisely, under the assumption that enough generic information about the measurement setup is available, the polynomial forward surrogate can be formed “offline,” that is, before the actual measurements are in hand. Thus, the efficiency of the “online” phase, or the least squares reconstruction algorithm, does not depend on how the parabolic forward problems were solved in the offline phase. For a related polynomial approximation approach to thermal tomography assuming the shape of the imaged body is known and without modeling the effects of thermal capacitance or the geometry of the heaters and sensors, see [16], where a (nonadaptive) Galerkin technique is used for forming the needed surrogates. To the best of our knowledge, apart from [16], there are no previous works considering polynomial surrogates as a computational tool in the framework of inverse parabolic boundary value problems.

In practical applications, the exterior shape of the imaged body is usually not the object of main interest, and the same definitely applies to the boundary heat transfer coefficient. Typically, the interesting quantities are the thermal conductivity and the heat capacity, or perhaps only one of those, e.g., if there exists a known functional dependence between the two of them. However, we demonstrate that ignoring the uncertainties in the object shape and the boundary heat transfer coefficient deteriorates the reconstructions of the thermal conductivity and the heat capacity and therefore we suggest that all four unknowns should be reconstructed simultaneously if possible. Although the main reason for only considering two-dimensional examples is their computational simplicity, there exists also some practical relevance since a two-dimensional model can be utilized if the object is cylindrically symmetric with suitable boundary conditions at the top and bottom boundaries.

A related shape estimation problem has previously been studied in the context of *electrical impedance tomography* (EIT) in [5, 6, 11, 17]. The fundamental difference between EIT and thermal tomography is that the former is modeled by an elliptic partial differential equation with one unknown physical field (conductivity), whereas thermal tomography corresponds to an (inverse) parabolic boundary value problem with two coefficients defining the properties of the imaged body (thermal conductivity and heat capacity). The algorithm of [17] is based on the so-called approximation error approach [13]: the error caused by mismodeling the measurement setup is treated as an additional noise process whose statistics are approximated via simulations based on prior probability models for the unknowns. The papers [5, 6] introduce Fréchet derivatives of the forward map of EIT with respect to the measurement geometry, which allows one to estimate the body shape as a part of a Bayesian output least squares algorithm, although the evaluation of the needed shape derivatives suffers from slight instability that originates from certain distributional boundary conditions in the elliptic boundary value problems defining the derivatives. Our approach is closely related to [11] where a polynomial surrogate is (nonadaptively) built for the forward operator of EIT. However, the efficiency gain induced by transferring computational burden to an offline phase is higher in thermal tomography than in EIT due to the possibility of completely avoiding time integration in the online phase when the reconstruction is formed by minimizing a least squares functional.

We denote integer vectors, or *multi-indices*, by bold symbols such as $\mathbf{k} := (k_1, \dots, k_N)$. In particular, the n th Euclidean basis vector is denoted by $\mathbf{e}_n$, having 1 as the n th element and 0 in the remaining $N - 1$ positions, where N is assumed to be clear from the context. The calligraphic font, e.g., $\mathcal{K}$, is used for sets of multi-

indices. Linear operators are denoted by scripted letters such as $\mathcal{Q}$ or $\mathcal{P}$. Symbols $\mathbf{R} = (-\infty, \infty)$ and $\mathbf{N} = \{1, 2, 3, \dots\}$ as well as $\mathbf{N}_0 = \mathbf{N} \cup \{0\}$ have their usual meanings. Furthermore, $\mathbf{P}_k$ denotes the space of real-valued polynomials in one real variable, having a degree $k \in \mathbf{N}_0$ or less.

This text is organized as follows. In section 2, the governing equations for heat flow are introduced and the dependence of the measurements on the finite-dimensional parameter vector is formulated. The pseudospectral approximation method and its adaptive version are described in section 3. Section 4 presents the numerical experiments that address both the accuracy of the forward surrogate and the feasibility of thermal tomography with an unknown exterior boundary. Finally, conclusions are drawn in section 5.

2. Thermal tomography. We start this section by mathematically formulating our forward model for thermal tomography following [19, 20]. Subsequently, a parametric extension is introduced for a certain two-dimensional measurement configuration.

2.1. Measurement model. Let $\Omega \subset \mathbf{R}^d$, where $d \in \{2, 3\}$, denote a bounded domain with a C^1 -boundary $\partial\Omega$ and the exterior unit normal $\hat{\mathbf{n}}$. Furthermore, let $I = (0, T)$ denote a time interval with $T > 0$. The initial/boundary value problem considered as a model for the temperature u in thermal tomography is

$$(1) \quad \begin{cases} b \partial_t u - \nabla \cdot (a \nabla u) = 0 & \text{in } \Omega \times I, \\ a \nabla u \cdot \hat{\mathbf{n}} = c(f - u) & \text{on } \partial\Omega \times I, \\ u = 0 & \text{in } \Omega \times \{0\}, \end{cases}$$

where $a, b \in L_+^\infty(\Omega)$ are the thermal conductivity and the heat capacity, respectively, and $c \in L_+^\infty(\partial\Omega)$ is the surface conductance or heat transfer coefficient. The time-dependent boundary source $f \in L^2(I; L^2(\partial\Omega))$ represents the external or outside temperature. All aforementioned function spaces are assumed to have $\mathbf{R}$ as the corresponding multiplier field and

$$L_+^\infty(D) := \{v \in L^\infty(D) : \text{ess inf } v > 0\}$$

for $D = \Omega$ or $\partial\Omega$. The weak formulation of (1) is to find $u: \Omega \times I \rightarrow \mathbf{R}$ such that

$$(2) \quad \partial_t(bu, v)_\Omega + (a \nabla u, \nabla v)_\Omega + (cu, v)_{\partial\Omega} = (cf, v)_{\partial\Omega} \quad \text{for all } v \in H^1(\Omega),$$

where $(\cdot, \cdot)_D$ denotes the (real) L^2 inner product on D and the time derivative as well as the initial condition from (1) are to be understood in the appropriate distributional sense (cf., e.g., [7]). Under the above assumptions, the variational problem (2) has a unique solution $u \in V$, where

$$V := L^2(I; H^1(\Omega)) \cap C(\bar{I}; L^2(\Omega));$$

see, e.g., [7].

We assume that the boundary is equipped with $J \in \mathbf{N}$ nonoverlapping heater elements $\{H_j\}_{j=1}^J$ that are identified with the open, nonempty, and connected parts of $\partial\Omega$ that they cover. In particular, it is natural to model the boundary heat source as

$$f = \sum_{j=1}^J g_j \chi_j,$$

where $g_j \in L^2(I)$ and $\chi_j: \partial\Omega \rightarrow \{0, 1\}$ is the characteristic function of H_j . Moreover, since our model for thermal tomography involves heating only one heater at a time and letting the target object cool down before activating the next heater, we actually consider solving (1) for a relatively short period of time but for J different functions

$$(3) \quad f = f_j := g \chi_j, \quad j = 1, \dots, J,$$

where $g \in L^2(I)$ is the chosen common time modulation for the heating patterns. Usually, the heat transfer coefficient c is large on those parts of the boundary that are covered by heaters and much smaller (but still positive) on the rest of the boundary (cf. [20]). Note that even though f vanishes outside the heaters, the heat still leaks to the surrounding air due to the nonvanishing Robin coefficient c .

It is assumed that the boundary temperature of Ω is measured at R pointlike sensors $\{s_r\}_{r=1}^R \subset \partial\Omega$ that are well separated from the heaters. We denote by $U \in \mathbf{R}^M$ a measurement vector that consists of the temperature values at the sensor locations for each measurement time $t_i \in \bar{I}$, $i = 1, \dots, M_T$. (To make such measurements well defined without further regularity assumptions for (1), a single measurement should be interpreted as the mean value of $u|_{\partial\Omega \times I} \in L^2(I; H^{1/2}(\partial\Omega)) \subset L^2(I; L^2(\partial\Omega)) \simeq L^2(\partial\Omega \times I)$ over a small neighborhood of the corresponding point (s_r, t_i) on $\partial\Omega \times \bar{I}$.) Assuming that for each active heater, i.e., for every f_j , all R sensor values are recorded at all measurement times, the total number of measurements becomes $M = JRM_T$.

