



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

Helin, Tapio; Hyvonen, Nuutti; Puska, Juha Pekka Edge-Promoting Adaptive Bayesian Experimental Design for X-ray Imaging

Published in: SIAM Journal on Scientific Computing

DOI: 10.1137/21M1409330

Published: 01/01/2022

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

Please cite the original version: Helin, T., Hyvonen, N., & Puska, J. P. (2022). Edge-Promoting Adaptive Bayesian Experimental Design for X-ray Imaging. SIAM Journal on Scientific Computing, 44(3), B506-B530. https://doi.org/10.1137/21M1409330

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

EDGE-PROMOTING ADAPTIVE BAYESIAN EXPERIMENTAL DESIGN FOR X-RAY IMAGING*

TAPIO HELIN[†], NUUTTI HYVÖNEN[‡], AND JUHA-PEKKA PUSKA[‡]

Abstract. This work considers sequential edge-promoting Bayesian experimental design for (discretized) linear inverse problems, exemplified by X-ray tomography. The process of computing a total variation-type reconstruction of the absorption inside the imaged body via lagged diffusivity iteration is interpreted in the Bayesian framework. Assuming a Gaussian additive noise model, this leads to an approximate Gaussian posterior with a covariance structure that contains information on the location of edges in the posterior mean. The next projection geometry is then chosen through A- or D-optimal Bayesian design, which corresponds to minimizing the trace or the determinant of the updated posterior covariance matrix that accounts for the new projection. Two- and three-dimensional numerical examples based on simulated data demonstrate the functionality of the introduced approach.

Key words. X-ray tomography, optimal projections, Bayesian experimental design, A-optimality, D-optimality, adaptivity, edge-promoting prior, lagged diffusivity

AMS subject classifications. 15A29, 62K05, 62F15, 65F10, 65F22, 65K10

DOI. 10.1137/21M1409330

SIAM J. SCI. COMPUT. Vol. 44, No. 3, pp. B506-B530

1. Introduction. Large-scale Bayesian inverse problems have rapidly gained popularity during the last two decades [31, 49]. While computational resources seem ever-increasing, data acquisition in a number of real-life inverse problems remains restricted or expensive. In consequence, there is a growing interest to develop computational methodologies for designing efficient data acquisition techniques or experimental setups to maximize the value of data in the solution process. Bayesian *optimal experimental design* (OED) provides a principled approach to such a task, and it has been widely adopted in the inverse problems community; see, e.g., [3] and references therein.

A Bayesian optimal design p^* maximizes the expected utility function U(p) over the design space D with respect to the data y and model parameters u according to

(1.1)
$$p^* = \underset{p \in D}{\arg \max} \mathbb{E}[U(p; u, y)]$$
$$= \underset{p \in D}{\arg \max} \int_{\mathcal{Y}} \int_{\mathcal{U}} U(p; u, y) \pi(u \mid p, y) \pi(y \mid p) \, \mathrm{d}u \, \mathrm{d}y.$$

Here $\pi(u \mid p, y)$ and $\pi(y \mid p)$ stand for the posterior distribution of the parameter u and the marginal distribution of the data y, respectively, under the design p. Moreover, \mathcal{Y} and \mathcal{U} are (subsets of) potentially high-dimensional real coordinate spaces to which

^{*}Submitted to the journal's Computational Methods in Science and Engineering section April 1, 2021; accepted for publication (in revised form) January 14, 2022; published electronically May 19, 2022.

https://doi.org/10.1137/21M1409330

Funding: The work of the first author was supported by the Academy of Finland through grants 320082 and 326961. The work of the second and third authors was supported by the Academy of Finland through grant 312124.

 $^{^{\}dagger}$ School of Engineering Science, LUT University, Lappeenranta, 53851 Finland (Tapio.Helin@lut.fi).

[‡]Department of Mathematics and Systems Analysis, Aalto University, FI-00076 Aalto, Finland (nuutti.hyvonen@aalto.fi, juha-pekka.puska@aalto.fi).

the data and the model parameters belong, respectively. The utility function can be devised in a number of ways; the two most common choices for U are arguably a *negative quadratic loss function* that measures the squared distance from u to a specific point estimator such as the posterior mean and the *expected information* gain, where U is defined by the Kullback–Leibler distance between the posterior and prior distributions.

The computational crux of (1.1) lies with the double integral over the potentially high-dimensional parameter and data spaces related to the considered imaging application. Moreover, if the set of possible designs is vast (e.g., p is a continuous parameter on a high-dimensional manifold), an exhaustive search may seem unfeasible. Still, significant progress has been made in the past working under conditions that allow closed form presentations for the above double integral. For the aforementioned two cost functions, the integrals in (1.1) are explicitly solvable when the posterior and data marginal distributions are Gaussian. In inverse problems, this occurs when the forward operator is linear and the prior and additive noise distributions are Gaussian [31]. In such a case, the double integral is essentially proportional to the trace and the determinant of the posterior covariance, respectively. In the literature, these are called the Bayesian A- and D-optimality criteria [18].

There has also been substantial effort to go beyond the conditions that enable explicit integration in (1.1). In this regard, important early work was carried out in [45, 27, 28] toward developing fast double loop Monte Carlo algorithms for tackling general inverse problems. More recent approaches have concentrated on improving efficiency of integral approximations by Laplace's method in the context of nonlinear inverse problems [37, 6, 20, 13]. Under well-designed approximation schemes, the computational complexity of such methods can be low in terms of the number of required forward solutions and scalable in the sense of being independent of the parameter and data dimensions [55, 56].

This paper grows out of the observation that the efficient use of non-Gaussian prior distributions in Bayesian OED for inverse problems has not been addressed in the literature. Indeed, successful solvers in imaging problems rely on well-designed prior information, which in variational regularization is often formulated in terms of nonquadratic penalty functionals [46]. Following the popularity of convex regularization in imaging, similar ideas have been successfully introduced to the Bayesian setting by formulating non-Gaussian priors in Banach spaces such as Besov spaces or functions of bounded variation; see, e.g., [54, 57, 52, 33, 1, 2, 38]. Motivated by these observations, our work contributes toward including non-Gaussian prior distributions in Bayesian OED practices for inverse problems and imaging.

1.1. Our contribution. This work introduces a computational method for performing greedy sequential OED for linear inverse problems with a *total variation* (TV) prior. The proposed algorithm is novel, as it does not utilize Laplace's approximation around the posterior mode or sampling schemes to tackle a non-Gaussian posterior distribution. Instead, its founding idea is based on the so-called lagged diffusivity approximation for TV introduced in [53]. At each step of the sequential algorithm, a lagged diffusivity iteration is employed to produce a sequence of Gaussian approximations for the TV prior, presumably with increasing accuracy close to the posterior mode. Assuming an additive Gaussian noise model, one of the two cost functions mentioned above, and replacing the TV prior by its final Gaussian approximation allow a closed form solution for the double integral in (1.1). This leads to a standard form A- or D-optimality criterion for choosing the subsequent measurement design.

B508 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

It should be emphasized that the Bayesian interpretation of the lagged diffusivity iteration is not new, but the main contribution of this work is its utilization in OED. Indeed, the Bayesian connection of lagged diffusivity iteration is detailed in [12, sect. 4.3.1] for the Laplace prior, with the associated line of reasoning readily applicable for the TV prior as well, and it was implicitly introduced already in [17]. Moreover, the iteration can be interpreted as a sequence of Laplace's approximations that depend on the current estimate for the posterior mode [51].

Like the lagged diffusivity approximation, our method could also be formulated for a large class of Gibbs prior measures. Moreover, it may be possible to extend some of our ideas to the framework of nonlinear inverse problems by combining them with Laplace's method. Be that as it may, in this work the proposed algorithm is only tested with a linear inverse problem and a TV prior.

We develop our method in the context of X-ray tomography, building upon our previous work [15] that considered efficiency and adaptivity of sequential OED in such a framework. X-ray tomography is particularly well-suited for a sequential approach to OED, as the radiation exposure (i.e., the number of projections) often needs to be minimized while maximizing the quality of the reconstruction in certain regions of interest, the locations of which may be unknown a priori.

We consider X-ray tomography in both two- and three-dimensional imaging setups with a narrow X-ray beam whose propagation angle and lateral position can be optimized. Our main hypothesis is that an approximate TV prior in Bayesian OED for X-ray tomography should promote designs that efficiently recover edges in the imaged target, which is a desired goal in X-ray tomography, as well as in many other imaging applications, because the imaged body often has piecewise constant absorption structure (consider, e.g., organs inside a human body); we refer to the textbooks [12, 16, 31] and the references therein for more information on other edge-promoting (Bayesian) reconstruction techniques. The presented numerical experiments, which are based on simulated data, demonstrate that our algorithm does indeed perform well for certain piecewise constant phantoms when compared with the use of equiangular full-width projections corresponding to an equivalent radiation dose.

This text is organized as follows. Section 2 introduces a discretized linear measurement model for X-ray tomography. In section 3 the basic ideas of lagged diffusivity iteration are interpreted in the Bayesian framework. The concepts of A- and D-optimality are recalled in section 4, and they are subsequently combined with the lagged diffusivity iteration to form a sequential OED algorithm in section 5. The numerical experiments are presented in section 6, and the concluding remarks are listed in section 7.

1.2. Literature review. Bayesian OED has gained substantial attention in large-scale inverse problems during the recent years. In addition to the works mentioned above, let us list [8, 11, 22, 23, 24, 25, 30, 32, 36, 37, 9] to name a few papers on this topic. In particular, there is an interesting line of research developing Bayesian OED for infinite-dimensional inverse problems [5, 4, 6, 7]. Here, we test our novel ideas in a sequential optimization strategy, which has previously been formalized for large-scale problems in [29] based on ideas from dynamical programming. For general references on Bayesian OED, we mention the review papers [18, 44] and the monograph [40].

Optimization of the imaging geometry in X-ray tomography has previously been considered in [43, 15]. The article [43] explored empirical A-optimal design in constrained problems based on training data by adopting sparse sensor-placing strategies

B509

and a gradient-based optimization scheme. The latter paper [15] introduced more degrees of freedom (lateral position of the source-receiver pair) to the problem, considered efficient evaluation of the A- and D-optimality target functions, and introduced adaptivity to the algorithm.

The idea of TV denoising was originally presented in [42], and the lagged diffusivity fixed point iteration for approximating TV regularization was introduced in [53]. The convergence of the algorithm has been considered, e.g., in [21, 19] for finite-dimensional image restoration problems, and its Bayesian interpretation was introduced in [12, 17]. Finally, let us remark that TV regularization is widely employed in computed tomography; see, e.g., [47, 35, 50].

2. Measurement model and its discretization. The X-ray measurements are modeled either as parallel beam or cone beam tomography, where multiple rays are directed into the object $D \subset \mathbb{R}^d$, d = 2 or 3, and the resulting intensities of the rays are measured at detectors [39]. Without severe loss of generality we assume that D is rectangular. The attenuation is described by the Lambert–Beer equation

(2.1)
$$I = I_0 \exp\left(-\int_L u \,\mathrm{d}s\right),$$

where L is the line along which the considered ray travels, I_0 is the intensity of the X-ray before entering the object, and $u: D \to \mathbb{R}_+$ is the absorption. Obviously, (2.1) can equivalently be given as

$$\log(I_0) - \log(I) = \int_L u \, \mathrm{d}s.$$

In particular, the difference between the logarithms of the emitted and measured intensities is typically considered as the available data when X-ray tomography is tackled mathematically.

We discretize the imaged domain into $n' \in \mathbb{N}$ pixels (or voxels in the threedimensional case) but assume the absorption distribution vanishes at the boundary pixels. The number of interior pixels is denoted by n < n'. The design parameter dependent forward operator, mapping the discretized absorption to the log-intensity measurements in a *single* projection image, can be approximated by a matrix $R \in \mathbb{R}^{m \times n}$, where m is the number of detectors in the considered measurement device whose geometric orientation can be altered between projection images (see, e.g., [48]). Typically the dimension of the unknown is higher than the number of pixels in a single projection image, i.e., $m \ll n$. In what follows, we abuse the notation by denoting with $u \in \mathbb{R}^n$, $n \in \mathbb{N}$, both the vector of pixel values defining the discretized (interior) absorption as well as a (smooth enough) function on D taking the given absorption values at the center points of the respective pixels. The correct interpretation should be clear from the context.

3. TV prior and lagged diffusivity. Let $u_{k-1} \in \mathbb{R}^n$ be the reconstruction after taking $k-1 \in \mathbb{N}_0$ X-ray projections, and assume that the *k*th projection image has just become available; section 4 below explains how the experimental design for this newest projection was chosen. In other words, *k* indicates the number of available projection images with *m* pixels each. Let us denote by

$$\mathbf{R}_{k} = \begin{bmatrix} R(p_{1}) \\ \vdots \\ R(p_{k}) \end{bmatrix} \in \mathbb{R}^{km \times n} \quad \text{and} \quad \mathbf{y}_{k} = \begin{bmatrix} y_{1} \\ \vdots \\ y_{k} \end{bmatrix} \in \mathbb{R}^{km}$$

the stacked X-ray matrix corresponding to all previous projections and the corresponding stacked noisy data vector, respectively. The vectors p_1, \ldots, p_k are the design parameters employed thus far. The measurements y_1, \ldots, y_k are modeled as realizations of the random variables

(3.1)
$$Y_j = R(p_j)U + N_j, \qquad j = 1, \dots, k_j$$

where U is the randomized discrete absorption and the noise N_j is assumed to follow a zero-mean Gaussian distribution $\mathcal{N}(0, \Gamma_{\text{noise}}^{(j)})$, where $\Gamma_{\text{noise}}^{(j)} \in \mathbb{R}^{m \times m}$ is symmetric and positive definite. The noise processes N_1, \ldots, N_k are assumed to be mutually independent.

The prior for the absorption U has an edge-promoting probability density of the form

(3.2)
$$\pi(u) \propto \exp(-\gamma \Phi(u)),$$

where $\gamma > 0$ is a free parameter and Φ is defined through

(3.3)
$$\Phi(u) = \int_D \varphi(|\nabla u|) \, \mathrm{d}x,$$

accompanied by the information that u vanishes at the pixels next to the boundary of D. In this work, we exclusively consider the smoothened TV prior [42] for which φ in (3.3) is a smooth approximation of the absolute value function,

(3.4)
$$\varphi(t) = \sqrt{t^2 + T^2} \approx |t|,$$

with T > 0 being a small parameter ensuring differentiability. However, it would also be possible to consider other edge-preferring priors of the exponential form (3.2); see, e.g., [12, 16, 31] and the references therein for more information on such priors.

According to the Bayes formula and assuming the measurement model (3.1), the posterior density for u thus reads

(3.5)
$$\pi(u \mid \mathbf{y}_k) \propto \pi(\mathbf{y}_k \mid u) \pi(u)$$
$$\propto \exp\Big(-\frac{1}{2}(\mathbf{y}_k - \mathbf{R}_k u)^{\mathrm{T}} (\mathbf{\Gamma}_{\mathrm{noise}}^{(k)})^{-1} (\mathbf{y}_k - \mathbf{R}_k u) - \gamma \Phi(u)\Big),$$

where $\mathbf{\Gamma}_{\text{noise}}^{(k)} := \text{diag}(\Gamma_{\text{noise}}^{(1)}, \ldots, \Gamma_{\text{noise}}^{(k)}) \in \mathbb{R}^{km \times km}$ is a block diagonal matrix defined by the noise covariance matrices for the previous measurements. Our leading idea is to iteratively approximate $\Phi(u)$ by quadratic terms in the spirit of the lagged diffusivity iteration [53]; see also [10, 26]. This results in an iterative algorithm for computing the reconstruction u_k , i.e., the mode of (3.5), after k measurements as well as forming a sequence of Gaussian approximations for (3.5) [12]. Although there definitely exist many other methods for computing the mode or the mean for (3.5), from the standpoint of Bayesian OED the intriguing aspect of the lagged diffusivity iteration is the possibility to use the corresponding final Gaussian approximation of (3.5) for choosing the next experimental design, which enables explicitly evaluating the double integral in (1.1).

3.1. First step: Gaussian approximation for the prior around u_{k-1} . Let $\{\phi_j\}_{j=1}^{n'}$ be a Lagrangian *finite element* (FE) basis for the rectangular dual mesh of the employed pixelification for D (or again, voxelification in three dimensions) numbered

B511

so that the first n basis functions correspond to the nodes at the midpoints of the interior pixels in D. In particular, the *j*th basis function ϕ_j takes value one at the midpoint of the *j*th pixel, i.e., at the *j*th node of the dual mesh, and vanishes in the closures of all rectangular elements defined by the dual mesh not intersecting the *j*th pixel of the original mesh; see, e.g., [14, pp. 68–69]. We denote the interior of the convex hull of the nodes in the dual mesh by $D' \subset D$. With this notation and due to the chosen numbering for the pixels in D, the first n FE basis functions ϕ_j, \ldots, ϕ_n belong to $H_0^1(D')$, i.e., the subspace of $H^1(D')$ consisting of the elements with a vanishing Dirichlet trace on $\partial D'$.

After identifying u with its interpolant in the FE basis and recalling that u is assumed to vanish at (the midpoints of) the boundary pixels numbered by $n+1, \ldots, n'$, through straightforward differentiation of $\Phi(u)$ one deduces that [10]

$$\nabla_u \Phi(u) = \Theta(u)u, \qquad u \in \mathbb{R}^n$$

where the matrix-valued function $\Theta : \mathbb{R}^n \to \mathbb{R}^{n \times n}$ is defined componentwise by

(3.6)
$$\Theta_{i,j}(w) := \int_{D'} \frac{1}{\sqrt{|\nabla_x w(x)|^2 + T^2}} \,\nabla\phi_i(x) \cdot \nabla\phi_j(x) \,\mathrm{d}x, \qquad i, j = 1, \dots, n_j$$

with $w \in \mathbb{R}^n$ being interpreted as an element of $H_0^1(D')$ via the introduced FE basis, excluding the boundary nodes of the dual mesh. In particular, for any fixed $w \in \mathbb{R}^n$, $\Theta(w) \in \mathbb{R}^{n \times n}$ is the FE system matrix with respect to the basis $\{\phi_j\}_{j=1}^n \subset H_0^1(D')$ for the partial differential operator

(3.7)
$$-\nabla \cdot \left(\frac{1}{\sqrt{|\nabla w|^2 + T^2}}\nabla(\cdot)\right) \colon H^1_0(D') \to H^{-1}(D');$$

see, e.g., [14] for more information on such discretizations. As the construction of $\Theta(w)$ includes a vanishing Dirichlet condition via the choice of the corresponding domain of definition $H_0^1(D')$ for (3.7), the matrix $\Theta(w)$ is positive definite and, in particular, invertible for any $w \in \mathbb{R}^n$.