After a *finite element* (FE) discretization, (2) results in a system of ordinary differential equations

$$(4) \quad \partial_t B \hat{u} + (A + C) \hat{u} = \hat{f},$$

where A , B , and C are symmetric matrices corresponding to the three integrals on the left-hand side of (2). Moreover, $\hat{u}$ and $\hat{f}$ are vectors carrying the coefficients of u and cf in the chosen FE basis. In practice, $\hat{f}$ and $\hat{u}$ can also be matrices whose J columns correspond to heating different heaters. The semidiscrete system (4) is solved in time by using the second-order Crank–Nicolson time integration method [15]. The measurement vector U is computed by evaluating the numerical solution at the specified measurement times and locations.

2.2. Parametric extension. In order to build a surrogate model that describes the dependence of the measurements on the fields a and b as well as on the coefficient c and the boundary shape $\partial\Omega$, we parametrize these quantities with a finite number of parameters. More precisely, after denoting by $N = N_a + N_b + N_c + N_\Omega$ the total number of parameters, we consider parameter vectors $\vartheta = (\vartheta^{(a)}, \vartheta^{(b)}, \vartheta^{(c)}, \vartheta^{(\Omega)}) \in \Xi^N$, where $\Xi = [-1/2, 1/2]$ and Ξ^N is the parameter hypercube with unit volume. The idea is that each fixed parameter vector defines one well-posed initial/boundary value problem of the form (1). That is, for each parameter vector there exists a unique solution u and a unique measurement vector U so that these quantities can be interpreted as mappings $u: \Xi^N \rightarrow V$ and $U: \Xi^N \rightarrow \mathbf{R}^M$, respectively. In the following, we only consider a specific two-dimensional parametrization, but emphasize that the extension to more complicated geometries or to three dimensions is conceptually straightforward.

Let us begin by considering a suitable representation for the shape of a domain $\Omega \subset \mathbf{R}^2$ as a perturbation of a disk. To this end, the boundary curve is parametrized in polar coordinates by two variables, the polar angle ϕ and the shape parameter $\vartheta^{(\Omega)} \in \Xi^{N_\Omega}$, that is,

$$(5) \quad \partial\Omega(\vartheta^{(\Omega)}) = \{(r(\phi; \vartheta^{(\Omega)}), \phi) : \phi \in [0, 2\pi)\}.$$

The radial coordinate is chosen to be of the form

$$r(\phi; \vartheta^{(\Omega)}) = \rho_0 + \sum_{i=1}^{N_\Omega} (\rho_+ - \rho_-) \vartheta_i^{(\Omega)} \psi_i(\phi),$$

where $\rho_0 = (\rho_+ + \rho_-)/2$ and $\rho_+ > \rho_-$ are the maximal and minimal radii of the parametrized domain, respectively. The local perturbations $\psi_i \in C^1([0, 2\pi))$ are 2π -periodic, nonnegative, and uniform quadratic B-splines that form a partition of unity on the unit circle. We also define a homeomorphism $\Phi(\cdot; \vartheta^{(\Omega)}): \overline{\Omega(\vartheta^{(\Omega)})} \rightarrow \overline{\Omega(0)}$ from a given perturbed domain to the unperturbed reference disk of radius ρ_0 via

$$\Phi((r', \phi); \vartheta^{(\Omega)}) = \left(\frac{\rho_0}{r(\phi, \vartheta^{(\Omega)})} r', \phi \right)$$

for $(r', \phi) \in \overline{\Omega(\vartheta^{(\Omega)})}$. It is assumed that the widths and starting polar angles (in the counterclockwise direction) of the heaters are fixed (i.e., known) and that the same applies to the polar angles of the pointlike heat sensors.¹ In other words, given the starting points $\{y_j(0)\}_{j=1}^J$ and the common width η of the heater patches $\{H_j(0)\}_{j=1}^J$ as well as the sensor locations $\{s_r(0)\}_{r=1}^R$ on the boundary of the reference disk $\Omega(0)$, the dependence of these geometric entities on the shape parameter $\vartheta^{(\Omega)}$ can be written as

$$H_j(\vartheta^{(\Omega)}) = \left\{ x \in \partial\Omega(\vartheta^{(\Omega)}) : 0 < \text{dist}\left(\Phi^{-1}(y_j(0); \vartheta^{(\Omega)}), x\right) < \eta \right\},$$

$$s_r(\vartheta^{(\Omega)}) = \Phi^{-1}(s_r(0); \vartheta^{(\Omega)}),$$

where $\text{dist}(y, x) = \text{dist}(y, x; \vartheta^{(\Omega)})$ denotes the distance between the points $y, x \in \partial\Omega(\vartheta^{(\Omega)})$ along the boundary in the counterclockwise direction. In particular, the boundary heat sources are still modeled by (3) under the interpretation that $\chi_j = \chi_j(\cdot; \vartheta^{(\Omega)})$ is the characteristic function of $H_j(\vartheta^{(\Omega)})$ for some $\vartheta^{(\Omega)} \in \Xi^{N_\Omega}$ that is clear from the context.

The parameter-dependent function-valued mappings $a, b: \Xi^N \rightarrow L_+^\infty(\Omega(\vartheta^{(\Omega)}))$ are defined through

$$(6) \quad \begin{aligned} a(x; \vartheta) &= \bar{a}(\Phi(x; \vartheta^{(\Omega)})) + 2\tilde{a} \sum_{i=1}^{N_a} \vartheta_i^{(a)} \varphi_i^{(a)}(\Phi(x; \vartheta^{(\Omega)})), \\ b(x; \vartheta) &= \bar{b}(\Phi(x; \vartheta^{(\Omega)})) + 2\tilde{b} \sum_{i=1}^{N_b} \vartheta_i^{(b)} \varphi_i^{(b)}(\Phi(x; \vartheta^{(\Omega)})), \end{aligned}$$

where $\bar{a}, \bar{b} \in L_+^\infty(\Omega(0))$ are the “default fields” and the bases $\{\varphi_i^{(a)}\}_{i=1}^{N_a}, \{\varphi_i^{(b)}\}_{i=1}^{N_b} \in L_+^\infty(\Omega(0))$ are assumed to form *partitions of unity* on $\Omega(0)$. The free parameters $\tilde{a}, \tilde{b} \in \mathbf{R}_+$ are to be tuned so that $a(\cdot; \vartheta)$ and $b(\cdot; \vartheta)$ really are elements of $L_+^\infty(\Omega(\vartheta^{(\Omega)}))$ for all $\vartheta \in \Xi^N$, that is, by construction they should satisfy

$$\tilde{a} < \text{ess inf } \bar{a} \quad \text{and} \quad \tilde{b} < \text{ess inf } \bar{b}.$$

¹It is natural to assume that the widths of the heaters are known, but the same does not necessarily apply to the (starting) polar angles of the heaters and the sensors. However, the precise nature of the available information on the positioning of the heaters and sensors is application specific, and thus we have decided to work with this relatively general and simple assumption.

The bases $\{\varphi_i^{(a)}\}_{i=1}^{N_a}$ and $\{\varphi_i^{(b)}\}_{i=1}^{N_b}$ could be formed using, e.g., B-splines or a piecewise linear FE basis, but in the numerical experiments of section 4 we simply resort to characteristic functions corresponding to pixelwise discretizations of the reference disk $\Omega(0)$.

The surface conductance could be parametrized in the same manner as a and b . However, as the (constant) value of the conductance underneath the heaters may be assumed known (cf. [20]) and its value between the heaters is not expected to vary considerably, we employ a simpler parametrization for $c: \Xi^N \rightarrow L_+^\infty(\partial\Omega(\vartheta^{(\Omega)}))$, that is,

$$(7) \quad c(x; \vartheta) = \begin{cases} c_0 & \text{if } x \in \bigcup_{j=1}^J H_j(\vartheta^{(\Omega)}), \\ \bar{c} + 2\tilde{c} \sum_{j=1}^J \vartheta_j^{(c)} \tilde{\chi}_j & \text{otherwise.} \end{cases}$$

Here $c_0 \in \mathbf{R}_+$ is the known (high) value of the surface conductance at the heaters, $\bar{c}$ is its default (low) value between the heaters, and $\{\tilde{\chi}_j\}_{j=1}^J$ are the characteristic functions of the J curve segments between the heaters. In particular, take note that $N_c = J$. To guarantee the positivity of c , it is assumed that $\bar{c} > \tilde{c}$.

We collect the assumptions on our parametrization in the following lemma.