The basic idea of the lagged diffusivity iteration can be interpreted as finding a sequence of quadratic approximations for the penalty term $\Phi(u)$ accurate to the first order around the current estimate for the posterior mode, which reduces the maximization of (3.5) to the minimization of a sequence of quadratic functions. To this end, let us introduce the quadratic penalty function

$$\Phi_{k-1}(u) = \frac{1}{2}u^{\mathrm{T}}\Theta(u_{k-1})u + \frac{1}{2}u_{k-1}^{\mathrm{T}}\Theta(u_{k-1})u_{k-1} + \int_{D'}\frac{T^2}{\sqrt{|\nabla_x u_{k-1}(x)|^2 + T^2}}\,\mathrm{d}x$$

One can easily see that

(3.8)
$$\Phi_{k-1}(u_{k-1}) = \Phi(u_{k-1})$$
 and $\nabla_u \Phi_{k-1}(u_{k-1}) = \nabla_u \Phi(u_{k-1}) = \Theta(u_{k-1})u_{k-1}$

meaning that the tangent planes for the graphs of $\Phi_{k-1} : \mathbb{R}^n \to \mathbb{R}_+$ and $\Phi : \mathbb{R}^n \to \mathbb{R}_+$ indeed coincide above the previous reconstruction u_{k-1} . Substituting Φ for Φ_{k-1} in (3.5), we have thus arrived at the approximate Gaussian posterior density

(3.9)
$$\pi^{(1)}(u | \mathbf{y}_k) \propto \exp\left(-\frac{1}{2}\left((\mathbf{y}_k - \mathbf{R}_k u)^{\mathrm{T}}(\mathbf{\Gamma}_{\mathrm{noise}}^{(k)})^{-1}(\mathbf{y}_k - \mathbf{R}_k u) + \gamma u^{\mathrm{T}}(\mathbf{\Gamma}_{k-1}^{(1)})^{-1}u\right)\right),$$

where $\Gamma_{k-1}^{(1)} := \Theta(u_{k-1})^{-1}$. Finding the mode of the density (3.9) can also be interpreted as a single step of a certain quasi-Newton scheme for maximizing the log-posterior defined by (3.5) starting from u_{k-1} [51].

B512 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

3.2. Second step: Iterating the argument. Building the initial Gaussian approximation (3.9) for the posterior (3.5) consists essentially of two steps: (i) assuming a reasonable estimate u_{k-1} for the solution of the studied inverse problem and (ii) forming the approximate prior covariance via $\Gamma_{k-1}^{(1)} = \Theta(u_{k-1})^{-1}$. Introducing the mean of the density (3.9) as a new, hopefully more accurate reconstruction and iterating the argument lead to a Bayesian interpretation [12] of the lagged diffusivity algorithm [53] for computing a reconstruction u_k after having k projection images in hand.

Define $u_{k-1}^{(0)} = u_{k-1}$. Assuming the availability of $u_{k-1}^{(j-1)}$, form an approximate prior covariance

(3.10)
$$\Gamma_{k-1}^{(j)} = \Theta(u_{k-1}^{(j-1)})^{-1}.$$

Introduce the corresponding posterior density

(3.11)
$$\pi^{(j)}(u | \mathbf{y}_k) \propto \Big(-\frac{1}{2} \Big((\mathbf{y}_k - \mathbf{R}_k u)^{\mathrm{T}} (\mathbf{\Gamma}_{\mathrm{noise}}^{(k)})^{-1} (\mathbf{y}_k - \mathbf{R}_k u) + \gamma u^{\mathrm{T}} (\mathbf{\Gamma}_{k-1}^{(j)})^{-1} u \Big) \Big),$$

and compute its mean

(3.12)
$$u_{k-1}^{(j)} = \Gamma_{k-1}^{(j)} \mathbf{R}_k^T \left(\mathbf{R}_k \Gamma_{k-1}^{(j)} \mathbf{R}_k^T + \gamma \mathbf{\Gamma}_{\text{noise}}^{(k)} \right)^{-1} \mathbf{y}_k;$$

see, e.g., [31, Thm. 3.7].

If the chosen stopping criterion is satisfied at j = J, one dubs $u_k := u_{k-1}^{(J)}$ the reconstruction after k projection images; the stopping criterion employed in our numerical tests is introduced and motivated in section 5, where the complete algorithm is formalized. The corresponding covariance matrix for the Gaussian density (3.11) with j = J is

(3.13)
$$\Gamma_{k} = \gamma^{-1} \big(\Gamma_{k-1}^{(J)} - \Gamma_{k-1}^{(J)} \mathbf{R}_{k}^{T} \big(\mathbf{R}_{k} \Gamma_{k-1}^{(J)} \mathbf{R}_{k}^{T} + \gamma \mathbf{\Gamma}_{\text{noise}}^{(k)} \big)^{-1} \mathbf{R}_{k} \Gamma_{k-1}^{(J)} \big);$$

see, e.g., [31, Thm. 3.7]. This covariance structure is then used for choosing the parameter vector p_{k+1} defining the next X-ray projection as explained in the following section.

Remark 3.1. The two steps (3.10) and (3.12) correspond to a lagged diffusivity iteration for maximizing the argument of the exponential in (3.5), that is, computing an approximation of the maximum a posteriori estimate for the absorption in D after the availability of k projection images. As the convergence of the lagged diffusivity iteration has been proven for denoising problems in [21, 19], it is not too far-fetched to hope that the above introduced iteration converges toward the mode of the posterior (3.5). For large enough j, the Gaussian density $\pi^{(j)}$ defined by (3.11) can thus be considered an approximation for the exact posterior (3.5) close to its mode.

Note also that the essential property of the regularization matrix $\Theta(w)$ is that it provides an approximation

(3.14)
$$v_1^T \Theta(w) v_2 \approx \int_D \frac{1}{\sqrt{|\nabla_x w(x)|^2 + T^2}} \nabla v_1(x) \cdot \nabla v_2(x) \, \mathrm{d}x,$$

where, on the left, v_1 and v_2 are interpreted as vectors in \mathbb{R}^n and, on the right, as functions over D with the indicated values at the midpoints of the n interior pixels and zeros as values at the boundary pixels. The use of FE basis functions above is motivated by the ease of notation as well as their direct applicability to nonrectangular

B513

meshes as well (cf. [10, 26]). However, on a rectangular mesh a first-order *finite* difference approximation for computing the gradients in (3.14) (cf., e.g., [51]) is the simplest approach to which we also resort to in our numerical studies, although on a regular mesh first-order FE and finite difference approximations for the right-hand side of (3.14) lead to similar discretizations.

4. A- and D-optimal designs. Let us assume that we have $k \in \mathbb{N}$ projection images of the imaged object D at our disposal. According to the construction in the previous section, this leads to the approximate, posterior probability distribution $\mathcal{N}(u_k, \Gamma_k)$ for the absorption U, with the mean and covariance defined via (3.12) and (3.13), respectively. Assuming the new X-ray projection obeys the same measurement model as the previous ones, i.e., (3.1), the Gaussian posterior covariance after the (k + 1)th projection reads

(4.1)
$$\Gamma_{\text{post}}^{(k+1)}(p) = \Gamma_k - \Gamma_k R(p)^T \left(R(p) \Gamma_k R(p)^T + \Gamma_{\text{noise}}^{(k+1)} \right)^{-1} R(p) \Gamma_k,$$

where p is the to-be-selected design parameter determining the (k+1)th projection.

The task at hand is now to choose the (k + 1)th projection or, more precisely, the corresponding design parameter p_{k+1} . In Bayesian OED, one often considers minimizing the expected squared distance of the unknown in a given (semi)norm around the posterior mean; see, e.g., [4, 18] for more details. In the considered simple setting, this leads to the so-called A-optimality criterion for choosing the (k + 1)th design parameter,

(4.2)
$$p_{k+1} = \operatorname*{arg\,min}_{p} \operatorname{tr} \left(A \Gamma_{\text{post}}^{(k+1)}(p) A^{T} \right),$$

with the employed seminorm induced by the positive semidefinite matrix $A^T A$ for a given $A \in \mathbb{R}^{n \times n}$. Another option is to maximize the expected information gain when replacing the prior with the posterior, which in the linear and Gaussian setting leads to the so-called D-optimality criterion [4, 18],

(4.3)
$$p_{k+1} = \operatorname*{arg\,min}_{n} \log\left(\det\left(\Gamma_{\mathrm{post}}^{(k+1)}(p)\right)\right),$$

if information at all pixels is considered equally valuable.

To solve the minimization problem (4.2) or (4.3) and to find the optimal design for the (k+1)th X-ray projection, we resort to the exhaustive optimization algorithm introduced in [15]. In our numerical experiments, the weight A for A-optimality is always the identity matrix $I \in \mathbb{R}^{n \times n}$; that is, we employ the (squared) $L^2(D)$ -norm as the distance measure and, in particular, consider the reconstruction accuracy equally important at all pixels. The details on how to tackle a setting where one is only interested in the accuracy of the reconstruction inside a certain region of interest can be found in [15] for both A- and D-optimality.

4.1. On computational cost and interpolation onto a sparser grid. In our numerical experiments, finding the optimal design parameter via (4.2) or (4.3)by resorting to the exhaustive search introduced in [15] is computationally more demanding than computing an edge-enhancing reconstruction using the lagged diffusivity ideas presented in section 3. The details on the algorithmic implementation for evaluating the A- and D-optimality target functions for a given p, i.e., the pivotal step in the corresponding exhaustive minimization algorithms, can be found for A- and D-optimality in sections 3.3.2 and 3.3.1 of [15], respectively. To summarize, as the region of interest in our tests is all of D, for both forms of optimality one needs to form a Cholesky decomposition

$$C(p)C(p)^T = R(p)\Gamma_k R(p)^T + \Gamma_{\text{noise}}^{(k+1)} \in \mathbb{R}^{m \times m}$$

for all p on the employed search grid. For D-optimality it is enough to evaluate the logarithm of the determinant of C(p), whereas for A-optimality one must form $C(p)^{-1}R(p)\Gamma_k \in \mathbb{R}^{m \times n}$ and compute its Frobenius norm. Depending on the sizes of mand n, the most time-consuming steps are the matrix multiplications involving the full high-dimensional matrix $\Gamma_k \in \mathbb{R}^{n \times n}$ or the formation of the Cholesky decomposition. In our numerical experiments, m is so small that the computational bottleneck is usually the former.

A straightforward idea for speeding up the optimization step is implementing it using a sparser discretization than the one employed for computing the actual reconstructions: Once the reconstruction $u_k \in \mathbb{R}^n$ corresponding to the first k projection images has become available, it is interpolated onto a sparser grid with $\tilde{n} \leq n$ interior pixels to obtain $\tilde{u}_k \in \mathbb{R}^{\tilde{n}}$. The corresponding covariance matrix $\tilde{\Gamma}_k$ is then formed as in (3.13) but with \mathbf{R}_k replaced by the analogous stacked X-ray projection matrix for the sparser discretization and with $\Gamma_{k-1}^{(J)}$ replaced by $\Theta(\tilde{u}_k)^{-1}$ also formed on the sparser mesh. The approximate posterior for the interpolated absorption $\tilde{\Gamma}_{\text{post}}^{(k+1)}(p)$ is then as in (4.1) but with Γ_k replaced by $\tilde{\Gamma}_k$ and R(p) with an X-ray projection matrix corresponding to the sparser discretization. Finally, the optimal design parameter is computed via (4.2) or (4.3) with $\Gamma_{\text{post}}^{(k+1)}(p) \in \mathbb{R}^{n \times n}$ replaced by $\tilde{\Gamma}_{\text{post}}^{(k+1)}(p) \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$, recalling that A is an identity matrix of an appropriate size in our numerical tests.

It should be emphasized that the evidence on the feasibility of the above described interpolation step is purely computational (cf. [15]) and there is no guarantee that it does not significantly affect the specifications of the experimental designs produced by the sequential optimization process, as demonstrated by our first numerical test in section 6.1.1. Even if one deduced bounds for the effect of the interpolation step on the values of the A- and D-optimality target functions, this would not yet provide means to analyze how the whole sequential optimization process is affected by the interpolation due to the potential existence of multiple local minima of almost equivalent quality (cf. [15]): If the interpolation error induces a major change in the specifications of the chosen design at a single step of the sequential algorithm, the following optimal designs are no longer comparable due to the greedy nature of the optimization process. Moreover, as the TV prior is known to suffer from a certain lack of discretization invariance [34], the comparison of results on different discretization levels is problematic as such.

5. Sequential edge-promoting optimization of projections. In this section, the above developments are summarized by combining the lagged diffusivity iteration and the sequential optimization of X-ray projections into a single concise algorithm. In the following it is assumed that the overall number of X-ray projections $K \in \mathbb{N}$ is known in advance, but in practice the user of the algorithm can stop the iteration as soon as the reconstruction is considered good enough, thus treating K as the maximum number of projection images.

ALGORITHM 1. Select the prior parameters T > 0 and $\gamma > 0$, a tolerance for the stopping criterion $\tau > 0$, and the number of iterations $K \in \mathbb{N}$, and set A = I. Initialization:

 \triangleright Set $u_0 = \mathbf{1} \in \mathbb{R}^n$. \triangleright Define $\Gamma_0 := \Theta(u_0)^{-1}$ according to (3.6)

Iteration:

for k = 1, ..., K do \triangleright Solve for p_k via (4.2) or (4.3) with $\Gamma_{\text{post}}^{(k)}(p)$ defined by (4.1) [15], resorting to interpolation onto a sparser grid if necessary (see section 4.1). \triangleright Form the projection matrix $R(p_k)$, and "measure" the data y_k . \triangleright Set j = 0, $u_{k-1}^{(0)} = u_{k-1}$, and $\Delta \Phi = \tau + 1$. while $\Delta \Phi > \tau$ do $\triangleright Set \ j \leftarrow j+1. \\ \triangleright Form \ \Gamma_{k-1}^{(j)} \ according \ to \ (3.10).$

$$\triangleright$$
 Compute $u_{k-1}^{(j)}$ according to (3.12).

 $\triangleright Compute \ \Delta \Phi = |\Phi(u_{k-1}^{(j-1)}) - \Phi(u_{k-1}^{(j)})| / \Phi(u_{k-1}^{(j)}).$

end while

$$\triangleright$$
 Set $J = j$. Define $u_k = u_{k-1}^{(J)}$ and Γ_k via (3.13).

end for

return u_K and Γ_K .

The stopping criterion for the interior loop is motivated by material in [10]: Apart from the case i = 0, the value of the (smoothened) TV functional Φ typically decreases monotonically during the lagged diffusivity iteration because the reconstruction gradually becomes better aligned with the prior information. The iteration is stopped once the relative convergence rate falls below a preselected tolerance $\tau > 0$.

In many of the following numerical examples, the deduction of the sequentially A- or D-optimal projections, i.e., the first step in the exterior loop of Algorithm 1, is performed on a sparser discretization of D consisting of $\tilde{n} < n$ interior pixels in order to speed up the computations. The modifications required by this accelerated algorithm are described in section 4.1. Consult [15] for more information on the exhaustive algorithm for defining the optimal projections.

6. Numerical experiments. Both two- and three-dimensional numerical examples are presented. In all tests, the free parameters in Algorithm 1 are chosen as $T = 10^{-6}$, $\gamma = 10^{-2}$, $\tau = 10^{-4}$, and A = I. The algorithm is not very sensitive to the chosen (reasonably small) value for T whose purpose is to assure the differentiability of φ in (3.4). If T is too large, φ is not an accurate approximation for the absolute value function, and one cannot consider (3.2) a smooth approximation of the TV prior. On the other hand, if T is very small, the matrix $\Theta(u)$ defined by (3.6) becomes ill-conditioned if u is constant over some regions in D, which, in fact, corresponds to the expected behavior of u. The values chosen for the other two parameters γ and τ do affect the numerical results: γ controls the strength of the prior, i.e., the weighting between the data fit and penalty terms in the exponent of the posterior (3.5), whereas τ defines the stopping criterion. The former should be chosen so that the reconstructions are in line with the a priori information on the unknown, whereas τ should be small enough to allow convergence to the vicinity of the mode of the posterior (3.5)but large enough to avoid unnecessary lagged diffusivity iterations. However, as our main aim is to compare reconstructions with and without sequential optimization of the projection geometries, we do not dwell on the selection of γ and τ . The choice of A reflects that we are equally interested in the reconstruction quality everywhere in Dwhen A-optimality is considered. The components of the additive zero-mean Gaussian noise contaminating the measurements are assumed to be mutually independent with a common standard deviation $\sigma > 0$ that may vary between the experiments. In other words, all noise covariance matrices appearing in sections 3 and 4 are assumed to be of the form $\sigma^2 I$, with I being an identity matrix of the appropriate size.

When the performance of Algorithm 1 is compared to reconstructions corresponding to, say, equiangular projections, the reference reconstructions are computed via a single lagged diffusivity iteration with the same, aforelisted values for the parameters T, γ , and τ . To be more precise, if **R** is the projection matrix corresponding to all employed reference geometries, **y** is the corresponding data vector, and $\mathbf{\Gamma}_{\text{noise}} = \sigma^2 \mathbf{I}$ is the assumed noise covariance, then one starts from the initial guess $u^{(0)} = \mathbf{1} \in \mathbb{R}^n$ and iterates the two steps

(6.1)
$$\Gamma^{(j)} = \Theta(u^{(j-1)})^{-1}, \qquad u^{(j)} = \Gamma^{(j)} \mathbf{R}^T \left(\mathbf{R} \Gamma^{(j)} \mathbf{R}^T + \gamma \boldsymbol{\Gamma}_{\text{noise}} \right)^{-1} \mathbf{y}$$

until

(6.2)
$$\frac{\left|\Phi(u^{(j-1)}) - \Phi(u^{(j)})\right|}{\Phi(u^{(j)})} < \tau,$$

after which $u^{(j)}$ is dubbed the reconstruction. In other words, one essentially runs the interior loop of Algorithm 1 assuming that all (equiangular) projection geometries and the associated data are available to start with.

Remark 6.1. We do not claim that the lagged diffusivity iteration is the best method for computing TV-type reconstructions in X-ray tomography. However, since the algorithm for deducing the optimal projection geometries is inherently connected to the lagged diffusivity ideas, we consider using the simplified version of Algorithm 1, detailed in (6.1)-(6.2), a well-motivated choice for computing the control reconstructions corresponding to nonoptimized projection geometries. In particular, this enables focusing solely on the effect of the optimal design when comparing the reconstructions.