LEMMA 2.1. *If the object boundary and the coefficients in (1) are parametrized according to (5), (6), and (7) under the associated assumptions, then (2) has a unique solution $u(\cdot; \vartheta) \in V$ for all $\vartheta \in \Xi^N$. In other words, the mapping $\Xi^N \ni \vartheta \mapsto u(\cdot; \vartheta) \in V$ is well defined.*

Proof. It is easy to see that for any fixed $\vartheta \in \Xi^N$, (5), (6), and (7) define a C^1 -domain $\Omega(\vartheta^{(\Omega)})$ as well as coefficient functions $a(\cdot; \vartheta), b(\cdot; \vartheta) \in L_+^\infty(\Omega(\vartheta^{(\Omega)}))$ and $c(\cdot; \vartheta) \in L_+^\infty(\partial\Omega(\vartheta^{(\Omega)}))$. The assertion thus follows from the standard theory on parabolic initial/boundary value problems; see, e.g., [7]. $\square$

The parametrizations in (6) and (7) could easily be made more general; see, e.g., [11], where a logarithmic parametrization for coefficient functions is used. Moreover, dropping the requirement on the bases for the partition of the unity property, the most common parametrization for a (random) coefficient in a partial differential equation is arguably the (exponential) Karhunen–Loève expansion (cf., e.g., [18]). In any case, it is imperative to choose the discretizations so that for each $\vartheta \in \Xi^N$ the functions a , b , and c stay positive.

3. Pseudospectral approximation. Let us assume that an accurate numerical forward solver for the parametrized measurement $U: \Xi^N \rightarrow \mathbf{R}^M$ is available. In other words, for any given vector $\vartheta \in \Xi^N = [-1/2, 1/2]^N$, we are able to compute the boundary measurements $\{U_m(\vartheta)\}_{m=1}^M$ according to the parametrized measurement setup and equations presented in section 2. The inaccuracies of the forward solver are neglected for the moment and, therefore, the same notation U is used for both the computed forward solution and the exact temperatures. In this section, our aim is to construct a polynomial surrogate $U^{(\mathcal{K})}: \Xi^N \rightarrow \mathbf{R}^M$ that satisfies $U^{(\mathcal{K})} \approx U$ in some appropriate sense. The approximation parameter $\mathcal{K} \subset \mathbf{N}_0^N$ is a multi-index set that is related to projection orders and polynomial degrees as explained shortly. The underlying motivation for using a polynomial surrogate is the possibility of faster evaluation of the forward solution and its derivatives when solving the inverse problem of thermal tomography by an iterative scheme. Moreover, the employment of a parametric model simplifies the handling of the unknown boundary shape (cf. [5, 6, 11]).

The construction of the surrogate is based on the SPAM introduced in [4] and, in particular, its adaptive version analyzed in [3]. We briefly outline the algorithm in what follows and refer the interested readers to, e.g., [9, 21, 22] for more details. For simplicity of exposition, in this section, a common approximation parameter $\mathcal{K}$ and the same set of basis polynomials are assumed for every measurement component. Alternatively, the surrogate could be constructed separately for each component of U . The latter is equivalent to considering M real-valued mappings U_m , each in turn, instead of one vector-valued mapping U . Thus, the case of componentwise approximations can be treated as a special case of the procedure presented here.

3.1. Sparse pseudospectral approximation method. The polynomial surrogate $U^{(\mathcal{K})}$ is written as linear combinations of orthonormal basis polynomials. More precisely,

$$(8) \quad U_m^{(\mathcal{K})}(\vartheta) = \sum_{\mathbf{i} \in \mathcal{P}(\mathcal{K})} \alpha_{m,\mathbf{i}} L_{\mathbf{i}}(\vartheta), \quad 1 \leq m \leq M, \quad \vartheta \in \Xi^N,$$

where $\alpha_{m,\mathbf{i}} \in \mathbf{R}$ are the expansion coefficients and $L_{\mathbf{i}}$ are suitably transformed multivariate Legendre polynomials with degrees defined by the multi-index set $\mathcal{P}(\mathcal{K}) \subset \mathbf{N}_0^N$ that depends on the approximation parameter via the chosen quadrature rule. The basis polynomials satisfy

$$\int_{\Xi^N} L_{\mathbf{i}}(\vartheta) L_{\mathbf{j}}(\vartheta) \, d\vartheta = \begin{cases} 1 & \text{if } i_n = j_n, \, 1 \leq n \leq N, \\ 0 & \text{otherwise} \end{cases}$$

for all $\mathbf{i}, \mathbf{j} \in \mathcal{P}(\mathcal{K})$ and they can be naturally constructed as products of univariate polynomials, i.e.,

$$L_{\mathbf{i}}(\vartheta) = \prod_{n=1}^N \ell_{i_n}(\vartheta_n),$$

where ℓ_i denotes the i th order orthonormal(ized) Legendre polynomial on $\Xi = [-1/2, 1/2]$. We refer to [8] for more details about orthogonal polynomials. Hereafter, the subscript in ϑ_n denotes the dimension (i.e., ϑ_n is the n th element of the vector ϑ), whereas superscripts are used to distinguish between different multidimensional nodes, i.e., vectors in Ξ^N .

The sparse pseudospectral approximation is built on sequential *full tensor projections*. Let us first discuss univariate quadratures and projections. A general quadrature rule of order $k \geq 0$ is a linear operator $\mathcal{Q}_k: C(\Xi) \rightarrow \mathbf{R}$, defined as²

$$(9) \quad \mathcal{Q}_k(F) := \sum_{i=0}^{q(k)} F(\vartheta^{(k,i)}) w^{(k,i)} \approx \int_{\Xi} F(\vartheta) \, d\vartheta,$$

where $q(k) + 1$ is the number of quadrature nodes and $\{(\vartheta^{(i,k)}, w^{(i,k)})\}_{i=0}^{q(k)} \subset \Xi \times \mathbf{R}$ are the nodes and weights. The rule is said to have an exactness of degree $s(k)$ if the approximation in (9) is an equality for all polynomials $F \in \mathbf{P}_{s(k)}$. The function $s: \mathbf{N}_0 \rightarrow \mathbf{N}_0$ is assumed to be nondecreasing. For $k \geq 0$, the k th-order *pseudospectral*

²A quadrature can naturally be applied to any function having well-defined values at the corresponding nodes, not only to continuous ones. The same is true for the associated projections.

projection onto the Legendre basis is $\mathcal{P}_k: C(\Xi) \rightarrow \mathbf{P}_{p(k)}$,

$$\mathcal{P}_k(F)(\vartheta) := \sum_{i=0}^{p(k)} \widehat{F}_i \ell_i(\vartheta), \quad \widehat{F}_i = \mathcal{Q}_k(F \ell_i),$$

where $p(k) = \lfloor s(k)/2 \rfloor$. It is easy to verify that $\mathcal{P}_k$ really is a projection, that is, an identity map (i.e., exact) on $\mathbf{P}_{p(k)}$. Finally, the *difference projection* operator for $k \geq 1$ is $\mathcal{D}_k: C(\Xi) \rightarrow \mathbf{P}_{p(k)}$,

$$\mathcal{D}_k := \mathcal{P}_k - \mathcal{P}_{k-1}$$

and for $k = 0$ we set $\mathcal{D}_0 = \mathcal{P}_0$. Notice that $\mathcal{D}_k$ is not actually a projection itself, but only a difference of two (unless $k = 0$).

The Gauss–Legendre quadrature admits the values $q(k) = k$ and $s(k) = 2k + 1$, yielding the convenient property $p(k) = k$. In contrast, the popular nested Clenshaw–Curtis rule grows exponentially, so that increasing the projection order by one gives many new polynomials. See, e.g., [3] for more details about different quadrature rules, including the Gauss–Patterson rule.