6.1. Two-dimensional parallel beam tomography. In our two-dimensional numerical experiments, the measurement setup is the same as described in [15]. That is, the domain $D = [0, 1]^2$ is discretized into $n = N^2$ square pixels, through which we take projections consisting of a number of parallel X-rays; see Figure 6.1. The individual X-rays are equally spaced and have a fixed width for a particular experiment. The width of the whole source-receiver pair satisfies $0 < w \leq 1$. The components of the design variable $p \in \mathbb{R}^2$ for a single projection geometry define the projection angle and the signed distance from the center of D to the median line of the source-receiver pair. The latter component of p is restricted within the interval [w - 1, 1 - w]/2.

We present three two-dimensional experiments. The first one exemplifies the general behavior of Algorithm 1 with a simple target. The effect of optimizing the projection geometries on a sparser grid than the one used for forming the reconstructions is also tested; see section 4.1. In the second test, the superiority of Algorithm 1 over the usage of equiangular full-width projections with an equivalent radiation dose is statistically demonstrated in the case of certain randomly selected phantoms. Finally, the third test applies Algorithm 1 to the Shepp–Logan phantom. The second test considers both A- and D-optimality, whereas the first and third tests concentrate solely on A-optimality.

6.1.1. 2D Test 1: Explicit example with a simple target. The aim of our first numerical experiment, which only considers the A-optimality criterion, is to



FIG. 6.1. Two-dimensional measurement setup.



FIG. 6.2. 2D Test 1. Left: Target with the rectangle, circle, and ellipse having absorption levels 1, 0.5, and 0.8, respectively. Right: Relative $L^2(D)$ errors for the reconstructions. The blue curve corresponds to projections A-optimized with the dense discretization for D, the red curve to projections A-optimized with the sparse discretization for D, and the black curve to the equiangular full-width reference projections. The horizontal axis indicates the number of projections with the beam width 0.25.

demonstrate the basic functioning of Algorithm 1. The target, shown in the left-hand image of Figure 6.2, consists of three simple shapes, each with a different uniform absorption level, placed randomly inside D. The absorption of the background is zero. The target has N = 100 pixels per edge, and the number of individual sensors in a full-width source-receiver pair is 51. The noise level is set to $\sigma = 10^{-3}$, which corresponds to a noise-to-signal ratio of at least 0.2% for all line integrals considered in the inversion. The beam width is chosen to be 0.25, which is a quarter of the maximal source-receiver pair width and corresponds to 13 individual X-rays. Algorithm 1 is run for a total of K = 16 iterations. In addition to considering the basic form of Algorithm 1, we also test speeding up the computations by performing the selection of the projection geometries on a considerably sparser discretization of the domain with only $\tilde{N} = 31$ pixels per edge; see section 4.1 for more details, and note, in particular, that the actual reconstructions are still formed on the denser grid with N^2 pixels. For comparison, we also compute reconstructions from equiangular full-width projections corresponding to equivalent radiation doses.

B518 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

The right-hand image of Figure 6.2 shows the relative $L^2(D)$ reconstruction errors after each step of Algorithm 1; the blue curve corresponds to optimizing the projection geometries on the reconstruction grid with 10⁴ pixels and the red curve to performing the optimization steps of the algorithm using the considerably sparser discretization with only $\tilde{N}^2 \approx 10^3$ pixels. The black line depicts the relative $L^2(D)$ errors resulting from the equiangular reference projections. Note that 1 projection with the maximal beam width of 1 approximately corresponds to the same amount of data, or equivalently the same radiation dose, as 4 projections with the beam width 0.25. As a consequence, the labels at 4, 8, 12, and 16 on the horizontal axis correspond to 1, 2, 3, and 4 equiangular reference projections, respectively.

According to Figure 6.2, the $L^2(D)$ reconstruction errors at equivalent radiation doses are lower for the sequentially optimized projection geometries with the quarterwidth source-receiver pair than for the equiangular full-width projections. This is not very surprising as the full-width projections (are forced to) waste radiation to image regions that contain nothing interesting, whereas the optimized projections concentrate on areas of interest; cf. Figure 6.3. On the other hand, deducing the optimal designs employing the sparser discretization for D does not seem to hamper the overall



FIG. 6.3. 2D Test 1. The sixteen A-optimized projection geometries and the corresponding reconstructions. The latest projection geometry is illustrated by the black lines.



FIG. 6.4. 2D Test 1. Reference reconstructions for 1, 2, 3, and 4 equiangular full-width projections.

performance of Algorithm 1, although the discretization level does affect the precise specifications of the individual optimized projection geometries (cf. section 4.1). After sufficiently many projections, the advantage of Algorithm 1 over the equiangular full-width projections becomes almost negligible.

Figure 6.3 shows the projection geometries and the corresponding reconstructions for all 16 iterations of Algorithm 1 without interpolation onto a sparser grid in the optimization steps. At least in the considered setup, the algorithm does indeed seem to have a tendency to concentrate the projections over areas where the reconstruction already shows quick variations. Occasionally other areas are also explored; cf. the 14th projection. After 14 iterations all target shapes are already clearly visible. For comparison, Figure 6.4 shows the four reconstructions corresponding to full-width equiangular reference projections.

Figure 6.5 presents the pixelwise (posterior) standard deviations corresponding to the reconstructions in Figure 6.3; that is, the square roots of the diagonal elements of the covariance matrix Γ_k are visualized as an image in Figure 6.5 at the same position where the corresponding reconstruction u_k is shown in Figure 6.3 for k = 1, ..., 16. Figure 6.5 clearly demonstrates that Algorithm 1 interprets the uncertainty in the reconstruction to be highest close to the already detected boundaries as well as in regions where previously undetected structures have recently emerged. Note, however, that the images in Figure 6.5 do not provide all information that is used when choosing the subsequent projection geometries, as the correlations between pixel values, described by the off-diagonal elements of Γ_k , also play a crucial in (4.2) and (4.3).

When computing the reconstructions in Figure 6.3, our nonoptimized MATLAB implementation spent almost 99.9% of its runtime performing the exhaustive optimization step in Algorithm 1. When the exhaustive search was performed using the sparser grid for D, this percentage decreased to about 90%. The number of lagged diffusivity steps required for forming the reconstructions in Figure 6.4 were 13, 13, 15, 10, 8, 11, 9, 6, 3, 5, 6, 5, 10, 14, 3, and 3, respectively; these numbers were not significantly altered by performing the optimization step on a sparser grid.

6.1.2. 2D Test 2: Average errors over random targets. In the second numerical experiment, the aim is to statistically demonstrate that Algorithm 1 has the potential to produce on average better reconstructions for a limited radiation dose than a straightforward approach with equiangular full-width projections as well as to compare the two optimality criteria (4.2) and (4.3). To this end, the algorithm is run with the beam widths of 0.25 and 0.5 for a set of random targets, and the average relative $L^2(D)$ reconstruction errors for A- and D-optimal designs are compared to those obtained by the equiangular full-width approach.



FIG. 6.5. 2D Test 1. The pixelwise (posterior) standard deviations associated to the corresponding reconstructions in Figure 6.3, i.e., the square roots of the diagonal elements of the covariance matrix Γ_k for the corresponding u_k in Algorithm 1 for k = 1, ..., 16. The latest projection geometry is illustrated by the black lines.

The targets consist of ellipses with constant absorption levels in a homogeneous nonabsorbing background. The number of ellipses is drawn from the uniform distribution over $\{2, 3, 4, 5\}$, their absorption levels from the uniform distribution over [0.5, 1.5], and their centers from the uniform distribution over the disk of radius 0.5 centered at the midpoint of D. Furthermore, the ellipses have (uniformly) random orientations, and their semimajor and semiminor axes are independently drawn from the uniform distribution over [0.05, 0.2]. In the regions where many ellipses overlap, the absorption level is defined to be the sum of those of the involved ellipses. An example of such a target is shown in Figure 6.6. In particular, note that the ellipses may extend over the domain boundary, which is not in line with the Dirichlet boundary condition for (3.7) but assures that any considered X-ray may pass through something interesting in a target.

The discretization of D is the same as in the previous example; that is, the reconstructions are formed on a uniform grid of $n = N^2 = 10^4$ pixels, and a full-width source-receiver pair corresponds to m = 51 individual X-rays (and the 0.25)



FIG. 6.6. 2D Test 2. Top left: Mean relative $L^2(D)$ reconstruction errors over 100 samples for A-optimal projection geometries and equiangular full-width projections with equivalent radiation doses. The red and blue curves show the errors with optimized projection geometries with beam widths 0.25 and 0.5, respectively, whereas the black curve shows the errors for equiangular projections with beam width 1. The error-bars show the confidence intervals of one standard deviation, and the horizontal axis indicates the number of projections with the narrowest beam width. Top right: Corresponding mean reconstruction errors for D-optimality. Bottom: Example of a random target composed of ellipses with randomly chosen shapes, sizes, positions, and absorption levels.

and 0.5 beam widths to 13 and 26 X-rays, respectively). However, encouraged by the observations in the previous test, the sequential optimization of the projection geometries is carried out on the sparser grid of $\tilde{n} = \tilde{N}^2 \approx 10^3$ pixels. The total number of considered random targets is 100, and the noise level is once again set to $\sigma = 10^{-3}$. To make the radiation doses comparable, the algorithm is run for 20 and 10 iterations for the beam widths of 0.25 and 0.5, respectively, and the corresponding relative $L^2(D)$ reconstruction errors are computed after each iteration. Analogously, the reference reconstructions and the corresponding relative $L^2(D)$ errors are computed for 1, 2, 3, 4, and 5 equiangular full-width projections.