Before introducing the multidimensional counterparts of the linear operators presented above, we need to fix some notation. First, $\mathbf{P}_{\mathbf{k}} := (\mathbf{P}_{k_1} \otimes \cdots \otimes \mathbf{P}_{k_N})$ denotes the space of N -variate polynomials with the given maximum univariate degrees $k_1, \dots, k_N$. Second, the vectorization $\mathbf{y}: \mathbf{N}_0^N \rightarrow \mathbf{N}_0^N$ of a function $y: \mathbf{N}_0 \rightarrow \mathbf{N}_0$ is defined via $\mathbf{y}(\mathbf{k}) := (y(k_1), \dots, y(k_N))$. Last, for a multi-index $\mathbf{k} \in \mathbf{N}_0^N$, the induced full tensor index set is

$$\mathcal{T}(\mathbf{k}) := \{\mathbf{i} \in \mathbf{N}_0^N : 0 \leq i_n \leq k_n, n = 1, \dots, N\},$$

which has the cardinality

$$(10) \quad |\mathcal{T}(\mathbf{k})| = \prod_{n=1}^N (k_n + 1).$$

A full tensor quadrature $\mathcal{Q}_{\mathbf{k}}: C(\Xi^N) \rightarrow \mathbf{R}$ of order $\mathbf{k} \in \mathbf{N}_0^N$ is

$$\begin{aligned} \mathcal{Q}_{\mathbf{k}}(F) &:= (\mathcal{Q}_{k_1} \otimes \cdots \otimes \mathcal{Q}_{k_N})(F) \\ &= \sum_{i_1=0}^{q(k_1)} \cdots \sum_{i_N=0}^{q(k_N)} F(\vartheta^{(k_1, i_1)}, \dots, \vartheta^{(k_N, i_N)}) w^{(k_1, i_1)} \cdots w^{(k_N, i_N)} \\ &= \sum_{\mathbf{i} \in \mathcal{T}(\mathbf{q}(\mathbf{k}))} F(\vartheta^{(\mathbf{k}, \mathbf{i})}) w^{(\mathbf{k}, \mathbf{i})}, \end{aligned}$$

where $\vartheta^{(\mathbf{k}, \mathbf{i})} := (\vartheta^{(k_1, i_1)}, \dots, \vartheta^{(k_N, i_N)}) \in \Xi^N$ and $w^{(\mathbf{k}, \mathbf{i})} := \prod_{n=1}^N w^{(k_n, i_n)} \in \mathbf{R}$ denote the tensorized quadrature nodes and weights, respectively. It is obvious from (10) that for a high dimension N , the full tensor quadratures are computationally feasible only if the order $\mathbf{k}$ is sparse, i.e., consists mostly of zeros. The full tensor pseudospectral projection $\mathcal{P}_{\mathbf{k}}: C(\Xi^N) \rightarrow \mathbf{P}_{\mathbf{p}(\mathbf{k})}$ of order $\mathbf{k} \in \mathbf{N}_0^N$ is defined as

$$\begin{aligned} \mathcal{P}_{\mathbf{k}}(F)(\vartheta) &:= (\mathcal{P}_{k_1} \otimes \cdots \otimes \mathcal{P}_{k_N})(F)(\vartheta) \\ &= \sum_{\mathbf{i} \in \mathcal{T}(\mathbf{p}(\mathbf{k}))} \widehat{F}_{\mathbf{i}} L_{\mathbf{i}}(\vartheta), \end{aligned}$$

where the pseudospectral coefficients are

$$\widehat{F}_{\mathbf{i}} := \mathcal{Q}_{\mathbf{i}}(FL_{\mathbf{i}}).$$

The tensorized difference projection operator $\mathcal{D}_{\mathbf{k}}: C(\Xi^N) \rightarrow \mathbf{P}_{\mathbf{p}(\mathbf{k})}$ is unsurprisingly

$$(11) \quad \mathcal{D}_{\mathbf{k}} := \mathcal{D}_{k_1} \otimes \cdots \otimes \mathcal{D}_{k_P} = \mathcal{P}_{\mathbf{k}} + \sum_{\mathbf{i} \in \mathcal{B}(\mathbf{k})} (-1)^{\|\mathbf{k}-\mathbf{i}\|_1} \mathcal{P}_{\mathbf{i}},$$

where

$$(12) \quad \mathcal{B}(\mathbf{k}) := \{\mathbf{i} \in \mathbf{N}_0^N : \|\mathbf{k} - \mathbf{i}\|_{\infty} = 1, i_n \leq k_n, 1 \leq n \leq N\}$$

is (a part of) the *backward neighborhood* of the multi-index $\mathbf{k} \in \mathbf{N}_0^N$ (see also [12, Proposition 1.7]). Note that $0 \leq |\mathcal{B}(\mathbf{k})| \leq 2^N - 1$.

In the univariate case, it is easy to see that

$$(13) \quad \mathcal{P}_k(F) = \sum_{i=0}^k \mathcal{D}_i(F)$$

due to the telescopic property of the sum. On the other hand, it is well known that the approximation $\mathcal{P}_k(F) \approx F$ is good if F is smooth enough and k is sufficiently large; see, e.g., [2] for proofs and details. The sparse pseudospectral approximation generalizes the idea of (13) to higher dimensions via Smolyak's construction [12]

$$(14) \quad F \approx F^{(\mathcal{K})} := \sum_{\mathbf{i} \in \mathcal{K}} \mathcal{D}_{\mathbf{i}}(F),$$

where $\mathcal{K} \subset \mathbf{N}_0^N$ is (usually) a sparse set, meaning that each multi-index $\mathbf{i} \in \mathcal{K}$ contains many zeros so that the workloads of the quadratures entangled in (11) stay feasible. It follows from (11) that $F^{(\mathcal{K})}$ is a nontrivial linear combination of the full tensor projections $\{\mathcal{P}_{\mathbf{i}}\}_{\mathbf{i} \in \mathcal{K}}$. Naturally, the projections can further be written as linear combinations of Legendre polynomials. The set of Legendre degrees in (14) is denoted by $\mathcal{P}(\mathcal{K})$ which (implicitly) depends on the chosen quadrature. For the Gauss–Legendre quadrature we have $\mathcal{P}(\mathcal{K}) = \mathcal{K}$.

The approximation (14) is valid only if the index set $\mathcal{K}$ is *Smolyak admissible*, which means that every $\mathbf{i} \in \mathcal{K}$ satisfies

$$(15) \quad \text{if } i_n > 0, \text{ then } \mathbf{i} - \mathbf{e}_n \in \mathcal{K}$$

for each $n = 1, \dots, N$. This makes the sum in (14) telescopelike. Clearly, full tensor index sets are Smolyak admissible, but a full tensor index set being sparse means that the approximation is poor for most of the dimensions. A common alternative to the full tensor index set is the total-order index set, which for $k \in \mathbf{N}_0$ is

$$(16) \quad \mathcal{T}_{\text{tot}}(k) := \left\{ \mathbf{i} \in \mathbf{N}_0^N : \sum_{n=1}^N i_n \leq k \right\}.$$

The cardinality of the total-order index set is

$$(17) \quad |\mathcal{T}_{\text{tot}}(k)| = \binom{N+k}{k} \sim N^k,$$

which for a large dimension N is significantly less than the cardinality $(k+1)^N$ of the full tensor index set based on a constant vector $(k, \dots, k) \in \mathbf{N}_0^N$.

3.2. Adaptive SPAM. The aim of the adaptive SPAM is to construct a sparse index set $\mathcal{K}$ in (14) so that the approximation is good while the cardinality of the set is kept as small as possible. In many applications, especially when the target function is expensive to evaluate, the main trouble with a too large index set is the huge number of quadrature nodes, since the target function has to be evaluated at all of them. In our application, however, there is also another incentive for having a small index set $\mathcal{K}$ and a small polynomial set $\mathcal{P}(\mathcal{K})$, namely, the evaluation cost of the surrogate and its derivatives when the inverse problem of thermal tomography is solved with an iterative method.

From now on, we consider the temperature measurement vector $U: \Xi^N \rightarrow \mathbf{R}^M$ instead of a generic real-valued function F , and thus we construct a polynomial approximation $U^{(\mathcal{K})}$ for U as in (8). The quadrature and projection operators directly generalize to the vector-valued case. As mentioned in the beginning of this section, it is also possible to construct the sets of basis polynomials individually for each component, but for the sake of generality we present the adaptive algorithm for the vector-valued function U .

The algorithm for constructing $\mathcal{K}$ is essentially the same as the adaptive quadrature presented in [9] and goes as follows. At the beginning, we choose $\mathcal{K} = \{\mathbf{0}\} \subset \mathbf{N}_0^N$. Then, at each iteration, we select one not-yet-selected *critical multi-index* $\mathbf{k} \in \mathcal{K}$ according to the selection rule specified shortly. All the *admissible forward neighbors*

$$\mathcal{K}^+(\mathbf{k}) := \{\mathbf{i} \in \mathbf{N}_0^N : \mathbf{i} = \mathbf{k} + \mathbf{e}_j, 1 \leq j \leq N, \mathcal{K} \cup \{\mathbf{i}\} \text{ is Smolyak admissible}\}$$

of the selected critical index are then added to $\mathcal{K}$ and this selection-addition procedure is continued until some suitable stopping criterion is satisfied. The determination of the critical indices encompasses pseudospectral projections so that when the algorithm terminates, we have not only computed the index set $\mathcal{K}$ but also all the necessary information that is needed when evaluating the polynomial $U^{(\mathcal{K})}$ for any given parameter vector $\vartheta \in \Xi^N$. The expansion coefficients in (8) correspond to a matrix $\alpha \in \mathbf{R}^{M \times |\mathcal{P}(\mathcal{K})|}$ and the polynomial degrees can be stored in a sparse matrix $\mathcal{P}(\mathcal{K}) \in \mathbf{N}_0^{|\mathcal{P}(\mathcal{K})| \times N}$.