The results, shown on the top row in Figure 6.6, indicate that the reconstruction errors for both A- and D-optimized projections decrease faster as functions of the radiation dose than that for the reference projections. As in the previous experiment, once enough projection data has been collected, this advantage starts to decrease. The A-optimal designs perform on average slightly better than their Doptimal counterparts, which is not surprising since the A-optimality criterion aims to minimize the expected squared reconstruction error, i.e., the performance metric considered in Figure 6.6. Performing the sequential experimental design with a narrower beam seems to be advantageous, presumably because the algorithm can

B522 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

concentrate on retrieving information on certain interesting local details in the target without "wasting radiation." However, this advantage comes with a fairly significant computational cost: In addition to having to run the algorithm for twice as many iterations, the search space is also much wider due to the increased number of possible lateral positions for the source-receiver pair. This poses a problem for our exhaustive optimization routine [15]. In addition, the overlapping confidence intervals in Figure 6.6 hint that the best approach is target-dependent: The A- and D-optimal designs produced by Algorithm 1 result on average in lower reconstruction errors than the equiangular full-width approach, but this does not apply to all individual targets.

6.1.3. 2D Test 3: Shepp–Logan phantom. In our third experiment, the target is the Shepp–Logan phantom shown in the top left image of Figure 6.7. The main aim is to compare the performance of Algorithm 1 with A-optimality and beam width



FIG. 6.7. 2D Test 3. Top left: Shepp-Logan phantom. Top right: Mean relative $L^2(D)$ reconstruction errors over 100 samples of noise realizations for sequentially A-optimized projection geometries (red: Algorithm 1, blue: Gaussian prior) and equiangular full-width projections with equivalent radiation doses (black). The error-bars show the confidence intervals of one standard deviation, and the horizontal axis indicates the number of projections with beam width 0.25. Middle row: Radiation dose corresponding to 5 full-width projections. Bottom row: Radiation dose corresponding to 10 full-width projections. Left column: Reconstructions for projection geometries A-optimized by Algorithm 1 with the latest projection depicted. Center column: Reconstructions for equiangular full-width projections. Right column: Reconstructions for projection geometries optimized based on a Gaussian prior with the latest projection depicted.

0.25 to reconstructions obtained from equiangular full-width reference projections. However, we also consider using in Algorithm 1 sequentially A-optimized quarterwidth projections corresponding to a Gaussian prior with a covariance matrix of the form

(6.3)
$$(\Gamma_{\text{prior}})_{i,j} = \eta^2 \exp\left(-\frac{|x_i - x_j|^2}{2\ell^2}\right).$$

Here $|\cdot|$ denotes the Euclidean norm, $\ell > 0$ is the so-called correlation length, $\eta > 0$ is the pixelwise standard deviation, and x_i denotes the center of the *i*th pixel. Under such a prior, the sequentially optimized projections do not depend on the measurements or the prior mean, and they can thus be computed in advance based on merely the covariance matrix (6.3) and the known structure of the additive Gaussian noise process; see [15] for more details. When employing a Gaussian prior with the covariance structure (6.3), we thus use in Algorithm 1 precomputed sequentially optimized design variables instead of determining the projection geometries adaptively as a part of the algorithm itself. However, the lagged diffusivity iteration is still employed in the computation of the reconstructions, as indicated by the interior loop of Algorithm 1.

We choose $\eta = 0.2$ and $\ell = 0.1$ in (6.3); the former is close to the pixelwise standard deviation of the Shepp–Logan phantom, whereas the latter simply seems to be in a relatively good agreement with the sizes of the areas with constant absorption in the top left image of Figure 6.7. All other parameters are the same as in the previous experiment. In particular, the optimization steps of Algorithm 1 are once again carried out on a sparser grid with $\tilde{N}^2 \approx 10^3$ pixels, and this same sparse discretization is also used for deducing the sequentially optimized projection geometries corresponding to the Gaussian prior with the covariance matrix (6.3). The test is run 100 times to examine how the measurement noise affects the reconstruction quality, mainly via changes in the adaptive optimal designs produced by Algorithm 1.

The top right image in Figure 6.7 shows the (mean) relative $L^2(D)$ reconstruction errors up to 40 and 10 projections for the two types of sequentially optimized geometries and the equiangular reference geometries, respectively. After a radiation dose that is equivalent to 1 full-width projection, both sets of reference projections, i.e., the full-width equiangular one and the 4 quarter-width ones based on the Gaussian prior, correspond to lower $L^2(D)$ errors than the reconstruction produced by Algorithm 1. This is likely due to the Shepp–Logan phantom covering most of the domain, which makes an initial full-width projection or 4 rather randomly distributed, nonadaptively chosen quarter-width projections sensible approaches. During the following 36 rounds of Algorithm 1, the edge-promoting sequentially optimal design first shows a clear advantage over the full-width reference reconstructions, but the advantage diminishes after enough projection data has been collected. On the other hand, the reference quarter-width projections based on the Gaussian prior perform almost as well as Algorithm 1 until about 10 projections, but subsequently the adaptive approach of Algorithm 1 leads to clearly superior results. It is also interesting to notice that the equiangular full-width heuristic starts to outperform the sequentially optimized projections corresponding to the Gaussian prior at about 25 quarter-width projections.

The final reconstructions after 40 sequentially optimized and with 10 full-width equiangular projections, as well as those after only 20 optimized and 5 equiangular projections, are presented on the bottom and middle rows of Figure 6.7. The reference reconstructions corresponding to the full-width projections in the middle column contain characteristic streaking artifacts of sparse-angle X-ray tomography, evenly spread

B524 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

around the target; this effect is particularly emphasized in the case of 5 full-width projections. For the 20 adaptively optimized projections in the left-hand column, some sections of the "head boundary" are reconstructed more accurately than in the corresponding reference reconstructions, and there is also less blurring in the interior of the phantom. However, with only 20 projections Algorithm 1 leads to a bad reconstruction of the bottom half of the phantom, as the optimized projections have not yet covered that region comprehensively. This exemplifies an obvious flaw in our approach: Even if the sequentially chosen projection geometries were locally optimal, their combination is no longer optimal after several rounds, and there is no guarantee that this nonoptimality could not be severe if a high number of projection geometries is considered. This unwanted behavior could presumably be controlled to a certain extent by the user of the algorithm via introducing, while the algorithm is running, a region of interest that guides the algorithm to retrieve data on the unexplored areas in the target (cf. [15]). Another potential remedy could be to ask the algorithm to occasionally choose a random design parameter or a design parameter furthest from the set of previously employed parameters.

The final reconstruction produced by Algorithm 1 after 40 quarter-width projections and the one corresponding to 10 equiangular full-width projections shown on the bottom row of Figure 6.7 are comparable in quality, whereas the one corresponding to the 40 quarter-width projections sequentially optimized based on the Gaussian prior does not reproduce the organ boundaries as accurately as the other two. These observations are in line with the information in the convergence plot of the top right image in Figure 6.7.

6.2. Three-dimensional cone beam tomography. In three dimensions, the unknown absorption distribution is located in the unit cube $D = [0,1]^3$ that is discretized into a uniform grid of $n = N^3$ voxels. We consider cone beam tomography, where a point-like source at $s \in \mathbb{R}^3$ sends X-rays to a two-dimensional receiver patch that occupies a "square" solid angle of the form $[\theta + \delta, \theta - \delta] \times [\phi + \delta, \phi - \delta]$ if the origin is transferred to s without affecting the orientation of the coordinate axes; see Figure 6.8. Here θ and ϕ denote the central polar and azimuthal angles of the detector, respectively. When considering full-aperture projections, the imaging system is always aligned so that the line between the source and the midpoint of the receiver



FIG. 6.8. Measurement setup of three-dimensional cone beam tomography.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

B525

passes through the center of the cube D. The receiver is discretized into a rectangular grid of $m = M^2$ detectors with respect to its polar and azimuthal angles in the coordinate system centered at s. To summarize, a single full-aperture projection geometry is defined by the central spherical angles of the detector (θ, ϕ) with respect to the source s (or the center of D), the corresponding opening angle δ , the distance d from the source to the center of D, and the number of pixels per edge M in the detector. Observe that the distance between the source and the detector does not play a role as long as the two are on opposite sides of D.

After assigning (fixed) values for d, δ , and M, a set of full-aperture projection geometries to be used in the exhaustive optimization algorithm of [15] is defined by choosing the corresponding central spherical angles (θ_j, ϕ_j) . Unlike in the twodimensional examples with parallel beam tomography, the projections are not symmetric with respect to reflections about the center of the object, and thus one cannot only focus on projections from one side of the object; i.e., one cannot exclude some projection directions as redundant by a symmetry argument. To simulate movement of a smaller detector in the lateral direction, it is possible to only consider some subset of detectors in a full-aperture receiver.

6.2.1. 3D Test with simple geometric shapes. Analogously to the first twodimensional experiment, our three-dimensional example only aims at demonstrating the basic operation of the algorithm for A-optimality. The target shown on the left in Figure 6.9 consists of two balls with radius 0.2 centered at (0.2, 0.2, 0.2) and (0.3, 0.6, 0.6), respectively, and a rectangular cuboid $[0.6, 0.8] \times [0.5, 0.9] \times [0.5, 0.9]$ in a homogeneous background with vanishing absorption. The common constant absorption level of the balls is 1, and that of the cuboid is 2. The target *D* is discretized into a grid with N = 50 voxels per edge, i.e., a total of $1.25 \cdot 10^5$ unknowns. For the optimization step of Algorithm 1, we interpolate once again onto a *significantly* sparser grid with $20^3 = 8000$ pixels to speed up the computations. The noise level is chosen to be $\sigma = 2 \cdot 10^{-3}$, the opening angle of the projection cones is $\delta = 0.24$ radians, and the distance from the source to the center of *D* is set to 2.5.