There are different selection criteria for the critical multi-index, but typically they are based on the norms of the difference projection operators. In this paper, we simply consider the norms

$$\epsilon(\mathbf{k}) := \|\mathcal{D}_{\mathbf{k}}(U)\|_{[L^2(\Xi^N)]^M} = \left(\sum_{m=1}^M \|\mathcal{D}_{\mathbf{k}}(U_m)\|_{L^2(\Xi^N)}^2 \right)^{1/2},$$

where

$$(18) \quad \mathcal{D}_{\mathbf{k}}(U_m) =: \sum_{\mathbf{i} \in \mathcal{T}(\mathbf{p}(\mathbf{k}))} \beta_{m,\mathbf{i}}^{(\mathbf{k})} L_{\mathbf{i}}$$

for some coefficients $\beta_{m,\mathbf{i}}^{(\mathbf{k})} \in \mathbf{R}$ that can be obtained by expanding the expression in (11). Among the not-yet-selected multi-indices, the one corresponding to the largest ϵ -norm has (loosely speaking) affected the approximation most and thus can be considered the critical index, whose admissible forward neighbors are added to the index set. By using the orthonormality of the Legendre polynomials, we can write

$$\|\mathcal{D}_{\mathbf{k}}(U_m)\|_{L^2(\Xi^N)} = \left(\sum_{\mathbf{i} \in \mathcal{T}(\mathbf{p}(\mathbf{k}))} \left(\beta_{m,\mathbf{i}}^{(\mathbf{k})} \right)^2 \right)^{1/2},$$

so that

$$\epsilon(\mathbf{k}) = \|\beta^{(\mathbf{k})}\|_F,$$

which is the Frobenius norm of the coefficient matrix in (18). According to Smolyak's construction (14), accumulating the coefficients $\beta^{(\mathbf{k})}$ computed during the algorithm results in the approximation $U^{(\mathcal{K})}$, and the polynomial degrees $\mathcal{P}(\mathcal{K})$ in (8) are just the union of the degrees appearing in the difference projections of the form (18).

In addition to the ϵ -norm introduced above, one may want to weigh different multi-index candidates based on, e.g., how much it costs to compute the associated projections or how many new polynomials one is going to obtain. See, e.g., [3] for work-considering algorithms. For the Gauss–Legendre quadrature, adding one multi-index $\mathbf{k}$ corresponds to adding precisely one new polynomial, namely, $L_{\mathbf{k}}$. In consequence, as we exclusively resort to the Gauss–Legendre rule in our numerical experiments, we employ ϵ as the selection criterion.

Algorithm 1: Adaptive sparse pseudospectral approximation.

```

1 function adaptiveSPAM (U);
   Input : Function  $U \in [C(\Xi^N)]^M$ 
   Output: Polynomial  $U^{(\mathcal{K})} \approx U$  of the form (8)
2  $U^{(\mathcal{K})} \leftarrow 0$ ; /* zero polynomial to start with */
3  $\mathcal{K} \leftarrow \emptyset$ ; /* indices already projected */
4  $\mathcal{J} \leftarrow \{\mathbf{0}\}$ ; /* indices to be projected next */
5 while  $|\mathcal{P}(\mathcal{K})| < \text{stopping criterion}$  do
6   for  $\mathbf{j} \in \mathcal{J}$  do
7     Compute and store  $\mathcal{P}_{\mathbf{j}}(U) \in \mathbf{P}_{\mathbf{p}(\mathbf{j})}$ ; /* forward solver needed */
8      $\mathcal{D}_{\mathbf{j}}(U) \leftarrow \mathcal{P}_{\mathbf{j}}(U) + \sum_{\mathbf{i} \in \mathcal{B}(\mathbf{j})} (-1)^{\|\mathbf{j}-\mathbf{i}\|_1} \mathcal{P}_{\mathbf{i}}(U)$ ; /* nothing to solve */
9     Compute and store  $\epsilon(\mathbf{j})$ ;
10     $U^{(\mathcal{K})} \leftarrow U^{(\mathcal{K})} + \mathcal{D}_{\mathbf{j}}(U)$ ; /* accumulate output ( $\mathcal{P}(\mathcal{K})$  and  $\alpha$ ) */
11  end
12   $\mathcal{K} \leftarrow \mathcal{K} \cup \mathcal{J}$ ;
13   $\mathbf{k} \leftarrow \arg \max_{\mathbf{i} \in \mathcal{K}} \epsilon(\mathbf{i})$ ; /* new critical index */
14   $\epsilon(\mathbf{k}) \leftarrow -\infty$ ; /* do not select again */
15   $\mathcal{J} \leftarrow \mathcal{K}^+(\mathbf{k})$ ; /* admissible forward neighbors */
16 end

```

The pseudocode of the method is presented in Algorithm 1. Note that $0 \leq |\mathcal{K}^+(\mathbf{k})| \leq N$ and the Smolyak admissibility implies that $\mathcal{B}(\mathbf{i}) \subset \mathcal{K}$ for each $\mathbf{i} \in \mathcal{K}^+(\mathbf{k})$; see (12) and (15). Thus, all the projections appearing on line 8 in Algorithm 1 have already been computed and the only line where the numerical forward solver for U is needed is line 7. Recall that by a forward solver we mean a computational function that returns the boundary measurement vector $U(\vartheta) \in \mathbf{R}^M$ for a given parameter vector $\vartheta \in \Xi^N$. In practice, such a solver only returns numerical approximations.

4. Numerical experiments. In this section, we present numerical examples demonstrating the feasibility of the proposed two-phase approach to thermal tomography:

1. In the offline phase, a polynomial surrogate is formed based (only) on generic information about the measurement setup, including approximate shape and

TABLE 1
Parameter values for the numerical experiments.

Parameter	Value	Notes
$N_a = N_b$	40 or 480	Discretization of the conductivity a and the capacity b
$J = R = N_c$	8	Number of heaters, sensors, and discretization of c
N_Ω	16	Discretization of the boundary
$\bar{a} = \bar{b}$	0.55	} $0.1 \leq a, b \leq 1$
$\tilde{a} = \tilde{b}$	0.45	
$\bar{c}$	0.11	} $0.02 \leq c \leq 0.2$ between the heaters
$\tilde{c}$	0.09	
c_0	10	$c = 10$ on the heaters
ρ_0	1	Radius of the reference domain $\Omega(0)$
ρ_-	0.8	Minimum radius
ρ_+	1.2	Maximum radius
T	2	Maximum time
M_T	6	Number of measurements in time
t_i	iT/M_T	Uniform measurement times ($i = 1, \dots, M_T$)
$g(t)$	$5t$	Time profile of the active heater
η	$\pi/8$	Heater width

size of the imaged body, the expected mean levels of the to-be-reconstructed parameters, and the number and size of the heaters and sensors. Because such computations can be performed prior to taking any measurements, one can assume the availability of considerable amounts of computation time and power for completing these tasks.

2. Once the measurements are in hand, the online phase consists merely of solving a least squares minimization problem that does not involve the forward solver but only its polynomial surrogate. Since the surrogate and its derivatives can be efficiently evaluated, this leads to fast reconstructions.

We first study the accuracy of the polynomial surrogate in subsection 4.1. More precisely, the values returned by the surrogate are compared with the values returned by the numerical solver upon which the surrogate is based. Naturally, the numerical solver itself returns only approximate measurement vectors for given parameter vectors, but since the surrogate can be constructed in the offline phase, one could, in principle, choose as accurate a solver as desired without affecting the efficiency of the actual reconstruction algorithm in the online phase. The errors in spatial and temporal discretizations are, however, taken into account when choosing the regularization or the Bayesian distributions for the inverse problem in subsection 4.2. On the other hand, since the number of employed basis polynomials affects the cost of evaluating the surrogate and its derivatives in the online phase [16], one should exploit adaptivity to keep the number of polynomials as low as possible without compromising the accuracy of the surrogate.

Subsection 4.2 addresses the inverse problem of thermal tomography by comparing some target functions a , b , c , and $\partial\Omega$ with the corresponding reconstructions; these considerations correspond to the online phase. The reconstructions are based on simulated measurement data, which are generated by a highly accurate numerical solver and further contaminated by a small amount of artificial noise in order to avoid any kind of an *inverse crime*. We also examine how reconstructions are affected if the uncertainties in the coefficient c and the boundary $\partial\Omega$ are not taken into account.

Table 1 lists the values defining the parametrized measurement setup used in all examples that follow. In particular, the zero vector $\vartheta = 0$ corresponds to the unit disk with constant fields $a \equiv b \equiv 0.55$ and the transfer coefficient $c = 0.11$

FIG. 1. *Parametrization of the measurement setup for the numerical experiments. The heaters are indicated by the bold boundary segments and the circles denote the sensors. The discretization of the conductivity and the capacity is shown for $N_a = N_b = 480$. Left: Unit disk $\Omega(0)$. Right: Domain that is obtained by perturbing two out of $N_\Omega = 16$ splines.*

between the heaters. Both the heater starting edges and the sensors are positioned in an equiangular fashion, and when the domain is the unit disk, the sensors are located exactly at the midpoints between the heaters. The geometry is illustrated in Figure 1; see also section 2 for more details on the parametrization. The total number of parameters is either $N = 104$ or $N = 984$ depending on the discretization of the fields a and b . The number of measurements is always $M = JRM_T = 384$.

4.1. Accuracy of the surrogate. We first focus on the accuracy of the polynomial surrogate which is constructed in three different ways. The parametrization of the measurement setup is such that $N_a = N_b = 40$ and thus $N = 104$ (this coarse discretization is visible in Figures 3 and 4). The first surrogate $U^{(\mathcal{K})}: \Xi^N \rightarrow \mathbf{R}^M$ is computed by choosing $\mathcal{P}(\mathcal{K}) = \mathcal{K} = \mathcal{T}_{\text{tot}}(2)$; see (16). That is, adaptivity is not used, but instead we fix the polynomials to those resulting from all full tensor projections having total order at most 2. (Recall that for Gauss–Legendre quadrature it holds that $\mathcal{P}(\mathcal{K}) = \mathcal{K}$, i.e., the polynomial degrees are precisely the projection orders.) The number of polynomials is $|\mathcal{P}(\mathcal{K})| = 5565$, which follows from (17). For the second surrogate, we apply the adaptive SPAM as described in subsection 3.2. The algorithm is run until $|\mathcal{P}(\mathcal{K})| \geq 5565$ and the same polynomial basis $\mathcal{P}(\mathcal{K})$ is constructed for all measurement components. The third surrogate is also constructed adaptively, but now separately for each measurement component, so that there are M different polynomial bases, each having cardinality of about 5565. In order to avoid too costly computations, the maximum degree for *univariate* Legendre polynomials is set to 4 in the third case.

The accuracy of these three surrogates is tested by drawing independent random parameter vectors $\{\vartheta^{(q)}\}_{q=1}^Q \subset \Xi^N$ from two different distributions. In the first case, all components of the random vector are independently and uniformly distributed on $\Xi = [-1/2, 1/2]$. The second distribution corresponds to (truncated) log-normal random fields for the conductivity a and the heat capacity b . More precisely, a and b are considered as discrete log-normal random fields with the (common) underlying normal distribution having the mean $\log(0.55)$ and the covariance function

$$(19) \quad K(x, y) = \varsigma^2 \exp\left(-\frac{\|x - y\|_2^2}{2\lambda^2}\right)$$

TABLE 2

Mean errors and variances for two random parameter distributions and three surrogates $U^{(\mathcal{K})}$: total-order polynomials of degree 2(T2), adaptively constructed polynomials with a common basis (adapt.), and adaptively constructed polynomials with separate bases (m-adapt.).

	Mean ($\mu_{\mathcal{K}} \cdot 10^2$)			Variance ($\sigma_{\mathcal{K}}^2 \cdot 10^2$)		
	T2	adapt.	m-adapt.	T2	adapt.	m-adapt.
Uniform	58.3	20.2	9.08	11.3	1.13	0.135
Log-normal	50.2	17.2	7.58	6.08	0.775	0.127

with $\varsigma = 0.5$ and $\lambda = 1/3$. In other words, $K(x, y)$ can be interpreted as the covariance matrix of an $N_a = N_b$ dimensional normal distribution when $x, y \in \Omega$ run through the centers of the “pixels” corresponding to the characteristic functions $\varphi_i^{(a)} = \varphi_i^{(b)}$ in the unit disk; cf. (6) and Figure 1. The fields a and b are drawn mutually independently as are the remaining $N_c + N_\Omega$ components of $\vartheta^{(q)}$, which again follow the uniform distribution. If a random draw for either of the two fields results in a pixel value that corresponds to a component $\vartheta_i^{(q)} < -1/2$ (resp., $\vartheta_i^{(q)} > 1/2$), then the field is truncated by redefining $\vartheta_i^{(q)} = -1/2$ (resp., $\vartheta_i^{(q)} = 1/2$).

For each surrogate $U^{(\mathcal{K})}$, we define the pointwise error, the mean error, and the (sample) variance as

$$e_{\mathcal{K}}^{(q)} := \|U(\vartheta^{(q)}) - U^{(\mathcal{K})}(\vartheta^{(q)})\|_2, \quad \mu_{\mathcal{K}} := \frac{1}{Q} \sum_{q=1}^Q e_{\mathcal{K}}^{(q)}, \quad \text{and} \quad \sigma_{\mathcal{K}}^2 := \frac{1}{Q-1} \sum_{q=1}^Q (e_{\mathcal{K}}^{(q)} - \mu_{\mathcal{K}})^2,$$

respectively. Here, U denotes our numerical forward solver that is based on piecewise linear FEs with meshes having a couple of thousands of nodes and the Crank–Nicolson time step of size $1/30$. We emphasize that the same solver is used for both the surrogate construction and the computation of the reference samples $\{U(\vartheta^{(q)})\}_{q=1}^Q$ since the aim at this point is to examine the accuracy of the polynomial approximation, not the accuracy of the numerical forward solver itself.

Table 2 lists the mean errors and variances corresponding to the two distributions for the parameter vector described above with sample size $Q = 1000$. It is evident that adaptivity improves the accuracy of the surrogate: the mean errors for the (commonly) adaptively constructed surrogate are about one third and the sample variances about one tenth of the corresponding values for the total-order surrogate with $\mathcal{P}(\mathcal{K}) = \mathcal{T}_{\text{tot}}(2)$. Because the adaptively constructed basis has approximately the same number of polynomials as the total-order basis, the evaluation and differentiation costs of these two surrogates are expected to be roughly the same. Although the complexity analysis in [16] considers only total-order bases, in practice there seems to be no significant difference in computational work between these two types of surrogates. In particular, for a given parameter vector, the evaluation of the surrogate still reduces to a simple matrix-vector product even if adaptivity is used. We refer to [16] for more detailed discussion on computational work involved when using the surrogate, and also remind the reader that in any case the cost of evaluating and differentiating the surrogate does not depend on the discretization of the forward solver that is used to form the surrogate. What is more, the adaptivity does not essentially slow down the offline phase, i.e., the construction of the surrogate, albeit some more evaluations of the forward solver may be needed.

Constructing the polynomial basis separately for each measurement component further enhances the surrogate, as seen in Table 2: the mean errors are only about 15%

FIG. 2. Mean errors for three different surrogates as functions of the expansion size.

of the corresponding numbers for $\mathcal{P}(\mathcal{K}) = \mathcal{T}_{\text{tot}}(2)$. However, in this case the surrogate construction algorithm has to be run several times, and the resulting surrogate cannot be evaluated as a simple matrix-vector product as in the case of a common basis (cf. [16]).

When the random realizations for a and b are smoother (i.e., the values of adjacent pixels are expected to be strongly correlated) as they are in the log-normal case, the mean errors seem to be slightly smaller compared to the case of random parameter vectors with independently and uniformly distributed components. For the zero parameter vector $\vartheta = 0$, the errors for the total order, adaptively constructed and adaptively/separately constructed surrogates are $e_{\mathcal{K}}^{(q)}(0) = 9.60 \cdot 10^{-2}$, $4.26 \cdot 10^{-2}$, and $1.41 \cdot 10^{-2}$, respectively. It is worth noticing that the mean errors for the surrogate that is constructed adaptively and separately for each measurement component are smaller than the pointwise error of the total-order surrogate at the origin $\vartheta = 0$ (cf. Table 2).