FIG. 6.9. 3D Test. Left: Target. Right: Relative $L^2(D)$ reconstruction errors for A-optimized quarter-aperture projection geometries and "equally spaced" full-aperture projections with equivalent radiation doses. The red curve depicts the errors for optimized projection geometries, whereas the black curve shows those for the equally spaced projections. The horizontal axis indicates the number of quarter-aperture projections.

B526 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

To define the set of (central) projection angles used in determining the search space for the exhaustive optimization algorithm from [15], we introduce 60 evenly spaced azimuthal angles ϕ_i over the interval $[0, 2\pi]$ and 3 polar angles $-\pi/4, 0, \pi/4$, with the 0 polar angle associated to directions parallel to the xy-plane. The total set of projection directions is then $[\phi_1, \ldots, \phi_{30}] \times [-\pi/4, 0, \pi/4]$. The detector is split into 4 quadrants, each with 10×10 detectors, to allow 4 quarter-aperture projection geometries for each projection direction. This construction results in a total number of $4 \times 3 \times 60 = 720$ available projection geometries for the exhaustive algorithm from [15]. In particular, note that the set of possible projection directions is both sparse and limited in the polar direction, which has a certain effect on the achievable reconstruction quality [41].

Algorithm 1 is run for a total of 40 rounds. For reference, we once again also consider reconstructions corresponding to "equally spaced" full-aperture projections of equivalent radiation dose. Unlike in two dimensions, there is no obvious methodology for choosing the directions for these reference projections: (i) There exists no fundamental way of uniformly sampling the available 180 directions, and (ii) it is obvious that projections from opposite directions contain similar yet not exactly the same information. Our heuristic for choosing the directions of the full-aperture projections is including in the computation of the reference reconstructions one by one more projection directions from the sequence: (0,0), $(0,2\pi/3)$, $(0,4\pi/3)$, $(\pi/4,\pi)$, $(\pi/4,0)$, $(-\pi/4,\pi/2)$, $(-\pi/4,3\pi/2)$, $(0,\pi/6)$, $(0,3\pi/2)$, $(0,5\pi/6)$. In particular, note that this construction does not even aim at globally optimal reference directions, as are arguably the equiangular directions in two dimensions, but the selection of the reference projection geometries is also sequential in the sense that all previously used projections are also included in the subsequent projection sets of higher cardinality.

The relative $L^2(D)$ reconstruction errors for equivalent radiation doses are shown on the right in Figure 6.9. For the optimized quarter-aperture projections, the reconstruction error initially starts to decrease before plateauing for iterations 3–6. At that point, the reconstruction error for the reference projections decreases faster, with the quality of the reference reconstructions being better for radiation doses equivalent to 3–8 quarter-aperture projections. However, between 7 and 10 iterations of Algorithm 1, the reconstruction error for the optimized quarter-aperture projections drops rapidly below the reference curve and stays there all the way until the limit of 40 quarter-aperture projections is reached. As in the two-dimensional experiments, once enough data has been collected the optimized quarter-aperture projections and the reference full-width projections result in roughly the same reconstruction errors for equivalent radiation doses.

The top row of Figure 6.10 shows three slices of the reconstruction parallel to the xy-plane after 5 rounds of Algorithm 1, while the bottom row illustrates the same cross-sections after 15 rounds. These images demonstrate that initially the optimization procedure focuses solely on the vicinity of the ball centered at (0.2, 0.2, 0.2), while the surroundings of the 2 other inclusions are left unexplored. This explains the rapid drop in the relative $L^2(D)$ reconstruction error over the first couple of iterations as one of the two balls is found and explored, but it also gives a reason for the slow convergence between 3 and 6 iterations: The algorithm prefers to first thoroughly investigate the detected ball, and it moves its focus on the other two objects only after an optimized projection accidentally passes through them. This demonstrates an inherent flaw in the algorithm: Areas with already detected distinguishable features are examined in depth, whereas other areas are left untouched until something interesting is found as a byproduct of the ongoing local exploration. This feature

ADAPTIVE EXPERIMENTAL DESIGN IN X-RAY IMAGING



FIG. 6.10. 3D Test. Slices parallel to the xy-plane of reconstructions produced by Algorithm 1 for quarter-aperture projections and A-optimality. Top row: 5 optimized projections. Bottom row: 15 optimized projections.

could possibly be mitigated, e.g., by initializing the algorithm with a low number of full-width projections that cover the entire target. We also refer to the possible remedies for the global nonoptimality of the designs produced by Algorithm 1 discussed in section 6.1.3.

7. Concluding remarks. In this work we studied sequential edge-promoting Bayesian experimental design for linear inverse problems and, in particular, for X-ray tomography. We introduced a novel greedy iterative method that aims at optimizing the measurement design when a TV-type prior is applied. The method is based on interpreting the so-called lagged diffusivity iteration [53] in the Bayesian framework; see [12, 17, 51]. Our two- and three-dimensional numerical examples based on simulated data suggest that the introduced approach promotes sequential designs that enhance recovery of edges in the target image.

Although we used X-ray tomography mainly as an example of a high-dimensional linear inverse problem with well-motivated reasons for resorting to TV-type priors, the presented computational results anyway demonstrate a potential to considerably reduce the radiation dose when the aim is to retrieve a certain amount of information on the imaged object via X-ray imaging. However, there are a couple of practical caveats: the assumption that the measurement noise is additive and Gaussian is essential for our analysis and the computational complexity of the optimization step in our algorithm may present an insurmountable obstacle in inherently threedimensional imaging setups. Moreover, the applicability of the presented approach definitely depends on the precise practical application of X-ray imaging.

There are a number of interesting avenues for future work. Due to the feedback from the data, our sequential algorithm often allocates subsequent projections to enhance already observed edges while a portion of the target image may remain uninvestigated. Such choices are not necessarily globally optimal, and we recorded reconstruction error plots that exhibit occasional jumps when previously unexplored

B528 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

objects are (accidentally) detected. Understanding the algorithmic balance between exploring new areas and improving already observed edges via the choice of the next design seems an interesting task.

The more straightforward questions are related to the performance of the algorithm for nonlinear inverse problems and its integration with more efficient optimization procedures than the exhaustive search employed here. Moreover, investigating whether the sequential designs obtained via the proposed approach approximate (at least asymptotically) the ones corresponding to the exact TV prior is also left for future studies.

Acknowledgments. We would like to thank an anonymous referee for invaluable contributions that significantly improved the quality of this article.

REFERENCES

- S. AGAPIOU, M. BURGER, M. DASHTI, AND T. HELIN, Sparsity-promoting and edge-preserving maximum a posteriori estimators in non-parametric Bayesian inverse problems, Inverse Problems, 34 (2018), 045002.
- [2] S. AGAPIOU, M. DASHTI, AND T. HELIN, Rates of contraction of posterior distributions based on p-exponential priors, Bernoulli, 27 (2021), pp. 1616–1642.
- [3] A. ALEXANDERIAN, Optimal experimental design for infinite-dimensional Bayesian inverse problems governed by PDEs: A review, Inverse Problems, 37 (2021), 043001.
- [4] A. ALEXANDERIAN, P. J. GLOOR, AND O. GHATTAS, On Bayesian A- and D-optimal experimental designs in infinite dimensions, Bayesian Anal., 11 (2016), pp. 671–695.
- [5] A. ALEXANDERIAN, N. PETRA, G. STADLER, AND O. GHATTAS, A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized l₀-sparsification, SIAM J. Sci. Comput., 36 (2014), pp. A2122–A2148.
- [6] A. ALEXANDERIAN, N. PETRA, G. STADLER, AND O. GHATTAS, A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems, SIAM J. Sci. Comput., 38 (2016), pp. A243–A272.
- [7] A. ALEXANDERIAN, N. PETRA, G. STADLER, AND I. SUNSERI, Optimal design of large-scale bayesian linear inverse problems under reducible model uncertainty: Good to know what you don't know, SIAM/ASA J. Uncertain. Quantif., 9 (2021), pp. 163–184.
- [8] A. ALEXANDERIAN AND A. K. SAIBABA, Efficient D-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems, SIAM J. Sci. Comput., 40 (2018), pp. A2956–A2985.
- N. ARETZ-NELLESEN, P. CHEN, M. A. GREPL, AND K. VEROY, A Sequential Sensor Selection Strategy for Hyper-Parameterized Linear Bayesian Inverse Problems, preprint, arXiv:2011.11391, 2020.
- [10] S. ARRIDGE, M. BETCKE, AND L. HARHANEN, Iterated preconditioned LSQR method for inverse problems on unstructured grids, Inverse Problems, 30 (2014), 075009.
- [11] A. ATTIA, A. ALEXANDERIAN, AND A. K. SAIBABA, Goal-oriented optimal design of experiments for large-scale Bayesian linear inverse problems, Inverse Problems, 34 (2018), 095009.
- [12] J. M. BARDSLEY, Computational Uncertainty Quantification for Inverse Problems, Comput. Sci. Engrg., 19, SIAM, Philadelphia, 2018.
- [13] J. BECK, B. M. DIA, L. F. ESPATH, Q. LONG, AND R. TEMPONE, Fast Bayesian experimental design: Laplace-based importance sampling for the expected information gain, Comput. Methods Appl. Mech. Engrg., 334 (2018), pp. 523–553.
- [14] D. BRAESS, Finite Elements, Cambridge University Press, Cambridge, UK, 1997.
- [15] M. BURGER, A. HAUPTMANN, T. HELIN, N. HYVÖNEN, AND J.-P. PUSKA, Sequentially optimized projections in X-ray imaging, Inverse Problems, 37 (2021), 075006.
- [16] D. CALVETTI AND E. SOMERSALO, Introduction to Bayesian Scientific Computing, Sur. Tutor. Appl. Math. Sci., Springer, New York, 2007.
- [17] D. CALVETTI AND E. SOMERSALO, Hypermodels in the Bayesian imaging framework, Inverse Problems, 24 (2008), 034013.
- [18] K. CHALONER AND I. VERDINELLI, Bayesian experimental design: A review, Stat. Sci., 10 (1995), pp. 273–304.
- [19] T. F. CHAN AND P. MULET, On the convergence of the lagged diffusivity fixed point method in total variation image restoration, SIAM J. Numer. Anal., 36 (1999), pp. 354–367.