Figure 2 shows the mean errors for the log-normal random realizations if only a subset of the computed polynomials is used in each surrogate. More precisely, for each surrogate, the polynomials are first sorted based on the norms of the coefficient vectors $(\alpha_{m,i})_{m=1}^M$ in the expansion (8). Then the polynomials having the largest magnitudes for their coefficients are retained while the others are discarded. For the separately constructed surrogate this filtering is performed componentwise, so that the norms above are just the absolute values. The values at the right end of the convergence plots in Figure 2 are those on the bottom line in Table 2. It is clear that even a prematurely stopped adaptive algorithm, say, at 500 polynomials, overwhelms the total-order surrogate, although the values in Figure 2 do not exactly

correspond to the mean errors one would obtain by stopping the algorithm after obtaining a certain number of polynomials. For the adaptively and separately formed surrogate this effect is even more apparent: already 100 polynomials per measurement component is enough for achieving higher accuracy than the total-order surrogate with $\mathcal{P}(\mathcal{K}) = \mathcal{T}_{\text{tot}}(2)$.

The error between our standard numerical forward solver for the surrogate construction and a highly accurate solver is $1.62 \cdot 10^{-2}$ for the case $\vartheta = 0$, measured in the same way as the error between the standard numerical solver and the corresponding surrogates. The highly accurate solver has tens of thousands of FEM nodes and a Crank–Nicolson time step of $1/50$. The discrepancy between the two forward solvers stems mainly from the spatial discretization, and it is expected to be higher for more complicated domains and nonconstant fields a and b . Roughly speaking, we anticipate that the discretization error of our standard numerical forward solver is of the same order as the surrogate error for the most accurate of the three surrogates.

4.2. Inverse problem. For the inverse examples that follow, we generate the (noiseless) boundary measurement data $\tilde{U} \in \mathbf{R}^M$ by solving the problem (2) with the highly accurate solver described at the end of the previous section. In general, the coefficient functions and boundary shapes used for simulating the data are not (exactly) representable by the parametrization upon which the surrogates are built. However, the circumference for all target objects is chosen to be $2\pi = |\partial\Omega(0)|$; it is not too far-fetched to assume that the approximate size of the imaged object is known in practice. A realization of a vector of zero mean Gaussian noise with independent components is added to the measurement, with the standard deviation of the m th component being $5 \cdot 10^{-3} \cdot \tilde{U}_m$. The resulting noisy measurement vector is denoted by $\tilde{U} \in \mathbf{R}^M$.

The inverse problem is treated as a nonlinear least squares minimization

$$(20) \quad \arg \min_{\vartheta \in \Xi^N} \left\{ \|\tilde{U} - U^{(\mathcal{K})}(\vartheta)\|_2^2 + \delta^2 \|G\vartheta\|_2^2 \right\},$$

where $G \in \mathbf{R}^{N \times N}$ is a block-diagonal regularization matrix. Its first two blocks of lengths $N_a = N_b$ are $\text{chol}(K^{-1})$, where $\text{chol}(\cdot)$ denotes the upper-triangular Cholesky factor of a matrix and K is formed by evaluating the covariance function (19) with $\varsigma^2 = 0.5$ and $\lambda = 1/3$ at all centerpoint pairs for the pixels in $\Omega(0)$. Therefore, the regularization operator prefers smooth fields a and b as in, e.g., [6, 20]. If the sum of the noise in $\tilde{U}$ and the surrogate and numerical errors in $U^{(\mathcal{K})}$ was assumed to be Gaussian with independent zero-mean components having standard deviation $\delta > 0$, then (20) would have a Bayesian interpretation of maximizing the posterior distribution for the parameter ϑ under the prior assumption that the fields a and b are mutually independent Gaussian random fields with pointwise variance ς^2 and correlation length λ . To be quite precise, the regularization matrix should be updated based on the current shape of the domain, but since the boundary perturbations are relatively small, it is considered to be enough to compute the corresponding covariances only in the reference disk. We employ no regularization (i.e., no informative prior) for the coefficient c or the shape of the object, meaning that the lower right-hand corner of G is actually empty. Such regularization would become necessary if the respective numbers of degrees of freedom, N_c and N_Ω , were increased significantly.

We use the generic value $\delta = 10^{-2}$ for the regularization parameter in all the following examples. Note that δ and ς couple in (20) in such a way that decreasing δ is equivalent to increasing ς . There is, however, a semiheuristic Bayesian justification

for our choices of δ and ς . Since the values of the functions a and b vary between 0.1 and 1, the choice $\varsigma^2 = 0.5$ seems reasonable. On the other hand, assuming (falsely) that the sums of the artificial noise and the surrogate and discretization errors at the sensors are independent mean-free Gaussians with the same distribution, one can choose $\delta \approx e/\sqrt{M}$, where $e \approx 20 \cdot 10^{-2}$ is a crude approximation for the magnitude of the total error in the Euclidean norm when the surrogate is based on the standard adaptive algorithm (cf. the second column of Table 2). Increasing δ would generally force the minimization algorithm to converge more robustly, but it would also result in reconstructions that are closer to the average values and not so informative. Moreover, increasing the amount of measurement noise would make the reconstructions less accurate. The choice of $\lambda = 1/3$ reflects our prior assumption on the characteristic length of variations (in a and b) inside Ω . If λ is increased, the reconstructions become more constant, whereas decreasing λ allows reproducing finer details but also leads to unwanted fluctuations in reconstructions of areas described by constant parameters.

We solve the problem (20) by using the `lsqnonlin` function in MATLAB without constraints starting from the initial guess $\vartheta = 0$. However, any nonlinear least squares approach should work, since computing the values and Jacobians of the to-be-minimized function only requires polynomial evaluations and applying the regularization makes the problem somewhat well behaving. Constraints that force the parameter vector to stay within Ξ^N could be implemented if necessary, but in our examples the algorithm seems to work without such. The reconstructed functions a , b , c , and the shape of $\partial\Omega$ can be deduced from the minimizing parameter vector ϑ by considering the mappings introduced in subsection 2.2. Observe that the parameter values in Table 1 are quite arbitrary and no big effort has been made to select a measurement configuration that yields the maximal amount of information about the target (notice that there is a lot of freedom in the choice of, e.g., the heating pattern g). This task of *optimal experimental design* is not addressed in this paper; see, e.g., [10] for more on that topic.

Let us first consider the (computationally) easier case $N_a = N_b = 40$ and investigate what happens if the uncertainties in c and $\partial\Omega$ are not taken into account. To this end, two additional “bad” surrogates are constructed. The first one assumes that $c = \bar{c} = 0.11$ between the heaters and so the number of parameters is effectively reduced by N_c . This surrogate is formed adaptively and a common polynomial basis of size $|\mathcal{P}(\mathcal{K})| \approx 5565$ is used for every measurement. Figure 3 shows the target configuration which corresponds to the case $\vartheta = 0$ except that the coefficient c takes two distinct values, namely, 0.02 on the first four boundary segments between the heaters and 0.2 on the remaining ones. The second and third rows of the figure demonstrate that optimizing (20) only for $\vartheta^{(a)}$, $\vartheta^{(b)}$, and $\vartheta^{(\Omega)}$ employing the new surrogate yields an inaccurate reconstruction, whereas using the full adaptively computed common-basis surrogate from subsection 4.1 with $|\mathcal{P}(\mathcal{K})| \approx 5565$ and optimizing the whole vector ϑ results in a far superior outcome.

The next target setup is illustrated in the top row of Figure 4; it corresponds to vanishing subvectors $\vartheta^{(a)}$, $\vartheta^{(b)}$, and $\vartheta^{(c)}$. The second bad surrogate is formed by assuming that the domain is the unit disk, effectively removing $\vartheta^{(\Omega)}$ from our parameter vector. This surrogate is again computed by the adaptive, common-basis method and involves $|\mathcal{P}(\mathcal{K})| \approx 5565$ polynomials. It is then employed in (20) to compute the corresponding reconstruction, which is shown on the second row of Figure 4. It is obvious that ignoring the uncertainties in the shape of Ω heavily deteriorates the reconstructions of the other parameters. The third row presents the result of using the full adaptive common-basis surrogate in (20): the reconstruction is clearly better

FIG. 3. Target configuration (top), reconstruction without reconstructing the coefficient c (middle), and full reconstruction (bottom). $N_a = N_b = 40$.

than when ignoring $\vartheta^{(\Omega)}$, albeit not as accurate as that on the third row of Figure 3 where the same surrogate was used for a different target.

Finally, let us consider a surrogate that corresponds to $N_a = N_b = 480$ as in Figure 1, resulting in $N = 984$ variables. One polynomial basis is adaptively constructed for two adjacent active heaters at a time, so that there are altogether four different bases. Due to certain symmetries in the parametrization, choosing to use such a four-bases surrogate does not actually increase the amount of offline computations. The maximal univariate Legendre degree is set to 4 and the adaptive algorithm is terminated after obtaining 485 605 polynomials, which is the same number as in the total-order set $\mathcal{T}_{\text{tot}}(2)$ for 984 parameters. The maximal total polynomial degree found by the algorithm is 11. About 81% of the polynomials are of the second order. The higher-order polynomials are mostly associated with the shape variables $\vartheta^{(\Omega)}$ and the parameters in $\vartheta^{(a)}$ defining the thermal conductivity a close to the exterior boundary. However, the parameter choices in Table 1 obviously have a significant impact on the way the adaptive algorithm tends to distribute the higher-order polynomials.

Figure 5 presents the result of the surrogate inversion for a smooth target configuration. Both the thermal conductivity and the heat capacity are relatively well reconstructed along with the object shape. The same applies to a somewhat lesser extent to the heat transfer coefficient. Figure 6 documents a more demanding reconstruction. In this case, the piecewise constant target functions for a and b are not compatible with the chosen smoothness regularization. However, the reconstructions of a and b arguably still carry useful qualitative information about the interior structure of the imaged object—in particular, far more information than one would obtain if the uncertainties in the exterior boundary and the heat transfer coefficient were simply ignored in the inversion.

FIG. 4. Target configuration (top), reconstruction without reconstructing the shape of Ω (middle) and full reconstruction (bottom). $N_a = N_b = 40$.

FIG. 5. Smooth target configuration (top) and its reconstruction using a surrogate with $N = 984$ variables (bottom).

The online phase of computing the reconstructions in Figures 5 and 6, i.e., solving the minimization problem (20) with $N = 984$ by the `lsqnonlin` function of MATLAB with a user-supplied Jacobian, took only roughly ten seconds on a standard desktop

FIG. 6. Piecewise constant target functions (top) and their reconstructions using a surrogate with $N = 984$ variables (bottom).

computer. Obviously, the efficiency depends on the implementation and on the precise settings of the minimization procedure. Note also that the coefficient matrix $\alpha \in \mathbf{R}^{M \times |\mathcal{P}(\mathcal{K})|}$ contains almost $2 \cdot 10^8$ elements so storing it requires quite a bit of memory.

5. Conclusions. We have extended the previous works on thermal tomography to include estimating the exterior shape of the imaged object in addition to reconstructing the thermal conductivity, heat capacity, and surface conductance. The presented algorithm is based on the adaptive pseudospectral approximation approach and simple output least squares minimization, resulting in a fast inversion method that requires only polynomial evaluation and differentiation while being independent of the discretization of the parabolic forward problem. Our numerical examples demonstrated that the shape estimation is indeed important if one wants to obtain reasonable reconstructions of the thermal parameters inside a physical domain whose shape is not accurately known.

REFERENCES

- [1] V. F. BAKIROV AND R. A. KLINE, *Diffusion-based thermal tomography*, J. Heat Transfer, 127 (2005), pp. 1276–1279, <https://doi.org/10.1115/1.2039115>.
- [2] C. CANUTO, M. Y. HUSSAINI, A. QUARTERONI, AND T. A. ZANG, JR., *Spectral Methods: Fundamentals in Single Domain*, Springer, Berlin, 2006.
- [3] P. R. CONRAD AND Y. M. MARZOUK, *Adaptive Smolyak pseudospectral approximations*, SIAM J. Sci. Comput., 35 (2013), pp. A2643–A2670, <https://doi.org/10.1137/120890715>.
- [4] P. G. CONSTANTINE, M. S. ELDERED, AND E. T. PHIPPS, *Sparse pseudospectral approximation method*, Comput. Methods Appl. Mech. Engrg., 229–232 (2012), pp. 1–12, <https://doi.org/10.1016/j.cma.2012.03.019>.
- [5] J. DARDÉ, N. HYVÖNEN, A. SEPPÄNEN, AND S. STABOULIS, *Simultaneous reconstruction of outer boundary shape and admittivity distribution in electrical impedance tomography*, SIAM J. Imaging Sci., 6 (2013), pp. 176–198, <https://doi.org/10.1137/120877301>.
- [6] J. DARDÉ, N. HYVÖNEN, A. SEPPÄNEN, AND S. STABOULIS, *Simultaneous recovery of admittivity and body shape in electrical impedance tomography: An experimental evaluation*, Inverse Problems, 29 (2013), 085004, <https://doi.org/10.1088/0266-5611/29/8/085004>.

- [7] R. DAUTRAY AND J.-L. LIONS, *Mathematical Analysis and Numerical Methods for Science and Technology, Vol. 5: Evolution Problems I*, Springer, Berlin, 1992.
- [8] W. GAUTSCHI, *Orthogonal Polynomials: Computation and Approximation*, Oxford University Press, Oxford, 2004.
- [9] T. GERSTNER AND M. GRIEBEL, *Dimension-adaptive tensor-product quadrature*, Computing, 71 (2003), pp. 65–87, <https://doi.org/10.1007/s00607-003-0015-5>.
- [10] X. HUAN AND Y. M. MARZOUK, *Simulation-based optimal Bayesian experimental design for nonlinear systems*, J. Comput. Phys., 232 (2013), pp. 288–317, <https://doi.org/10.1016/j.jcp.2012.08.013>.
- [11] N. HYVÖNEN, V. KAARNIOJA, L. MUSTONEN, AND S. STABOULIS, *Polynomial collocation for handling an inaccurately known measurement configuration in electrical impedance tomography*, SIAM J. Appl. Math., 77 (2017), pp. 202–223, <https://doi.org/10.1137/16M1068888>.
- [12] V. KAARNIOJA, *Smolyak Quadrature*, Master's thesis, University of Helsinki, Helsinki, 2013.
- [13] J. P. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer, New York, 2005.
- [14] V. KOLEHMAINEN, J. P. KAIPIO, AND H. R. B. ORLANDE, *Reconstruction of thermal conductivity and heat capacity using a tomographic approach*, Int. J. Heat Mass Transf., 50 (2007), pp. 5150–5160, <https://doi.org/10.1016/j.ijheatmasstransfer.2007.07.015>.
- [15] M. G. LARSON AND F. BENGTZON, *The Finite Element Method: Theory, Implementation, and Applications*, Springer, Heidelberg, 2013.
- [16] L. MUSTONEN, *Numerical study of a parametric parabolic equation and a related inverse boundary value problem*, Inverse Problems, 32 (2016), 105008, <https://doi.org/10.1088/0266-5611/32/10/105008>.
- [17] A. NISSINEN, V. KOLEHMAINEN, AND J. P. KAIPIO, *Reconstruction of domain boundary and conductivity in electrical impedance tomography using the approximation error approach*, Int. J. Uncertain. Quantif., 1 (2011), pp. 203–222, <https://doi.org/10.1615/Int.J.UncertaintyQuantification.v1.i3.20>.
- [18] CH. SCHWAB AND C. J. GITTTELSON, *Sparse tensor discretizations of high-dimensional parametric and stochastic PDEs*, Acta Numer., 20 (2011), pp. 291–467, <https://doi.org/10.1017/S0962492911000055>.
- [19] J. M. TOIVANEN, V. KOLEHMAINEN, T. TARVAINEN, H. R. B. ORLANDE, AND J. P. KAIPIO, *Simultaneous estimation of spatially distributed thermal conductivity, heat capacity and surface heat transfer coefficient in thermal tomography*, Int. J. Heat Mass Transf., 55 (2012), pp. 7958–7968, <https://doi.org/10.1016/j.ijheatmasstransfer.2012.08.024>.
- [20] J. M. TOIVANEN, T. TARVAINEN, J. M. J. HUTTUNEN, T. SAVOLAINEN, H. R. B. ORLANDE, J. P. KAIPIO, AND V. KOLEHMAINEN, *3D thermal tomography with experimental measurement data*, Int. J. Heat Mass Transf., 78 (2014), pp. 1126–1134, <https://doi.org/10.1016/j.ijheatmasstransfer.2014.07.080>.
- [21] J. WINOKUR, P. CONRAD, I. SRAJ, O. KNIO, A. SRINIVASAN, W. C. THACKER, Y. MARZOUK, AND M. ISKANDARANI, *A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database*, Comput. Geosci., 17 (2013), pp. 899–911, <https://doi.org/10.1007/s10596-013-9361-3>.
- [22] J. WINOKUR, D. KIM, F. BISETTI, O. P. LE MAÎTRE, AND O. M. KNIO, *Sparse pseudo spectral projection methods with directional adaptation for uncertainty quantification*, J. Sci. Comput., 68 (2016), pp. 596–623, <https://doi.org/10.1007/s10915-015-0153-x>.