- [20] B. CRESTEL, A. ALEXANDERIAN, G. STADLER, AND O. GHATTAS, A-optimal encoding weights for nonlinear inverse problems, with application to the Helmholtz inverse problem, Inverse problems, 33 (2017), 074008.
- [21] D. C. DOBSON AND C. R. VOGEL, Convergence of an iterative method for total variation denoising, SIAM J. Numer. Anal., 34 (1997), pp. 1779–1791.
- [22] J. FOHRING AND E. HABER, Adaptive A-optimal experimental design for linear dynamical systems, SIAM/ASA J. Uncertain. Quantif., 4 (2016), pp. 1138–1159.
- [23] E. HABER, L. HORESH, AND L. TENORIO, Numerical methods for experimental design of largescale linear ill-posed inverse problems, Inverse Problems, 24 (2008), 055012.
- [24] E. HABER, L. HORESH, AND L. TENORIO, Numerical methods for the design of large-scale nonlinear discrete ill-posed inverse problems, Inverse Problems, 26 (2009), 025002.
- [25] E. HABER, Z. MAGNANT, C. LUCERO, AND L. TENORIO, Numerical methods for A-optimal designs with a sparsity constraint for ill-posed inverse problems, Comput. Optim. Appl., 52 (2012), pp. 293–314.
- [26] L. HARHANEN, N. HYVÖNEN, H. MAJANDER, AND S. STABOULIS, Edge-enhancing reconstruction algorithm for three-dimensional electrical impedance tomography, SIAM J. Sci. Comput., 37 (2015), pp. B60–B78.
- [27] X. HUAN, Accelerated Bayesian Experimental Design for Chemical Kinetic Models, Ph.D. thesis, Massachusetts Institute of Technology, 2010.
- [28] X. HUAN AND Y. M. MARZOUK, Simulation-based optimal Bayesian experimental design for nonlinear systems, J. Comput. Phys., 232 (2013), pp. 288–317.
- [29] X. HUAN AND Y. M. MARZOUK, Sequential Bayesian Optimal Experimental Design via Approximate Dynamic Programming, preprint, arXiv:1604.08320, 2016.
- [30] N. HYVÖNEN, A. SEPPÄNEN, AND S. STABOULIS, Optimizing electrode positions in electrical impedance tomography, SIAM J. Appl. Math., 74 (2014), pp. 1831–1851.
- [31] J. KAIPIO AND E. SOMERSALO, Statistical and Computational Inverse Problems, Appl. Math. Sci. 160, Springer, New York, 2005.
- [32] M. KHODJA, M. PRANGE, AND H. DJIKPESSE, Guided Bayesian optimal experimental design, Inverse Problems, 26 (2010), 055008.
- [33] M. LASSAS, E. SAKSMAN, AND S. SILTANEN, Discretization-invariant Bayesian inversion and Besov space priors, Inverse Probl. Imaging, 3 (2009), pp. 87–122, https://doi.org/10.3934/ ipi.2009.3.87.
- [34] M. LASSAS AND S. SILTANEN, Can one use total variation prior for edge-preserving Bayesian inversion?, Inverse Problems, 20 (2004), pp. 1537–1563.
- [35] Y. LIU, J. MA, Y. FAN, AND Z. LIANG, Adaptive-weighted total variation minimization for sparse data toward low-dose X-ray computed tomography image reconstruction, Phys. Med. Biol., 57 (2012), 7923.
- [36] Q. LONG, M. MOTAMED, AND R. TEMPONE, Fast Bayesian optimal experimental design for seismic source inversion, Comput. Methods Appl. Mech. Engrg., 291 (2015), pp. 123–145.
- [37] Q. LONG, M. SCAVINO, R. TEMPONE, AND S. WANG, Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations, Comput. Methods Appl. Mech. Engrg., 259 (2013), pp. 24–39.
- [38] D. LV, Q. ZHOU, J. K. CHOI, J. LI, AND X. ZHANG, Nonlocal TV-Gaussian prior for Bayesian inverse problems with applications to limited CT reconstruction, Inverse Probl. Imaging, 14 (2020), pp. 117–132.
- [39] F. NATTERER AND F. WÜBBELING, Mathematical Methods in Image Reconstruction, Monogr. Math. Model. Comput. 5, SIAM, Philadelphia, 2001.
- [40] F. PUKELSHEIM, Optimal Design of Experiments, SIAM, Philadelphia, 2006.
- [41] E. T. QUINTO, Singularities of the X-ray transform and limited data tomography in ℝ² and ℝ³, SIAM J. Math. Anal., 24 (1993), pp. 1215–1225, https://doi.org/10.1137/0524069.
- [42] L. I. RUDIN, S. OSHER, AND E. FATEMI, Nonlinear total variation based noise removal algorithms, Phys. D, 60 (1992), pp. 259–268.
- [43] L. RUTHOTTO, J. CHUNG, AND M. CHUNG, Optimal experimental design for inverse problems with state constraints, SIAM J. Sci. Comput., 40 (2018), pp. B1080–B1100, https://doi. org/10.1137/17M1143733.
- [44] E. G. RYAN, C. C. DROVANDI, J. M. MCGREE, AND A. N. PETTITT, A review of modern computational algorithms for Bayesian optimal design, Int. Stat. Rev., 84 (2016), pp. 128– 154.
- [45] K. J. RYAN, Estimating expected information gains for experimental designs with application to the random fatigue-limit model, J. Comput. Graph. Statist., 12 (2003), pp. 585–603.
- [46] O. SCHERZER, M. GRASMAIR, H. GROSSAUER, M. HALTMEIER, AND F. LENZEN, Variational Methods in Imaging, Springer, New York, 2009.

B530 TAPIO HELIN, NUUTTI HYVÖNEN, AND JUHA-PEKKA PUSKA

- [47] E. Y. SIDKY AND X. PAN, Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization, Phys. Med. Biol., 53 (2008), 4777.
- [48] S. SILTANEN, V. KOLEHMAINEN, S. JÄRVENPÄÄ, J. P. KAIPIO, P. KOISTINEN, M. LASSAS, J. PIRTTILÄ, AND E. SOMERSALO, Statistical inversion for medical x-ray tomography with few radiographs: I. General theory, Phys. Med. Biol., 48 (2003), pp. 1437–1463.
- [49] A. M. STUART, Inverse problems: A Bayesian perspective, Acta Numer., 19 (2010), pp. 451– 559, https://doi.org/10.1017/S0962492910000061.
- [50] Z. TIAN, X. JIA, K. YUAN, T. PAN, AND S. B. JIANG, Low-dose CT reconstruction via edgepreserving total variation regularization, Phys. Med. Biol., 56 (2011), 5949.
- [51] F. URIBE, J. M. BARDSLEY, Y. DONG, P. C. HANSEN, AND N. A. B. RIIS, A Hybrid Gibbs Sampler for Edge-Preserving Tomographic Reconstruction with Uncertain View Angles, preprint, arXiv:2104.06919v1, 2021.
- [52] S. VÄNSKÄ, M. LASSAS, AND S. SILTANEN, Statistical X-ray tomography using empirical Besov priors, Int. J. Tomogr. Statist., 11 (2009), pp. 3–32.
- [53] C. R. VOGEL AND M. E. OMAN, Iterative methods for total variation denoising, SIAM J. Sci. Comput., 17 (1996), pp. 227–238.
- [54] Z. WANG, J. M. BARDSLEY, A. SOLONEN, T. CUI, AND Y. M. MARZOUK, Bayesian inverse problems with l₁ priors: A randomize-then-optimize approach, SIAM J. Sci. Comput., 39 (2017), pp. S140–S166.
- [55] K. WU, P. CHEN, AND O. GHATTAS, A Fast and Scalable Computational Framework for Large-Scale and High-Dimensional Bayesian Optimal Experimental Design, preprint, arXiv:2010.15196, 2020.
- [56] K. WU, P. CHEN, AND O. GHATTAS, A Fast and Scalable Computational Framework for Goal-Oriented Linear Bayesian Optimal Experimental Design: Application to Optimal Sensor Placement, preprint, arXiv:2102.066271, 2021.
- [57] Z. YAO, Z. HU, AND J. LI, A TV-Gaussian prior for infinite-dimensional Bayesian inverse problems and its numerical implementations, Inverse Problems, 32 (2016), 075006.